

DE GRUYTER

Gabriel Frahm

RATIONAL CHOICE AND STRATEGIC CONFLICT

THE SUBJECTIVISTIC APPROACH TO GAME
AND DECISION THEORY

BUSINESS & ECONOMICS



Gabriel Frahm

Rational Choice and Strategic Conflict

Gabriel Frahm

Rational Choice and Strategic Conflict

The Subjectivistic Approach to Game and Decision Theory

DE GRUYTER

Mathematics Subject Classification 2010

Primary: 91A35, 91A40, 91A80; Secondary: 60A05

Author

Prof. Dr. Gabriel Frahm
Helmut Schmidt University
Faculty of Economics and Social Sciences
Department of Mathematics/Statistics
Chair of Applied Stochastics and Risk
Management
Hamburg
Germany
frahm@hsu-hh.de

ISBN 978-3-11-059515-4

e-ISBN (PDF) 978-3-11-059610-6

e-ISBN (EPUB) 978-3-11-059615-1

Library of Congress Control Number: 2019939833

Bibliographic information published by the Deutsche Nationalbibliothek

The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available on the Internet at <http://dnb.dnb.de>.

© 2019 Walter de Gruyter GmbH, Berlin/Boston

Cover image: Steve McAlister / Photolibrary / gettyimages

Typesetting: VTeX UAB, Lithuania

Printing and binding: CPI books GmbH, Leck

www.degruyter.com

“The logical roots of game theory are in Bayesian decision theory. Indeed, game theory can be viewed as an extension of decision theory (to the case of two or more decision-makers), or as its essential logical fulfillment. Thus, to understand the fundamental ideas of game theory, one should begin by studying decision theory.”

Myerson (1991, p. 5)

“Deliberation about what to do in any context requires reasoning about what will or would happen in various alternative situations, including situations that the agent knows will never in fact be realized.”

Stalnaker (1996)

Preface

Genesis and Structure of This Book

Some time ago, I began to doubt whether the predominant solution concept of game theory, the equilibrium doctrine, can really be justified on the grounds of decision theory. After holding my lectures on game and decision theory for a couple of years, it became increasingly harder for me to explain why Nash equilibrium should be the only rational solution of a strategic conflict, and even why *any* Nash equilibrium should be a rational solution at all. This is the main reason for writing this book.

Traditional game theory is based on the foundation of objectivistic decision theory. I quickly realized that objectivistic decision theory is quite limited when it comes to predict human behavior. It turned out that subjectivistic decision theory is much better suited in order to describe and explain what rational subjects do in situations of conflict. Unfortunately, the subjectivistic approach to rational choice, and even expected-utility theory per se, has become discredited during the last decades.

In my personal opinion, this is without much substantial justification. Human behavior that is often considered irrational can very well be explained by subjectivistic decision theory if we do not treat subjective probabilities like objective ones and make use of counterfactual reasoning. This is really essential and enables us to solve Ellsberg's paradox, which seems to be the major reason why subjectivistic decision theory has been banished. One of my purposes is to show that the subjectivistic approach is much better than its reputation.

Hence, in the first part of this book, I recapitulate the mathematical foundation of the subjectivistic approach to rational choice, with a special emphasis on Bayesian rationality, in order to develop a unified framework for game and decision theory. To be more precise, I concentrate on Savage's seminal work. Savage is the first who provides a complete axiomatic theory of rational choice under uncertainty. Moreover, his approach is brilliant. Nonetheless, it lacks one important element: counterfactual reasoning. We will see that decision theory, let it be subjectivistic or objectivistic, unfolds its full effect only by counterfactual reasoning.

The second part of this book deals with subjectivistic game theory, which can be considered a straightforward application of subjectivistic decision theory to the case in which two or more rational decision makers meet each other in a situation of conflict. I come to conclusions that differ essentially from those that are usually obtained in traditional game theory under the assumption of strategic independence, which claims that the choice of one player does not depend on the choice of any other player. After dropping the strategic-independence assumption, the results that are obtained in the subjectivistic framework differ even more from those that are generally taught in lectures on traditional game theory.

<https://doi.org/10.1515/9783110596106-201>

In the subjectivistic framework, I do not treat strategic games, i.e., situations of conflict, essentially different from games against Nature, i.e., decision problems, which are typically analyzed in decision theory. Why should a rational player in a strategic game do *not* try to choose any optimal action (or strategy) in the same way as he aims at maximizing his expected utility in a game against Nature? It turns out that strategic conflicts can very well be solved by means of subjectivistic decision theory, irrespective of whether the strategic-independence assumption is satisfied or violated. It seems quite odd that traditional game theory has developed into a completely different direction, but I think that the reasons are historical.

Hence, in the subjectivistic framework, each player can be considered a rational decision maker who treats his opponents like Nature. In most practical applications we can solve dynamic games by backward induction, which is a typical instrument of decision theory. In order to provide a better understanding of the entire solution concept, I distinguish between action, reaction, and interaction. Action and reaction describe the typical situations in static and dynamic games, whereas interaction occurs in coherent games. The latter category is fully ignored in traditional game theory and seems to be novel.

In the third part of this book, I discuss alternative concepts of game theory. Of course, I do not ignore the main contributions of traditional game theory, i.e., Nash equilibrium, Harsanyi's Bayesian Nash equilibrium, Aumann's correlated equilibrium, and von Neumann's minimax solution. This is done not only in order to provide a broader picture of game theory but also to highlight the differences between subjectivistic and objectivistic game theory. I discuss also a relatively new branch of game theory, namely the epistemic approach, which goes very much in the same direction as the subjectivistic approach. However, its underlying assumptions about the strategic reasoning of the players are by far more intricate. I close my analysis by discussing the theory of moves, an alternative branch of game theory that deals with strict ordinal games in which the players have perfect information.

Throughout this book, I give my very best to demonstrate the subjectivistic approach by practical examples and by referring to real-life situations. However, game and decision theory cannot exist without some degree of abstraction, which means that a rigorous mathematical treatment of the subject matter is indispensable. Readers who are not familiar with the basic methods, or who do not wish to go into the mathematical details, may skip Chapter 1, Section 6.6, Section 7.4, and Section 7.5. Nonetheless, I would not deny that the formal details are essential if somebody wants to get a comprehensive picture (not only) of subjectivistic game and decision theory. Further, it could be difficult to follow the explanations after skipping some sections that refer to methodological issues at discretion. In any case, the reader should be aware of Savage's representation theorem, which is treated in Chapter 1, since this is the main pillar of the subjectivistic framework.

Acknowledgments

First of all, I would like to thank Herbert Hax, who passed away in 2005. I had the great opportunity to attend his splendid lectures in finance at the University of Cologne in the 1990s, where I came into contact with decision theory. I am very grateful also for being part of the Institute of Econometrics and Statistics at the University of Cologne, where I worked for Karl Mosler and Friedrich Schmid. I had some nice discussions with Karl and his former assistant Rainer Dyckerhoff (not only) about Ellsberg's paradox, which represents an important topic of this book.

In 2015 I came into contact with Steven Brams. It seems that he was the first who did not consider my thoughts about Nash equilibrium preposterous, and since then I discovered that I was not alone with my opinion. I can say that Steve had the most significant impact on developing my, somewhat heretical, ideas and he was there countless times whenever I had something on my mind concerning game theory. Actually, he was also the one who eventually brought me to the idea of writing this book. Thus, for many reasons, I am indebted to Steven Brams.

I would like to thank also Klaus Beckmann, who gave me the opportunity to hold my lectures on game theory at the Bundeswehr Command and Staff College. Every time I hold these lectures it becomes more and more clear to me why game theory should try to solve *real-life* strategic conflicts. What else should game theory be all about? Does it make any sense to propagate models in which the players act in a way that is completely detached from reality? I think that this does not help anyone, and after our many discussions about the subject matter, I am convinced that Klaus would agree with me. Probably, the same holds true for my excellent colleagues at that college, i.e., Stefan Bayer and Burkhard Meißner.

Before writing this book, I tried very hard to convey my ideas in a premature manuscript that I sent to Robert Aumann. To be honest, I made no great attempt to conceal my doubts about the equilibrium doctrine and, more specifically, about correlated equilibrium. Despite my provocative opinion, his answer was calm and polite. He indicated that his personal view of game theory turned (even more) into the subjectivistic direction during the last decades. Thus, I am not sure how much he adheres to the equilibrium doctrine these days. In any case, I think that without reading his beautiful publications, I would never have started to make any effort for a deeper understanding of game theory.

Many thanks belong also to Andrés Perea, who was very nice and patient when answering my probing questions about epistemic game theory. Indeed, this branch is fascinating and anything but trivial. He clarified that game theory is based, to a large extent, on the standard assumption of strategic independence. This is the reason why the concept of best response and the dominance principle play such a big role. However, our discussions revealed that best response and dominance must be treated

<https://doi.org/10.1515/9783110596106-202>

with great caution if the strategic-independence assumption is violated, which is a typical phenomenon in dynamic and in coherent games.

I still have to thank many others for discussing game and decision theory with me, e.g., Giacomo Bonanno, Uwe Jaekel, Gert Mittring, Markus Riedinger, Mathias Risse, Ulrich Schwalbe, Wolfgang Spohn, and my family. I thank also my friends and colleagues at the Helmut Schmidt University, who attended my lengthy seminar talks, and my students for giving me a constant inspiration for a critical investigation of prevailing opinions. Finally, I would like to thank my editor Stefan Giesen from de Gruyter. He was always nice and relaxed during the entire project. I guess this cannot be taken for granted and so I was very fortunate. Last but not least I thank also Michael Kudisch, who did a great job proofreading the manuscript.

Gabriel Frahm
Hamburg, August 2, 2019

Contents

Preface — VII

Acknowledgments — IX

Part I: Rational Choice

- 1 The Subjectivistic Approach to Rational Choice — 3**
 - 1.1 States, Consequences, and Acts — 3
 - 1.2 The Preference Relation — 6
 - 1.2.1 Binary Relations — 6
 - 1.2.2 Orders — 8
 - 1.3 Savage's Postulates of Rational Choice — 10
 - 1.4 The Representation Theorem — 18
 - 1.4.1 Subjective Probability — 18
 - 1.4.2 Expected Utility — 19
 - 1.5 Bayesian Rationality — 22
 - 1.6 Conclusion — 26
- 2 How the Subjectivistic Approach Works — 27**
 - 2.1 Procedural Aspects of Deliberation — 27
 - 2.2 Classification of Decision Theories — 28
 - 2.2.1 Descriptive vs. Prescriptive — 28
 - 2.2.2 Normative vs. Positive — 29
 - 2.3 Counterfactual Reasoning — 29
 - 2.3.1 Substantive Conditionals — 30
 - 2.3.2 Event Tree vs. Decision Tree — 35
 - 2.3.3 Strategy and Scenario — 37
 - 2.3.4 Simple Decision Problems — 42
 - 2.4 Deliberation Crowds Out Prediction — 49
 - 2.4.1 Levi's Postulate — 49
 - 2.4.2 Wald's Maximin Rule — 51
 - 2.4.3 Savage's Minimax-Regret Rule — 53
 - 2.5 Knowledge and Belief — 55
 - 2.5.1 Waking Up in the Morning — 57
 - 2.5.2 Newcomb's Paradox — 58
 - 2.5.3 Communication vs. Information — 60
 - 2.6 The Dominance Principle — 61
 - 2.6.1 Dominance — 61

2.6.2	Superdominance —	65
2.6.3	Quintessence —	66
2.7	Ellsberg's Paradox —	67
2.7.1	Ellsberg's Thought Experiments —	67
2.7.2	The Solution to the Paradox —	70
2.7.3	Quintessence —	74
2.8	Backward Induction —	74
2.8.1	Regular Decision Problems —	74
2.8.2	Irregular Decision Problems —	78
2.8.3	The Time-Inconsistency Paradox —	80
2.9	The Oil-Wildcatting Problem —	83
2.9.1	Decision Tree —	83
2.9.2	Decision Matrix —	84
2.9.3	Final Remarks —	85
2.10	Conclusion —	87

Part II: Strategic Conflict

3	The Subjectivistic Approach to Strategic Conflict —	91
3.1	Games against Nature vs. Strategic Games —	91
3.2	The Subjectivistic Framework —	91
3.2.1	What Is a Rational Solution? —	91
3.2.2	Strategic and Social Interaction —	92
3.2.3	Cooperative vs. Noncooperative Games —	93
3.3	Static and Dynamic Games —	94
3.3.1	Static Games —	94
3.3.2	Dynamic Games —	97
3.4	Coherent Games —	101
3.5	Normal vs. Extensive Form —	105
3.6	The Dominance Principle —	107
3.6.1	Weak Dominance —	107
3.6.2	Strict Dominance —	108
3.7	Conclusion —	109
4	Action —	111
4.1	Strategic Independence —	111
4.2	The Subjectivistic Solution Concept —	113
4.2.1	The Entrepreneurship —	114
4.2.2	Rebel Without a Cause —	116
4.2.3	The Game Show —	117
4.3	Typical 2×2 Games —	123

4.3.1	Anti-Coordination Games —	124
4.3.2	Discoordination Games —	128
4.3.3	Coordination Games —	134
4.4	Zero-Sum Games —	140
4.4.1	Rock-Paper-Scissors —	140
4.4.2	The Mallorca Game —	142
4.5	The Prisoners' Dilemma —	144
4.5.1	The Classic Dilemma —	144
4.5.2	Split or Steal —	148
4.5.3	Nuclear Threat —	150
4.6	Conclusion —	151
5	Reaction —	153
5.1	The Ultimatum Game —	153
5.2	Nuclear Threat —	156
5.3	Penalty Shoot-Out —	158
5.4	The Game of Chess —	160
5.4.1	Zermelo's Chess Theorem —	161
5.4.2	The Subjectivistic Explanation of Chess —	162
5.4.3	Backward Induction —	165
5.5	The Iterated Prisoners' Dilemma —	170
5.5.1	The Axelrod-Hamilton Experiment —	170
5.5.2	General Form of the Dilemma —	172
5.5.3	Forward Deduction —	174
5.6	Conclusion —	180
6	Interaction —	183
6.1	Coherent vs. Quasicoherent Games —	183
6.2	Private Information and Interaction —	186
6.3	Solving Coherent Games —	191
6.3.1	Response Diagrams —	191
6.3.2	Rational Solutions —	194
6.3.3	Refinement —	197
6.4	Typical 2×2 Games —	199
6.4.1	Anti-Coordination Games —	199
6.4.2	Discoordination Games —	203
6.4.3	Coordination Games —	208
6.5	The Prisoners' Dilemma —	211
6.5.1	The Classic Dilemma —	211
6.5.2	Split or Steal —	214
6.6	Coherent n -Person Games —	215
6.6.1	Preliminary Remarks —	216

- 6.6.2 Main Theorems — 220
- 6.7 Conclusion — 227

Part III: Alternative Concepts of Game Theory

7 The Traditional Approach — 231

- 7.1 Transparency of Reason — 231
 - 7.1.1 Complete Information and Common Knowledge — 232
 - 7.1.2 What Does Common Knowledge Tell Us? — 233
 - 7.1.3 Some Examples — 234
- 7.2 Nash Equilibrium — 238
 - 7.2.1 Basic Model — 238
 - 7.2.2 Some Examples — 240
 - 7.2.3 Bacharach's Transparency of Reason — 245
 - 7.2.4 Do We Make Our Choices at Random? — 247
 - 7.2.5 Subjectivistic Nash Equilibrium — 249
- 7.3 Bayesian Nash Equilibrium — 252
 - 7.3.1 Basic Model — 253
 - 7.3.2 Some Examples — 254
- 7.4 Correlated Equilibrium — 259
 - 7.4.1 Basic Model — 260
 - 7.4.2 Aumann's Notion of Strategy — 261
 - 7.4.3 A Simple Example — 262
 - 7.4.4 Controversy — 265
- 7.5 Minimax Solution — 267
 - 7.5.1 Von Neumann's Minimax Theorem — 267
 - 7.5.2 2-Person Zero-Sum Games — 268
 - 7.5.3 Some Examples — 273
- 7.6 Refinement — 275
 - 7.6.1 Payoff Efficiency — 276
 - 7.6.2 Risk Dominance — 277
 - 7.6.3 Perfectness — 281
- 7.7 Conclusion — 282

8 The Epistemic Approach — 287

- 8.1 Rationalizability — 288
 - 8.1.1 Basic Idea — 288
 - 8.1.2 Solution Concept — 289
 - 8.1.3 Some Examples — 291
 - 8.1.4 General Remarks — 293
 - 8.1.5 The Subjectivistic Interpretation — 296

8.2	Reasonability —	297
8.2.1	Common Belief —	297
8.2.2	Stochastic Dependence —	301
8.2.3	Optimality vs. Rationality —	303
8.2.4	Reasonable Solutions —	304
8.2.5	Belief Hierarchies —	305
8.3	Conclusion —	311
9	Theory of Moves —	315
9.1	Motivation —	315
9.2	Basic Rules of TOM —	316
9.3	Standard Example —	318
9.3.1	The Solution According to TOM —	318
9.3.2	The Coherent Solution —	321
9.4	TOM vs. Coherent Games —	323
9.5	Conclusion —	326

Bibliography — 327

Games — 333

Index — 335

Part I: Rational Choice

1 The Subjectivistic Approach to Rational Choice

This work builds on the foundation of subjectivistic decision theory, i.e., of rational choice under uncertainty. Hence, it seems worth recapitulating the basic theory for those who are not familiar with the subject matter. Nonetheless, readers who are not so much interested in a rigorous treatment of the foundation may skip or, at least, skim through the mathematical details of this chapter. However, in any case, one should not miss Section 1.4 because it contains Savage's celebrated representation theorem, which can be considered the main result of subjectivistic decision theory.

The basic foundation of the subjectivistic approach to rational choice is laid by Borel (1924), de Finetti (1937), and Ramsey (1926), but to the best of my knowledge, Savage (1954) is the first who provides a complete axiomatic theory. Many other (subjectivistic) theories have been developed afterwards. See, e.g., Anscombe and Aumann (1963), Bolker (1967), Chai et al. (2016), Domotor (1978), Fishburn (1967, 1969), Jeffrey (1981, 1983), Kopylov (2007), Luce (2000), Luce and Krantz (1971), Pfanzagl (1967, 1968), Pratt et al. (1964), Schmeidler (1989), Suppes (1969), Wakker (1990), and Wakker (1993b). Of course, the given list of contributions is far from exhaustive.¹

Readers who are particularly interested in an exhaustive mathematical treatment of subjectivistic decision theory should read Fishburn (1970) and Wakker (1989). These authors provide a very detailed account of the subjectivistic approach to rational choice and point to some interesting aspects of Savage's representation theorem. According to Fishburn (1970, p. 191), Savage's seminal work is "the most brilliant axiomatic theory of utility ever developed." I present Savage's postulates of rational choice but ignore the proof of his representation theorem. Instead, I discuss the typical assumptions about the structure of the decision problem and the decision maker's preferences. The following exposition is mainly based on Fishburn (1970) and Savage (1972), i.e., the second revised edition of Savage (1954).

1.1 States, Consequences, and Acts

Savage's theory of rational choice consists of three primitives: states, consequences, and acts. These building blocks constitute the structure of the decision problem.

According to Savage (1972, p. 9), the world is "the object about which the person is concerned" and a state of the world is "a description of the world, leaving no relevant aspect undescribed." The state space of the decision maker is denoted by Ω . It is assumed that Ω is nonempty and that the elements of Ω are mutually exclusive. There

¹ A splendid, though not contemporary, overview of subjectivistic decision theories can be found in Fishburn (1981). Further, a nice and modern overview of decision theories, which goes far beyond the subjectivistic approach, is given by Wakker (2010).

exists one and only one true state of the world, that is “the state that does in fact obtain, i.e., the true description of the world” (Savage, 1972, p. 9). Throughout this book, ω denotes any state of the world, i.e., an element of Ω , whereas the *true* state of the world is always symbolized by ω_0 .

The state space, Ω , precisely contains those descriptions of the world, i.e., states, that are not *known*, a priori, by the decision maker to be impossible.² Put another way, he knows that the true state of the world, ω_0 , belongs to Ω , but he does not know more than that. Hence, in this sense, Ω must be exhaustive (Machina, 2003), and Savage’s postulates of rational choice, which will be elaborated in Section 1.3, require also that Ω is at least (countably) infinite (Wakker, 1989, p. 90).

Let $\mathcal{F} = 2^\Omega$ be the power set of Ω . Hence, \mathcal{F} is a σ -algebra on Ω , i.e.,

1. $\Omega \in \mathcal{F}$;
2. $F \in \mathcal{F} \Rightarrow \Omega \setminus F \in \mathcal{F}$, and
3. $F_1, F_2, \dots \in \mathcal{F} \Rightarrow \bigcup_{i=1,2,\dots} F_i \in \mathcal{F}$.³

Each element of \mathcal{F} is referred to as an event. I say that the event $F \in \mathcal{F}$ happens if and only if $\omega_0 \in F$.

Fishburn (1981, p. 141) notes that “states [...] lead to specific consequences that depend on the course of action adopted by the individual” and he mentions that “the occurrence of one consequence precludes the occurrence of any other consequence.” Thus, let C be a general (nonempty) set of consequences. According to Fishburn (1981, p. 141), each element of C “provides a complete description of everything the individual may be concerned about [...]” Further, let $\mathcal{S} = C^\Omega$ be the set of all functions from Ω to C . The elements of \mathcal{S} are referred to as Savage acts (Fishburn, 1981, pp. 143, 160).⁴ Savage (1972, Chapter 2.5) points out that “If two different acts had the same consequences in every state of the world, there would [...] be no point in considering them two different acts at all.” Thus, two functions $s, t \in \mathcal{S}$ are considered identical if and only if $s(\omega) = t(\omega)$ for all $\omega \in \Omega$.

Figure 1.1 illustrates two Savage acts $s, t \in \mathcal{S}$. The state space, Ω , is given by the horizontal axis, whereas the set of consequences, C , coincides with the vertical axis. Each graph represents a Savage act. Savage (1972, p. 13) provides an example in order to demonstrate his concept: A cook has already broken five good eggs into a bowl in order to make an omelette. Now, there is an unbroken egg left, which can either be good or rotten. The cook considers three alternatives:

- (a) Break the egg into the bowl that already contains the other five eggs.
- (b) Break it into a saucer for inspection.
- (c) Throw it away without inspection.

² The distinction between knowledge and belief is fundamentally important. I will come back to this crucial point in Section 2.5.

³ By contrast, a Boolean algebra requires only that $F_1, F_2, \dots, F_n \in \mathcal{F} \Rightarrow \bigcup_{i=1,2,\dots,n} F_i \in \mathcal{F}$.

⁴ The reason why the elements of \mathcal{S} are called “acts” will become clear below.




Figure 1.1: State space, Ω , set of consequences, C , and Savage acts $s, t \in \mathcal{S}$.

This leads us to the decision matrix in Table 1.1. Savage calls Good and Rotten “states.” Since the state space, Ω , is infinite (Wakker, 1989, p. 90), these descriptions of the world actually represent *events* rather than states, but the chosen terminology shall not bother us here. Anyway, each act leads to another consequence, which depends essentially on the true state of the world. For example, if the cook decides to break the egg into the bowl that already contains the other five eggs, he has a six-egg omelette if the egg is good, but no omelette and five good eggs destroyed if it is rotten. Put another way, each decision leads to a function from Ω to C . For this reason, every element of \mathcal{S} is said to be a (Savage) “act.”

Table 1.1: Savage’s omelette example.

Act	“State”	
	Good	Rotten
<i>a</i>	Six-egg omelette	No omelette and five good eggs destroyed
<i>b</i>	Six-egg omelette and a saucer to wash	Five-egg omelette and a saucer to wash
<i>c</i>	Five-egg omelette and one good egg destroyed	Five-egg omelette

Savage assumes that the decision problem is well-defined. This means that the decision maker knows the potential consequences of his available actions. For example, suppose that the cook is not sure about whether or not a rotten egg spoils a six-egg omelette, given that the other five eggs are good. In this case, the decision problem is not well-defined because the consequence of breaking the egg into the bowl is ambiguous: It could either be “No omelette and five good eggs destroyed” or “Six-egg omelette.” However, this problem can easily be solved by splitting each “state” apart until we are able to assign each available action an unambiguous consequence (Savage, 1972, p. 15):

1. Good egg and a rotten egg spoils a six-egg omelette.
2. Good egg and a rotten egg does not spoil a six-egg omelette.
3. Rotten egg and a rotten egg spoils a six-egg omelette.
4. Rotten egg and a rotten egg does not spoil a six-egg omelette.

This looks like as if we have split two different “states,” namely Good and Rotten, each into two new states, but in my opinion this interpretation is somewhat misleading. A state of the world represents an atom in Ω and so it is not divisible. However, since Good and Rotten are events rather than states, we have just chosen a finer σ -algebra on Ω , and this is perfectly possible. According to Savage (1972, p. 15), we should always be able to cut each event into pieces, i.e., to dissect every potential source of ambiguity, until the resulting σ -algebra is fine enough in order to assign each combination of action and event one and *only one* consequence. Hence, we may consider each action a Savage act.

Although the verbal definition of Ω provided by Savage remains vague, the reader should understand the basic idea: In general, the true state of the world, ω_0 , is unknown to the decision maker. That is, if he decides to choose some Savage act $s \in S$, he cannot say, beforehand, which consequence occurs. Fishburn (1970, S. 164) points out also that “states [...] should be formulated in such a way that the state that obtains does not depend on the act selected.” This remark is essential. However, it poses some conceptual problems, which are solved in Section 2.3.

1.2 The Preference Relation

In the previous section, I described the world in which the decision maker takes place. Now, I concentrate on his individual preferences among the set S of all Savage acts. Preferences are usually expressed by using binary relations, more specifically, orders. Hence, we should have a basic understanding of these terms.

1.2.1 Binary Relations

A binary relation R on any set X is a subset of $\{(x, y) : x, y \in X\}$.

I write xRy to express that x is related to y , i.e., $(x, y) \in R$, whereas $\neg(xRy)$ means that x is not related to y , i.e., $(x, y) \notin R$. Moreover, $xRy \wedge yRx$ means that x is related to y and y is related to x , whereas $xRy \vee yRx$ shall indicate that x is related to y or y is related to x .⁵

Definition 1 (Basic properties). A binary relation R on a set X is said to be

- reflexive if and only if xRx ;
- irreflexive if and only if $\neg(xRx)$;
- symmetric if and only if $xRy \Rightarrow yRx$;
- asymmetric if and only if $xRy \Rightarrow \neg(yRx)$;

⁵ Throughout this book, whenever I say “ A or B ,” I mean that A or that B is true in the nonexclusive sense, whereas the statement “either A or B ” shall be understood in the exclusive sense.

- antisymmetric if and only if $(xRy \wedge yRx) \Rightarrow x = y$;
- transitive if and only if $(xRy \wedge yRz) \Rightarrow xRz$;
- negatively transitive if and only if $xRz \Rightarrow (xRy \vee yRz)$;
- complete if and only if $xRy \vee yRx$,⁶ and
- connex if and only if $x \neq y \Rightarrow (xRy \vee yRx)$

for all $x, y, z \in X$.

Note that connexity is *almost* the same as completeness. The former requires that $x, y \in X$ must be related to one another whenever x and y are *distinct*, whereas the latter does not make any requirement on the two elements of X .

Lemma 1. *Let R be a binary relation on a set X .*

- (a) *R is both reflexive and irreflexive if and only if $R = X = \emptyset$.*
- (b) *R is asymmetric if and only if R is irreflexive and antisymmetric.*
- (c) *R is both symmetric and antisymmetric if and only if $R \subseteq \{(x, x) : x \in X\}$.*
- (d) *Suppose that R is transitive. Then R is irreflexive if and only if R is asymmetric.*
- (e) *R is negatively transitive if R is asymmetric, transitive, and connex.*
- (f) *R is transitive if R is both asymmetric and negatively transitive.*
- (g) *R is empty if R is irreflexive, symmetric, and transitive.*
- (h) *R is both reflexive and connex if R is complete.*

Proof.

- (a) If R is reflexive and irreflexive, we have that xRx and $\neg(xRx)$ for all $x \in X$, which is possible only if $X = \emptyset$ and thus $R = \emptyset$. The converse is trivial.
- (b) If R is asymmetric, we have that $\neg(xRx)$ if xRx . Hence, xRx cannot be true for any $x \in X$, i.e., R is irreflexive. Moreover, the implication $(xRy \wedge yRx) \Rightarrow x = y$ is true whenever its antecedent is wrong. Since R is asymmetric, this is always the case and so R is antisymmetric. Conversely, let R be irreflexive and antisymmetric. In the case of $x \neq y$, the antisymmetry of R precludes $xRy \wedge yRx$. Otherwise, since R is irreflexive, we have that $\neg(yRx)$. Hence, $xRy \Rightarrow \neg(yRx)$ for all $x, y \in X$, which means that R is asymmetric.
- (c) If R is symmetric and antisymmetric, xRy implies that yRx and thus $x = y$. Hence, R must be a subset of $\{(x, x) : x \in X\}$. Conversely, if R is a subset of $\{(x, x) : x \in X\}$, it is clearly symmetric. Moreover, $xRy \wedge yRx$ can be true only if $x = y$, which means that R is also antisymmetric.
- (d) If R is transitive, it holds that $\neg(xRx) \Rightarrow (\neg(xRy) \vee \neg(yRx))$. Thus, if R is irreflexive, we have that $\neg(yRx)$ whenever xRy , which means that R is asymmetric. Conversely, (b) implies that R is irreflexive if it is asymmetric.
- (e) We must show that $(\neg(xRy) \wedge \neg(yRz)) \Rightarrow \neg(xRz)$ if the given conditions are satisfied. Since R is asymmetric, it is also irreflexive. Hence, if $y = x$ or $y = z$, the

⁶ A complete relation is called also “total” or “connected.”

implication is trivial. By contrast, if $y \neq x, z$, the connexity of R implies that yRx and zRy whenever $\neg(xRy) \wedge \neg(yRz)$. From the transitivity of R we conclude that zRx and the asymmetry of R yields $\neg(xRz)$.

- (f) Let R be asymmetric and negatively transitive. Suppose that xRy and yRz . From xRy we may conclude that $xRz \vee zRy$. However, since R is asymmetric, we cannot have that zRy . Thus, it holds that xRz and thus R is transitive.
- (g) If R is irreflexive, symmetric, and transitive, we conclude from (d) that it is also asymmetric. A symmetric and asymmetric relation is empty.
- (h) This assertion is trivial. □

1.2.2 Orders

In mathematics we often deal with (binary) relations that are referred to as orders, which are typically denoted by “ \leq ” or “ $<$.” For example, let X be the set of all living people, where $x < y$ shall indicate that x is shorter than y (Fishburn, 1970, p. 10). The relation $<$ is irreflexive, asymmetric, and negatively transitive. Hence, $<$ is also anti-symmetric and transitive. If we assume that two different people cannot have equal height, $<$ is also connex but incomplete because neither $x < y$ nor $y < x$ for $x = y$. The symbol \leq usually indicates a nonstrict order, whereas $<$ represents a strict order. Whenever the attribute “nonstrict” or “strict” is missing, it is implicitly assumed that the given order is nonstrict.

An order is always reflexive, whereas a strict order is always irreflexive. A partial order can be incomplete, whereas both a total order and a weak order are always complete. Further, a weak order need not be antisymmetric, whereas both a partial order and a total order are always antisymmetric.

Definition 2 (Orders). Let R be a binary relation on a set $X \neq \emptyset$.

- (a) R is said to be a partial order if and only if it is
 - 1. transitive;
 - 2. antisymmetric, and
 - 3. reflexive.
- (b) R is said to be a strict partial order if and only if it is
 - 1. transitive and
 - 2. asymmetric.
- (c) R is said to be a total order if and only if it is
 - 1. transitive;
 - 2. antisymmetric, and
 - 3. complete.
- (d) R is said to be a strict total order if and only if it is
 - 1. transitive;
 - 2. asymmetric, and

3. connex.
- (e) R is said to be a weak order if and only if it is
 1. transitive and
 2. complete.
- (f) R is said to be a strict weak order if and only if it is
 1. negatively transitive and
 2. asymmetric.
- (g) R is said to be an equivalence if and only if it is
 1. transitive;
 2. symmetric, and
 3. reflexive.

Now, by making use of Lemma 1, the reader can easily verify the following implications:

- An order is both partial and weak if and only if it is total.
- A strict total order is a strict weak order is a strict partial order.

Let \leq be a weak order on S that reflects the individual preferences of some decision maker. Of course, his preferences are inherently subjective. The statement $s \leq t$ is understood as “ s is not preferred to t .” Given the preference relation \leq , we may define the statements

- $t \geq s$ by $s \leq t$;
- $s < t$ by $\neg(t \leq s)$;
- $t > s$ by $s < t$, and
- $s \sim t$ by $s \leq t \wedge s \geq t$.

Here, $s < t$ means that “ s is less preferred than t ” and $s \sim t$ shall indicate that the decision maker is indifferent among the Savage acts s and t .

Since \leq is complete, we have that $s \leq t \vee t \geq s$ and thus $\neg(s > t \wedge t > s)$ for all $s, t \in S$. This means that $<$ is asymmetric. Moreover, since \leq is transitive, we have that $(s \leq t \wedge t \leq u) \Rightarrow s \leq u$, which is equivalent to $s > u \Rightarrow (s > t \vee t > u)$ for all $s, t, u \in S$. That is, $<$ is negatively transitive and thus a strict weak order. However, $<$ is not (necessarily) a strict total order.

The preference relation \sim is an equivalence and it may happen that the decision maker is indifferent among two *different* Savage acts s and t . In this case, we have that $s \neq t$ but neither $s < t$ nor $t < s$, which means that $<$ is not connex. In the context of decision theory, we usually deal with strict weak orders that are not strict total.

If we start our investigation with a strict weak order $<$ rather than a weak order \leq but define the statement $s \leq t$ by $\neg(t < s)$, it turns out that \leq is a weak order. Hence, \leq is a weak order if and only if $<$ is a strict weak order. Note that for all $s, t \in S$, the assertions $s < t$, $s \sim t$, and $s > t$ are mutually exclusive, whereas $s \leq t$ and $s \geq t$ may hold true together. Whether we should begin with \leq or $<$ is just a matter of taste.

Savage (1972) starts with \preceq , whereas Fishburn (1981) favors $<$. In this book, I use $<$, $>$, \preceq , \succeq , and \sim at discretion.

1.3 Savage's Postulates of Rational Choice

Throughout this section, capital Roman letters denote events, small Roman letters indicate Savage acts, and small Greek letters symbolize consequences.

Savage's subjectivistic theory of rational choice is normative.⁷ The decision maker is said to be rational if and only if he satisfies a number of postulates concerning his individual preference relation \preceq . It is a nice feature of the subjectivistic approach that we need not worry about any other aspect of "rationality."

Savage's first postulate (Savage, 1972, p. 18) goes like this:

P1. The preference relation \preceq is complete and transitive.

P1 just states that \preceq is a weak order. This means that for all Savage acts $s, t \in S$, the decision maker should be able to say whether $s \preceq t$ or $s \succeq t$. It is also possible that $s \preceq t$ and $s \succeq t$, in which case he is indifferent among s and t . Further, a weak order is always transitive. Thus, if the decision maker does not prefer s to t and t to u , then he must not prefer s to u , too. Similarly, for the strict weak order $<$ that is associated with \preceq we have that $s < t$ or $s > t$ or $s \sim t$ for all $s, t \in S$, and from $s < t$ and $t < u$ it follows that $s < u$.

For example, a person is sitting in a restaurant and studies the menu. The restaurant offers burgers and seafood. He considers to eat either a cheeseburger or shrimps. Now, he must be able to say whether he prefers the cheeseburger or the shrimps. Alternatively, he can be indifferent among the two dishes. In any case, he must be able to compare every dish with (itself and) each other. Moreover, if he prefers shrimps to cheeseburger and cheeseburger to salmon, he must also prefer shrimps to salmon.

Fishburn (1991) provides the following example, which illustrates a situation in which the preferences of a decision maker are not transitive: A professor thinks about moving to another university. He takes only two criteria into account: salary and prestige. If two offers are far apart on salary, the professor considers prestige not important at all. Otherwise, his decision is (also) based on the prestige of the university. Suppose that the professor has the three job offers given in Table 1.2.

Since a and b are far apart on salary, he prefers a to b , i.e., $a > b$. The difference in salary between b and c is not too big and thus he prefers b to c , i.e., $b > c$, since b has the higher prestige compared with c . Finally, since a and c are not far apart on salary, too, due to the same reason, he feels that $c > a$. Obviously, his preferences are

⁷ In Section 2.2, I will explain the meaning of "normative" in more detail.

Table 1.2: Job offers of the professor.

Offer	Criterion	
	Salary	Prestige
<i>a</i>	\$65,000	Low
<i>b</i>	\$50,000	High
<i>c</i>	\$58,000	Medium

not transitive because $a \succ b$ and $b \succ c$ but not $a \succ c$. Hence, the professor is irrational in the sense of subjectivistic decision theory.

Let $F \in \mathcal{F}$ with $F \neq \Omega$ be some event and $\neg F := \Omega \setminus F$ be the complement of F . Further, let s be some Savage act and suppose that the function $s_F : F \rightarrow C$ is such that $s_F(\omega) = s(\omega)$ for all $\omega \in F$. That is, s_F corresponds to s on F but it remains undefined on the complement $\neg F$. Hence, s_F represents a *restricted* Savage act (Fishburn, 1981, pp. 143, 160).

Now, we come to the next postulate, which is a key element of subjectivistic decision theory (Savage, 1972, p. 23):

P2. Suppose that $F \in \mathcal{F}$ and $s, t, s', t' \in \mathcal{S}$ such that $s_F = s'_F$ and $t_F = t'_F$ but $s_{\neg F} = t_{\neg F}$ and $s'_{\neg F} = t'_{\neg F}$. Then we have that $s \leq t \Rightarrow s' \leq t'$.

P2 can be understood as follows: When the decision maker compares the Savage acts s and t , he takes only that part of Ω into consideration on which s and t differ from one another. The other part of Ω is simply irrelevant because there the Savage acts s and t lead to the same consequences.

We can express the quintessence of **P2**, equivalently, by using composite acts. Consider some Savage acts $s, t \in \mathcal{S}$. Then $(s_F, t_{\neg F}) \in \mathcal{S}$ is the composite act that coincides with s on F but with t on $\neg F$. Hence, **P2** states that

$$(s_F, t_{\neg F}) \preceq (t_F, t_{\neg F}) \iff (s_F, u_{\neg F}) \preceq (t_F, u_{\neg F})$$

for all $F \in \mathcal{F}$ and $s, t, u \in \mathcal{S}$. This basic principle is illustrated in Figure 1.2.




Figure 1.2: If s is not preferred to t (left), the same must hold true for s' and t' (right).

Savage (1972) makes use also of restricted preferences (Savage, 1972, p. 22):

D1. Suppose that $F \in \mathcal{F}$ and $s, t \in \mathcal{S}$. We say that s_F is not preferred to t_F , i.e., $s_F \preceq t_F$, if and only if $s' \preceq t'$ for some $s', t' \in \mathcal{S}$ such that $s'_F = s_F$, $t'_F = t_F$, and $s'_{-F} = t'_{-F}$.

D1 can be expressed, equivalently, in terms of composite acts as follows:

$$s_F \preceq t_F : \Longleftrightarrow (s_F, u_{-F}) \preceq (t_F, u_{-F})$$

for any $u \in \mathcal{S}$. In fact, due to **P2**, the choice of u does not matter: If we have that $(s_F, u_{-F}) \preceq (t_F, u_{-F})$ for *some* Savage act u , the same holds true for *all* $u \in \mathcal{S}$. For example, s_F is not preferred to t_F if and only if it holds that $s \preceq t$ or, equivalently, $s' \preceq t'$ for the Savage acts s, t, s', t' that are depicted in Figure 1.2.

At this point we must be somewhat careful: I am not speaking about another preference relation but still refer to \preceq . The reader can easily verify that restricted preferences are complete and transitive. Thus, for every $F \in \mathcal{F}$, \preceq represents a weak order on the set of all Savage acts that are restricted to F .

D1, **P1** and **P2** imply Savage's so-called sure-thing principle (Savage, 1972, Section 2.7), which is a central element of subjectivistic decision theory: Consider the Savage acts $s = (s_F, s_{-F})$ and $t = (t_F, t_{-F})$. Suppose that $s_F \preceq t_F$ as well as $s_{-F} \preceq t_{-F}$. Then it must hold that $s \preceq t$. This can be seen as follows:

$$s = (s_F, s_{-F}) \preceq (t_F, s_{-F}) \preceq (t_F, t_{-F}) = t.$$

This result can be generalized by the following theorem (Savage, 1972, p. 24). Its proof is based on the previous argument and thus it can be skipped. Recall that a finite set $\{A_1, A_2, \dots, A_n\}$ of subsets of a nonempty set A is said to be a partition of A if and only if A_1, A_2, \dots, A_n are mutually disjoint and their union equals A .

Theorem 1 (Sure-thing principle, Part I). *Let $\{F_1, F_2, \dots, F_n\}$ be a partition of $F \in \mathcal{F}$ and suppose that $s, t \in \mathcal{S}$.*

- *If $s_{F_i} \preceq t_{F_i}$ for $i = 1, 2, \dots, n$, then $s_F \preceq t_F$.*
- *Moreover, if $s_{F_j} \prec t_{F_j}$ for some $j \in \{1, 2, \dots, n\}$, then $s_F \prec t_F$.*

For example, suppose that the decision maker thinks about going to the theatre or staying home. He does not know whether it will rain or shine. Let us consider the decision matrix in Table 1.3. If the decision maker prefers a to b and b to c , he must also prefer a to c , i.e., going to the theatre (a) to staying home (c).

Note that the Savage act b does not make any *physical* sense because the decision maker cannot choose to miss the play and keep dry if it rains but to enjoy the play and keep dry if not. However, the given approach to rational choice does not require any reasonable interpretation of Savage acts. In fact, most Savage acts have no particular meaning at all.

Table 1.3: Some Savage acts of the decision maker.

Act	Event	
	Rain	Shine
a	Enjoy the play but get wet	Enjoy the play and keep dry
b	Miss the play and keep dry	Enjoy the play and keep dry
c	Miss the play and keep dry	Miss the play and keep dry

The sure-thing principle can be explained also without the (somewhat dubious) Savage act b . For this purpose, consider the restricted Savage acts a_{Rain} , c_{Rain} , a_{Shine} , and c_{Shine} . If the decision maker prefers a_{Rain} to c_{Rain} and a_{Shine} to c_{Shine} , he must also prefer a to c .

From now on, I write “ $a \equiv \alpha$ ” in order to indicate a constant act $a(\omega) = \alpha \in C$ for all $\omega \in \Omega$. Correspondingly, “ $a_F \equiv \alpha$ ” is a constant act that is restricted to $F \in \mathcal{F}$, i.e., $a_F(\omega) = \alpha$ for all $\omega \in F$.

We proceed with the following definitions (Savage, 1972, pp. 24–25):

D2. The statement $\alpha \leq \beta$ means that $a \leq b$ with $a \equiv \alpha \in C$ and $b \equiv \beta \in C$.

D3. An event $F \in \mathcal{F}$ is said to be null if and only if $s_F \leq t_F$ for all $s, t \in \mathcal{S}$.

D3 implies that the decision maker is indifferent among all Savage acts that agree outside F . That is, he considers the null event F negligible. However, F cannot be dispensed of because the true state of the world can still belong to F . Put another way, even though the decision maker might *think* that F is negligible, this event can very well happen.

The set of all null events shall be denoted by $\mathcal{N} \subset \mathcal{F}$. Each event outside \mathcal{N} is said to be essential. When applying Theorem 1, we may neglect all null events in the partition $\{F_1, F_2, \dots, F_n\}$ because $s_{F_i} \leq t_{F_i}$ is always satisfied if F_i is null. By contrast, $s_{F_j} < t_{F_j}$ cannot be satisfied if F_j is null.

The next proposition (Savage, 1972, p. 26) constitutes the second part of the sure-thing principle and is illustrated in Figure 1.3:

P3. Suppose that $F \in \mathcal{F}$ is not null and consider some constant restricted acts $a_F \equiv \alpha \in C$ and $b_F \equiv \beta \in C$. Then we have that $a_F \leq b_F \Leftrightarrow \alpha \leq \beta$.

Note that **P3** requires the event F to be *essential* from the perspective of the decision maker, i.e., $F \notin \mathcal{N}$. In fact, it makes no sense at all to conclude anything about his preference between α and β on the basis of some null event.

Theorem 1 refers to arbitrary Savage acts and their corresponding restrictions to some pieces of Ω . By contrast, the following result, which can be found in Savage (1972, p. 26), refers to Savage acts that are piecewise constant on Ω .




Figure 1.3: The restricted Savage act a_F is not preferred to b_F (left) if and only if the constant act $a \equiv \alpha$ is not preferred to the constant act $b \equiv \beta$ (right).

Corollary 1 (Sure-thing principle, Part II). *Let $\{F_1, F_2, \dots, F_n\}$ be a partition of $F \in \mathcal{F}$ and suppose that $a, b \in \mathcal{S}$ are such that $a_{F_i} \equiv \alpha_i \in C$ and $b_{F_i} \equiv \beta_i \in C$ for $i = 1, 2, \dots, n$.*

- *If $\alpha_i \leq \beta_i$ for $i = 1, 2, \dots, n$, then $a_F \leq b_F$.*
- *Moreover, if $\alpha_j < \beta_j$ for some $j \in \{1, 2, \dots, n\}$ such that $F_j \notin \mathcal{N}$, then $a_F < b_F$.*

Proof. The first part is trivial for $F \in \mathcal{N}$. Otherwise, we may conclude from **P3** that $\alpha_i \leq \beta_i \Rightarrow a_{F_i} \leq b_{F_i}$ for $i = 1, 2, \dots, n$ and $\alpha_j < \beta_j \Rightarrow a_{F_j} < b_{F_j}$ for all $F_j \notin \mathcal{N}$. The rest of the proof is an immediate consequence of Theorem 1. \square

Thus, Corollary 1 is a straightforward implication of **P3** and Theorem 1. We already know that we may neglect all null events in the partition $\{F_1, F_2, \dots, F_n\}$ when applying Theorem 1. However, in contrast to Theorem 1, the second part of Corollary 1 explicitly requires F_j to be *not* null, since $\alpha_j < \beta_j$ can very well be true if $F_j \in \mathcal{N}$. Thus, when applying Savage's sure-thing principle, we not only may, we even *must* neglect all null events in $\{F_1, F_2, \dots, F_n\}$!

Table 1.4: Simple decision problem.

Act	Event	
	Rain	Shine
a	\$100	\$100
b	\$0	\$0
c	\$0	\$100

For example, let us inspect the decision matrix in Table 1.4 and suppose that the decision maker does not consider the event Rain null. If he prefers \$100 to \$0, i.e., $a > b$, he must also prefer a to c . By contrast, if the decision maker considers Rain null, we must not conclude that $a > c$ but $a \sim c$.⁸

⁸ Since the event Rain is null and the consequences of a and c are identical for Shine, Theorem 1 implies that $a \geq c$ and $a \leq c$, i.e., $a \sim c$.

In objectivistic decision theory, the content of Corollary 1 is referred to as the dominance principle: If some act dominates another act, i.e., if its consequences are not worse for all elements of some partition of Ω but even *better* for (at least) one *essential* event, the decision maker should refuse the dominated act. Note that we must neglect all null events in the partition when applying the dominance principle. Otherwise, our conclusions could be highly misleading. I will come back to this crucial point in Section 2.6.

I guess that most readers agree that Savage's sure-thing principle appears to be quite plausible and intuitive, at least from a normative point of view. However, some authors argue that, in real life, people often violate the sure-thing principle because they are ambiguity averse (see, e.g., Camerer and Weber, 1992; Ellsberg, 1961; Gilboa, 1987; Schmeidler, 1989; Slovic and Tversky, 1974; Wakker, 2001, 2010, Chapter 11). I think that the arguments raised by those authors do not fully exploit the possibilities of subjectivistic decision theory. Indeed, if we choose the state space, Ω , in an appropriate way,⁹ it is not difficult to show that ambiguity aversion is not only compatible with Savage's axioms of rational choice—we are even able to *explain* why a rational subject might be ambiguity averse. I will discuss this important issue in Section 2.7.

Next, we should try to provide a formal definition of qualitative probability before we start to derive a (quantitative) probability measure on \mathcal{F} . If somebody says that “it is more probable that Trump will be reelected in 2020 than that Yeti exists,” he primarily expresses a qualitative probability. Note that, in the subjectivistic framework, events like “Trump is reelected as President” and “Yeti exists” are propositions, but “Yeti exists” cannot be an event in the objectivistic, i.e., frequentistic, sense. In fact, Trump could be elected one more time, but the Yeti does either exist or not.

Suppose that the decision maker has the following choices:

- (a) He wins \$100 if it rains and \$0 if it shines.
- (b) He wins \$0 if it rains and \$100 if it shines.

If he prefers a to b , he considers the event Rain more probable than Shine.

Now, assume that the decision maker has the following additional choices:

- (c) He wins \$200 if it rains and \$100 if it shines.
- (d) He wins \$100 if it rains and \$200 if it shines.

If he prefers a to b he should also prefer c to d , i.e., his qualitative probability should not depend on the level of the payoffs.

This basic principle is formulated in a more general setting by the following postulate (Savage, 1972, p. 31), which is illustrated in Figure 1.4:

P4. Suppose that $F, G \in \mathcal{F}$ as well as $\alpha, \beta, \gamma, \delta \in C$ with $\alpha \succ \beta$ and $\gamma \succ \delta$. Further, assume that $a \equiv \alpha$, $b \equiv \beta$, $c \equiv \gamma$, and $d \equiv \delta$. Then we have that

$$(a_F, b_{-F}) \preceq (a_G, b_{-G}) \implies (c_F, d_{-F}) \preceq (c_G, d_{-G}).$$

⁹ The meaning of “appropriate” will be elaborated in Section 2.3.




Figure 1.4: If s is not preferred to t (left), then u must not be preferred to v (right).

The next proposition is an obvious nontriviality axiom (Savage, 1972, p. 31):

P5. There exist some $\alpha, \beta \in C$ such that $\alpha \succ \beta$.

This proposition implies that C is not a singleton.

Now, we are ready to make the following definition (Savage, 1972, p. 31):

D4. Suppose that $F, G \in \mathcal{F}$ and $\alpha, \beta \in C$ such that $\alpha \succ \beta$. We say that F is not more probable than G , i.e., $F \leq G$, if and only if $(a_F, b_{-F}) \leq (a_G, b_{-G})$ with $a \equiv \alpha$ and $b \equiv \beta$.

Note that in **D4** we make use of the symbol \leq in order to indicate a qualitative probability. There should be no danger of confusion because the objects F and G represent events rather than Savage acts.

We conclude that **P4**, **P5**, and **D4**—together with all other postulates and definitions—guarantee the existence of a qualitative probability. This probability is inherently subjective. It is a building bridge between the sure-thing principle and our final destination, i.e., the quantitative probability of the decision maker. In order to derive the latter we need two additional postulates.

The first one can be found in Savage (1972, pp. 39–40):

P6. Let $s, t \in S$ be such that $s < t$ and consider a constant act $a \equiv \alpha \in C$. Then there exists a partition $\{F_1, F_2, \dots, F_n\}$ of Ω such that

- $(a_{F_i}, s_{-F_i}) < t$ and
 - $s < (a_{F_i}, t_{-F_i})$
- for $i = 1, 2, \dots, n$.

This continuity axiom excludes situations in which the (strict) preference order of the decision maker among two Savage acts changes after “contaminating” an arbitrarily small part of a Savage act. Put another way, there is no consequence that is “infinitely desirable” or “infinitely undesirable” (Fishburn, 1981, p. 161). This postulate is illustrated in Figure 1.5.

The next definition introduces preferences among Savage acts and consequences (Savage, 1972, p. 72):




Figure 1.5: If t is preferred to s (left), then t is still preferred to $(a_{F_i}, s_{\neg F_i})$ (right), where F_i is an arbitrary element of some partition $\{F_1, F_2, \dots, F_n\}$ of Ω .

D5. Suppose that $F \in \mathcal{F}$, $s \in \mathcal{S}$, $\alpha \in C$, and consider the constant act $a \equiv \alpha$. We say that

- s_F is not preferred to α , i.e., $s_F \leq \alpha$, if and only if $s_F \leq a_F$ and that
- α is not preferred to s_F , i.e., $\alpha \leq s_F$, if and only if $a_F \leq s_F$.

This definition can be used, in the usual way, to define the statements $s_F > \alpha$, which means that s_F is preferred to α , and $\alpha > s_F$, i.e., that α is preferred to s_F .

Our final postulate goes like this (Savage, 1972, p. 77):

P7. Suppose that $F \in \mathcal{F}$ and $s, t \in \mathcal{S}$.

- If $s_F \leq t(\omega)$ for all $\omega \in F$, then $s_F \leq t_F$.
- If $s(\omega) \leq t_F$ for all $\omega \in F$, then $s_F \leq t_F$.

Hence, if the worst consequence of the Savage act t is not worse than the Savage act s , then s must not be preferred to t . Similarly, if the best consequence of s is not better than t , then s must not be preferred to t either. This postulate is illustrated in Figure 1.6 for the case of $F = \Omega$. An immediate consequence of **P7** is that there must be some $\omega_1 \in \Omega$ such that $s \geq t(\omega_1)$ and some $\omega_2 \in \Omega$ such that $s(\omega_2) \geq t$ whenever $s > t$. Put another way, $s > t$ implies that s is bounded below by some consequence of t and t is bounded above by some consequence of s .

A decision maker is said to be Savage rational if and only if he satisfies **P1–P7**.




Figure 1.6: If s is not preferred to any constant act that goes through the blue area, it must not be preferred to t .

1.4 The Representation Theorem

1.4.1 Subjective Probability

A real-valued function P on \mathcal{F} is said to be a probability measure if and only if

1. $P(\Omega) = 1$;
2. $P(F) \geq 0$ for all $F \in \mathcal{F}$, and
3. $P(F \cup G) = P(F) + P(G)$ for all $F, G \in \mathcal{F}$ such that $F \cap G = \emptyset$.

Every probability measure is (finitely) additive. That is, if the events F_1, F_2, \dots, F_n are mutually disjoint, we have that

$$P\left(\bigcup_{i=1}^n F_i\right) = \sum_{i=1}^n P(F_i).$$

Moreover, P is said to be σ -additive (or countably additive) if and only if

$$P\left(\bigcup_{i=1}^{\infty} F_i\right) = \sum_{i=1}^{\infty} P(F_i)$$

whenever the events F_1, F_2, \dots are mutually disjoint.

Definition 3 (Subjective probability measure). A probability measure P on \mathcal{F} is said to be subjective if and only if $P(F) > P(G) \Leftrightarrow F \succ G$ for all $F, G \in \mathcal{F}$.

We will see that every Savage-rational subject has a subjective probability measure, and this probability measure is even *unique*.

Moreover, for each $F \in \mathcal{F}$ and $0 < \lambda < 1$ we can find some $G \subseteq F$ such that $P(G) = \lambda P(F)$ (Fishburn, 1981, p. 161). Put another way, the probability measure P is atomless, which implies that $P(\{\omega\}) = 0$ for all $\omega \in \Omega$, and by setting $F = \Omega$ it follows that we can find an event G such that $P(G) = \lambda$ for every real number λ between 0 and 1. Hence, the elements of \mathcal{F} are continuously divisible. However, this does *not* imply that the state space Ω is uncountable, as is often suggested in the literature (see, e.g., Fishburn, 1981, p. 161). Wakker (1993a) shows that we can very well construct situations in which the power set of Ω is uncountable, but Ω itself is only countably infinite. This phenomenon can be best understood by realizing that Cantor's theorem requires $2^{\mathbb{N}}$, i.e., the power set of \mathbb{N} , to be uncountable although \mathbb{N} is countable.¹⁰

The subjective probability measure need not be σ -additive (Savage, 1972, p. 40).¹¹ In most real-life applications, this is not a serious obstacle. A probability measure that is only finitely additive can even be more flexible than a σ -additive one. For example, it is well-known that the power set of $[0, 1]$, i.e., the set $2^{[0,1]}$ of all subsets of $[0, 1]$, has no Lebesgue measure,¹² but we can find at least a finitely additive probability measure on

¹⁰ Here, it is implicitly assumed that the Continuum Hypothesis is true.

¹¹ Wakker (1993a) even observes that it *cannot* be σ -additive if the Continuum Hypothesis is true.

¹² Every power set is a σ -algebra and, by definition, a Lebesgue measure is always σ -additive.

$2^{[0,1]}$ that coincides with the Lebesgue measure for each Lebesgue-measurable subset of $[0, 1]$ (Savage, 1972, p. 41).

In any case, the subjective probability measure satisfies the elementary axioms of probability (Kolmogoroff, 1933, Section I.1). This implies that the probabilities of a rational subject must be consistent in the usual sense of probability theory. Thus, if we observe that some person behaves in a way that is inconsistent with Kolmogoroff's axioms, this person cannot be rational in the sense of **P1–P6**.

The subjective probability measure must not be confounded with the objective one. Even if we assume that the objective probability measure exists in a real-life situation,¹³ P can very well be different! Consider the following striking example: Roulette. Suppose that Black has occurred five times in a row. People who are not trained in probability theory often think that, in the next round, Red must obtain with higher probability than Black. Students are usually taught that this is nonsense, since the outcomes of the roulette wheel are (serially) independent and identically distributed. More precisely, the (objective) probability of Red is either $\frac{18}{37}$ or $\frac{18}{38}$, depending on whether there is only one zero pocket on the wheel or two.

Well, from a frequentistic point of view, the latter argument is completely right but not from a subjectivistic perspective. A rational subject might very well think that Red has a higher (or lower) probability than Black—and his opinion may even change from one round to another. Hence, the subjectivistic understanding of probability is inherently descriptive, which is precisely the reason why it is able to explain the real-life behavior of (rational) subjects.

In principle, we now have already achieved our goal to provide a normative framework that guarantees the existence of a (unique) subjective probability measure. However, there is one essential point missing—the so-called utility function of the decision maker. The latter plays an essential role in decision theory. Without a utility function, we could have contented ourselves with the preference relation of the decision maker, in which case **P1** would already suffice.

1.4.2 Expected Utility

Our final goal is to establish Savage's *representation theorem*, i.e.,

$$s \succ t \iff E(u(s)) > E(u(t))$$

for all $s, t \in S$. Here, u is a real-valued function on C and E denotes the expectation of the decision maker.¹⁴ More precisely, u is the decision maker's utility function and the

¹³ The objective probability measure is always unique by definition.

¹⁴ Hence, the symbol u shall not represent a Savage act.

expectation E is based on his subjective probability measure P . Hence, the representation theorem states that the decision maker prefers s to t if and only if the expected utility of s is greater than the expected utility of t .

As already mentioned, the probability measure P need not be σ -additive and so the expectations expressed by the representation theorem cannot be interpreted in the usual sense, i.e., as Lebesgue integrals. For this reason, we need **P7** and an appropriate definition of “expectation” in order to present the main result of subjectivistic decision theory, i.e., Savage’s representation theorem, in its most general form. This definition can be found in Fishburn (1970, Section 10.3).

Note that the function $f : \omega \mapsto u(s(\omega))$ is always \mathcal{F} -measurable because \mathcal{F} represents the power set of Ω . The function f is said to be simple if and only if its range, $\{f(\omega) : \omega \in \Omega\}$, is finite. Similarly, a Savage act s is said to be simple if and only if it takes on a finite number of values in C .¹⁵ In this case, also the associated function f is simple. In most practical applications of decision theory, we assume that s and thus f is simple.

Let s be a simple Savage act and $\{u_1, u_2, \dots, u_n\}$ be the range of f . Further, let $\{F_1, F_2, \dots, F_n\}$ be a partition of Ω such that $u_i = u(s(\omega)) = f(\omega)$ for all $\omega \in F_i$ and $i = 1, 2, \dots, n$. Finally, let $p_i := P(F_i)$ be the probability that the event F_i happens. Now, the expected utility of the (simple) Savage act s is defined as

$$E(u(s)) \equiv E(f) := \sum_{i=1}^n p_i u_i.$$

For most real-life examples of (subjectivistic or objectivistic) decision theory, this definition of expected utility is fully adequate.

Now, suppose that we have found some probability measure P on \mathcal{F} and a real-valued function u on C such that $s \succ t \Leftrightarrow E(u(s)) > E(u(t))$ for all simple $s, t \in \mathcal{S}$. Then P is a subjective probability measure. This can be seen like this: Assume that $F \succ G$ for some $F, G \in \mathcal{F}$, which means that $(a_F, b_{-F}) \succ (a_G, b_{-G})$ for any $a \equiv \alpha \in C$ and $b \equiv \beta \in C$ with $\alpha \succ \beta$. It follows that $E(u((a_F, b_{-F}))) > E(u((a_G, b_{-G})))$, i.e.,

$$P(F)u(\alpha) + P(\neg F)u(\beta) > P(G)u(\alpha) + P(\neg G)u(\beta). \quad (1.1)$$

From $\alpha \succ \beta$ we conclude that $u(\alpha) > u(\beta)$ and thus $P(F) > P(G)$. Conversely, suppose that $P(F) > P(G)$ and consider some $\alpha, \beta \in C$ such that $\alpha \succ \beta$. Hence, we have that $u(\alpha) > u(\beta)$ and so (1.1) holds true. This means that $(a_F, b_{-F}) \succ (a_G, b_{-G})$, i.e., $F \succ G$. We have just shown that $F \succ G \Leftrightarrow P(F) > P(G)$. Put another way, P is a subjective probability measure.

The following result can be considered a precursor of Savage’s representation theorem. It characterizes the individual preferences of a rational subject for *simple* acts

¹⁵ Savage (1972, Section 5.2) calls a simple act a gamble.

by means of expected utility. Its proof can be found in Fishburn (1970, Section 14.3 and Section 14.4) and Savage (1972, Section 3.3 and Section 5.3).

Proposition 1. *If the postulates P1–P6 are satisfied, the subjective probability measure P exists and is unique. Moreover, there exists a real-valued function u on C such that*

$$s \succ t \iff E(u(s)) > E(u(t))$$

for all simple $s, t \in S$, which is unique up to a positive affine transformation.

This proposition guarantees that u is unique—but only up to a positive affine transformation. This means that if u is a utility function, then any other function $v = a + bu$ with $a \in \mathbb{R}$ and $b > 0$ can serve as a utility function as well.

According to Fishburn (1970, p. 135), a sequence f_1, f_2, \dots of simple functions is said to converge uniformly from below to some (not necessarily simple) function f if and only if, for all $\omega \in \Omega$,

1. $f_1(\omega) \leq f_2(\omega) \leq \dots$;
2. $f(\omega) = \sup \{f_j(\omega) : j = 1, 2, \dots\}$, and
3. for every $\varepsilon > 0$, there exists a positive integer m (which may depend on ε but not on ω) such that $f_m(\omega) \geq f(\omega) - \varepsilon$.

We just say that f_1, f_2, \dots converges from below to f if the sequence satisfies the first two conditions. Only the last one constitutes uniform convergence. In fact, that condition requires that we are able to find some simple function f_m such that $f_m(\omega) \geq f(\omega) - \varepsilon$ for *all* and not just for some $\omega \in \Omega$.

The function f is said to be bounded if and only if there exist some numbers $a, b \in \mathbb{R}$ such that $a \leq f(\omega) \leq b$ for all $\omega \in \Omega$. For every *bounded* function f there exists a sequence of simple functions that converges uniformly from below to f (Fishburn, 1981, p. 135). It is clear that a simple function f is always bounded and if u is bounded, then f must be bounded, too.

Now, we are ready to define expected utility for arbitrary Savage acts: Suppose that f is bounded and consider any sequence f_1, f_2, \dots that converges uniformly from below to f . Then

$$E(u(s)) \equiv E(f) := \sup \{E(f_j) : j = 1, 2, \dots\}. \quad (1.2)$$

For all bounded utility functions u , the expected utility $E(u(s))$ is finite and unique (Fishburn, 1970, p. 136).

If \mathcal{F} were a σ -algebra on Ω such that P is σ -additive, it would not be necessary to assume that f_1, f_2, \dots converges *uniformly* from below to f , and f could even be unbounded. That is, provided that f is measurable, we could just use the Lebesgue integral in order to define expected utility for more general Savage acts. However, in our context, the σ -additivity of P is questionable, and thus we must assume that the utility function is bounded or restrict to simple Savage acts.

After adding **P7** to **P1–P6**, we can drop the assumption that u is bounded and need no longer restrict to simple Savage acts. According to Fishburn (1970, Section 14.5) and Fishburn (1981, p. 161), **P7** implies that u is a bounded utility function and so the expected utility is always well-defined according to Equation 1.2. This leads us to Savage's representation theorem, which is the main theorem of subjectivistic decision theory:

Theorem 2 (Savage's representation theorem). *If the postulates **P1–P7** are satisfied, the subjective probability measure P exists and is unique. Moreover, there exists a real-valued function u on C such that*

$$s \succ t \iff E(u(s)) > E(u(t))$$

for all $s, t \in S$, which is bounded and unique up to a positive affine transformation.

The proof can be found in Fishburn (1970, Chapter 14). If the utility function u is clear from the context, I write $EU(s)$ instead of $E(u(s))$ for notational convenience.

It is worth emphasizing that there exist many variants of the representation theorem. In fact, depending on our structural assumptions about Ω , \mathcal{F} , and C , on which subset of S we take into consideration, and on our basic requirements regarding the preference relation \preceq , we may come to different conclusions about the subjective probability measure P and the utility function u (Fishburn, 1981).

For example, after an appropriate modification of our postulates of rational choice, the utility function u need no longer be bounded (Wakker, 1993b). We can even find a σ -additive, not only a finitely additive, probability measure P , provided that we do no longer use the power set 2^Ω for \mathcal{F} but restrict to a Boolean algebra on Ω (Wakker, 1989, Theorem V.6.1). However, irrespective of the chosen axioms, the representation theorem remains valid. Hence, we are always able to characterize the individual preferences of a rational subject by his *expected utility*. This is the quintessence of subjectivistic decision theory and the main pillar of our subsequent analysis. Since we deal only with simple acts in most practical applications, the measure-theoretic subtleties shall no longer bother us.

1.5 Bayesian Rationality

The tuple (Ω, \mathcal{F}) represents a measurable space. **P1–P6** guarantee that we are able to extend (Ω, \mathcal{F}) to a (unique) probability space (Ω, \mathcal{F}, P) . The event $F \in \mathcal{F}$ is said to happen if and only if the true state of the world, ω_0 , belongs to F . Savage (1972) implicitly assumes that the decision maker is informed about the fact that $\omega_0 \in \Omega$. Put another way, he has the trivial information Ω .

In the Bayesian context, the probability measure P is referred to as a prior. Now, assume that the decision maker is equipped with a private information set, i.e., a par-

tion $\mathcal{I} = \{I_1, I_2, \dots, I_n\}$ of Ω .¹⁶ This means that he is informed about the event $I \in \mathcal{I}$ that is such that $\omega_0 \in I$ and thus, a posteriori, he realizes that each element of $\neg I$ is impossible. Simply put, he knows which event in his (private) information set happens.¹⁷ The event $I \ni \omega_0$ is said to be the decision maker's information or evidence.

Hence, we have to distinguish between two situations of the decision maker:

1. His situation *a priori*, in which his information set is trivial, i.e., $\mathcal{I} = \{\Omega\}$, and so he knows only that $\omega_0 \in \Omega$.
2. His situation *a posteriori*, in which his information set corresponds to $\mathcal{I} = \{I_1, I_2, \dots, I_n\}$, so that he knows the element of \mathcal{I} that contains $\omega_0 \in \Omega$.

Savage's postulates of rational choice hold true a priori as well as a posteriori. Consequently, we may distinguish between two preference relations:

1. The preference relation of the a-priori situation, i.e., \preceq , and
2. the preference relation of the a-posteriori situation, which is denoted by \preceq_I .¹⁸

In the same way as we can derive the preference relations \succeq , $<$, $>$, and \sim from the prior preference relation \preceq (see Section 1.2.2), we may derive the preference relations \succeq_I , $<_I$, $>_I$, and \sim_I from the posterior preference relation \preceq_I .

Both \preceq and \preceq_I refer to the set \mathcal{S} of all (unrestricted) Savage acts. That is, we still consider Ω our state space and do not leave Savage's model of rational choice. Moreover, it should be clear that \preceq does not depend on the specific location of $\omega_0 \in \Omega$ and, correspondingly, \preceq_I does not depend on $\omega_0 \in I$. I will come back to this point in Section 2.5.

After the previous arguments, the following postulate seems obvious:

P8. A posteriori, the decision maker considers $\neg I$ null.

This is referred to as the Bellman postulate. It states that the decision maker neglects the complementary event $\neg I$ when he is informed about I . More precisely, he *ignores* all impossible states, i.e., all descriptions of the world that lie outside I , and thus considers each event outside I negligible.

Note that **P8** does not require $P(I) > 0$. That is, the decision maker can very well be surprised, a posteriori, by the occurrence of I ! However, the Bellman postulate guarantees that, a posteriori, the decision maker ignores the consequences that occur at any state outside I . This enables us to solve sequential decision problems by backward induction. I will come back to this point in Section 2.8.

The (unique) subjective probability measure based on the information I is referred to as his posterior and it is denoted by P_I . Moreover, we may conclude from Theorem 2

¹⁶ The trivial partition of Ω is $\{\Omega\}$.

¹⁷ Then he knows also that $\omega_0 \in J$ for any other event $J \supset I$.

¹⁸ The preference relation \preceq_Ω is identical with \preceq .

that

$$s \succ_I t \iff E_I(u_I(s)) > E_I(u_I(t))$$

for all $s, t \in S$, where the expectation E_I is based on the posterior P_I and u_I is the posterior utility function of the decision maker. In particular, we have that $P_\Omega \equiv P$ and $u_\Omega \equiv u$. The decision maker is said to be Bellman rational if and only if he satisfies **P1–P8**.

The next axiom is said to be the Bayes postulate:

P9. If I is not a priori null, then $s \preceq_I t \iff s_I \preceq t_I$ for all $s, t \in S$.

This axiom requires that the Savage act s is not preferred to the Savage act t , a posteriori, if and only if s_I is not preferred to t_I , a priori, provided that I is not negligible. Note that when comparing s with t , a posteriori, as well as when comparing s_I with t_I , a priori, the decision maker neglects both the part of s and the part of t that lies outside I .

We will see in Section 2.8 that the Bayes postulate guarantees that solving a sequential decision problem by backward induction cannot lead to an optimal strategy that is considered suboptimal by forward deduction, i.e., by using a decision matrix rather than a decision tree.

In the case in which $P(I) > 0$, the conditional probability of $F \in \mathcal{F}$ is defined as

$$P(F \mid I) := \frac{P(F \cap I)}{P(I)}$$

and it is clear that $P(F \mid I) = 0$ whenever $F \cap I = \emptyset$. By contrast, if I is negligible, probability theory tells us that we can use any number between 0 and 1 as a probability of F conditional on I . The conditional expectation $E(u(s) \mid I)$ is defined, in the usual way, by using the conditional probability measure $P(\cdot \mid I)$ instead of the unconditional one, i.e., P . Similarly, in the case in which I is negligible, any real number can serve as an expectation of $u(s)$ conditional on I .

The next theorem generalizes Savage's representation theorem and it will play a crucial role throughout this book. I call it Bayes theorem:

Theorem 3 (Bayes theorem). *Let P be the prior and u be the utility function of a subject who satisfies the postulates **P1–P9**. Further, assume that its private information I is not negligible. Then the following statements hold true:*

- *The posterior P_I exists and is uniquely determined by $P(\cdot \mid I)$.*
- *Moreover, we have that*

$$s \succ_I t \iff E(u(s) \mid I) > E(u(t) \mid I)$$

for all $s, t \in S$.

Proof. Let s and t be any Savage acts. From Theorem 2 we conclude that the posterior P_I exists and is unique. Further, if $s \succ_I t$, **P9** reveals that $s_I \succ t_I$, i.e., $(s_I, r_{-I}) \succ (t_I, r_{-I})$ for any $r \in \mathcal{S}$. Now, Theorem 2 implies that $E(u((s_I, r_{-I}))) > E(u((t_I, r_{-I})))$ and thus

$$P(I)E(u(s) \mid I) + P(\neg I)E(u(r) \mid \neg I) > P(I)E(u(t) \mid I) + P(\neg I)E(u(r) \mid \neg I),$$

i.e., $E(u(s) \mid I) > E(u(t) \mid I)$. Conversely, if $E(u(s) \mid I) > E(u(t) \mid I)$, then $P(I) > 0$ implies that $E(u((s_I, r_{-I}))) > E(u((t_I, r_{-I})))$ for any $r \in \mathcal{S}$. That is, $s_I \succ t_I$ and **P9** tells us that $s \succ_I t$. To sum up, $s \succ_I t \Leftrightarrow E(u(s) \mid I) > E(u(t) \mid I)$ for all $s, t \in \mathcal{S}$, which means that $P(\cdot \mid I)$ is the posterior P_I . \square

A subject is said to be Bayes rational, or just “rational,” if and only if he satisfies the postulates **P1–P9**. The Bayes theorem states that the posterior of a rational subject, who is equipped with some essential information, corresponds to his prior *conditional* on the given information. Moreover, the utility function of a rational decision maker does not change a posteriori, i.e., after receiving some essential information. The entire hierarchy of rationality is depicted in Figure 1.7.

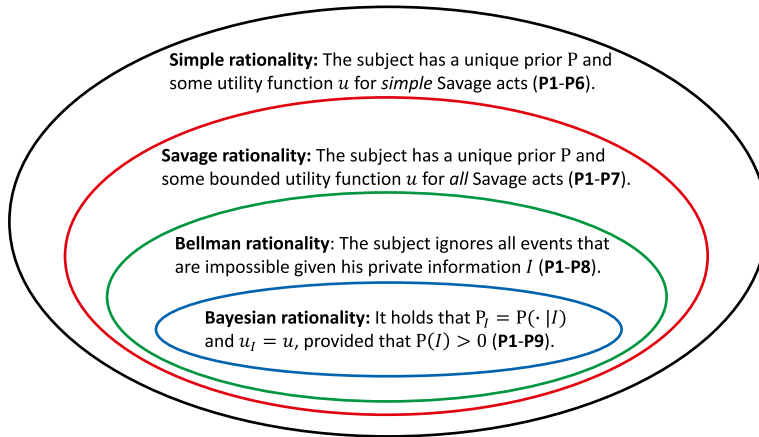


Figure 1.7: Hierarchy of rationality.

It is worth emphasizing that Bayesian rationality does not require us to apply the well-known Bayes rule:

$$P(F \mid I) = \frac{P(I \mid F) P(F)}{P(I)}.$$

This rule is just another way to calculate the posterior of a rational decision maker, provided that his private information I is essential.

1.6 Conclusion

Savage (1972) uses the notion of state, consequence, and act in order to formalize the concept of subjective probability. In my opinion, Savage's subjectivistic theory of rational choice is by far more than a decision theory. It can be considered a self-contained probability theory, in which the existence (and uniqueness) of some utility function is obtained as a by-product. However, expected utility plays a major role in this book, which is based on Savage's representation theorem.

A decision maker is said to be rational if and only if his preference relation among acts obeys Savage's postulates of rational choice and two additional axioms. The Bellman postulate guarantees that, at least in most practical applications, we can solve a sequential decision problem by backward induction. Further, we will see that the Bayes postulate guarantees that a strategy that is considered optimal by backward induction turns out to be optimal also by forward deduction.

2 How the Subjectivistic Approach Works

Chapter 1 might appear somewhat technical. A rigorous mathematical foundation of decision theory is indispensable, but nonetheless we should not forget to ask whether or not the presented approach is able to explain, or even predict, human behavior in a meaningful way. This will be demonstrated by a large number of examples, which shall bring Savage's abstract notion of state, consequence, and act to life.

2.1 Procedural Aspects of Deliberation

In many textbooks on decision theory one can find criteria of rationality that go far beyond our postulates of rational choice. If we leave the present framework, there is no general agreement about the necessary and sufficient conditions for "rationality." The criteria suggested in the literature are typically motivated by a prescriptive point of view. For example, according to Eisenführ and Weber (1994, Section 1.2.2), a rational decision maker should satisfy the following list of conditions:

1. **Effectiveness:** He should take all *effective* actions into consideration and ignore ineffective ones.
2. **Diligence:** He should invest the time and effort that is necessary in order to accumulate information. Moreover, he should process the given information in a *diligent*, i.e., reasonable and systematic, way.
3. **Objectivity:** His (conditional) expectations should be based on the *objective*, i.e., frequentistic, probability measure. Moreover, he should use only relevant information in order to solve his decision problem.
4. **Awareness:** He should be *aware* about his individual goals and preferences. He should also try to avoid any kind of cognitive bias and self-deception.

Obviously, these conditions refer to the procedural aspects of deliberation and they are very ambitious. The reader should judge for himself whether or not any known person satisfies the conditions above.

The subjectivistic approach does *not* require anything that is related to the procedure of decision making. It is concerned only with the individual preferences of the decision maker. More precisely, it presumes that the preferences are *consistent* in the sense of the postulates **P1–P9**, which implies that the subjective probabilities are consistent in the usual sense of probability theory (Kolmogoroff, 1933).

Throughout this book, I often write that the decision maker "thinks" or that he "believes" that some event F happens with probability $P(F)$, and sometimes I mention that he is "convinced" about F , etc. These expressions should only help us to understand his individual preferences. However, I do not presume that they are a result of any cognitive process. Further, I do not claim that the preferences of the decision maker are reasonable in any way, and they may very well be subject to instinctive or

<https://doi.org/10.1515/9783110596106-002>

intuitive factors. Also emotions like hope and fear may have an impact on the individual preferences of a rational subject and thus determine his subjective probabilities.

In principle, I even do not require that a rational subject *performs* an optimal action. Hence, I do not make any behavioral requirements either. Nonetheless, if I argue as an outside observer in order to explain the result of a decision problem or a strategic conflict, I usually assume that the subject(s) choose(s) an optimal action or strategy, but this assumption is not made for methodological reasons. It just simplifies the overall presentation and underpins the fact that subjectivistic decision theory tries to predict the actual behavior of human beings.

2.2 Classification of Decision Theories

2.2.1 Descriptive vs. Prescriptive

Savage's representation theorem states that the individual preferences of a rational subject can be understood *as if* he were comparing expected utilities. However, we need not assume that the decision maker *de facto* calculates any expected utility. The representation theorem guarantees that there exists some probability measure P and utility function u that can be used in order to characterize the individual preferences of a rational subject, but our intention is not to dictate any specific measure or utility function. We only try to describe the preferences of the subject in terms of expected utility. That is, we do not want to stipulate the decision maker's behavior. Hence, the subjectivistic approach to rational choice is descriptive.

By contrast, the objectivistic approach to rational choice holds that a rational decision maker should use the objective probability measure, and it is typically assumed that this probability measure is the frequentistic one.¹⁹ Of course, this requires (i) the existence and (ii) the uniqueness of such a probability measure. Moreover, as the name "objective" suggests, (iii) all (rational) subjects should agree that this probability measure is the one and only. Moreover, in most applications of objectivistic decision theory, the utility function is supposed to satisfy certain conditions. For example, it is typically assumed that the set of consequences, C , is a subset of \mathbb{R} and that u is strictly concave, which means that the decision maker is risk averse. In the economics literature, we can often find additional assumptions like constant absolute (or relative) risk aversion (Arrow, 1971; Pratt, 1964), etc. All these assumptions constitute a substantial limitation of "rationality." Thus, we may conclude that the objectivistic approach to rational choice is prescriptive.

¹⁹ This viewpoint can be found in von Neumann and Morgenstern (1953, p. 19), which is mentioned also by Savage (1972, p. 3). For a broad philosophical treatment of probability see, e.g., Burdzy (2016).

2.2.2 Normative vs. Positive

The essential difference between the subjectivistic and the objectivistic approach is that the former considers the probability measure endogenous, whereas the latter presumes that it is exogenous. However, both approaches are still normative because they assume that the decision maker is rational. By contrast, the behavioral approach (Kahneman and Tversky, 1979) holds that decision makers violate some, if not all, postulates of rational choice. Then the existence of both a probability measure and a utility function cannot be guaranteed at all. A decision theory in which the subjects are assumed to be irrational is called positive.

In the literature about decision theory, the words “normative” and “prescriptive” are often used synonymously, but I think that this is somewhat misleading. In fact, the subjectivistic approach is normative but nonetheless descriptive. It can be viewed as a compromise between the objectivistic and the behavioral approach, but my personal opinion is that the subjectivistic approach comes much closer to the behavioral one, since it tries to explain the real-life behavior of (rational) subjects. The classification of decision theories is illustrated in Table 2.1.

Table 2.1: Classification of decision theories.

	Descriptive	Prescriptive
Normative	Subjectivistic	Objectivistic
Positive	Behavioral	—

2.3 Counterfactual Reasoning

This book is based on a general philosophical principle that goes beyond Savage’s approach to rational choice: counterfactual reasoning. Thus, I will first explain this topic before I go into the details of subjectivistic decision theory.

In his famous omelette example, which has been discussed in Section 1.1, Savage presumes that the cook has no influence on the state of the unbroken egg. For example, it cannot happen that the unbroken egg is good if he breaks it into a saucer but rotten if he throws it away. Indeed, in that specific example we can hardly imagine any other possibility. However, in many real-life situations the decision of a person indeed may *have* an influence on his environment or, more generally, on the world by itself. That is, the subject can “change the world.”

Suppose that a person considers two possibilities to get to work, namely by car or by foot. From his individual perspective, the only relevant question regarding the world is whether he will be punctual or not. Thus, he takes only the two events “I will

arrive on time” and “I will not arrive on time” into consideration. It is clear that whether he will arrive on time or not depends essentially on his own decision.

Here is another example: An undercover agent is informed about the fact that a group of bank robbers is going to raid Fort Knox. He fears that the robbery will take place if he does not report the plan to his superior. Otherwise, he expects that the attempt will be foiled by the authorities. Obviously, his own decision has a substantial impact on the two descriptions “Raid” and “No raid.”

The two situations described above are prototypical. They demonstrate that, in general, each decision of a person can change the world. However, as Fishburn (1970, S. 164) points out, the state space, Ω , should be formulated in such a way that the true state of the world, ω_0 , does not depend on the action of the decision maker! Hence, at first glance, it seems that we have to concentrate on a small and insignificant set of decision problems.

To me it is an open question whether or not Savage was aware of this problem: His famous omelette example presumes that the cook has no influence on the unbroken egg and the reader can find many other examples in his book that suggest that the decision maker has (almost) no influence on the course of events. The same can be observed in Savage’s earlier work on decision theory (Savage, 1951). Hence, some authors answer the question in the negative (Chai et al., 2016).

However, that problem is not a serious one because it can easily be solved by the principle of counterfactual reasoning. This principle seems to become more and more important in game and decision theory during the last decades (see, e.g., Aumann, 1995; Bicchieri and Dalla Chiara, 1992; Bonanno, 2015; Harper et al., 1981; Rodriguez Marine, 1995; Samet, 2005; Stalnaker, 1981a,b, 1996; Zambrano, 2004) and it plays a major role also in this book. In my opinion, both decision theory and game theory would make no sense at all without counterfactual reasoning.

2.3.1 Substantive Conditionals

Consider this statement: “If Peter is in Canada, then he is in Africa.” How can we understand this sentence from a logical perspective? Obviously, it is a conditional statement, but is it an implication (“ \Rightarrow ”) in the usual sense of logic?

Implications are also referred to as material conditionals. Let A and C be some (logical) propositions. The material conditional $A \Rightarrow C$ is a composite proposition. It asserts that the consequent, C , is true if the antecedent, A , is true. However, it does not require anything about C in the case in which A is false. Thus, $A \Rightarrow C$ is true whenever A is false. To sum up, the implication $A \Rightarrow C$ just states that it cannot be true that A is true and C is false (see Table 2.2).

Hence, if the statement “If Peter is in Canada, then he is in Africa” is understood as a material conditional, then it would be true whenever Peter is in Europe! The reason is because the antecedent “Peter is in Canada” is false if Peter is in Europe and so the

Table 2.2: Truth table of a material conditional.

A	C	$A \Rightarrow C$
true	true	true
true	false	false
false	true	true
false	false	true

given conditional is true, irrespective of whether the consequent “Peter is in Africa” is true or false. Certainly, this is not what we expect in everyday life when we say, colloquially, “If Peter is in Canada, then he is in Africa.” Thus, we need another kind of conditional in order to express what we actually mean.

Let W be a nonempty set of possible worlds. The letter w always indicates an element of W , i.e., a possible world, where $w_0 \in W$ shall be the actual world. In general, we do not know the exact location of w_0 in W .

Now, let P be any proposition. Formally, I do not distinguish between the proposition P and the set $P \subseteq W$ of all possible worlds in which P is true. Thus, I say that P is true in $w \in W$ if and only if $w \in P$. Correspondingly, P is false in $w \in W$ if and only if $w \in \neg P \equiv W \setminus P$. Moreover, I call the proposition P empty if and only if $P = \emptyset$. For example, the proposition “Ridley Scott’s Alien truly exists” is empty unless the reader thinks that we should consider a world possible in which Alien appears in reality.

Throughout this section, capital Roman letters denote propositions.²⁰ Now, let us define the notion of substantive conditional (Aumann, 1995).

Definition 4 (Substantive conditional). A substantive conditional is a statement $A \subseteq C$ with $A, C \subseteq W$ and $A \neq \emptyset$.

Note that the antecedent of a substantive conditional must always be *nonempty*. Otherwise, both $A \subseteq C$ and $A \subseteq \neg C$ would be vacuously true for all $A = \emptyset$, but Definition 4 implies that $A \subseteq \neg C$ is false whenever $A \subseteq C$ is true. However, if $A \subseteq C$ is false, we must not conclude that $A \subseteq \neg C$ is true, since we can very well have that neither $A \subseteq C$ nor $A \subseteq \neg C$ is true. By contrast, if $A \Rightarrow C$ is true, then also $A \Rightarrow \neg C$ can be true, but if $A \Rightarrow C$ is false, then $A \Rightarrow \neg C$ must be true.

Further, note that the truth value of a material conditional depends on $w \in W$, whereas a substantive conditional is either universally true or universally false. We can see in Figure 2.1 that the substantive conditional $A \subseteq C$ is true and that $B \subseteq C$ is false. Further, the material conditional $B \Rightarrow C$ is true in w_1, w_2 , and w_3 , but it is false in w_4 , whereas $A \Rightarrow C$ is true in every $w \in W$.

We conclude that a material conditional $A \Rightarrow C$ with nonempty antecedent, A , is true in every possible world if and only if $A \subseteq C$ is true.

²⁰ The trivial proposition W , i.e., “This world is possible,” is true in every $w \in W$.




Figure 2.1: Propositions and possible worlds.

In everyday life, the statement “If Peter is in Canada, then he is in Africa” is understood as a substantive conditional. Of course, Peter can be in Canada but not in Africa and so this substantive conditional is false. If we ignore any form of extraterritoriality such as a foreign embassy or military base, Peter even *cannot* be in Africa while being in Canada, but this is not necessary at all in order to conclude that the substantive conditional “If Peter is in Canada, then he is in Africa” is false. On the contrary, the substantive conditional “If Mary is in Germany, then she is in Europe” is true because Mary cannot be in Germany without being in Europe.

What about the statement “If Mary *were* in Germany, then she *were* in Europe”? This is a special form of a substantive conditional, which is called counterfactual. Counterfactuals are typically studied in philosophy (Lewis, 1973), but as already mentioned above, during the last decades they have become increasingly popular in game and decision theory. I make use of counterfactual reasoning throughout this book in order to develop a consistent and meaningful theory of rational choice.

Definition 5 (Counterfactual). A counterfactual is a substantive conditional $A \subseteq C$ with $w_0 \in \neg A$.

Hence, a counterfactual is a substantive conditional whose antecedent is false in the actual world. For example, the statement “If Mary *were* in Germany, then she *were* in Europe” represents a (true) counterfactual, since it implies that the person addressed is, *actually*, not in Germany.

If we leave open whether the antecedent is true or false in the actual world, the substantive conditional is said to be an indicative conditional. The following example shall clarify how we distinguish between indicative conditionals and counterfactuals in our natural language: “If you went to school, then you know that $1 + 1$ equals 2” is an indicative conditional, but “If you had gone to school, then you would know that $1 + 1$ equals 2” represents a counterfactual.

Here is another example, which I have found in Dalla Chiara (1992). Assume that we are in the 1990s, in which Karol Wojtyla served as Pope, and consider the following statement: “If Wojtyla were a communist, then he would not be the pope.” It is clear that Wojtyla *is* the pope in w_0 and it is well-known also that he is not a communist. However, the question is about whether or not Wojtyla *would* still be the pope if he were a communist. This question refers to a situation that, de facto, never happened in real life, which shall be illustrated in Figure 2.2.




Figure 2.2: The set of possible worlds in the 1990s.

The real world, w_0 , can be found in the middle of the diagram. In this world, Wojtyla is Pope but not a communist. The statement “If Wojtyla were a communist, then he would not be the pope” refers to some alternative world that belongs to the set C , i.e., “Wojtyla is a communist,” which is marked red in Figure 2.2. Since the antecedent C is false in reality, the given statement represents a counterfactual.

Now, consider the proposition P , i.e., “Wojtyla is the pope,” which is marked green in Figure 2.2. In $w_1 \in C \cap \neg P$, Wojtyla is a communist but not the pope. However, in $w_2 \in C \cap P$ it turns out that he is both a communist and the pope. Hence, the counterfactual $C \subseteq \neg P$, i.e., “If Wojtyla were a communist, then he would not be the pope,” is false because there exists a possible world that belongs to C but not to $\neg P$. It would be true if and only if $C \cap P = \emptyset$.

The truth value of a substantive conditional depends essentially on W , i.e., on the given set of possible worlds. For example, a person in the 1990s who says “If Wojtyla were a communist, then he would not be the pope” renders any world in which Wojtyla is both a communist and the pope impossible. That is, given his set of possible worlds, the counterfactual “If Wojtyla were a communist, then he would not be the pope” is true. However, another person might possess another W and consider a world in which Wojtyla is a communist and the pope possible, in which case the same counterfactual turns out to be false from the latter’s viewpoint. Of course, that subject may very well

believe that Wojtyla would not be the pope if he were a communist. Nonetheless, the given counterfactual is still false.²¹

Here is another example: Heracles is standing in front of two closed doors, each one leading to the underworld. He expects to encounter Cerberus, the hound of Hades, when passing through a door. However, Heracles does not really know whether Cerberus is waiting behind any door. Thus, from his own perspective, “If Heracles goes through Door A, he will meet Cerberus” represents a false indicative conditional because it can happen that Cerberus is not waiting behind Door A. Analogously, “If Heracles goes through Door A, he will not meet Cerberus” is false, too, and the indicative conditionals are still false if we replace “Door A” with “Door B.” By contrast, Cerberus knows that he is waiting, say, behind Door A. For this reason, the indicative conditionals “If Heracles goes through Door A, he will meet Cerberus” and “If Heracles goes through Door B, he will not meet Cerberus” are true, whereas the opposite statements are false from Cerberus’ point of view.

Now, suppose that Heracles passes through Door A and meets Cerberus. In this moment, Heracles knows that Cerberus was waiting behind Door A, but does this mean that Heracles would *not* have met Cerberus if he had gone instead through Door B? Astonishingly enough, the answer is “No”! In fact, from Heracles’ perspective, it would have been possible to meet Cerberus at Door B, although he *actually* met Cerberus at Door A.²² Thus, Heracles considers the counterfactual “If Heracles had gone through Door B, he would not have met Cerberus” false. By contrast, Cerberus knows that he is waiting behind Door A and so his own set of possible worlds, W , differs essentially from Heracles’ set of possible worlds. To be more precise, it contains no world in which Cerberus is not waiting behind Door A. Hence, from his point of view, the counterfactual “If Heracles had gone through Door B, he would not have met Cerberus” is true!

Thus, whenever we deal with substantive conditionals, we must fix the set W . Otherwise, our conclusions can be erroneous, in particular if we analyze two or more substantive conditionals. For example, let A be “Trump did not offer himself as a presidential candidate in 2016” and B be “Trump has not won the presidential election in 2016.” The counterfactual $A \subseteq B$ is clearly true. Now, suppose that C is “The Republicans have lost the presidential election in 2016.” At first glance, the counterfactual $B \subseteq C$ seems to be true, but it is not!

The first counterfactual, $A \subseteq B$, is based on some set of possible worlds, W , in which it can happen that Trump did not offer himself as a presidential candidate in 2016. Thus, we must take this possibility into account also when deciding whether the

²¹ The essential difference between knowledge and belief will be discussed in Section 2.5.

²² We could imagine that Hades is able to predict Heracles’ decision and thus to take Cerberus to the right place just before Heracles passes through the door.

second counterfactual, $B \subseteq C$, is true or false, and we can imagine a world in which the Republicans have won the presidential election in 2016 *without* Trump.

According to Definition 4, substantive conditionals are transitive. Hence, $A \subseteq B$ and $B \subseteq C$ implies that $A \subseteq C$.²³ Moreover, substantive conditionals are monotonic. That is, from $A \subseteq C$ and $A \cap B \neq \emptyset$ it follows that $A \cap B \subseteq C$. For example, “If Peter is in Canada, then he is in North America” is a true indicative conditional. Further, “Peter is in Canada and he is visiting a theatre” is a nonempty proposition. Thus, “If Peter is in Canada and he is visiting a theatre, then he is in North America” is a true indicative conditional. Finally, it is clear that substantive conditionals are also reflexive, i.e., $A \subseteq A$ for each (nonempty) proposition A .

2.3.2 Event Tree vs. Decision Tree

Substantive conditionals can be illustrated by event trees. Let us come back to the 1990s, in which we know that Wojtyla is the pope but not a communist. Our status quo can be found on the upper right of the event tree in Figure 2.3. The circles in the event tree are referred to as event nodes and each outgoing branch represents an event. The given event tree indicates that the counterfactual “If Wojtyla were a communist, then he would not be the pope” is true, since Chance Node 3 has only one outgoing branch, namely “Wojtyla is not the pope.” This means that it is impossible for Wojtyla to be the pope if he is a communist. By contrast, if Chance Node 3 would have also the outgoing branch “Wojtyla is the pope,” then the counterfactual would be false.




Figure 2.3: Event tree in the 1990s.

Nonetheless, even if Chance Node 3 has the two outgoing branches “Wojtyla is the pope” and “Wojtyla is not the pope,” we might still *believe* that Wojtyla would not be

²³ Examples that try to abolish the transitivity of substantive conditionals are often based on erroneous conclusions like that one of the presidential election in 2016.

the pope if he were a communist. Put another way, we could assign the event “Wojtyla is not the pope” behind Chance Node 3 probability 1 and thus consider the complementary event “Wojtyla is the pope” at the same chance node null. However, the latter event remains *possible*, which means that this branch must not be eliminated from the event tree.

Now, what about the indicative conditional “If Wojtyla is not a communist, then he is the pope”? This statement suggests that we do not know whether or not Wojtyla is a communist. However, it is clearly possible that Wojtyla (as well as any other person in the 1990s) is anything else than the pope if he is not a communist. This is indicated by the two outgoing branches of Chance Node 2 in Figure 2.3. We conclude that the indicative conditional “If Wojtyla is not a communist, then he is the pope” is false. Note that its truth value does not depend on whether or not it turns out later on that Wojtyla, in fact, *is* the pope.

Let us reconsider Heracles’ situation after he passes through Door A and meets Cerberus, which can be described by the decision tree in Figure 2.4. The square at the beginning of the decision tree is said to be a decision node and each outgoing branch represents an action. Heracles’ state of mind, i.e., his status quo, can be seen on the upper right of that decision tree. It indicates that both the counterfactual “If Heracles had gone through Door B, he would have met Cerberus” and the counterfactual “If Heracles had gone through Door B, he would not have met Cerberus” are false from Heracles’ perspective.




Figure 2.4: Heracles’ decision tree.

By contrast, Cerberus’ event tree is depicted in Figure 2.5. Cerberus knows, right from the start, that he is waiting behind Door A and so the given event tree indicates that the counterfactual “If Heracles had gone through Door B, he would not have met Cerberus” is true from his point of view, whereas “If Heracles had gone through Door B, he would have met Cerberus” is clearly false.




Figure 2.5: Cerberus' event tree.

An event tree just describes the possible course of events and does not refer to any decision at all. However, in this book we typically deal with actions. For example, Heracles has to *decide* whether to go through Door A or through Door B. The main point is that a decision maker *changes the world* by making some decision. More precisely, he changes the place of w_0 in W .²⁴ Thus, his actions can change the course of events, and this is precisely what decision theory is all about.

Let us step into Heracles' shoes. We can reformulate the counterfactuals that are given in Figure 2.4, in a more personalized way, as “If I had gone through Door B, I would have met Cerberus” and “If I had gone through Door B, I would not have met Cerberus.” A decision tree differs from an event tree in that it allows the subject to perform some actions and thus to change the course of events. I will explain in Section 2.4 why actions and events are very different things. Beforehand, I would like to concentrate on the fundamental notion of strategy and scenario.

2.3.3 Strategy and Scenario

In order to understand the overall concept, the reader should take a look at the stylized decision tree in Figure 2.6. As already mentioned in the previous section, the squares in the decision tree are said to be decision nodes and their outgoing branches represent the decision maker's available actions. Further, the circles are referred to as chance nodes and each branch behind a circle represents an event. Finally, the triangles at the end of the decision tree are called end nodes.

The following definitions are crucial throughout this book.

²⁴ It is worth pointing out that W is not a state space, i.e., Ω . It is a set of possible worlds. In Section 2.4 we will see why we need to distinguish between W and Ω .




Figure 2.6: Stylized decision tree.

Definition 6 (Strategy). A strategy is an exhaustive plan of moves from each decision node to the next node in the decision tree.

Definition 7 (Scenario). A scenario is an exhaustive plan of moves from each chance node to the next node in the decision tree.

Hence, a scenario can be considered a “strategy of Nature.” Each node in the decision tree represents a certain position. A strategy (or scenario) is a rule that specifies how to *change* the position in the decision tree when it is the decision maker’s (or Nature’s) turn. Changing the position means to perform some action or to generate an event. Each position reflects all actions and events that have been made or realized up to the corresponding node. Whenever we combine some strategy with a scenario, we end up in one and only one end node, which contains the associated consequence of the decision maker’s strategy.

A decision problem is said to be simple if and only if the decision tree contains only one decision node. Otherwise, it is said to be sequential. Further, a decision problem is called finite if and only if the number of strategies is finite. In a simple decision problem, the available actions represent (simple) strategies. By contrast, as soon as the decision tree contains two or more decision nodes, strategies are no longer identical with actions. This holds true, e.g., for the (stylized) decision problem that is depicted in Figure 2.6. Obviously, this decision problem is finite.²⁵

How many strategies and scenarios are contained in the decision tree that is shown in Figure 2.6? In order to count the number of strategies and scenarios, we start at the end nodes and run through the decision tree from back to front.

First of all, we count the number of strategies. The counting scheme is depicted in Figure 2.7. We go from right to left. Each end node counts one. At each chance node

²⁵ Almost all sequential decision problems that are considered in this book are finite.




Figure 2.7: Counting the number of strategies.

we *multiply* the numbers at the next nodes. By contrast, at each decision node we *add* the numbers at the next nodes. When we arrive at the beginning of the decision tree, we find the number of strategies. Hence, there are 6 strategies:

1. Choose Action *a*. If Nature chooses Event *A*, choose Action *d*, otherwise choose Action *f*.
2. Choose Action *a*. If Nature chooses Event *A*, choose Action *e*, otherwise choose Action *f*.
3. Choose Action *a*. If Nature chooses Event *A*, choose Action *d*, otherwise choose Action *g*.
4. Choose Action *a*. If Nature chooses Event *A*, choose Action *e*, otherwise choose Action *g*.
5. Choose Action *b*.
6. Choose Action *c*.

Now, we count the number of scenarios (see Figure 2.8). For this purpose, we consider Nature a decision maker. Once again, each end node counts one, but at the chance nodes we *add* and at the decision nodes we *multiply* the numbers that are assigned to the next nodes. After we have finished this procedure, we come to the conclusion that there are 8 scenarios:

1. If the subject chooses Action *a*, choose Event *A*, and if he chooses Action *e*, choose Event *E*. If the subject chooses Action *b*, choose Event *C*.
2. If the subject chooses Action *a*, choose Event *A*, and if he chooses Action *e*, choose Event *F*. If the subject chooses Action *b*, choose Event *C*.
3. If the subject chooses Action *a*, choose Event *B*, and if he chooses Action *g*, choose Event *G*. If the subject chooses Action *b*, choose Event *C*.




Figure 2.8: Counting the number of scenarios.

4. If the subject chooses Action a , choose Event B , and if he chooses Action g , choose Event H . If the subject chooses Action b , choose Event C .
5. If the subject chooses Action a , choose Event A , and if he chooses Action e , choose Event E . If the subject chooses Action b , choose Event D .
6. If the subject chooses Action a , choose Event A , and if he chooses Action e , choose Event F . If the subject chooses Action b , choose Event D .
7. If the subject chooses Action a , choose Event B , and if he chooses Action g , choose Event G . If the subject chooses Action b , choose Event D .
8. If the subject chooses Action a , choose Event B , and if he chooses Action g , choose Event H . If the subject chooses Action b , choose Event D .

Finally, we are able to represent the given decision problem by means of a decision matrix, which can be found in Table 2.3. Each scenario represents an event, i.e., a subset of Ω . When the decision problem has been solved, i.e., when “the game is over,” we know both the decision maker’s strategy and Nature’s scenario. Put another way, we know the row in Table 2.3 that has been chosen by the decision maker and the column that has been “chosen” by Nature. The intersection of that row and that column leads to some consequence, i.e., an element of C , which is symbolized by a Greek letter.

Assume that different Greek letters indicate different consequences and consider the situation of the decision maker before he has made any decision. Table 2.3 reveals that there are 15 different indicative conditionals. For example, “If I choose Strategy 1, the consequence is α .” This indicative conditional is false unless the decision maker knows, a priori, that ω_0 , i.e., the true state of the world, belongs to Scenario 1, 2, 5, or 6. Further, “If I choose Strategy 2, the consequence is β ” is false unless he knows that ω_0

Table 2.3: Stylized decision matrix.

Strategy	Scenario							
	1	2	3	4	5	6	7	8
1	α	α	δ	δ	α	α	δ	δ
2	β	γ	δ	δ	β	γ	δ	δ
3	α	α	ε	ζ	α	α	ε	ζ
4	β	γ	ε	ζ	β	γ	ε	ζ
5	η	η	η	η	θ	θ	θ	θ
6	ι	ι	ι	ι	ι	ι	ι	ι

is part of Scenario 1 or 5, etc. By contrast, “If I choose Strategy 6, the consequence is ι ” is a true indicative conditional.

When all actions have been made and all events have happened, 14 out of the 15 indicative conditionals turn into counterfactuals, where each counterfactual is of the form “If I had performed Strategy x , then the consequence would have been y .” Deliberation is nothing other than thinking about all possible consequences *before* we come to our final conclusion. However, typically most counterfactuals are false because, in general, the decision maker has imperfect information. Thus, he does not know, *a priori*, which consequence occurs after performing some strategy.²⁶ This holds true *a posteriori*. That is, even when the decision problem is solved, the decision maker usually cannot say which consequence would have occurred if he had performed another strategy.

It is important to distinguish between the decision maker’s situation before he has made some choice and the situation when the decision is finished. I call the former situation *ex ante* and the latter *ex post*. This distinction is made not only when analyzing decision problems but also in the context of game theory.²⁷ Note that the attributes “*ex ante*” and “*ex post*” refer to the decision maker’s actions, whereas the attributes “*a priori*” and “*a posteriori*” refer to the course of events, and thus to the information flow of the decision maker.

The vacuous action “do nothing” represents a decision. Further, not making any decision does not mean to decide to do nothing. In some cases it can happen that doing nothing changes the status quo, whereas doing something preserves it. For example, suppose that we see that a free climber is hanging over an abyss and is losing his grip. We can now decide either to do nothing or to come to his rescue. If we decide to do nothing, the climber will inevitably fall into the abyss, but if we decide to come to his rescue, we will save his life.

Counterfactual reasoning allows us to think about how a decision maker can change the course of events by performing some strategy. The subjectivistic approach

²⁶ The precise meaning of “knowledge” will be clarified in Section 2.5.

²⁷ I will come back to this issue in Section 7.4.4.2.

to rational choice enables us to take counterfactuals into account that make no physical sense at all. This can be illustrated by Savage's omelette example. As already mentioned in Section 2.3, Savage suggests that the cook has no influence on the state of the unbroken egg. Counterfactual reasoning is less restrictive. For example, it allows the cook to think that the unbroken egg is good if he breaks it into a saucer but rotten if he throws it away. This feature is extremely important and it is precisely the reason why we can use subjectivistic decision theory in order to explain human behavior. A particularly remarkable example will be discussed in Section 2.7.

If we do not allow the decision maker to change the world by performing some action, as is implicitly done by Savage (1972), then we end up with a brilliant theory of subjective probability. However, at least in my opinion, this theory has not much to do with decision making. Indeed, Savage's subjectivistic approach allows us to conclude that a rational subject assigns each branch of an *event tree* a unique probability, but this is not enough in order to guarantee that the subject is able to assign also each event in a *decision tree* one and only one probability—unless we assume that the decision maker has no influence on Nature's "response." This shortcoming is remedied by counterfactual reasoning.

2.3.4 Simple Decision Problems

2.3.4.1 The Homeowner's Problem

The basic principle of counterfactual reasoning shall be clarified by a simple example: A homeowner is thinking about protecting his house against burglary. His action set A consists only of the following two options:

- (a) Doing nothing and
- (b) making the windows burglarproof.

Whenever I say in the following that the homeowner "changes his windows," I mean that he makes his windows burglarproof.

The homeowner is interested only in two possible descriptions of the world, namely "No break-in" and "Break-in." His available actions are "Do nothing" and "Change windows." Thus, we can express his situation by the scenario matrix in Table 2.4. The scenario set $S = \{s_1, s_2, s_3, s_4\}$ is a partition of Ω , i.e., the events s_1, s_2, s_3, s_4 are such that $s_i \cap s_j = \emptyset$ for all $i \neq j$ and $\bigcup_{i=1}^4 s_i = \Omega$. In most textbooks on decision theory, the columns of the scenario matrix are called "states." I refrain from doing so because in Savage's framework the state space Ω is infinite.

The action of the homeowner might have an influence on Nature's "response." Thus, in order to analyze the given decision problem, we must take all possible assignments of "No break-in" and "Break-in" to the available actions of the homeowner into account. Each scenario, i.e., column of the scenario matrix, represents a possible assignment, and it can be considered a strategy of Nature.

Table 2.4: Scenario matrix of the homeowner.

Action	Scenario			
	s_1	s_2	s_3	s_4
Do nothing	No break-in	Break-in	No break-in	Break-in
Change windows	No break-in	No break-in	Break-in	Break-in

Hence, the basic idea of counterfactual reasoning can be stated like this: The scenario set is $S := P^A$, i.e., the set of all functions from A to P , where $A \neq \emptyset$ is the (finite) action set of the decision maker and $P \neq \emptyset$ is a finite set of propositions. In the homeowner's example, the action set is $A = \{\text{Do nothing, Change windows}\}$ and the set of propositions is $P = \{\text{No break-in, Break-in}\}$. Each element of S represents a column of the scenario matrix in Table 2.4. This elementary approach to counterfactual reasoning is already discussed by Gibbard and Harper (1981), who attribute the basic idea to Jeffrey (1965b).

A scenario in which the propositions depend on the decision maker's action is said to be heterogeneous, whereas all other scenarios are called homogeneous. For example, Scenario s_2 and Scenario s_3 are heterogeneous, whereas Scenario s_1 and Scenario s_4 are homogeneous. If we ignore all heterogeneous scenarios of the homeowner, we obtain the reduced scenario matrix in Table 2.5. Reduced scenario (and decision) matrices play an important role throughout this book because they occur whenever the decision maker knows or, at least, believes that he has no influence on Nature's "response." However, for the time being, we shall proceed further with the (full) scenario matrix of the homeowner.

Table 2.5: Reduced scenario matrix of the homeowner.

Action	Scenario	
	s_1	s_4
Do nothing	No break-in	Break-in
Change windows	No break-in	Break-in

The homeowner does not know which scenario happens. However, if he is rational, he assigns each scenario a subjective probability. For example, we could have the probability distribution given in Table 2.6. Obviously, the homeowner considers the third scenario, i.e., "No break-in if I do nothing and break-in if I change the windows," null. If he would think that his action has no influence at all on Nature's "response," then we would also have that $P(s_2) = 0$. Similar situations are considered by Savage (1972).

However, the homeowner assigns Scenario s_2 , i.e., "Break-in if I do nothing and no break-in if I change the windows," probability 15%. Thus, he expects a positive

Table 2.6: Scenario matrix with subjective probabilities.

Action	Scenario			
	80 %	15 %	0 %	5 %
	s_1	s_2	s_3	s_4
Do nothing	No break-in	Break-in	No break-in	Break-in
Change windows	No break-in	No break-in	Break-in	Break-in

effect of making his windows burglarproof. This phenomenon, namely that a decision can have an impact on Nature's "response," occurs in many, if not even most, real-life situations, which is precisely the reason why we need counterfactuals in decision theory.

Now, we can assign each combination of action and scenario a consequence. Suppose that changing the windows costs \$5,000 and if somebody breaks into the house, the homeowner loses \$50,000. This leads us to the decision matrix in Table 2.7. The decision matrix reveals that each action can be considered a Savage act. In fact, the action "Do nothing" is a function on Ω that assigns s_1 and s_3 the consequence \$0, whereas s_2 and s_4 lead to the consequence $-\$50,000$. By contrast, the action "Change windows" assigns s_1 and s_2 the consequence $-\$5,000$, whereas s_3 and s_4 lead to the consequence $-\$55,000$.

Table 2.7: Decision matrix of the homeowner with consequences.

Action	Scenario			
	80 %	15 %	0 %	5 %
	s_1	s_2	s_3	s_4
Do nothing	\$0	$-\$50,000$	\$0	$-\$50,000$
Change windows	$-\$5,000$	$-\$5,000$	$-\$55,000$	$-\$55,000$

Since the homeowner is rational, he has also some utility function u , which assigns each dollar amount a real number. Suppose that the utilities are given by the decision matrix in Table 2.8.²⁸ This leads us to the expected utilities

$$EU(\text{Do nothing}) = 0.8 \cdot 0 + 0.2 \cdot (-9) = -1.8$$

and

$$EU(\text{Change windows}) = 0.95 \cdot (-1) + 0.05 \cdot (-10) = -1.45.$$


We conclude that the homeowner prefers to change his windows.

²⁸ Throughout this book, consequences that can be found in tables and figures are written without parentheses, whereas the corresponding utilities are given in parentheses.

Table 2.8: Decision matrix of the homeowner with utilities.

Action	Scenario			
	80 %	15 %	0 %	5 %
	s_1	s_2	s_3	s_4
Do nothing	(0)	(-9)	(0)	(-9)
Change windows	(-1)	(-1)	(-10)	(-10)

We can express the entire problem in a more convenient way by the decision tree in Figure 2.9. The probabilities in the decision tree can easily be calculated on the basis of the scenario matrix in Table 2.6. The end nodes contain the potential consequences of the homeowner's actions as well as the associated utilities in parentheses. Correspondingly, the numbers in parentheses at the chance nodes are the expected utilities of the homeowner, given that he decides to do nothing or to change the windows.

**Figure 2.9:** Homeowner's decision tree.

The first branch behind the first circle (from above) means that nobody will break in if the homeowner decides to do nothing. Hence, this branch represents the event $s_1 \cup s_3$. Further, the second branch behind the first circle means that there is a break-in if he decides to do nothing and so this is the complementary event $s_2 \cup s_4$. Similarly, the first branch behind the second circle is $s_1 \cup s_2$ and the second branch behind the second circle corresponds to the event $s_3 \cup s_4$. Thus, each event in the decision tree is also an event in the usual sense of probability theory, i.e., an element of \mathcal{F} . For example, we have that $P(s_1 \cup s_3) = 0.8$, but in the subjectivistic framework I will never write “ $P(\text{No break-in} \mid \text{Do nothing}) = 0.8$.” Such a statement makes no sense because “Do nothing” is a *decision*, but the decision of a subject cannot be an event, i.e., a subset of his own state space Ω . I will come back to this crucial point in Section 2.4.

Recall that we can interpret the descriptions “No break-in” and “Break-in” as Nature's possible “responses” to the homeowner's action. Thus, a scenario represents an exhaustive plan of Nature's “reactions” to the homeowner's action. Simply put,

Nature is prepared for each possible action of the decision maker. For example, Scenario s_2 states that Nature “sends” a burglar to break in if the homeowner decides to do nothing, but it does not send anybody if he decides to change his windows. This scenario is marked bold in Figure 2.10. Hence, a scenario is a *combination* of events. Of course, in the same way, we could depict any other scenario by choosing the right combination of events.




Figure 2.10: Homeowner's decision tree with Scenario s_2 marked bold.

Although the given example is very simple, it clearly reveals that it is by far more convenient (and perhaps less prone to error) to use a decision tree rather than a decision matrix in order to find an optimal choice of the decision maker. This will become manifest, in particular, when dealing with sequential decision problems.²⁹ The same conclusion can be drawn for strategic conflicts, which are treated in Part II of this book.

Before we proceed further, I would like to discuss an important issue: A scenario assigns each available action of the decision maker, i.e., each element of A , a proposition, i.e., a description of the world, not a consequence! Indeed, depending on the chosen action, the same proposition can lead to different consequences. For example, it makes an essential difference whether the homeowner has done nothing or, instead, has changed his windows if nobody breaks in, since we may suppose that making the windows burglarproof does not come for free.

In contrast to Karni (2017), I refrain from using C^A as a state space Ω or as a scenario set S . Karni's approach goes very much in the same direction as $S := P^A$. However, it is not the same, and choosing C^A as a state space can lead to some conceptual difficulties. A straightforward argument why C^A cannot be a state space in Savage's framework has already been mentioned before: The state space Ω must be infinite. Otherwise, we cannot derive a unique (quantitative) probability measure for each rational decision maker. Nonetheless, beyond this trivial observation, it turns out that C^A cannot serve as a scenario set either. I would like to explain my opinion here and in the following section.

²⁹ I will discuss some examples in Section 2.8.

The homeowner's set of consequences is $C = \{-\$55,000; -\$50,000; -\$5,000; \$0\}$ or it is even larger than that. Hence, C^A would lead us not only to $2^2 = 4$ but to $4^2 = 16$ different scenarios, and some of them are impossible. For example, the scenario "The homeowner loses \$55,000 irrespective of whether he did nothing or made the windows burglarproof" cannot happen. More precisely, the decision maker *knows* that he cannot lose \$55,000 no matter what he does.³⁰ It is clear that a (Bellman) rational decision maker assigns each impossible scenario probability 0. However, the mere fact that those scenarios remain part of C^A implies that they are still considered *possible* by the subject.

2.3.4.2 The Horse Race

There is another essential difficulty: The state space $\Omega = C^A$ precludes constant acts, which are crucially important in Savage's framework. For example, consider a horse race. If the decision maker decides to bet on his favorite horse "Lucky," he can either win \$100 or lose \$100, depending on whether Lucky prevails or not. By contrast, if he decides not to bet, he goes away empty-handed. This means that the decision maker's action set is $A = \{\text{No bet, Bet}\}$ and his set of consequences reads $C = \{-\$100, \$0, \$100\}$.

Now, let us assume that C^A is the state space. Then each row in the decision matrix must contain all elements of C and so "No bet" cannot be a constant act! This is illustrated in Table 2.9. No bet would be a constant act only after removing all states but $\omega_2, \omega_5, \omega_8$ from C^A . Actually, we would have to remove also ω_5 because this state is impossible, too. Hence, C^A cannot be the state space, and the same arguments hold true if we propose C^A as a scenario set.

Table 2.9: Decision matrix of the horse race if " $\Omega = C^A$."

	Ω								
	ω_1	ω_2	ω_3	ω_4	ω_5	ω_6	ω_7	ω_8	ω_9
No bet	-\$100	\$0	\$100	-\$100	\$0	\$100	-\$100	\$0	\$100
Bet	-\$100	-\$100	-\$100	\$0	\$0	\$0	\$100	\$100	\$100

The conceptual difficulties discussed above can easily be solved by choosing

1. $P = \{\text{Lucky wins, Lucky loses}\}$ as a set of propositions;
2. $S = P^A$ as a scenario set, and
3. assigning each combination of action and scenario a consequence.

³⁰ Of course, it can very well happen that the homeowner loses \$55,000 in stock trading, but this loss is not a result of the decision at hand.

This method has already been applied to the homeowner's problem and it is illustrated for the horse race in Table 2.10.

Table 2.10: Scenario and decision matrix of the horse race if $S = P^A$.

	Scenario			
	s_1	s_2	s_3	s_4
No bet	Lucky wins	Lucky loses	Lucky wins	Lucky loses
	\$0	\$0	\$0	\$0
Bet	Lucky wins	Lucky wins	Lucky loses	Lucky loses
	\$100	\$100	−\$100	−\$100

2.3.4.3 The Meeting

Both in the homeowner's problem and in the horse race we have chosen P^A as a scenario set. Of course, this is only one possibility of how to construct a scenario set. We have already seen in Section 2.3.3 how to create scenarios for decision problems that are (much) more difficult than the homeowner's problem and the horse race. Then, in general, the scenario set can no longer be defined by P^A . However, we are always able to deduce the scenario set, in a quite simple way, from the decision tree of the given problem. This will be demonstrated once again in Section 2.9. Our previous choice, i.e., $S := P^A$, is just a special case of the general procedure based on decision trees, which has already been illustrated in Section 2.3.3.

Now, consider the following example in order to see that P^A can be an *inappropriate* choice for S even if the decision problem is very simple: A manager is sitting in his home office, just being struck by the fact that he is going to have an important meeting in one hour. Unfortunately, there is no possibility to do a video or phone conference. He can decide either to stay at home, in which case he will certainly miss the meeting, or hit the road and try to arrive on time. Hence, his action set is $A = \{\text{Stay}, \text{Go}\}$ and the set of propositions is $P = \{\text{On time}, \text{Too late}\}$. However, it makes no sense to consider P^A a scenario set because the manager already *knows* that he cannot arrive on time if he decides to stay at home. This means that S is smaller than P^A . The correct scenario matrix is given in Table 2.11.

Table 2.11: Scenario matrix of the meeting.

	Scenario	
	s_1	s_2
Stay	Too late	Too late
Go	On time	Too late

2.4 Deliberation Crowds Out Prediction

2.4.1 Levi's Postulate

Philosophers have long argued about whether or not a decision can be an event in the sense of probability theory. This question turns out to be essential when modeling strategic conflicts by means of decision theory because each action represents an individual decision. Jeffrey (1965a) tries to abolish Savage's trinity of states, consequences, and acts. He holds that events, consequences, and actions are nothing other than logical propositions. According to his point of view, choosing an action means to choose some event and thus to make some proposition be true.

Shin (1992) points to some peculiarities, which naturally arise from that understanding of "action." A similar argument can be found in Burdzy (2016, p. 128), who refers to Lewis' (1981) causal decision theory. Nonetheless, I have the impression that most philosophers still follow Jeffrey's monotheistic approach. Some well-known opponents of Jeffrey are, e.g., Levi (1997, 2008) and Spohn (1977). For a nice discussion on that topic see also Briggs (2017) as well as Liu and Price (2018).

This work is based on Levi's (1997) famous postulate: "Deliberation crowds out prediction." This phrase can be reformulated in an equivalent manner by saying that "prediction crowds out deliberation." Hence, the true state of the world, $\omega_0 \in \Omega$, is not optional—irrespective of whether we argue from the perspective of a single decision maker or a player in a game against each other.³¹ That is, a decision maker cannot make any event happen and so he cannot create his own evidence. Here, the term "event" shall be understood in the usual sense of probability theory. Thus, it makes no sense for the decision maker to assign his *own* action a probability (Spohn, 1977). This contradicts Jeffrey's monotheistic argument.

I support Savage's polytheistic argument and thus distinguish between deliberation and prediction. If deliberation would not crowd out prediction, the entire probabilistic approach would become meaningless. In my opinion, this assertion holds true not only in the subjectivistic but also in the objectivistic framework. We can say that each element of Ω , i.e., state of the world, is a list of (logical) propositions (Aumann, 1999; Savage, 1972, p. 3) and a subject might very well *know* whether or not some proposition is true. However, he cannot *decide* upon its truth value. Otherwise, it would make no sense to consider any proposition part of some state space.

Levi's postulate shall be motivated by the following example: Suppose that Peter may choose between Heads and Tails. Now, he considers his own choice a probabilistic event and calls himself "Nature." Counterfactual reasoning leads us to the conclusion that there are four scenarios, which can be seen in Table 2.12. The scenarios s_1 , s_2 , and s_4 are clearly impossible. More precisely, they are not only improbable but also

31 However, a subject can very well change, deliberately, the actual world $w_0 \in W$ (see Section 2.3.2).

Table 2.12: Peter’s scenario matrix.

Action	Scenario			
	s_1	s_2	s_3	s_4
Heads	Heads	Tails	Heads	Tails
Tails	Heads	Heads	Tails	Tails

logically inconsistent. I guess that most readers would consider Peter schizophrenic if he would assign the scenarios s_1 , s_2 , or s_4 a positive probability. At least, we could say that Peter is not Bellman rational. Only Scenario s_3 makes sense, but this is a trivial tautology.

Levi’s postulate holds true also in the objectivistic framework. For example, assume that somebody is asked to bet on one of the six sides of a dice. If the dice roll generates “1,” he wins \$100, if the result is “2,” he wins \$200, etc. Each side of the dice represents an event, and we typically consider the dice roll a Laplace experiment. That is, we assign each side the (objective) probability $\frac{1}{6}$. However, the given probabilities do not matter at all. In any case, the situation of the decision maker is illustrated on the left-hand side of Figure 2.11.




Figure 2.11: The dice roll can be considered either an event (left) or an action (right).

Now, suppose that the decision maker were able to make any event happen. This means that he could place any side of the dice facing upwards. Well, in this case, we must replace the chance node on the left-hand side of Figure 2.11 with a decision node. This is depicted on the right-hand side of Figure 2.11. Put another way, deliberation crowds out prediction. On the contrary, if the decision maker is not able to control the dice roll, we must substitute the decision node with a chance node, which means that prediction crowds out deliberation.

Hence, Levi's postulate simply states that each node in a decision tree cannot be both a decision node and a chance node. This is not to say that a decision maker cannot consider his own decision stochastic. However, this is possible only at some *preceding* node in the decision tree.³² Of course, each action, event, and consequence may be considered a proposition, but this does not relieve us of making a clear distinction between those elements of a decision tree. Hence, I warn against neglecting the fact that actions, events, and consequences have quite different meanings in decision theory.

Since deliberation crowds out prediction, the state space must be constructed in such a way that the decision maker cannot choose between the elements of Ω . For example, after getting up in the morning, Bob can decide whether to put on red or black socks. This decision cannot be described by a state of the world, i.e., it must not be part of his *own* state space. By contrast, Ann cannot decide upon Bob's socks. Hence, Bob's possible choices can very well be part of her state space.

Suppose that Ann sits in the next room and that she is rational. Although she might be uncertain about the color of his socks, she is able to assign Red and Black each a subjective probability. If Ann has access to a surveillance camera that has been installed in Bob's room, she even *knows* the color of his socks. However, she is still not able to *control* his choice of dress. That is, knowledge and deliberation are totally different things. These arguments are crucial in situations of strategic conflict, which will be described in Part II of this book.

Now, the reader hopefully understands why each person must have its *own* state space, i.e., Ω is subjective. Assuming that the state space is common to all people in a situation of conflict can lead to substantial contradictions and artificial solutions.³³ Allowing each subject to have his own state space resolves the problem of deliberation and determination, which is often discussed in philosophy. By the way, also the set of consequences, C , is subjective because each element of C shall describe everything that the *individual* is concerned about (Fishburn, 1981, p. 141).

We conclude that the elements of \mathcal{F} are not optional to the decision maker. Readers who are familiar with probability theory or statistics might wonder why I discuss this issue in so much detail. I guess that most of us assume, implicitly, that a probabilistic event cannot be optional. However, I emphasize this point because some authors do *not* agree with me, as I already mentioned at the beginning of this section, and I think that it should be up to the reader to form his own opinion.

2.4.2 Wald's Maximin Rule

Peter is faced with the possibility to make a bet on the weather tomorrow. If he decides to bet, he wins \$200 in case it shines, whereas he loses \$100 if it rains. By contrast, if he

³² This is a typical situation in sequential decision problems, which will be discussed in Section 2.8.

³³ Some of these solutions will be described in Chapter 7.

refuses to bet, he cannot win or lose anything. His decision problem can be described by the scenario matrix and the decision matrix in Table 2.13.

Table 2.13: Peter's scenario matrix (left) and decision matrix (right).

Action	Scenario			
	s_1	s_2	s_3	s_4
Bet	Rain	Shine	Rain	Shine
Don't bet	Rain	Rain	Shine	Shine

Action	Scenario			
	s_1	s_2	s_3	s_4
Bet	−\$100	\$200	−\$100	\$200
Don't bet	\$0	\$0	\$0	\$0

Objectivistic decision theory ignores the heterogeneous scenarios s_2 and s_3 . More precisely, it ignores the scenarios in which the natural events Rain and Shine depend on Peter's choice because this makes no physical sense. This leads us to the reduced decision matrix in Table 2.14. A reduced decision matrix is just the counterpart of a reduced scenario matrix (see Section 2.3.4.1).

Wald's famous maximin rule (Wald, 1950) is prescriptive and thus it belongs to objectivistic decision theory. It goes like this: If Peter decides to bet, the worst case is Rain. By contrast, if he refuses to bet, Rain and Shine lead to the same consequence, \$0, and so those cases are equally "worse." Now, according to Wald, Peter should choose an optimal action by taking only the worst cases into account. Thus, he should *not* bet.

Table 2.14: Peter's reduced decision matrix.

Action	Scenario	
	Rain	Shine
Bet	−\$100	\$200
Don't bet	\$0	\$0

At first glance, Wald's line of argument seems plausible, but a closer look reveals that it violates Levi's postulate. The problem is that Peter has no influence on the true state of the world. That is, the events that are given in Table 2.14 are not optional to him. However, whenever we search for a worst case, we implicitly assume that Nature may change its "opinion" with each action of the decision maker. This means that Peter is able to control the true state of the world!

In the context of subjectivistic decision theory, this problem can easily be solved by counterfactual reasoning. Wald's maximin rule can be motivated like this: Being extremely pessimistic means to believe that Nature has conspired against oneself—according to Murphy's law: "Anything that can go wrong will go wrong!" That is, one

assigns the set of all worst cases probability 1. More precisely, Peter’s subjective probability of the event $s_1 \cup s_3$ (see Table 2.13) amounts to 1, which means that he is convinced that it rains if he decides to bet. Well, then it is clear that Peter should not decide to bet. This can be better illustrated by the decision tree in Figure 2.12.




Figure 2.12: Peter’s decision tree if he applies the maximin rule.

Thus, we come to the same conclusion as the maximin rule, but our arguments are essentially different. In contrast to the objectivistic approach, we do not require the heterogeneous scenarios s_2 and s_3 to be null. This enables us to consider extreme pessimism by fading out all scenarios that do not represent worst cases. To put it more simply: We search for the worst outcomes in the decision tree and neglect all other branches of the chance nodes. The maximin rule turns out to be the right one for an utterly pessimistic, not to say paranoid, decision maker. This rule makes sense only for the most fearful subjects and, in general, it leads to very conservative decisions.

2.4.3 Savage’s Minimax-Regret Rule

Savage goes one step further and suggests to transform Peter’s payoffs into losses, which are contained in Table 2.15: If it rains, Peter would have done better not to bet, since if he has decided to bet, he loses \$100 compared with “Don’t bet.” By contrast, if it shines he would have done better to bet. In this case, he loses \$200 compared with “Bet,” given that he has decided not to bet. The given losses can be interpreted as regret because they quantify the cognitive dissonance that the decision maker suffers from when comparing the best choice with his actual chose, *a posteriori*, i.e., after the true state of the world has been revealed to him.

Now, Peter should choose an action that has the smallest maximum loss, which leads us to Savage’s minimax-regret rule (Savage, 1951, 1972, Section 9.4). Savage allows Peter to randomize his actions, i.e., he may choose Bet with probability p and

Table 2.15: Peter's reduced loss matrix.


Action	Scenario	
	Rain	Shine
Bet	\$100	\$0
Don't bet	\$0	\$200

Don't bet with probability $1 - p$. However, this does not change our overall conclusion.³⁴ Thus, let us ignore any randomization. It turns out that Peter should decide to bet, which contradicts Wald's maximin rule!

The minimax-regret rule suffers from the same problem as the maximin rule: It violates Levi's postulate. Again, this problem can be solved by counterfactual reasoning. Consider the (full) loss matrix in Table 2.16, in which Peter is utterly pessimistic, since he assigns Scenario s_3 probability 1. The corresponding decision tree is given in Figure 2.13. Of course, in this case, he should definitely decide to bet.

Table 2.16: Peter's full loss matrix if he is utterly pessimistic.

Action	Scenario			
	0 %	0 %	100 %	0 %
	s_1	s_2	s_3	s_4
Bet	\$100	\$0	\$100	\$0
Don't bet	\$0	\$0	\$200	\$200

**Figure 2.13:** Peter's decision tree if he applies the minimax-regret rule. The consequences behind each end node represent losses.

The decision tree reveals also that randomization makes no sense at all: If Peter performs a "mixed strategy," he substitutes the decision node in Figure 2.13 with a chance

³⁴ I will come back to this point below.

node whose branches possess the probabilities p and $1 - p$ ($0 < p < 1$). Well, in this case he would risk to lose \$200. Thus, it is better for Peter to choose a *definite* action and not to apply a random generator.

Once again, by counterfactual reasoning, we come to the same conclusion as the minimax-regret rule and the underlying arguments are not essentially different compared with Wald's maximin rule: Savage's minimax-regret rule is the right one for an utterly pessimistic decision maker who focuses on regret rather than payoff. Nonetheless, Savage's minimax-regret rule turns out to be by far less conservative than Wald's maximin rule because it favors actions whose potential payoffs are high compared with the other actions that are available to the decision maker.

2.5 Knowledge and Belief

It is impossible to construct a subjectivistic theory of rational choice without having a clear understanding of knowledge and belief. As I already mentioned in Section 1.1, the state space, Ω , precisely contains those descriptions of the world that are not known, a priori, by the decision maker to be impossible. This means that the decision maker knows that the true state of the world, ω_0 , belongs to Ω , but he does not know that $\omega_0 \in F$ for any proper subset F of Ω .

This shall be illustrated by the following example: A good friend promised to give you back some book that you have lent him long ago. You are just meeting him on Times Square in New York City and you wonder whether or not he has brought the book with him. Some possible descriptions of the world are thus: "The friend is now on Times Square and has brought the book with him" and "The friend is now on Times Square and has not brought the book with him." By contrast, since it is evident to you that your friend is on Times Square, "The friend is now on Wallstreet and has brought the book with him" is clearly impossible and thus must be ignored when forming your state space Ω .

Hence, even though we call the information Ω "trivial," it should be clear that the state space reflects the prior knowledge of a subject and so each decision maker may have his own state space.³⁵ The Bellman postulate just guarantees that we can treat the information that is *a posteriori* available to the decision maker like the trivial one. Put another way, after receiving some new information, a Bellman-rational decision maker ignores all states of the world that now turn out to be impossible, which is precisely the reason why we can apply backward induction when solving a sequential decision problem.³⁶

³⁵ This holds true not only due to informational asymmetry but also because of the mere fact that deliberation crowds out prediction (see Section 2.4). Hence, the action of a decision maker cannot be part of his *own* state space, but it may very well belong to the state space of another subject.

³⁶ This will be demonstrated in Section 2.8.

Further, the subjective probability measure, P , of a rational decision maker, i.e., his prior, does not depend on the true state of the world $\omega_0 \in \Omega$. Savage presumes that the decision maker has only the trivial information Ω and P is just a result of his individual preferences, which do not vary with ω_0 if the decision maker has only the trivial information. Bayesian decision theory goes one step further.

Let F be any event and suppose that the decision maker is equipped with some information set $\mathcal{I} = \{I, \neg I\}$, where I shall be a nontrivial information. In case the true state of the world, ω_0 , belongs to I , he knows that I happens and otherwise he knows that $\neg I$ happens. Let P_I be his posterior based on the information I and $P_{\neg I}$ the posterior given $\neg I$. His prior probability $P(F)$ still does not depend on ω_0 , but his posterior probabilities $P_I(F)$ and $P_{\neg I}(F)$, in general, depend on whether ω_0 belongs to I or $\neg I$. However, the posterior probabilities of F , once again, do *not* depend on the specific location of ω_0 in I or $\neg I$, respectively. This holds true for the same reason why the prior probability $P(F)$ does not depend on the specific location of ω_0 in Ω if the decision maker has only the (trivial) information Ω .

Definition 8 (Belief). A rational subject believes that $F \in \mathcal{F}$ happens if and only if $P_I(F) = 1$ with $\omega_0 \in I \in \mathcal{I}$.

Assume that $\omega_0 \in I$, which means that the subject receives the information I . Then he assigns F probability 1, i.e., $P_I(F) = 1$, if and only if he believes that F happens. This means that he is *convinced* about F . Of course, then the decision maker must consider the complementary event $\neg F$ null, i.e., we have that $P_I(\neg F) = 0$. That is, given his particular information I , he believes that $\neg F$ does not happen.

At this point, I should mention some linguistic peculiarities: “Believing” that F happens or, equivalently, being “convinced” about F , means that $P_I(F) = 1$. If I say that the decision maker “doubts” or that he does *not* believe that F happens, I mean that $P_I(F) < 1$. By contrast, if the decision maker believes that F does *not* happen, then we have that $P_I(F) = 0$. Hence, believing that F does not happen is much stronger than not believing that F happens.

Now, suppose that the decision maker is convinced about F , i.e., $P_I(F) = 1$, and F truly happens, i.e., $\omega_0 \in F$. Does he also *know* that F happens? How can we properly distinguish between “knowledge” and “belief” or, equivalently, between “evidence” and “conviction”? This question is quite philosophical. However, its answer is very simple but crucially important if we want to understand the basic principles of rational choice.

Definition 9 (Knowledge). A subject knows that $F \in \mathcal{F}$ happens if and only if $I \subseteq F$ with $\omega_0 \in I \in \mathcal{I}$.

Hence, in order to know that F happens, two conditions must be satisfied:

1. The subject must receive any information I , i.e., $\omega_0 \in I \in \mathcal{I}$, and
2. it must be *impossible* that F does not happen if I happens, i.e., $I \subseteq F$.

Note that, in contrast to the definition of belief, the definition of knowledge does not require any probability measure and thus a *rational* subject. Knowledge is only a matter of the information set of the decision maker.

The overall concept of knowledge is illustrated in Figure 2.14: The subject is equipped with the information set $\mathcal{I} = \{I_1, I_2, I_3\}$. We can see that $I_1 \subseteq F$. Hence, if he receives the information I_1 , i.e., if $\omega_0 \in I_1$, he knows that the event F happens. If he receives the information I_2 , he does not know whether F or $\neg F$ happens because neither $I_2 \subseteq F$ nor $I_2 \subseteq \neg F$. Finally, if he receives the information I_3 , he knows that the complementary event $\neg F$ happens, since it holds that $I_3 \subseteq \neg F$.




Figure 2.14: State space Ω (black line), the information set $\mathcal{I} = \{I_1, I_2, I_3\}$ (blue line), the event F , and the complementary event $\neg F$ (red line).

Suppose that the decision maker is Bellman rational and knows that F happens. Then we have that

$$P_I(F) = P_I(F \cap I) + \underbrace{P_I(F \cap \neg I)}_{=0} = \underbrace{P_I(F \cap I)}_{=I} = 1.$$

We conclude that knowledge implies belief. However, the converse is not true.

2.5.1 Waking Up in the Morning

The essential difference between knowledge and belief shall be illustrated like this: A student wakes up in the morning. All windows and curtains are closed. Suppose that the house is built in such a way that he cannot hear raindrops clattering on a window. That is, he cannot see or hear what is going on outside. He wonders whether or not it is raining and comes to the conclusion that it is raining. After standing up and taking a look outside, it turns out that he was right. We may assume that the weather has not changed in the meantime, i.e., after the student has come to his conclusion. However, did he really *know* that it is raining? Of course, the answer is “No”!

Let F be the event Rain and $\neg F$ be the event Shine. Further, assume that $I \in \mathcal{F}$ is the student's available information. For example, this could be the current time and date, the appearance of the room, his own physical condition, etc. The problem is that $I \not\subseteq F$. Put another way, it is possible that it shines outside although the student is convinced that it is raining. The fact that it rains, i.e., $\omega_0 \in F$, does not change anything. Indeed, he would have been still convinced that it rains if it were, in fact, dry outside, since his posterior, P_I , does not depend on the specific location of $\omega_0 \in I$. This situation is illustrated in Figure 2.15.




Figure 2.15: Student's information set (blue line) if the curtains are closed.

By contrast, if the curtains were open, he would have known that it is raining. More precisely, then the weather condition would have been *evident* to him. In this case, his private information set, \mathcal{I} , were at least as fine as $\{F, \neg F\}$ and so it would hold that $I \subseteq F$ or $I \subseteq \neg F$ for each $I \in \mathcal{I}$ (see Figure 2.16). Then, by the very definition of “knowledge,” he would know if it rains and also if it shines.




Figure 2.16: Student's information set (blue line) if the curtains are open.

These insights will play a major role later on when analyzing strategic conflicts. In my opinion, we cannot understand the meaning of action, reaction, and interaction without a proper distinction between knowledge and belief.

2.5.2 Newcomb's Paradox

Newcomb's paradox (Nozick, 1969) represents a touchstone in decision theory and can be stated like this: There is a rational subject, a predictor, and two boxes designated A and B. Box A is transparent and contains \$1,000, which is visible to the subject. Box B is opaque and the subject knows that the predictor has put either nothing or \$1 million into the box. Now, the subject has the following choices:

1. Box A and Box B together or
2. Box B alone.

He knows that the predictor is an omniscient being. More precisely, the subject is aware of the following: If the predictor foresees that the subject chooses Box A and Box B together, then he puts nothing into Box B. By contrast, if he predicts that the subject chooses Box B alone, then he decides to fill that box with \$1 million.

Which choice is optimal for the subject? The subjectivistic approach comes to a clear answer: Box B alone!

The solution to Newcomb's paradox turns out to be trivial in our context. This can be seen immediately by drawing a decision tree, which is done in Figure 2.17. Since the

subject *knows* the strategy of the predictor, each chance node has only one outgoing branch. It is evident that Box B is optimal.




Figure 2.17: Decision tree of the subject in Newcomb’s paradox.

The fact that the content of Box B has already been set before the subject makes his choice poses no real challenge at all in the subjectivistic framework. Although the chance nodes appear after the decision node in Figure 2.17, the corresponding event, i.e., the prediction, does not take place *after* the subject has made his decision. This example nicely demonstrates that the particular order of nodes in a decision tree need not have any chronological meaning and we should not confuse cause and effect, i.e., the subject’s choice and the omniscient being’s prediction. Hence, Newcomb’s paradox represents an educational example in which the effect takes place *before* its cause.

Prediction crowds out deliberation. Hence, if somebody predicts the action of another person, the former cannot also decide upon the choice of the latter. The omniscient being only *predicts* the choice of the subject—he is not responsible for his decision. This can be explained also by counterfactual reasoning: Actually, the predictor decides to put \$1 million into Box B, but if the subject had chosen Box A and Box B together, the omniscient being would have decided to put nothing into that box. That is, the predictor’s decision depends on the subject’s, not vice versa.

The subjectivistic solution to Newcomb’s paradox becomes more transparent if we realize that the predictor’s strategy represents a scenario from the subject’s viewpoint. The subject *knows* the strategy of the predictor, i.e., he is equipped with evidence. That is, he simply knows that each other scenario is impossible and thus it goes beyond his state space Ω . In fact, the given decision matrix is trivial (see Table 2.17), and the subject clearly prefers Action B to Action A+B.

Table 2.17: Subject’s decision matrix of Newcomb’s paradox.

Action	Scenario
A+B	\$1,000
B	\$1 million

2.5.3 Communication vs. Information

In general, communication does not lead to the desired information. For example, suppose that your best friend claims that he is going to be promoted to head of some corporation. Well, in the first place you only know that he is *claiming* to be promoted to head, but do you really know that his statement is true? Of course, if you were at the shareholders' meeting, you would have the necessary information in order to judge whether your friend is right or wrong.³⁷ Otherwise, it is hard to imagine why his statement should provide you with *evidence*.

This can be seen as follows: You know that your friend tells you that he is going to be promoted to head. This information may very well have a substantial impact on your posterior probabilities. Let us assume that, based on this information, you are convinced that his assertion is true. Is it still possible that your friend is wrong? The answer is "Yes"! Thus, you do not know that he is going to be promoted. This can be illustrated in Figure 2.15, where we just have to substitute "Rain" with "He is going to be promoted to head" and "Shine" with "He is not going to be promoted to head." Further, the information *I* stands for "He tells me that he is going to be promoted to head" and thus $\neg I$ means that "He does not tell me that he is going to be promoted to head."

By contrast, if you were at the shareholders' meeting, you would be equipped with more information. This can be illustrated in Figure 2.16, once again, after substituting "Rain" with "He is going to be promoted to head" and "Shine" with "He is not going to be promoted to head." In this case it would be evident to you whether or not your friend is right, which means that you would certainly *not* believe that he is going to be promoted to head if he, in fact, is not. By "certainly" I presuppose that you are Bellman rational. Hopefully, this demonstrates that evidence is not the same as conviction.

Nonetheless, we can imagine some situations in which communication, in fact, leads to evidence. For example, a warehouse service provider seeks for a forklift truck operator. Suppose that we live in a world in which a forklift licence cannot be forged. This means that it is *impossible* that an applicant possessing a license has not attended a driver's instruction. Hence, showing the license when applying for the job is a form of communication that leads to evidence. In the economics literature this is called signaling (Spence, 1973).

Indeed, since the event "The applicant has not attended a driver's instruction, but he has a license" cannot happen, our specific criterion for knowledge is fulfilled as soon as the applicant brings a license. Without any license, the employer does not know whether or not the applicant has attended a driver's instruction. We can expect that the employer assigns the event that the applicant has got some instruction if he

³⁷ Here, I implicitly rule out the possibility of being deceived by a staged performance of the company's board, i.e., the decisions made at the shareholders' meeting are binding.

comes without any license a very small probability. The overall situation is illustrated in Figure 2.18.




Figure 2.18: Signaling.

Here is another example: A husband wants to know whether his wife is at home. He calls her at the telephone and she answers. Well, for the sake of simplicity, let us suppose that we live in a world in which it is technically impossible to answer a phone call at home while being somewhere else. Then the husband knows that his wife is at home. This can very well be illustrated, once again, in Figure 2.18, where we have to substitute “License” with “She answers,” “No license” with “She does not answer,” “Applicant has attended a driver’s instruction” with “She is at home,” and “Applicant has not attended a driver’s instruction” with “She is not at home.” If his wife does not answer, the husband still does not know whether or not she is at home. However, at least he knows that she does not answer the phone call, which is more than not knowing whether she answers or not. The fact that she does not answer can have a substantial (negative) impact on his subjective probability that his wife is at home.

2.6 The Dominance Principle

The sure-thing principle is a main pillar of subjectivistic decision theory. This shall be demonstrated by using Corollary 1, i.e., Part II of the sure-thing principle. In objectivistic decision theory, the content of this theorem is commonly referred to as the dominance principle, which is considered the most basic principle of rational choice.³⁸ However, the dominance principle seems to be often misunderstood and a naive application can lead to wrong or paradoxical conclusions. This shall be demonstrated by a number of examples below.

2.6.1 Dominance

I say that Action a is better than Action b in Scenario s if and only if the consequence of Action a in Scenario s is preferred to the consequence of Action b in Scenario s .

³⁸ The close connection between those principles is stressed also by Savage (1972, p. 114).

Dominance. Neglect all scenarios with probability 0.

1. If Action a is not worse than Action b in each scenario and
 2. Action a is better than Action b at least in one scenario;
- then the decision maker prefers Action a to Action b .

It does not matter whether we refer to consequences or to utilities when applying the dominance principle. In fact, Savage's representation theorem just states that a consequence is better than another if and only if its utility is greater than the utility of the other consequence. Thus, if we know whether the decision maker considers a given consequence better, equivalent, or worse than another, we need not specify the utility function of the decision maker in order to apply the dominance principle. However, it is worth emphasizing that, when searching for dominant actions, we must not take scenarios with probability 0 into account.

The decision matrix in Table 2.18 shall clarify a typical mistake. Does Action a dominate Action b ? We could be inclined to say "Yes", but actually the answer is: "We don't know"! It depends essentially on the probabilities of the decision maker: If the subject considers Scenario s_2 null, he is indifferent among Action a and Action b . Otherwise, we can apply the dominance principle and come to the conclusion that Action a is preferred to Action b , since the former dominates the latter.

Table 2.18: Decision matrix without probabilities.

Action	Scenario		
	s_1	s_2	s_3
a	9	3	2
b	9	2	2

Now, consider the decision matrices in Table 2.19. At first glance, we could be tempted to ignore the dominance principle because Action a is better than Action b in Scenario s_2 , but in Scenario s_3 it is worse. However, this is a mistake, too. We cannot search for dominated actions without knowing the null scenarios of the decision maker! If a rational subject has the probabilities on the left-hand side, he prefers Action a to Action b . By contrast, the subjective probabilities on the right-hand side reveal that Action b is preferred to Action a .

Table 2.19: Decision matrices with probabilities.

Action	Scenario		
	40 %	60 %	0 %
	s_1	s_2	s_3
a	9	3	1
b	9	2	2

Action	Scenario		
	40 %	0 %	60 %
	s_1	s_2	s_3
a	9	3	1
b	9	2	2

Now, suppose that the homeowner from Section 2.3.4.1 considers both Scenario s_2 and Scenario s_3 null. Thus, he believes that his decision about either to do nothing or to make his windows burglarproof has no influence at all on whether or not somebody breaks into his house. In this case, we could have the decision matrix in Table 2.20. Since we must neglect the null scenarios s_2 and s_3 when searching for a dominated action, we can use also the reduced decision matrix in Table 2.21.

Table 2.20: Full decision matrix of the homeowner.

Action	Scenario			
	80 %	0 %	0 %	20 %
	s_1	s_2	s_3	s_4
Do nothing	\$0	\$0	−\$50,000	−\$50,000
Change windows	−\$5,000	−\$55,000	−\$5,000	−\$55,000

Table 2.21: Reduced decision matrix of the homeowner.

Action	Scenario	
	80 %	20 %
	No break-in	Break-in
Do nothing	\$0	−\$50,000
Change windows	−\$5,000	−\$55,000

The reduced decision matrix is based on the reduced scenario matrix of the homeowner, which can be found in Table 2.5. This explains why the scenario s_1 is labelled “No break-in,” whereas s_4 is associated with “Break-in.” The utility function of the homeowner tells us that he prefers more money to less. Obviously, since “Do nothing” dominates “Change windows,” the homeowner prefers to do nothing. This result is not very surprising: Why should the homeowner invest some money in making his windows burglarproof if he doubts that this will have any (positive) impact? By contrast, if he believes that making his windows burglarproof can improve his situation, there is no dominance at all. To me it seems much more convincing to assume that (the homeowner assumes that) changing the windows *has* some positive impact, i.e., that Scenario s_3 is not null.

Now, consider another problem, which once again reveals that the dominance principle must be applied with care³⁹: Suppose that there are two urns, where Urn I contains only one ball and Urn II contains two balls. The ball in Urn I can either be red or black, whereas one ball in Urn II is red and the other ball is black. A decision maker can choose between Urn I and Urn II. If he chooses Urn I, the ball in that urn is revealed to him. Otherwise, he draws a ball at random from Urn II. If Red occurs, he

³⁹ Similar examples will be discussed in Section 2.7.

wins \$200 if he has chosen Urn I, but he wins only \$100 if he has opted for Urn II. If Black occurs, he goes away empty-handed in either case.

Let us take a look at the decision matrix in Table 2.22. At first glance, it seems clear that a rational subject should prefer Urn I. Indeed, after a naive application of the dominance principle, we might come to the conclusion that Urn I dominates Urn II because Urn I is better than Urn II if Red occurs but not worse if Black occurs. However, this is a fallacy. In fact, the decision matrix in Table 2.22 is *reduced*. That is, by using that decision matrix we implicitly presume that $P(s_2) = P(s_3) = 0$. Moreover, the given decision matrix suggests that $P(s_1) = P(s_4) = 0.5$, which is not necessarily true from a subjectivistic point of view.

Table 2.22: Reduced decision matrix of the two-urns problem.

Action	Scenario	
	50 %	50 %
	•	•
Urn I	\$200	\$0
Urn II	\$100	\$0

There is no reason why both Scenario s_2 and Scenario s_3 should be considered null by a rational decision maker, and he even need not consider Red and Black equally probable. More precisely, even if some objective probabilities *are* available, which de facto holds true at least for Urn II, they need not correspond to the subjective probabilities of the decision maker!

Table 2.23 contains the (full) scenario matrix and the associated decision matrix with an exemplary distribution of subjective probabilities. The problem is that the decision maker assigns Scenario s_2 , in which Urn I turns out to be worse than Urn II, positive probability. Thus, Urn I does *not* dominate Urn II.

Table 2.23: Scenario (left) and decision matrix (right) of the two-urns problem.

Action	Scenario			
	20 %	30 %	20 %	30 %
	s_1	s_2	s_3	s_4
Urn I	•	•	•	•
Urn II	•	•	•	•

Action	Scenario			
	20 %	30 %	20 %	30 %
	s_1	s_2	s_3	s_4
Urn I	\$200	\$0	\$200	\$0
Urn II	\$100	\$100	\$0	\$0

The decision matrix in Table 2.23 can be expressed, equivalently, by the decision tree in Figure 2.19. It is clear that we can find some utility function u such that the decision maker prefers Urn II and not Urn I. For example, assume that $u(\$0) = 0$, $u(\$100) = 9$, and $u(\$200) = 10$, in which case we have that

$$EU(\text{Urn I}) = 0.6 \cdot 0 + 0.4 \cdot 10 = 4$$

and

$$EU(\text{Urn II}) = 0.5 \cdot 0 + 0.5 \cdot 9 = 4.5.$$




Figure 2.19: Decision tree of the two-urns problem.

Now, let us come back to Newcomb's paradox (see Section 2.5.2). According to Nozick (1969), the paradox comes from a typical misapplication of the dominance principle. In order to understand this point, the reader may take a look at Table 2.24, which contains the reduced decision matrix of the paradox. This decision matrix suggests that the subject has no influence on the predictor's choice, which does not appropriately reflect the given information. In fact, the subject *knows* that the two (homogeneous) scenarios that are expressed in Table 2.24 cannot happen! More precisely, it is impossible that the predictor chooses either nothing or \$1 million, *irrespective* of whatever the subject does. This means that the (subjective) probabilities that are associated with the homogeneous scenarios are zero and for this reason we cannot apply the dominance principle to that decision matrix at all.

Table 2.24: Reduced decision matrix of Newcomb's paradox.

Subject's action	Predictor's action	
	Nothing	\$1 million
A+B	\$1,000	\$1,001,000
B	\$0	\$1,000,000

2.6.2 Superdominance

The following criterion is stronger than dominance and does not require us to specify the null scenarios of the decision maker:

Superdominance. If the worst consequence of Action a is better than the best consequence of Action b , then the decision maker prefers Action a to Action b .

For example, reconsider the two-urns problem, which has been discussed in the last section, and suppose that the decision maker can now receive the payoffs in Table 2.25. We can see that Urn I superdominates Urn II: The worst consequence of Urn I is \$200, whereas the best consequence of Urn II is \$100. The former is better than the latter and so Urn I superdominates Urn II. This can be verified by the (full) decision matrix in Table 2.26. However, in order to find superdominant actions, it suffices to use the *reduced* decision matrix.

Table 2.25: Reduced decision matrix of the two-urns problem with alternative payoffs.

Action	Scenario	
	•	•
Urn I	\$300	\$200
Urn II	\$100	\$0

Table 2.26: Decision matrix of the two-urns problem with alternative payoffs.

Action	Scenario			
	s_1	s_2	s_3	s_4
Urn I	\$300	\$200	\$300	\$200
Urn II	\$100	\$100	\$0	\$0

The superdominance of Urn I can be illustrated by the decision tree in Figure 2.20. Since the expected utility of an action is a convex combination of the smallest and the largest utility of the action, the expected utility of Urn I must be greater than the expected utility of Urn II. Put another way, the superdominant action, Urn I, must have a greater expected utility, irrespective of whether or not the subjective probabilities of the decision maker depend on his choice between Urn I and Urn II. For this reason, we cannot make any mistake by searching for superdominance in the reduced decision matrix. Recall that this does not hold true for dominance!

2.6.3 Quintessence

The quintessence of this section is that a naive application of the dominance principle is very dangerous. Although the dominance principle is the most basic principle of rational choice, it can lead to wrong conclusions and paradoxical results. The typical mistakes are that




Figure 2.20: Superdominance in the two-urns problem.

1. we do not neglect null scenarios or
2. we neglect scenarios that have positive probability.

These observations are all the more important when applying the dominance principle to strategic conflicts. In fact, the dominance principle plays a major role in game theory and I will return to this issue in Section 3.6.

2.7 Ellsberg's Paradox

The subjectivistic approach to decision theory is badly affected by two famous thought experiments that are made by Ellsberg (1961). He argues that most decision makers are ambiguity averse. This means that they prefer situations in which objective probabilities are available to situations without those probabilities. According to Knight (1921), the former class of situations is typically referred to as “risk,” whereas the latter class of situations is called “uncertainty.” Ellsberg concludes that ambiguity aversion contradicts Savage’s postulates of rational choice.

In this section, I provide a solution to Ellsberg’s paradox. I do not question ambiguity aversion at all in order to solve the paradox. By contrast, I just demonstrate that ambiguity aversion can very well be explained by Savage’s postulates of rational choice. The solution is astonishingly simple. It nicely reflects the basic idea of counterfactual reasoning and subjective probability. We will see that Ellsberg’s paradox evaporates as soon as subjective probabilities are no longer treated like objective ones.

2.7.1 Ellsberg’s Thought Experiments

Ellsberg’s two-colors experiment is discussed from p. 650 to p. 653, whereas the three-colors experiment is investigated from p. 653 to p. 655 in Ellsberg (1961). Here, I recapitulate Ellsberg’s thought experiments not only for convenience but also in order to provide a better illustration of the solution to the paradox.

2.7.1.1 The Two-Colors Experiment

Consider two urns, each one containing 100 balls. Each ball can either be red or black. The number of red balls and the number of black balls in Urn I are unknown to the decision maker. There can even be 100 red balls or 100 black balls in Urn I. By contrast, he knows that Urn II contains exactly 50 red and 50 black balls. Now, he can place one of the following bets:

1. Red_I : Red in Urn I
2. Black_I : Black in Urn I
3. Red_{II} : Red in Urn II
4. Black_{II} : Black in Urn II

After placing his bet, the decision maker draws a ball at random from the respective urn. In case he proves correct, he wins \$100 and otherwise he goes away empty-handed. There are no other costs or benefits.

Ellsberg's arguments on the two-colors experiment are only verbal. That is, he does not use any scenario (or decision) matrix. However, I think that his arguments are better understood with a decision matrix, which is given in Table 2.27. The reason why I call each column of the decision matrix a scenario, not a state, has been already explained in Section 1.1. However, this does not alter Ellsberg's line of argument at all. Moreover, the specific choice of the indices, "1, 4, 13, 16," will be clarified in Section 2.7.2.

Table 2.27: Ellsberg's decision matrix of the two-colors experiment.

	Scenario			
	s_1	s_4	s_{13}	s_{16}
Urn I	•	•	•	•
Urn II	•	•	•	•
Red_I	\$100	\$0	\$100	\$0
Black_I	\$0	\$100	\$0	\$100
Red_{II}	\$100	\$100	\$0	\$0
Black_{II}	\$0	\$0	\$100	\$100

The first row of colors represents Urn I and the second row represents Urn II. For example, Scenario s_4 means that the decision maker draws a black ball from Urn I, given that he has placed a bet on Red or Black in Urn I, but a red ball from Urn II, provided that his bet refers either to Red or to Black in Urn II, etc.

Ellsberg (1961, p. 651) conjectures that most people prefer Red_{II} to Red_I and Black_{II} to Black_I , since the number of red (or black) balls in Urn I is ambiguous. Those deci-

sion makers apparently violate Savage's postulate **P2**⁴⁰: Red_I agrees with Black_{II} and Black_I agrees with Red_{II} on $s_4 \cup s_{13}$, whereas Red_I agrees with Red_{II} and Black_I agrees with Black_{II} on $s_1 \cup s_{16}$. **P2** implies that a rational decision maker who prefers Red_{II} to Red_I should also prefer Black_I to Black_{II}. This result seemingly contradicts the idea of ambiguity aversion.

Ellsberg assumes also that most people are indifferent among Red_I and Black_I as well as among Red_{II} and Black_{II}. If this is true, Savage's postulates imply that those people cannot prefer any bet at all, i.e., Red_I ~ Black_I ~ Red_{II} ~ Black_{II}.⁴¹ Ellsberg concludes that ambiguity aversion cannot be expressed by Savage's subjectivistic approach to rational choice.

2.7.1.2 The Three-Colors Experiment

Now, there is only one urn, which contains 90 balls. More precisely, it contains 30 red balls and 60 other balls, each one being either black or yellow. However, the distribution of black and yellow balls is unknown to the decision maker. He can make the following bets:

- I: Red
- II: Black
- III: Red or Yellow
- IV: Black or Yellow

Ellsberg's decision matrix of the three-colors experiment (Ellsberg, 1961, p. 654) is reproduced in Table 2.28. It turns out to be simpler than the decision matrix of the two-colors experiment. This is because now there is only one urn, whereas the two-colors experiment involves two urns. Once again, the particular choice of the indices, "1, 41, 81," will become clear in Section 2.7.2.

Table 2.28: Ellsberg's decision matrix of the three-colors experiment.

Urn	Scenario		
	s_1	s_{41}	s_{81}
	•	•	•
I	\$100	\$0	\$0
II	\$0	\$100	\$0
III	\$100	\$0	\$100
IV	\$0	\$100	\$100

⁴⁰ Ellsberg (1961, p. 649) suggests that ambiguity aversion violates the sure-thing principle (Savage, 1972, Section 2.7), but de facto he refers to **P2**, which is an essential ingredient of that principle.

⁴¹ We can assume without loss of generality that $u(\$0) = 0$ and $u(\$100) = 1$. Then it follows that $P(s_1) + P(s_{13}) = P(s_4) + P(s_{16})$ and $P(s_1) + P(s_4) = P(s_{13}) + P(s_{16})$, which implies that $P(s_4) = P(s_{13})$.

Ellsberg argues that most people prefer I to II and IV to III, and the reason is ambiguity aversion: The decision maker knows that the urn contains 30 red balls, but he does not know the number of black balls. Thus he prefers I to II. Moreover, he knows that the total number of black and yellow balls is 60, but he does not know the total number of red and yellow balls. For this reason, he prefers IV to III. This is a seeming contradiction to **P2**: I agrees with III and II agrees with IV on $s_1 \cup s_{41}$. Further, I agrees with II and III agrees with IV on s_{81} . Hence, if a rational decision maker prefers I to II he should also prefer III to IV.

It seems that decision makers who are ambiguity averse cannot be rational in the sense of Savage (1972), and Ellsberg (1961, p. 655) claims that, at least for them, it is even impossible to infer qualitative probabilities. He comes to the conclusion that their behavior cannot be explained by expected-utility theory. In the next section, I will show that this conclusion is wrong.

2.7.2 The Solution to the Paradox

The solution to Ellsberg's paradox is based on the simple observation that the scenario sets of the decision problems that are (implicitly) used by Ellsberg are not properly specified. The problem is that the decision matrices, in fact, contain (much) more columns, but they have been neglected in his thought experiments. Put another way, the associated probabilities have been implicitly assumed to be zero. We have already seen similar examples in Section 2.6.

This can be best illustrated by using decision trees. Let us start with the two-colors experiment (see Figure 2.21). Counterfactual reasoning tells us that the subjective probabilities for Red and Black may depend on the choice of the decision maker: He believes that Red occurs with probability 30 % if he chooses Red_I , but his probability of Red *changes* to 70 % if he chooses Black_I . By contrast, his probability of Red is 50 % whenever he chooses Red_{II} or Black_{II} .

Actually, even for Urn II the decision maker's *subjective* probability of Red need not coincide with its *objective* probability, which corresponds to 50%. We may still suppose that the decision maker does not believe that his particular choice of color regarding Urn II has any influence on the outcome. Nonetheless, he might very well prefer Red to Black in Urn II, although he knows that the numbers of red and of black balls are equal in that urn. In this case, his subjective probability to draw a red ball from Urn II must be greater than 50%.⁴²

In his thought experiment, Ellsberg (1961, p. 651) precisely describes the situation depicted by the decision tree in Figure 2.21, i.e., we have that

- $\text{Red}_I \sim \text{Black}_I$ and $\text{Red}_{II} \sim \text{Black}_{II}$ but
- $\text{Red}_{II} > \text{Red}_I$ and $\text{Black}_{II} > \text{Black}_I$.

⁴² A similar observation has already been made in Section 1.4.1 regarding Roulette.

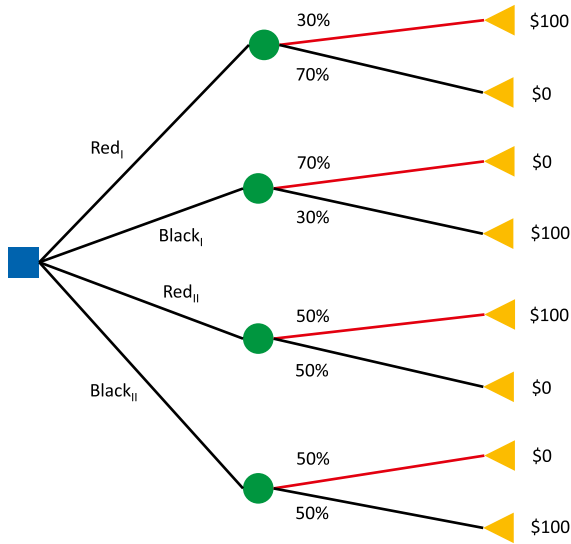


Figure 2.21: Decision tree of Ellsberg's two-colors experiment.

We conclude that the decision maker, in fact, is ambiguity averse. Nonetheless, this phenomenon can be simply explained by means of subjectivistic decision theory. In order to understand this point we must allow for counterfactual reasoning and not treat subjective probabilities like objective ones.

We have already seen in Section 2.3.4.1 that the probabilities in a decision tree may depend on the actions of the decision maker. However, it might seem obscure to the reader that we allow the probability of Red to depend on whether the decision maker chooses Red_I or Black_I . One could think that the decision maker ponders on the possible number of red (or black) balls in Urn I before making his decision. After a while he comes to the conclusion that Urn I contains 30 red and 70 black balls. Of course, if this is actually the way he creates his subjective probabilities, it cannot happen that the probability of Red depends on whether he chooses Red_I or Black_I . In fact, the number of red and the number of black balls are fixed, i.e., they cannot be influenced by the decision maker. Then he should place his bet on Black rather than Red in Urn I.

However, this way of thinking is objectivistic. It is based on a physical model and goes back to the frequentistic interpretation of probability. Simply put, the objectivistic approach suggests that the preferences of a rational subject should be a result of the (objective) probabilities, but the subjectivistic approach goes the other way around: The given probabilities are just a result of the preferences of the decision maker! Since subjectivistic decision theory does not scrutinize the reasoning of the decision maker, it may always happen that his probabilities change with each single action. Subjective probabilities are not bound to physical laws—they simply reflect the decision maker's individual preferences among (Savage) acts.

The probabilities in the decision tree in Figure 2.21 describe the whole idea of ambiguity aversion, namely that the decision maker is pessimistic regarding Urn I. Irrespective of whether he chooses Red or Black, as long as he is betting on a color in Urn I, he believes that the probability of winning is only 30%. This is because he does not know how much red and black balls are in Urn I. Thus, he feels uncomfortable when placing a bet on this urn. By contrast, he knows the physical distribution of red and black balls in Urn II, and so his (subjective) probability of winning amounts to 50 % if he is betting on a color in Urn II. Hence, ambiguity aversion translates into the subjective probabilities of the decision maker. They could even be used in order to *quantify* the decision maker's ambiguity aversion—provided he obeys Savage's postulates of rational choice.

We can apply the same arguments to the three-colors experiment. For example, we could obtain the subjective probabilities in the decision tree in Figure 2.22. Now, we have that $I > II$ but $III < IV$, in accordance with Ellsberg's (1961) hypothesis on p. 654. We conclude that ambiguity aversion can be expressed by the decision tree of the three-colors experiment as well.




Figure 2.22: Decision tree of Ellsberg's three-colors experiment.

Now, how can we see that Savage's postulate **P2** is neither violated in the two-colors nor in the three-colors experiment by referring to decision matrices? For this purpose, we have to count the number of scenarios in each decision problem, i.e., we must take all branches in the decision tree into account (see Section 2.3.3). In the two-colors experiment there are $2^4 = 16$ scenarios, whereas the three-colors experiment contains

$3^4 = 81$ scenarios. The exponent, 4, quantifies the number of actions, whereas the base, 2 or 3, is the number of branches behind each chance node. Each scenario represents a strategy of Nature. More precisely, it is a complete list of (actual and hypothetical) “reactions” of Nature to every action that is available to the decision maker. In this way, we are able to assign each action another *subjective* probability of Red and Black in the two-colors experiment, or Red, Black, and Yellow in the three-colors experiment.

The full scenario matrix of the two-colors experiment is given in Table 2.29. Scenario s_1 means that Red occurs, irrespective of whatever the decision maker does. Scenario s_2 describes that the outcome is Black if the decision maker chooses Red_I , whereas otherwise he obtains Red, etc. An ambiguity averse decision maker considers $s_2 \cup s_6 \cup s_{10} \cup s_{14}$ more probable than $s_3 \cup s_7 \cup s_{11} \cup s_{15}$. The Ellsberg paradox is simply based on the fact that all scenarios except for s_1, s_4, s_{13} , and s_{16} are neglected. Thus, an essential part of the state space is supposed to be null.⁴³ Of course, if we eliminate some scenarios ad libitum, we cannot properly explain the behavior of a decision maker who is rational in the sense of Savage—unless he really *considers* the eliminated scenarios null. The same arguments apply to the three-colors experiment, where Ellsberg’s paradox comes from the fact that all scenarios except for s_1, s_{41} , and s_{81} are neglected (see Table 2.28).

Table 2.29: Full scenario matrix of Ellsberg’s two-colors experiment.

Action	Scenario															
	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8	s_9	s_{10}	s_{11}	s_{12}	s_{13}	s_{14}	s_{15}	s_{16}
Red_I	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Black_I	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Red_{II}	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Black_{II}	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•

Considering ambiguity aversion a natural part of the subjectivistic approach to rational choice might seem odd to the reader. Why should the heterogeneous scenarios, in which the decision maker’s own choice can change the outcome of the experiment, be possible at all? Subjectivistic decision theory is not based on any physical law and it does not presume that the choice of the decision maker *causes* the outcome. It can very well be the other way around, i.e., that the outcome causes the choice of the decision maker! In any case, these considerations are purely subjective, i.e., they take place only in the mind of the given subject and need not make any real sense.

43 After reducing the set of scenarios to $\{s_1, s_4, s_{13}, s_{16}\}$, the scenario matrix in Table 2.29 turns into the upper part of Ellsberg’s decision matrix in Table 2.27.

For example, you saw that some person has placed a coin on the back of his hand. He is covering that coin with the other hand and now you are asked to bet either on heads or tails. Of course, you know that your choice has no influence on the outcome. Nonetheless, you can very well assume that you are an unlucky fellow, which means that you always back the wrong horse. Put another way, you believe that if the outcome is Heads, you will choose Tails and if it is Tails, you will choose Heads. Perhaps the reader remembers that we have already considered this form of extreme pessimism when discussing Wald's maximin rule in Section 2.4.2 and Savage's minimax-regret rule in Section 2.4.3.

2.7.3 Quintessence

Ellsberg vividly demonstrates that, in real-life situations, there can be different degrees of uncertainty, and it seems obvious that people prefer less ambiguous alternatives to more ambiguous ones. However, ambiguity aversion can be readily explained by subjectivistic decision theory.

The solution to Ellsberg's paradox is astonishingly simple and underpins the basic idea of counterfactual reasoning as well as subjectivistic decision theory: Subjective probabilities reflect the individual preferences of a rational subject, and his preferences may very well be affected by ambiguity aversion. The paradox disappears as soon as we begin to analyze the decision problem by counterfactual reasoning and stop treating subjective probabilities like objective ones.

2.8 Backward Induction

2.8.1 Regular Decision Problems

Our assumption that the decision maker is Bayes rational is particularly helpful for solving sequential decision problems. This shall be demonstrated by a simple example: Peter is (Bayes) rational and thinks about what to do on the weekend. He could either go climbing or visit the local swimming pool, in which case he probably meets his girlfriend. Peter will choose one of these options only if it is sunny. By contrast, if it is cloudy, he will stay at home and read a book.

Figure 2.23 contains the decision tree of the weekend problem. We can see that there are three stages:

1. Nature chooses to be sunny or cloudy;
2. Peter decides between swimming or climbing if it is sunny; otherwise he decides to read a book, and
3. Nature reveals Peter's girlfriend or not, provided it has chosen to be sunny.

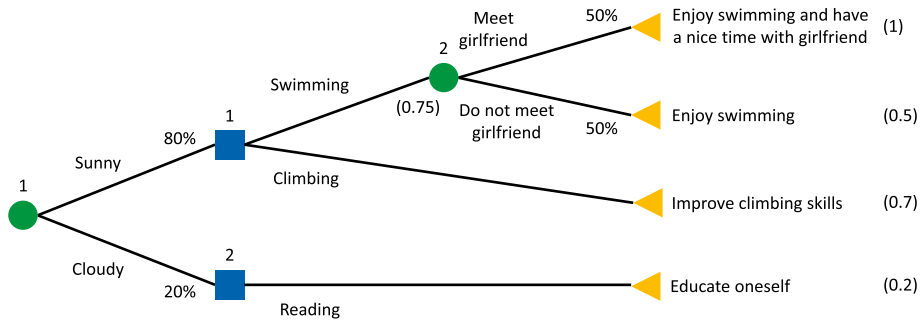


Figure 2.23: Decision tree of the weekend problem.

A strategy is an exhaustive plan of actions and reactions to Nature’s “actions,” i.e., to the events that are contained in the decision tree. Thus, Peter can choose between two strategies:

1. Go swimming if it is sunny and otherwise read a book.
2. Go climbing if it is sunny and otherwise read a book.

A scenario is an exhaustive plan of events, which can be interpreted as Nature’s “actions” and “reactions” to Peter’s actions in the decision tree. We can find three scenarios in Figure 2.23:

1. Be sunny and let Peter meet his girlfriend if he goes swimming.
2. Be sunny and do not let Peter meet his girlfriend if he goes swimming.
3. Just be cloudy.

The weekend problem represents a sequential decision problem. Its possible consequences can be found in the descriptions without the parentheses that are placed behind the end nodes of the decision tree. Each combination of strategy and scenario leads to one and only one consequence. Hence, every strategy can be considered a Savage act, i.e., a function from Ω to C , whereas each scenario is an event, i.e., a subset of Ω .⁴⁴ Theorem 2 guarantees that we can assign each consequence some utility, which is given within the parentheses after the descriptions behind the end nodes of the decision tree.

Obviously, Peter’s best consequence is to enjoy swimming and have a nice time with his girlfriend, whereas he considers educating himself by reading a book worst. Savage’s representation theorem asserts that a rational decision maker has an optimal strategy, i.e., a strategy that maximizes his expected utility.⁴⁵ Peter’s expected utility

⁴⁴ The Savage acts are constant on the given subsets of Ω .

⁴⁵ The set of optimal strategies need not be a singleton.

of Strategy 1 is

$$EU(1) = \underbrace{0.8 \cdot 0.5}_{=0.4} \cdot 1 + \underbrace{0.8 \cdot 0.5}_{=0.4} \cdot 0.5 + 0.2 \cdot 0.2 = 0.64, \quad (2.1)$$

whereas his expected utility of Strategy 2 amounts to

$$EU(2) = 0.8 \cdot 0.7 + 0.2 \cdot 0.2 = 0.6.$$

Thus, Peter prefers Strategy 1.

Note that calculating the expected utility of Strategy 1 by using the decision tree in Figure 2.23 involves *conditional* probabilities. The probability that Peter meets his girlfriend is conditional on the event that it is sunny, and the same holds true for the probability that Peter does not meet his girlfriend if it is sunny. These conditional probabilities can be found in the decision tree after Chance Node 2.

By contrast, the decision matrix of the weekend problem, which is given in Table 2.30, contains the *unconditional* probabilities of the three scenarios mentioned above. Put another way, these are the prior probabilities of Peter. The existence of these subjective probabilities is guaranteed by Theorem 2. Hence, in order to solve the weekend problem by the decision matrix, it suffices to assume that Peter is Savage rational—he need not be Bayes rational. In fact, he even need not be Bellman rational. However, if he violates the Bellman postulate, calculating the prior probabilities by the decision tree might be impossible.

Table 2.30: Decision matrix of the weekend problem.

Strategy	Scenario		
	40 %	40 %	20 %
	1	2	3
1	Enjoy swimming and have a nice time with girlfriend (1)	Enjoy swimming (0.5)	Educate oneself (0.2)
2	Improve climbing skills (0.7)	Improve climbing skills (0.7)	Educate oneself (0.2)

Knowing that it is sunny means to know that $\omega_0 \in s_1 \cup s_2$ and, for this reason, the conditional probabilities that can be found in the decision tree behind Chance Node 2 are given by $0.4/(0.4 + 0.4) = 0.5$. It is worth emphasizing that these probabilities are *not* conditional on Peter's action because his action is not an event in the probabilistic sense (see Section 2.4).

Since Peter is Bayes rational, the conditional probabilities in the decision tree correspond to his *posterior* probabilities. By contrast, if the Bayes postulate were violated, it could happen that Peter has any other subjective probabilities as soon as he arrives at Decision Node 1. If also the Bellman postulate were violated, it could even happen

that Peter does not ignore the event Cloudy at Decision Node 1 although it is *evident* to him that it is not cloudy but sunny. Using the decision tree in order to solve the weekend problem would make no sense at all in this case.

Hence, when arriving at Decision Node 1, Peter thinks that he will meet his girlfriend with probability 50 % if he decides to go swimming. This means that the conditional expected utility of Swimming amounts to $0.5 \cdot 1 + 0.5 \cdot 0.5 = 0.75$, which can be found in parentheses at Chance Node 2. Swimming is better than Climbing, which yields only a utility of 0.7. We conclude that Peter will decide to go swimming if he arrives at Decision Node 1.

Further, the probability that Peter arrives at Decision Node 1 is 80 % and so the (unconditional) expected utility of Strategy 1 can be calculated by

$$EU(1) = 0.8 \cdot \underbrace{(0.5 \cdot 1 + 0.5 \cdot 0.5)}_{= 0.75} + 0.2 \cdot 0.2 = 0.64. \quad (2.2)$$

As we can see, the given result is the same as in Equation 2.1.

By comparing Equation 2.2 with Equation 2.1, and using the distributive law of arithmetic, we observe that Peter prefers Strategy 1 to Strategy 2 right from the start because the conditional expected utility of Swimming, i.e., 0.75, is greater than the conditional expected utility of Climbing, i.e., 0.7. Here, the given condition refers to Decision Node 1, in which Nature has already decided to be sunny. Thus, we can solve the decision problem more efficiently by thinking about what Peter will do (or, at least, prefer) in the future.

Searching for Peter's optimal strategy by Equation 2.1 is referred to as forward deduction, whereas the method used in Equation 2.2 is called backward induction. Forward deduction means that we calculate the expected utility of each available *strategy* by running through the decision tree from left to right. In this case, we can deduce the optimal strategy only after we have calculated the expected utilities of all available strategies. In fact, this is precisely what we are doing when using a decision matrix. Here, it is implicitly assumed that the decision tree contains the conditional probabilities and prior utilities of the subject, which may differ from his posterior probabilities and the corresponding utilities if the decision maker violates the Bayes postulate. This interesting phenomenon will be discussed in Section 2.8.3.

By contrast, backward induction means that we calculate the expected utility of each available *action* by running through the decision tree from right to left and eliminating all suboptimal actions concomitantly. This method enables us to find the optimal strategy without calculating the expected utility of any other strategy. It is clear that this is possible only if we use a decision tree, not a decision matrix, unless the given problem is simple.

In fact, the weekend problem is a very simple example of a sequential decision problem. Nonetheless, I hope that the reader can understand the basic principle without a precise mathematical treatment: We can solve any finite decision problem by backward induction, provided that the decision maker is, at least, Bellman rational.

For this purpose, we start at the end of the decision tree and delete all suboptimal actions at each decision node by moving from right to left. An action is suboptimal if and only if there is another action at the same decision node that provides a greater expected utility. This well-known optimality principle goes back to Bellman (1957), which explains the chosen terminology. Note that the expected utilities at the decision nodes are always calculated by the *posterior* probabilities of the decision maker, which may differ from his *conditional* probabilities if he is not Bayes rational.

However, if the decision maker is Bayes rational, his posterior probabilities, which are given in the decision tree, correspond to the conditional probabilities.⁴⁶ In this case, every strategy that is found to be optimal in the decision tree, i.e., by backward induction, may be considered optimal also in the decision matrix, i.e., by forward deduction. Conversely, every strategy that is considered optimal in the decision matrix turns out to be optimal in the decision tree, too.

We conclude that solving a (sequential) decision problem by the decision tree, on the basis of the conditional but not the posterior probabilities of the subject, makes sense only if the decision maker is Bayes rational. If the Bayes postulate is violated, it can happen that a strategy that is considered optimal by backward induction is suboptimal by forward deduction, and a strategy that is considered optimal by forward deduction might be suboptimal by backward induction. I will come back to this point in Section 2.8.3.

2.8.2 Irregular Decision Problems

A decision problem is said to be irregular if and only if its decision tree contains some null event that leads to a decision node. Obviously, the decision problem discussed in the previous section is regular.

Now, let us assume that Peter is convinced that it will be cloudy on the weekend, i.e., Cloudy has probability 1. Then his decision problem turns out to be irregular and, a priori, it does not make any difference to him whether he decides to go swimming or climbing at Decision Node 1, i.e., both strategies appear to be optimal in the decision matrix. However, is this actually true, *a posteriori*?

Peter is (Bayes) rational and so he is Bellman rational, too. Although he assigns Sunny probability 0, Nature *can* choose to be sunny. Put another way, it is *possible* to be sunny. The Bellman postulate guarantees that Peter will update his subjective probabilities when this null event happens. This means that he will possess some posterior probabilities for the events Meet girlfriend and Do not meet girlfriend.⁴⁷ Moreover,

⁴⁶ Most examples of backward induction that can be found in the literature implicitly presume that the decision maker is Bayes rational.

⁴⁷ In fact, since Peter is Bayes rational, his posterior probabilities can still be considered conditional probabilities, although the event Sunny is null, a priori (see Section 1.5).

since he is Bellman rational, too, he will ignore the event Cloudy when he realizes that it is sunny. Obviously, his posterior probabilities for the subsequent events Meet girlfriend and Do not meet girlfriend must differ from the prior probabilities, which are zero, a priori, but cannot be zero together, a posteriori.

Correspondingly, Peter assigns each consequence that appears after Decision Node 1 some posterior utility. Although he is Bayes rational, his posterior utility function, u_{Sunny} , need not coincide with his prior utility function u .⁴⁸ However, since Peter is Bellman rational, he will consider some action, i.e., Swimming or Climbing, optimal. The posterior probabilities and (expected) posterior utilities are marked red in Figure 2.24. It reveals that Peter prefers to go swimming when he is surprised by the sun. By contrast, he was indifferent among swimming and climbing *before* the weekend started, since he was firmly convinced that the sun would not show up.




Figure 2.24: Decision tree if Peter is a priori convinced that it will be cloudy.

To sum up, we can very well apply Savage's representation theorem at each single decision node of the decision tree. For this purpose, we have to use the posterior probabilities of the decision maker. However, if the decision problem is irregular, it can happen that we are not able to find a *unique* optimal strategy by forward deduction, whereas backward induction leads us to a unique optimal solution. The problem is that a strategy can be optimal beforehand, but it need not be optimal in the course of time if the decision problem is irregular. This holds true even if the decision maker is Bayes rational. However, a Bellman-rational decision maker rules out all strategies that turn out to be suboptimal when he is surprised by the occurrence of an unexpected event.

Whenever we express a decision problem by a decision tree, we can solve it by backward induction. By contrast, if we express the decision problem by a decision

⁴⁸ Recall that the posterior utility function must coincide with the prior utility function only after the decision maker has received some *essential* information (see Section 1.5).

matrix, we ignore the sequential structure of the problem. If the decision problem is irregular, the set of optimal strategies that is obtained by backward induction can be smaller than the set of strategies that are considered optimal by forward deduction. However, it does never exceed the set of strategies that appear to be optimal in the decision matrix. We could say, equivalently, that the posterior set of optimal strategies is a subset of the prior set of optimal strategies.

The decision matrix just represents the situation of the decision maker *a priori*. It fails to consider situations in which he is surprised by some null event, in which case he must revise his subjective probabilities of forthcoming events. By contrast, the decision tree depicts all situations of the decision maker *a posteriori* and enables us to derive his actual preferences at each point in time or, more generally, at each single decision node by backward induction. In general, this gives us a much better impression of the decision maker's actual choices, which is one of the reasons why I highly prefer to use decision trees rather than decision matrices.

However, it is worth emphasizing that solving an irregular decision problem by forward deduction does not lead to wrong conclusions. The only problem is that we could obtain a multitude of optimal strategies, in which case we cannot say which one the decision maker actually performs. Hence, the decision matrix can be ineffective if we want to predict the strategy that a (rational) decision maker is going to perform in the course of time. Of course, in some cases also backward induction can lead to multiple solutions and then we are still left with ambiguity. However, solving a decision problem by backward induction usually leads to a much smaller set of optimal strategies. Ideally, this set is a singleton.

2.8.3 The Time-Inconsistency Paradox

If the Bayes postulate is violated, but the Bellman postulate still holds true, we can solve the decision problem by backward induction. However, then the posterior probabilities of the decision maker need no longer coincide with the conditional probabilities, and his utility function may change at each subsequent decision node. This is referred to as the time-inconsistency paradox (see, e.g., Kydland and Prescott, 1977; Loewenstein and Prelec, 1992; Strotz, 1955).

In the last section, we have already observed that a strategy can be optimal beforehand, but if the decision problem is irregular, the same strategy need not be optimal in the course of time. However, as long as the decision maker is Bayes rational, the posterior set of optimal strategies is a subset of the prior set. This statement no longer holds true if we drop the assumption of Bayesian rationality. Hence, irrespective of whether the decision problem is regular or not, the violation of the Bayes postulate can lead to situations in which a strategy that is considered suboptimal by forward deduction turns out to be optimal by backward induction. The prior and posterior sets of optimal strategies can even be disjoint.

We already know that Peter’s optimal strategy in the weekend problem is to go swimming if it is sunny and otherwise to read a book. This conclusion was based on the decision tree in Figure 2.23, which contains Peter’s *conditional* probabilities behind Chance Node 2. Now, suppose that Peter is not Bayes rational (but still Bellman rational) and consider Figure 2.25. The red numbers in parentheses after Decision Node 1 represent Peter’s posterior probabilities and (expected) utilities. As we can see, these numbers differ from those in Figure 2.23.




Figure 2.25: Decision tree if Peter is not Bayes rational.

Beforehand, Peter plans to go swimming if it is sunny, but he changes his mind at Decision Node 1. Figure 2.25 reveals that Peter’s posterior expected utility of Swimming amounts to $0.4 \cdot 1 + 0.6 \cdot 0.3 = 0.58$, whereas Climbing has a posterior expected utility of 0.6. This means that his individual preferences at this stage are not congruent with those at the beginning. Simply put, he turns into another person as soon as it is sunny. The problem is that the person that we consider before Decision Node 1 pursues different interests. That is, we mix up two different decision makers in one decision tree. It is clear that applying backward induction in such a case can lead to paradoxical results.

The time-inconsistency paradox can be described also by a famous mythological narrative: Odysseus and the Sirens. The Sirens were dangerous creatures trying to lure sailors with their fantastic music and beautiful voices to make them shipwreck on the rocky coast of their island. Figure 2.26 contains the decision tree of an ancient sailor, who must decide either to head for the island or to pass by. Before the sailor hears the Sirens singing, he clearly prefers to pass by in order to save himself and his crew from shipwreck—irrespective of whether the Sirens will sing or not. However, as soon as he captures their music, he is charmed by its beauty. Thus, he decides to head for the island and willingly accepts shipwrecking.

At Decision Node 1 he just turns into another person, whose utility function differs completely from that of the person before, i.e., the person who did not hear the Sirens singing. This is illustrated in Figure 2.26 by the red numbers in parentheses and the resulting strategy is marked bold. Is this behavior irrational? Well, this depends on the




Figure 2.26: Decision tree of an ancient sailor.

perspective! Of course, the sailor *before* Decision Node 1 considers the given strategy suboptimal, but the sailor *at* Decision Node 1 is happy with his decision.

Odysseus was curious about the Sirens' singing. However, he was forewarned by Circe and so he had the following idea: Why not let his crew tie him to the mast and have all of his men plug their ears with beeswax? Moreover, they should leave him tied—no matter how much he would beg and plead. Since Odysseus was aware of the danger that he *will not be himself* at Decision Node 1, we must substitute the corresponding square in Figure 2.26 with a circle and obtain the decision tree in Figure 2.27. Note that whether the Sirens sing or not is no longer important at all for Odysseus, i.e., he is interested only to avoid shipwrecking and not to enjoy their singing.

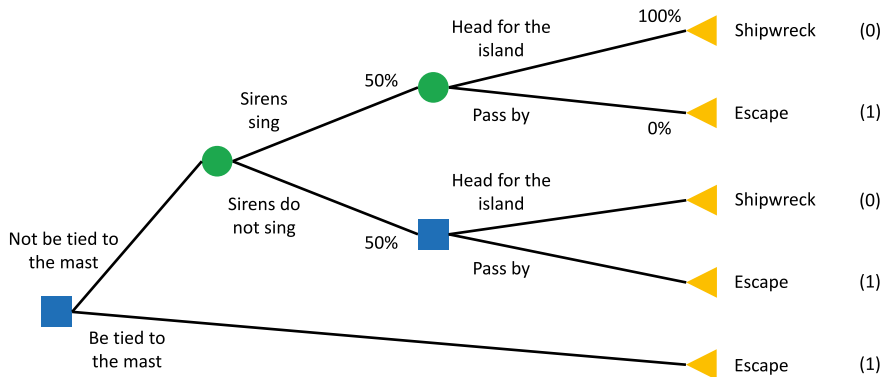


Figure 2.27: Decision tree of Odysseus.

We have solved the paradox by substituting a decision node with a chance node. The new decision tree depicts the situation of Odysseus at the beginning of his journey to the island and does not mix up two different persons or state of minds. Odysseus may choose between the following three strategies:

1. Not be tied to the mast and head for the island if the Sirens do not sing.
2. Not be tied to the mast and pass by if the Sirens do not sing.
3. Be tied to the mast.

Now, both forward deduction and backward induction lead us to the same optimal choice, namely to be tied to the mast. In fact, this is the only optimal strategy because the event that the Sirens will sing is not negligible from Odysseus' point of view. Otherwise, he would have been indifferent among Strategy 2 and Strategy 3. However, he knew from Circe what he was letting himself in for, and the myth tells us that Odysseus chose Strategy 3. Nowadays, that kind of behavior is commonly referred to as self-commitment.

2.9 The Oil-Wildcatting Problem

The subjectivistic solution concept shall now be demonstrated by the so-called oil-wildcatting problem. This sequential decision problem is frequently used in decision theory as a typical example of backward induction (Eisenführ and Weber, 1994; Raiffa, 1968): The rational CEO of an oil company thinks about drilling for oil on a particular spot. He considers three options:

1. Apply a seismic test before entering the venture.
2. Drill without seismic test.
3. Do not drill at all.

If the CEO decides not to drill, the company does not have to pay anything. Without the seismic test, the (subjective) probability of finding oil is 60%. If the company finds oil, its profit amounts to \$3 million. Otherwise, it loses \$1 million because of the drilling costs. Further, if the CEO decides to apply the seismic test, the company has to pay \$300,000 extra. He expects to get a positive test result with 80 % probability. Moreover, he thinks that

- the probability of finding oil is 95 % if the test result is positive, whereas
- it is just 5 % in the case in which the test result is negative.

He may quit the project immediately after a positive or a negative test result.

2.9.1 Decision Tree

The decision tree of the oil-wildcatting problem is given in Figure 2.28, where the (expected) utilities of the CEO can be found in parentheses.

Fortunately, the oil-wildcatting problem is a regular decision problem, which can be solved in a simple way by backward induction:

- At Decision Node 2, the CEO decides to drill because drilling yields an expected utility of

$$0.95 \cdot 0.99 + 0.05 \cdot 0 = 0.9405.$$

By contrast, not drilling gives only a utility of 0.8.

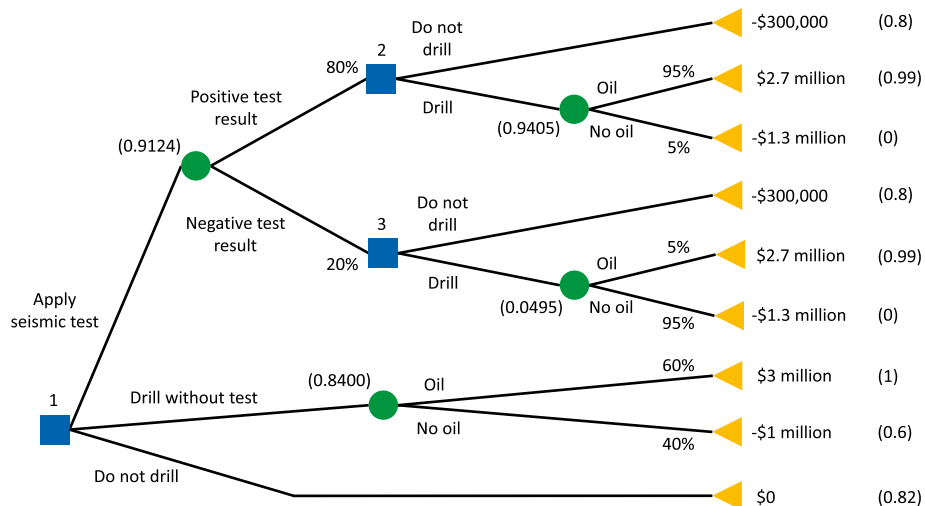


Figure 2.28: Decision tree of the oil-wildcatting problem.

- At Decision Node 3, he decides *not* to drill, which yields an expected utility of 0.8 vs.

$$0.05 \cdot 0.99 + 0.95 \cdot 0 = 0.0495.$$

- At Decision Node 1, the expected utility of applying the seismic test amounts to

$$0.8 \cdot 0.9405 + 0.2 \cdot 0.8 = 0.9124.$$

If the CEO decides to drill without the test, he receives the expected utility

$$0.6 \cdot 1 + 0.4 \cdot 0.6 = 0.84.$$

Finally, the utility of not drilling at all is 0.82. Hence, he decides to apply the seismic test.

We conclude that the decision maker favors the following strategy: Apply the seismic test. Drill if the result is positive and do not drill if it is negative.

2.9.2 Decision Matrix

In order to provide a complete picture, I should mention that the oil-wildcatting problem can be represented, equivalently, by a decision matrix. For this purpose, we have to count the number of strategies and the number of scenarios. In fact, this has already

been done in Section 2.3.3, where the decision tree of the oil-wildcatting problem occurs in stylized form in Figure 2.6.

We come to the conclusion that the CEO considers 6 strategies:

1. Apply the seismic test and do not drill, irrespective of the test result.
2. Apply the seismic test. Drill if the test result is positive and do not drill if it is negative.
3. Apply the seismic test. Do not drill if the test result is positive and drill if it is negative.
4. Apply the seismic test and drill, irrespective of the test result.
5. Drill without seismic test.
6. Do not drill at all.

Moreover, there are 8 scenarios, which can be described like this:

1. Oil and positive test result → oil.
2. Oil and positive test result → no oil.
3. Oil and negative test result → oil.
4. Oil and negative test result → no oil.
5. No oil and positive test result → oil.
6. No oil and positive test result → no oil.
7. No oil and negative test result → oil.
8. No oil and negative test result → no oil.

The corresponding decision matrix is sketched in Table 2.31, which provides the potential consequences and the associated utilities in parentheses.⁴⁹ There we can find also the prior probabilities of the CEO, which can be deduced from the decision tree in Figure 2.28.⁵⁰ Now, we could solve the oil-wildcatting problem in the usual way, i.e., by maximizing the expected utility of each strategy. However, trying to solve such a sequential decision problem by a decision matrix is a tedious task and, to be honest, I cannot see any benefit compared with backward induction.

2.9.3 Final Remarks

The solution procedure based on backward induction is quite nice and simple. It can be applied to much more complicated problems. We need not assume that the CEO actually sits down and draws a decision tree. Moreover, we even do not require that the CEO himself considers the given decision problem sequential, in which case he

⁴⁹ Each row of the decision matrix contains a *strategy*, not an action, and each column represents a *scenario*, neither a single event in the decision tree nor a state of the world, i.e., an element of Ω .

⁵⁰ Here, it is implicitly assumed that the CEO is Bayes rational.

Table 2.31: CEO's decision matrix (amounts in \$1 million).

Strategy	Scenario			
	45.6 %	2.4 %	...	7.6 %
	1	2	...	8
1	−0.3 (0.8)	−0.3 (0.8)	...	−0.3 (0.8)
2	2.7 (0.99)	−1.3 (0)	...	−0.3 (0.8)
⋮	⋮	⋮		⋮
6	0 (0.82)	0 (0.82)	...	0 (0.82)

would have to think ahead what to do if the seismic test leads to a positive result and what to do if it leads to a negative one, etc. That is, we do not scrutinize the way he comes to his conclusions.

The subjective probabilities of the CEO need not coincide with the objective probabilities either. For example, the objective probability of exploiting oil in the particular region, without applying a seismic test, could be 20 % instead of 60%. The fact that the CEO's subjective probability is much higher than the objective probability conveys that he is overconfident. In the subjectivistic framework, this poses no problem at all. We are just trying to explain the decision maker's behavior but not to judge whether or not his preferences are in any sense reasonable.

Hence, the subjectivistic approach *essentially* differs from the objectivistic one. I guess that most objectivists would say that the CEO is irrational because he thinks that the probability of finding oil is 60 %, while the "true" probability is only 20%. The problem is that objective probabilities are not accessible in most real-life situations and in many cases they do not even exist. They require us to propagate a frequentistic model, and each rational subject should agree that the given model is correct. However, then we still have to specify the decision maker's information because a probability, irrespective of whether it is subjective or objective, is always *conditional* on some (trivial or nontrivial) information. Even such a simple example like a coin toss or a dice roll can be considered a Laplace experiment only if we assume that the decision maker's information is trivial. Otherwise, the frequentistic model becomes quite complicated and I doubt that the subject is able to calculate (or compute) the objective probabilities in any practical application.

Thus, it seems helpful to take a descriptive rather than a prescriptive approach in order to explain rational behavior in real-life situations. This is not because we capitulate in the face of the big challenges of the objectivistic approach. The reason is that, at least in my opinion, the objectivistic approach is not able to describe the actual preferences or to predict human behavior in everyday life.

2.10 Conclusion

Procedural aspects of deliberation play no role in the subjectivistic framework. This is in direct contrast to prescriptive decision theories, which are much more restrictive than the subjectivistic approach. The latter requires only the decision maker's preference relation to be consistent with our postulates of rational choice. By contrast, prescriptive approaches assume that the decision maker satisfies a given list of conditions that refer to procedural aspects of deliberation. For example, he should be effective, diligent, objective, and aware, which implies that his individual expectations should be based on the objective probability measure or, at least, on empirical data in order to make use of the law of large numbers. It is well-known that these conditions are hardly met in real life.

Subjectivistic decision theory is normative. Nonetheless, it is descriptive because it accepts the individual preferences of a rational subject just as they are. This means that it does not tell people what they *should* prefer. It just tries to describe and predict the actual behavior of decision makers. Although it assumes that human beings are rational, i.e., that their individual preferences obey some minimal conditions, it comes much closer to the behavioral approach than to the objectivistic approach.

Decision theory makes not much sense without counterfactual reasoning, i.e., without thinking about what would have happened if the decision maker had performed a strategy that is different from the strategy that he has *actually* performed. The subjectivistic approach considers Nature a subject who can “respond” to the decision maker's actions. Hence, the course of events may depend on each available action of the decision maker. However, Nature's “responses” need not make any physical sense—they just take place in the mind of the decision maker.

Counterfactual reasoning enables us to solve several paradoxes of decision theory, like Newcomb's paradox and Ellsberg's paradox. An essential ingredient of our theory of rational choice is Levi's postulate, which states that deliberation crowds out prediction. Simply put, a subject cannot consider his own decisions stochastic. Another, not less important, pillar of subjectivistic decision theory is a proper distinction between knowledge and belief. This will become all the more important when analyzing human behavior in situations of conflict.

In principle, we can solve every (finite) sequential decision problem by a decision matrix. However, a decision tree proves to be much better suited than a decision matrix because the latter ignores the sequential structure of the problem. For this reason, forward deduction can be ineffective when trying to solve an irregular decision problem. Moreover, if the Bayes postulate is violated, solving a sequential decision problem by forward deduction is practically impossible because then, in general, the prior set of optimal strategies does not encompass the posterior set. In this case, backward induction still leads us to the actual strategy, provided that the decision maker is Bellman rational and that we use his posterior, not his conditional, probabilities as well as his posterior utilities in the decision tree.

To sum up, the subjectivistic approach is powerful enough in order to describe and predict the actual behavior of human beings, even if they do not satisfy all postulates of rational choice. Many paradoxes and shortcomings that are typically attributed to subjectivistic decision theory can easily be solved by counterfactual reasoning. The given solutions nicely reflect human behavior that is commonly considered irrational from an objectivistic point of view but turns out to be rational in the subjectivistic framework. I do not claim that the subjectivistic approach is universally valid. Of course, some of us may violate Savage's axioms of rational choice. However, the axiomatic foundation of subjectivistic decision theory is fairly broad. I think that subjective expected utility makes sense in most practical applications of decision theory. Hence, the subjectivistic approach seems to be a good compromise between the objectivistic and the behavioral approach. In the next part of this book, we will see that subjectivistic decision theory can be applied, as well, in order to solve strategic conflicts, which leads us to the area of game theory.

Part II: **Strategic Conflict**

3 The Subjectivistic Approach to Strategic Conflict

3.1 Games against Nature vs. Strategic Games

In the context of decision theory, we are concerned with a rational subject that tries to solve some decision problem. There is no adversary who pursues his own interests. This means that the decision maker is faced only with Nature. By contrast, in a strategic conflict there are, at least, two rational subjects and each one tries to act in a way that is favorable for himself.

Despite the different circumstances, Aumann and Dreze (2009) distinguish between “games against Nature” and “strategic games.” This suggests that the basic principles of rational choice, which have been elaborated in Part I of this book, hold true irrespective of whether we suppose that there is only one single decision maker, faced with Nature, or a number of players competing with each other in a situation of conflict. Hence, according to this view, every decision maker may be considered a player and vice versa.

A game against Nature represents a decision problem, whereas a strategic game is a situation of conflict between two or more decision makers.⁵¹ The unique feature of the subjectivistic approach is that strategic conflicts are not treated differently from decision problems. A game with n players is considered a set of n decision problems. Thus, we can solve a strategic conflict just by solving each single decision problem, separately. This simple and obvious idea substantially differs from traditional game theory (see Chapter 7), but it is reconsidered in epistemic game theory (see Chapter 8).

In this book, I use the term “game” for any kind of situation that involves two or more rational subjects. That is, a game shall be understood as a strategic conflict, not as a decision problem, and I frequently assume that the number of players is $n = 2$. This assumption is made only for the sake of simplicity but without loss of generality. Indeed, the solution concept presented here can easily be extended to games with an arbitrary number of players. This would require us to complicate the notation, but the associated benefits usually do not seem to outweigh the costs. However, there still exist some cases in which a result that holds true for $n = 2$ is no longer applicable for $n > 2$. It can happen also that some matters become essentially more complicated if there are at least 3 players. Whenever I am aware of such a case, I point out and clarify that issue.

3.2 The Subjectivistic Framework

3.2.1 What Is a Rational Solution?

The solution of a game simply describes what the players do. To be more precise, it is just an n -tuple of actions or strategies that are performed by the players. This is meant

⁵¹ In some cases, the term “conflict” can be somewhat misleading, as is shown in Section 4.3.3.

only in the descriptive but not in the prescriptive sense. Hence, I do not claim that the players *should* behave in any specific way. By contrast, I try to explain their *actual* behavior in order to predict the solution of a game.

Definition 10 (Rational solution). A solution is said to be rational if and only if all players choose an optimal action or strategy, i.e., maximize their expected utilities.

Solving a strategic game, or a game against Nature, just means to apply methods of game theory or of decision theory in order to discover its set of rational solutions. In this book I try to explain and to predict the actual solution of a game in which the players can be considered rational in the sense of subjectivistic decision theory.

Each player in a (strategic) game is faced with rational adversaries. He can expect that his opponents will do their very best in order to maximize their *own* expected utilities, i.e., they try to behave in a strategically reasonable way. However, the same holds true also for a (rational) decision maker in a game against Nature. Decision theory is a theory of rational choice. Thus, we can readily apply that theory in order to explain the behavior of rational subjects in a game, too. Indeed, as is pointed out by Kadane and Larkey (1982), it is an astonishing fact that this route has not been taken by traditional game theory. This book tries to contribute filling this gap.

3.2.2 Strategic and Social Interaction

It is often argued that game theory is about strategic or social interaction of rational subjects, whereas decision theory ignores any interaction. This argument is wrong and it fails to mention that a major part of traditional game theory deals with strategic independence—a typical assumption of noncooperative game theory. I would even say that noncooperative game theory is *not* about interaction at all. It typically presumes that the action of each player has no influence on the action of any other or, at least, that the players *believe* that their actions are independent. However, we can imagine situations in which the players are able to adapt their actions to one another. More precisely, they can react to each other, which means that they can *interact*. Here, I do not speak about the action-reaction scheme that can typically be observed in dynamic games, in which the players act one after another. What I mean is interaction in a one-shot game. It is clear that the strategic-independence assumption is heavily violated in this case. To the best of my knowledge, this form of strategic dependence has not yet been considered in game theory.

In game theory we can often observe arguments like “Ann thinks that Bob performs Action *a* because she thinks that he thinks that she performs Action *b* because she thinks that he thinks that she thinks that he performs Action *c*, etc.” That form of “interaction” refers to the individual *belief* of each player, which is typically an-

alyzed in epistemic game theory.⁵² By contrast, interaction in the sense of this book is based on counterfactual reasoning and implies that the players *know* the action of each other, which might substantially depend on their own actions, etc. This form of interaction plays a major role in this book.

If each player in a game or member of a society is rational, why should decision theory not be able to describe strategic or social interaction? The subjectivistic approach does not require us to neglect any form of interaction. A player can very well have an impact on what each other player knows about him. Indeed, by performing his own action, he could reveal information to the others and, based on their evidence, they might be able to react accordingly. In Part I we have seen that the same holds true in a game against Nature, where Nature can be considered “Player 0” without any self-interest, whereas “Player 1” is the decision maker.

Each player aims at maximizing his own expected utility, which can lead to quite complicated interactions that typically do not appear in games against Nature. This distinguishes decision problems from strategic conflicts, i.e., decision theory from game theory. Nonetheless, interaction poses no real challenge for us. We will see that the subjectivistic approach is perfectly able to cope with action, reaction, and interaction. Here, the terms “action,” “reaction,” and “interaction” have no chronological meaning. They just paraphrase the strategic opportunities that are available to the players, which essentially depend on the question of whether or not they *know* each other’s action or strategy.

3.2.3 Cooperative vs. Noncooperative Games

Cooperative game theory presumes that players form coalitions in order to pursue a common interest, whereas noncooperative game theory claims that the players have different goals and thus act out of self-interest. The pathbreaking work of von Neumann and Morgenstern (1944) is a classic in cooperative game theory and Nash (1951) marks a cornerstone of noncooperative game theory. For a nice overview of cooperative and noncooperative game theory see, e.g., Peters (2015).

I make no distinction between cooperative and noncooperative games. This distinction seems unnecessary to me and it can even be misleading. According to Selten (2001), every noncooperative game can lead to cooperation and thus it may turn into a cooperative game, provided that this is in the interest of each player and that the rules of the game allow for cooperation.⁵³ Conversely, we may doubt that cooperation

⁵² For example, Board (2004), Brandenburger (1992) and Geanakoplos (1992) refer to interaction, but it seems that epistemic game theory has dropped this terminology in the meanwhile.

⁵³ Similar conclusions can be found in the pioneering work of Schelling (1980).

would ever take place in a cooperative game if it were *not* in the self-interest of all players who form a coalition—unless some player can be forced to make an unfavorable decision.

Hence, in contrast to the classical paradigm of game theory, I think that cooperation should be viewed as a possible *result* instead of a *prerequisite* (Selten, 2001). Put another way, whether players cooperate or not should follow from the formal description of the game and not be determined from the outset. This has forced me to develop a methodological framework that is general enough to allow for both cooperative and noncooperative behavior. I believe that this (quite ambitious) goal has been accomplished with this book.

It is well-known that cooperation is a desirable solution of many games that are supposed to be noncooperative. A main conclusion of this work is that cooperation is not only a question of the payoffs or the order of moves in a (dynamic) game. We even need not assume that the players are able to make a binding agreement or that they are altruistic. In some games, cooperation occurs if the players trust in each other or if they are equipped with evidence. In any case, cooperation always turns out to be a result of the individual interests of all players. This means that members of a society cooperate only if it is *optimal* for them to cooperate.⁵⁴ Hence, subjectivistic game theory does not presume that cooperative players are altruistic. In fact, whenever cooperation occurs, it is a consequence of the rational behavior of *egoistic* players, whose choices are based on evidence and belief.

Counterfactual reasoning is an essential ingredient of the given approach to rational choice. Without the possibility of a decision maker to “change the world” by his own actions, decision theory would be restricted to quite artificial problems, like Savage’s omelette example (see Section 1.1). We have already seen in Section 2.6 and Section 2.7 how counterfactual reasoning can solve many questions that naturally arise when working with subjective probabilities. Moreover, it enables us to solve games in which the players are able (and willing) to interact. Eventually, this leads us to games that are usually considered “cooperative” in traditional game theory.

I conclude that there is no essential difference between game and decision theory—at least from a subjectivistic point of view. Subjectivistic game theory explains why rational players behave in a way that is typically not predicted by traditional game theory. This shall be demonstrated in the following sections.

3.3 Static and Dynamic Games

3.3.1 Static Games

Game theory distinguishes between static and dynamic games. A static game is often viewed as a strategic conflict in which the players choose their actions simultaneously,

⁵⁴ For a discussion about social choice and cooperation see, e.g., Moulin (1988).

after which the game is over. Thus, static games are often referred to as simultaneous games or one-shot games. For example, Odds and Evens and Rock-Paper-Scissors are static games. However, the terms “simultaneous” and “one shot” can be misleading. In a static game, the players do not always have to make their choices at the same time and we could even allow them to make more than one decision during the game. Actually, what we are trying to say is that their choices are *independent*.⁵⁵ Static games are treated in Chapter 4.

Traditional game theory typically makes use of payoff matrices in order to solve static 2-person games. A so-called payoff matrix usually contains the *utilities* of both players instead of the potential consequences of their actions. Hence, the term “pay-off” must not be understood in its literal sense, which means that the given entries shall not be interpreted as cash flows. However, payoff has become an established term in game theory. In this book, I write utilities in parentheses and consequences without parentheses. Hence, the reader should always be able to see whether a given payoff matrix contains utilities or consequences. In particular, monetary consequences are preceded by “\$” or “-.\$”.

Let us consider the game Odds and Evens, whose payoff matrix is given in Table 3.1. Odds and Evens each must show either one or two fingers, simultaneously. Equivalently, we could assume that Odds and Evens hide their choices from the other, in which case the decisions need not be simultaneous at all. In any case, Odds wins if the sum of fingers is odd, whereas Evens wins if it is even.

Table 3.1: Payoff matrix of Odds and Evens.

Odds	Evens	
	One	Two
One	(-1, 1)	(1, -1)
Two	(1, -1)	(-1, 1)

Subjectivistic game theory considers the two players *separately*. More precisely, each player treats his counterpart like Nature. Table 3.2 contains the composite scenario and decision matrix of Odds on its left-hand side, whereas the composite scenario and decision matrix of Evens can be found on its right-hand side. Odds’ scenarios are indicated by $s_{O1}, s_{O2}, s_{O3}, s_{O4} \subset \Omega_O$, whereas Evens’ scenarios are given by $s_{E1}, s_{E2}, s_{E3}, s_{E4} \subset \Omega_E$. Note that each player has his *own* state space.

It is true that Odds and Evens make their choices independently. However, we need not assume that they *believe* or that they even *know* that the other’s decision is independent. Indeed, in some games Player 1 can very well assume that Player 2 is able

⁵⁵ Perhaps, it would be better to speak about parallel instead of simultaneous games.

Table 3.2: Composite scenario and decision matrix of Odds (left) and Evens (right).

Action	Scenario			
	s_{01}	s_{02}	s_{03}	s_{04}
One	One (-1)	Two (1)	One (-1)	Two (1)
Two	One (1)	One (1)	Two (-1)	Two (-1)

Action	Scenario			
	s_{E1}	s_{E2}	s_{E3}	s_{E4}
One	One (1)	Two (-1)	One (1)	Two (-1)
Two	One (-1)	One (-1)	Two (1)	Two (1)

to make his own action dependent on the action of Player 1 or vice versa.⁵⁶ Nonetheless, in Odds and Evens we can readily justify the assumption that both players think that their actions are independent. Hence, we may focus on their reduced decision matrices, which are provided in Table 3.3.

Table 3.3: Reduced decision matrices of Odds (left) and Evens (right).

Action	Scenario	
	s_{01}	s_{04}
One	(-1)	(1)
Two	(1)	(-1)

Action	Scenario	
	s_{E1}	s_{E4}
One	(1)	(-1)
Two	(-1)	(1)

The scenarios s_{01} and s_{04} in Odds' decision matrix, i.e., on the left-hand side of Table 3.3, represent Evens' actions One and Two. Correspondingly, the scenarios s_{E1} and s_{E4} in Evens' decision matrix, i.e., on the right-hand side of Table 3.3, represent Odds's actions One and Two. If both players think that their actions are independent, we can simply combine both decision matrices in order to obtain the payoff matrix in Table 3.1.

The payoff matrix is a standard tool of traditional game theory. However, we should keep in mind that it is an appropriate instrument only if we assume, at least, that the players consider all heterogeneous scenarios null, which means that they *neglect* those scenarios. Traditional game theory goes even further. It presumes that each player knows that the heterogeneous scenarios are impossible. Hence, the players *ignore* the scenarios s_{02} and s_{03} as well as s_{E2} and s_{E3} .⁵⁷ The subjectivistic approach does not require this implicit assumption and thus it is more general than the objectivistic, i.e., traditional, approach. In principle, we could assume that a player assigns some heterogeneous scenario a positive probability. This will be very useful in the following analysis.

⁵⁶ I will come back to this point in Chapter 4.

⁵⁷ A rational decision maker neglects an event if and only if he assigns it probability 0, whereas he ignores the event if and only if he is informed about the fact that it cannot happen (see Section 1.5).

Even in such a simple game like Odds and Evens, the players have more scenarios than actions. The action set of Evens corresponds to a proper subset of the scenario set of Odds. More precisely, Odds' set of scenarios is $\{s_{O1}, s_{O2}, s_{O3}, s_{O4}\}$ and the actions of Evens are associated with the subset $\{s_{O1}, s_{O4}\}$. Analogously, $\{s_{E1}, s_{E2}, s_{E3}, s_{E4}\}$ is Evens' scenario set and Odds' action set corresponds to the subset $\{s_{E1}, s_{E4}\}$. If and only if the players consider the scenarios $s_{O2}, s_{O3}, s_{E2}, s_{E3}$, i.e., the strategies that are unavailable to the other player, null, we can analyze this game by its payoff matrix. For this purpose, it suffices to assume that the players know the action set of the other and that they know that the game is static, i.e., that their actions are, de facto, independent.⁵⁸

Of course, similar conclusions can be drawn for any n -person game, but then the notion of payoff "matrix" makes no sense. In the case of $n > 2$ we could speak about a payoff tensor. However, as already mentioned in Section 3.1, in most parts of this book I concentrate on 2-person games in order to keep things simple.

3.3.2 Dynamic Games

A dynamic game is a strategic conflict in which the players choose their actions in an alternating way until the game comes to an end. These games are also referred to as "sequential games." Well-known examples are Go and Chess. We can usually think about a chronological order in which the players make their moves one after another. Nonetheless, it is not necessary to require a chronological order. In fact, we want to express only that the choice of one player can *depend* on the choice of the other.⁵⁹ Dynamic games are discussed in Chapter 5.

For example, consider a strategic conflict between Ann and Bob in which Ann's action set is $A = \{\text{Up}, \text{Down}\}$ and Bob's action set is $B = \{\text{Left}, \text{Right}\}$. Each player is asked by a referee to make two moves, i.e., to choose two times an element from his own action set. Ann must make her two choices immediately, but Bob can start with his first choice only. No player is aware of the other's choice(s). The referee takes Bob's (first) move and places Ann's first move after Bob's. That is, he reveals Ann's first move to Bob. Now, Bob has to make his second move, after which the referee places Ann's second move. Then the game is over and the players receive their payoffs according to the rules of the game. We conclude that this dynamic game consists of four moves, where three of them are made, in any chronological order, at the beginning of the game, whereas one is made thereafter.

In traditional game theory, such a game is usually illustrated by a game tree (see Figure 3.1). Bob starts making his first move. Ann makes the next move, but she does

⁵⁸ This implicitly rules out pessimism and optimism (see Section 2.4.2, Section 2.4.3, and Section 2.7).

⁵⁹ If we call static games parallel, we can call dynamic games serial. However, I maintain the usual terminology, i.e., distinguish between static and dynamic games, in order to avoid any confusion.

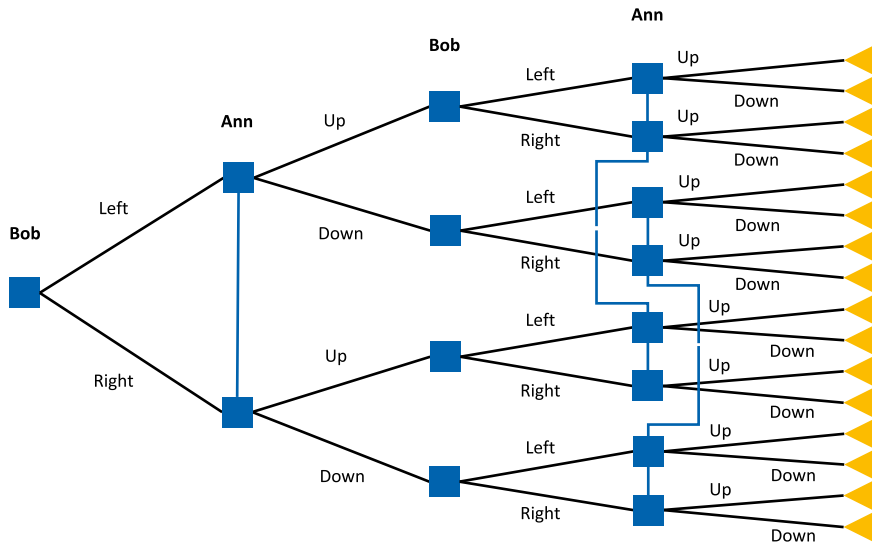


Figure 3.1: Game tree of Ann's and Bob's conflict.

not know Bob's first move, which is indicated by the dashed line between her decision nodes. This line connects all nodes in the game tree that are indistinguishable for Ann. Traditional game theory calls the connected nodes an information "set." This shall not be confounded with the meaning of information set that has been developed in Section 1.5. To be more precise, the connected nodes represent Ann's information at this stage, which is an *element* of her private information set but not the set itself. In any case, Ann cannot vary her action at any node that is connected with other nodes. For example, if Ann goes up if Bob goes left, she must also go up if he goes right.

Now, it is Bob's turn. He knows Ann's first move and so we need not connect any nodes at that stage. Further, Ann knows her first move, but she does not know Bob's first and second move. Hence, in the next step we must connect those nodes that descend from Up and those nodes that descend from Down, where "Up" and "Down" shall be the first move of Ann. Once again, Ann's second choice must be the same at each node that is connected by a dashed line. When she is ready, the game is finished at the end nodes.

Traditional game theory thus makes use of only one (game) tree in order to analyze the strategic conflict of Ann and Bob. Hence, it considers the two players *together*. By contrast, how does subjectivistic game theory deal with Ann's and Bob's situation? It considers each player *separately* and makes use of decision trees instead of a game tree. Once again, this approach is more general than the traditional approach. Moreover, it is more easy to handle. I would like to demonstrate this below.

Ann's decision tree is depicted in Figure 3.2. She does not know Bob's choices and thus her decision nodes come first. The chance nodes behind Ann's decision nodes




Figure 3.2: Ann's decision tree.

represent Bob's decisions. Nonetheless, Bob's first decision may take place even *before* Ann's decisions. Indeed, at the beginning of the game, the referee could have asked Bob first to make his (first) choice. However, in Ann's decision tree, we must place his first move *after* Ann's decision nodes because she is not aware of his actions when making her own choices.

Figure 3.3 contains Bob's decision tree. Bob starts with his first move. Then the referee reveals Ann's first move to him, after which Bob is able to choose his second move. Since he knows Ann's move when making up his mind for the second move, the corresponding decision nodes come after the first chance nodes, which represent Ann's first move. However, Bob still does not know Ann's second moves and thus the next chance nodes are placed at the end of his decision tree.

Once again, the given placement of nodes does not reflect any chronological order, since Ann's decisions in fact may have taken place, on the timeline, before Bob's decisions. In any case, we need not make use of dashed lines or information sets in the decision trees. The placement of nodes in the decision trees precisely reflects the information flow of each player.

Another example, which demonstrates that a decision tree need not reflect any chronological order, is Newcomb's paradox (see Section 2.5.2). The decision tree of the subject is given in Figure 2.17. We may consider the predictor an adversary, who acts *beforehand* by putting either nothing or \$1 million into Box B. Nonetheless, we must place the chance nodes *after* the decision node, since the choice of the omniscient being depends on the decision of the subject.

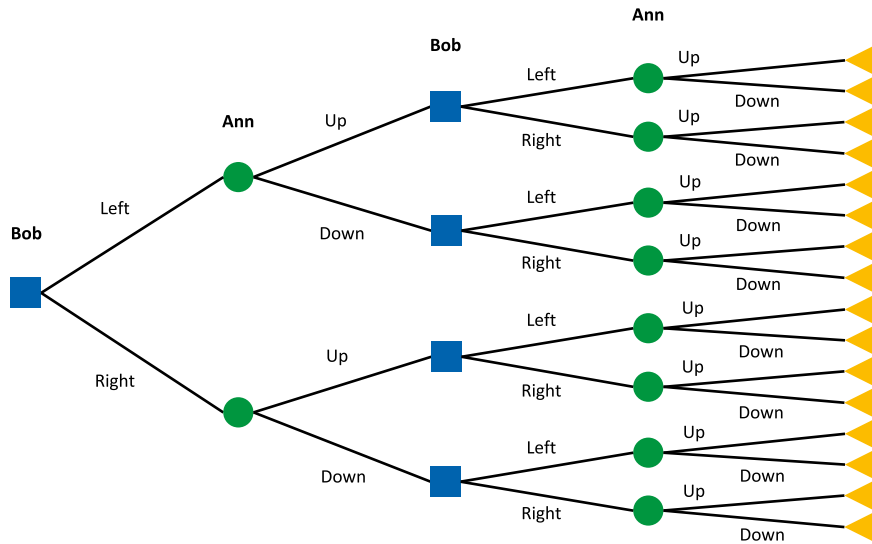


Figure 3.3: Bob's decision tree.

Let us turn back to Ann and Bob. The counting scheme presented in Section 2.3.3 reveals that Ann has 4 strategies and 256 scenarios, whereas Bob has 8 strategies and 64 scenarios. Thus, 8 out of the 256 scenarios of Ann represent Bob's available strategies and 4 out of the 64 scenarios of Bob represent Ann's available strategies.

Under which circumstances does it make sense at all to represent a (2-person) game by a game tree? The answer is analogous to static games: We can use a game tree if and only if both players consider the unavailable strategies of their opponent null! For example, we could assume that each player knows the private information flow of the other. In traditional game theory this is typically justified by the assumption that the players know the rules of the game. Hence, they know which actions the other is able to perform and also which information he receives before or during the game. Put another way, each player knows the available strategies of the other and so it is assumed that the players ignore all unavailable strategies.

Suppose that the referee asks Bob to make also his second move beforehand. That is, both Ann and Bob must make all their choices before the referee places one move after another. Well, in this case the game is *static* because the players have to make their choices independently. It does not matter at all in which chronological order the referee asks the players to make their choices and it does not make any difference that the players have to make two moves instead of one. However, it is obvious that this game is not one shot.

Ann's decision tree does not change at all, but Bob possesses a new decision tree, which can be seen in Figure 3.4. In principle, this decision tree is the same as Ann's. We only have to substitute "Ann" with "Bob," "Up" with "Left," and "Down" with "Right."

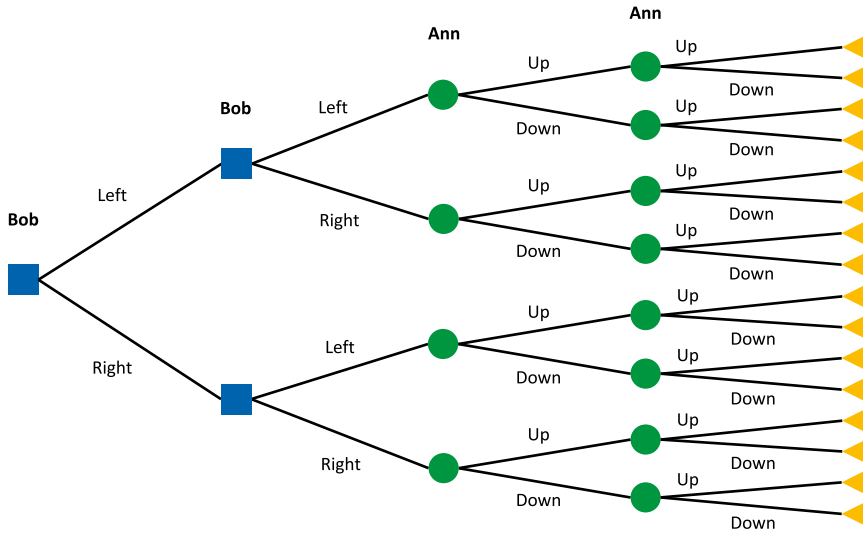


Figure 3.4: Bob's new decision tree.

Now, both Ann and Bob have 4 strategies and 256 scenarios. More precisely, Ann's available strategies are

1. Up and Up;
2. Down and Up;
3. Up and Down;
4. Down and Down.

By contrast, Bob's available strategies read

1. Left and Left;
2. Right and Left;
3. Left and Right;
4. Right and Right.

If both players know that their opponent must make his choice in advance, we may assume that they consider $256 - 4 = 252$ scenarios null. More precisely, Ann and Bob neglect all strategies that are unavailable to the other player. Then we can represent the game by a 4×4 payoff matrix, where each row is an available strategy of Ann and each column is an available strategy of Bob. However, throughout this book, I call the rows and columns of a payoff matrix actions rather than strategies. It does not harm to focus on one-shot games when speaking about static games.

3.4 Coherent Games

A main reason for writing this book was to point out that there exist games that can neither be solved by a payoff matrix nor a game tree. These games are characterized

by the strange fact that the players can act and react to one another at the same time. Hence, each player knows the action of his adversary when making his own choice and so the players are able to interact. I call such games coherent.

Table 3.4: Payoff matrix of a coherent game.

Ann	Bob	
	Left	Right
Up	(2, 1)	(9, 0)
Down	(1, -1)	(9, 8)

Table 3.5: Composite scenario and decision matrix of Ann (left) and of Bob (right).

Action	Scenario			
	s_{A1}	s_{A2}	s_{A3}	s_{A4}
Up	Left (2)	Right (9)	Left (2)	Right (9)
Down	Left (1)	Left (1)	Right (9)	Right (9)

Action	Scenario			
	s_{B1}	s_{B2}	s_{B3}	s_{B4}
Left	Up (1)	Down (-1)	Up (1)	Down (-1)
Right	Up (0)	Up (0)	Down (8)	Down (8)

Consider the payoff matrix in Table 3.4 and assume that Ann and Bob make their choices evident to one another.⁶⁰ The composite scenario and decision matrices of Ann and Bob are given in Table 3.5. Since each player can make his own choice dependent on the choice of the other player, the strategies s_2 and s_3 are no longer negligible or even impossible. For example, suppose that Ann performs Strategy 2. In this case, she chooses Down if Bob chooses Left, whereas she chooses Up if he chooses Right. Hence, Ann has four available strategies:

1. Choose always Up.
2. Choose Down if Bob chooses Left and choose Up otherwise.
3. Choose Up if Bob chooses Left and choose Down otherwise.
4. Choose always Down.

Similarly, the four available strategies of Bob read:

1. Choose always Left.
2. Choose Right if Ann chooses Up and choose Left otherwise.
3. Choose Left if Ann chooses Up and choose Right otherwise.
4. Choose always Right.

⁶⁰ This payoff matrix has some interesting properties and will often appear throughout this book.

Can we analyze such a game by a payoff matrix? Well, the “payoff matrix” in Table 3.6 reveals that this cannot be done because there exist some combinations of strategies that are impossible or that lead to multiple solutions. In these cases we cannot assign a unique consequence or utility to each player.

Table 3.6: “Payoff matrix” of Ann and Bob if their choices are evident to one another.

Ann	Bob			
	1	2	3	4
1	(2, 1)	(9, 0)	(2, 1)	(9, 0)
2	(1, -1)	either (1, -1) or (9, 0)	no solution	(9, 0)
3	(2, 1)	no solution	either (2, 1) or (9, 8)	(9, 8)
4	(1, -1)	(1, -1)	(9, 8)	(9, 8)

A game tree must have a single initial node, i.e., root, and it does not allow any path to get back to the root. Since the players are able to act and react, at the same time, it is not possible to represent their situation by a game tree either. This is illustrated in Figure 3.5. The problem is that Ann’s action can have an influence on Bob’s action, but his action can have an influence on her action, too. This leads to interdependencies that cannot be described by a game tree. In Chapter 6 we will see how to solve coherent games in a relatively simple way by using response diagrams.




Figure 3.5: “Game tree” of Ann and Bob if their choices are evident to one another.

We can use response diagrams also in order to provide a better explanation why the payoff matrix given in Table 3.6 leads to no solution at all or to multiple solutions. For example, if Ann performs Strategy 2 while Bob performs Strategy 3, she will choose Down if he chooses Left and otherwise she will choose Up, whereas he will choose Left if she chooses Up and otherwise he will choose Right. This is illustrated by the

response curves on the left-hand side of Figure 3.6. Response curves shall not be confounded with *best-response curves*, which are often used in traditional game theory in order to solve 2-person games.⁶¹ Best-response curves occur also in the subjectivistic framework when solving *static* games, which will be demonstrated in Chapter 4.




Figure 3.6: The left response diagram contains no possible solution and the right one contains two possible solutions (green points).

The red response curve represents Ann's strategy, whereas the black response curve is Bob's strategy. Correspondingly, the red bullets are the potential responses of Ann and the black bullets are the potential responses of Bob. The solid lines have no particular meaning. They are given only for a better illustration of the strategies.

The two response curves on the left-hand side of Figure 3.6 have no bullet in common and so that choice of strategies, where Ann chooses Strategy 2 and Bob chooses Strategy 3, is impossible. By contrast, the response curves on the right-hand side of Figure 3.6 indicate that both Ann and Bob perform Strategy 2. This means that Ann chooses Down if and only if Bob chooses Left and Bob chooses Left if and only if Ann chooses Down. Now, the given response curves have two points in common, which indicate the possible solutions of this game. The first one can be found on the lower left and leads to the payoffs 1 and -1 . The second one is on the upper right and leads to the payoffs 9 and 0. In order to make a final conclusion, we have to specify the actions of Ann and Bob. If we assume that Ann and Bob are rational, she will choose Up and he will choose Right.

Hence, this work provides a solution concept for strategic conflicts in which the protagonists are able to act, to react, or to interact. It is worth emphasizing that interaction has nothing to do with the typical action-reaction scheme that is depicted by a game tree. Players who are able to interact can choose their actions in a *coherent* way. This phenomenon seems to be widely ignored by traditional game theory. In many situations of strategic conflict, also in games that are usually considered "non-cooperative," interaction leads to cooperation. We will see that this holds true even in

⁶¹ I will come back to this point in Section 7.2.

the prisoners' dilemma, although this serves as a prime example of a noncooperative game in traditional game theory.

3.5 Normal vs. Extensive Form

Traditional game theory distinguishes between the normal and the extensive form of a game. The normal form is also referred to as the strategic form.⁶² In order to represent a (2-person) game in normal form, traditional game theory typically makes use of a payoff matrix or a best-response diagram, whereas a game tree is the tool of choice if the game shall be represented in extensive form.⁶³

Usually, static games are represented in normal form, whereas game theorists prefer to use the extensive form for dynamic games. However, von Neumann and Morgenstern (1953, Section 11) claim that we are able to analyze *every* game in normal form. Hence, the terms “normal-form game” and “extensive-form game,” which can often be observed in the literature, are somewhat misleading. The normal form and the extensive form are only two possible *representations* of a game. A so-called “normal-form game” just represents a static game, whereas an “extensive-form game” refers to a dynamic game.

According to von Neumann and Morgenstern (1953), each row and column of the payoff matrix represents an action or a strategy. If the game is static, we may call it action and if it is dynamic, we can say that it is a strategy. For example, Chess is a dynamic game and thus it is usually represented in extensive form. Nonetheless, it could be represented also in normal form, where each row and column in the payoff matrix is an exhaustive plan of moves and countermoves in the game tree. However, representing Chess in normal form is possible only because the players know the rules of the game and so they know the strategic opportunities of their adversary. Thus, if we follow traditional game theory, we can assume that the players ignore all unavailable strategies of their opponent. Put another way, each scenario that is taken into consideration by a player corresponds to an available strategy of the other player.

The formal expression of a game does not contain any time dimension. This holds true both for the normal form and for the extensive form. Hence, in a purely formal sense, actions, reactions, and their consequences cannot be associated with any point in time. Of course, in a game tree we typically *interpret* the moves (and countermoves) of the players chronologically, and in many games it is obvious or even necessary to refer to time. However, in order to formalize a game, we need no time dimension at all, but it is often convenient to think in chronological terms.

The fact that one player in a dynamic game is able to react to another implies that the former *knows* the action of the latter. Conversely, a player who does not know the

⁶² In their original work, von Neumann and Morgenstern (1953, p. 85) call this form “normalized.”

⁶³ Best-response diagrams are treated in Chapter 4.

action of another cannot be able to react. The extensive form precisely describes what the players know (and what they do not know) at each stage of the game. Thus, we have to clarify only what the players know and what they do not know at each stage of the game but not at which *time* they make their decisions.

Despite the arguments raised by von Neumann and Morgenstern (1953), I think that analyzing a dynamic game in normal form is unnecessary, counterintuitive, and dangerous. Thus, I strongly recommend against using payoff matrices or best-response diagrams in order to solve dynamic (2-person) games. Searching for the rational solutions of a dynamic game in the normal form can be difficult, ineffective, and erroneous because we are not able to specify the private information sets of the players in a proper way. This means that in the normal form it is not possible to describe what the players know and what they do not know at each stage of the game. As a result, serious confusion and fallacies are inevitable.

Instead, I highly prefer to use the extensive form in order to solve dynamic games. However, I do not mean game trees. As already explained above, game trees are quite restrictive because they implicitly presume that the players know the available strategies of their opponents and that they ignore each unavailable strategy. Moreover, since the moves and countermoves of all players are represented in one single tree, this tool can sometimes be cumbersome. By contrast, it is quite easy to solve a dynamic game by decision trees if the game is finite, which means that both the number of players and the number of their available strategies are finite (see Section 2.3.3). In this case, we can simply apply backward induction. In some special cases, it is possible to solve even a dynamic game with an infinite number of strategies by backward induction.

Hence, whenever I focus on the subjectivistic approach, I represent dynamic games in the extensive form by using decision trees, whereas game trees refer to the traditional approach. In subjectivistic game theory, static games are generally solved by decision matrices or by decision trees, depending on whether one prefers the normal or the extensive form. If we assume that the players know (or, at least, believe) that they have no influence on the action of their adversary, a static (2-person) game can easily be solved by a payoff matrix or a best-response diagram, which seems to be more convenient in this special case. Nonetheless, we must be quite careful when applying the dominance principle to a payoff matrix. This will be explained in the next section.

To sum up, a 2-person game in normal form can be solved, in principle, by using a payoff matrix, decision matrices, a best-response diagram, or a response diagram, whereas it can be solved in the extensive form by applying either a game tree or decision trees. I use the term “action” whenever I focus on a simple decision problem, a static, or a coherent (one-shot) game. Without loss of generality, I speak about actions also when discussing a dynamic game in normal form, notwithstanding the fact that such games typically involve strategies but not actions.

By contrast, in case I represent a sequential decision problem or a dynamic game in extensive form, I use the term “strategy.” Nonetheless, I speak about actions when-

ever I refer to the outgoing branches of a decision node, irrespective of whether we consider a sequential decision problem or a dynamic game. In any case, an action represents a trivial strategy in a static or a coherent one-shot game, but an ordinary strategy is an exhaustive plan of moves and countermoves in a dynamic game.

3.6 The Dominance Principle

The dominance principle plays an important role in decision theory and it has already been discussed in Section 2.6. We have seen that a rational decision maker must not consider an action optimal that is dominated, or even superdominated, by another action. However, we must be careful when applying the dominance principle. A naive application to decision problems can lead to serious misunderstandings and the same holds true, all the more, in the context of game theory.

In order to understand this issue, we have to distinguish between weak and strict dominance. Although I argue from the subjectivist's perspective, the reader will see that the following statements are valid also in the objectivistic framework.

3.6.1 Weak Dominance

Consider an n -person game in normal form. Let $a, b \in A$ be two possible actions of a player with some action set A . Action a is said to dominate Action b in the *weak* sense if and only if the following conditions are satisfied:

1. The payoffs of Action a are greater or equal to the payoffs of Action b for all actions of every opponent.
2. The payoff of Action a is greater than the payoff of Action b for some action of an opponent.

For example, consider the following payoff matrix of a 2-person game:

$$\begin{bmatrix} (1, 2) & (3, 0) \\ (1, 3) & (2, 5) \end{bmatrix}.$$

The second row is weakly dominated by the first row.

The reader might wonder why I do not distinguish between weak and strict dominance in Section 2.6. This distinction does not make much sense when dealing with games against Nature: We already know that we (i) *must* neglect scenarios with probability 0 and (ii) *must not* neglect scenarios with positive probability when applying the dominance principle. A (simple) decision problem is typically solved by making use of a decision matrix, which contains the probabilities of each scenario. Hence, we can readily apply the dominance principle by fading out all negligible scenarios, i.e., scenarios with probability 0, and in most practical applications those scenarios have already been eliminated when we start applying the dominance principle.

A payoff matrix contains two decision matrices, i.e., the reduced decision matrix of Player 1 and the reduced decision matrix of Player 2. The fact that we represent a (static 2-person) game by a payoff matrix implies that the players believe that their actions have no influence on the action of the other player. Hence, when using a payoff matrix, we cannot make the mistake of neglecting some essential scenario. However, we are still exposed to the danger of taking negligible scenarios into account, which can very well lead us to erroneous conclusions!

This shall be demonstrated by the 2×2 payoff matrix given above: Should the row player prefer the first row to the second row? We already observed that the first one weakly dominates the second one. However, the answer is “No”! The row player prefers the first row to the second *only* if his subjective probability for the second column is positive. Otherwise, he is indifferent among the first and the second row because both actions lead to the same (expected) utility. Thus, we cannot say that the row player prefers the first row to the second without knowing whether or not he considers the second column null.

To sum up, applying weak dominance in order to eliminate some action requires us to know the zero probabilities of the players.

3.6.2 Strict Dominance

Action a is said to dominate Action b in the *strict* sense if and only if the payoffs of Action a are greater than the payoffs of Action b for all actions of every opponent. It is clear that strict dominance implies weak dominance. For example, consider now the following payoff matrix:

$$\begin{bmatrix} (1, 2) & (3, 0) \\ (0, 3) & (2, 5) \end{bmatrix}.$$

The second row is strictly dominated by the first row.

An action that is strictly dominated should never be preferred to the dominating action. Hence, a strictly dominated action cannot be optimal for a rational player. This holds true irrespective of the subjective probability distribution of the player. It is just a simple consequence of the representation theorem: The utilities of the dominating action are always greater than the utilities of the dominated action. Hence, the expected utility of the former must be greater than that of the latter. However, the dominating action could be strictly dominated by another action or just be suboptimal per se.

Applying strict dominance can very well be used in order to find the set of rational solutions, but in most cases we are not able to eliminate all but one action for each player. Then we must proceed further with other solution concepts. Nonetheless, it is worth noting that the elimination of strictly dominated actions plays a big role in epistemic game theory. This will be discussed in more detail in Chapter 8.

To sum up, applying strict dominance in order to eliminate some action does *not* require us to know the zero probabilities of the players.

3.7 Conclusion

In the subjectivistic framework, the distinction between games against Nature, i.e., decision problems, and strategic games, i.e., situations of conflict between two or more rational subjects, is merely taxonomic. In contrast to traditional game theory, a strategic conflict is not treated differently from a decision problem. A rational subject always tries to perform an optimal action or strategy, irrespective of whether he is faced only with Nature or with rational adversaries. The simple fact that Nature does not have any self-interest, whereas his opponents act strategically, too, does not force us to change our basic methodology.

In contrast to common belief, both strategic interaction and social interaction can very well be described by means of subjectivistic decision theory. Interaction is a phenomenon that occurs if (but not *only* if) the corresponding players know the action of each other. This phenomenon seems to be widely ignored by traditional game theory. Later on, I will show that players are often willing to cooperate—even in games that are typically considered noncooperative—if they are able to interact. In any case, cooperation is a *choice* and it appears only if the players in a game, or members of a society, have an incentive to cooperate.

By contrast, traditional game theory a priori distinguishes between cooperative and noncooperative games. It presumes that players in a cooperative game are able and willing to cooperate, whereas they are not able or willing to cooperate in a noncooperative game. Subjectivistic game theory makes no distinction between cooperative and noncooperative games. Cooperation can and *should* be viewed as a possible result of a game, not as a prerequisite. This enables us to explain cooperation and to create situations in which the players are willing to cooperate.

Traditional game theory distinguishes between the normal and the extensive form of a (noncooperative) game. A dynamic 2-person game in normal form is typically solved by a payoff matrix, whereas its solution is based on a game tree if the game is represented in extensive form. Solving a dynamic game by a payoff matrix can be highly misleading and cumbersome. Even a game tree is not always the best choice because it presumes that the players know the available strategies of one another and ignore all unavailable strategies. Further, we can neither use a payoff matrix nor a game tree to solve coherent games, i.e., games in which the players can interact.

I highly recommend to use decision trees rather than a game tree in order to solve a dynamic game. Of course, this fits in well with our overall subjectivistic approach, which considers each player in a (dynamic) game a rational decision maker. However, decision trees should not be used for a coherent game. The same holds true for pay-

off and decision matrices. Coherent games can be solved by making use of response diagrams, which will be shown later on.

Finally, game theory forces us to distinguish between weak and strict dominance. This is not necessary in decision theory because a decision matrix already contains the subjective or objective probabilities of each scenario. By contrast, in order to solve a (static 2-person) game, we typically make use of a payoff matrix, which is usually not equipped with any probability. However, applying weak dominance requires us to know the zero probabilities of the players. Otherwise, we might come to erroneous conclusions, which cannot happen to us with strict dominance.

4 Action

4.1 Strategic Independence

The basic hypothesis of this chapter is called strategic independence: The action of each player has no influence on the action of the other. Simply put, the players cannot react to one another. In this case, the players cannot interact either. This is a standard assumption of noncooperative game theory.

The name of this chapter is “Action,” which shall indicate that all games that we consider at this point satisfy the strategic-independence assumption, i.e., they are static. By contrast, the following chapters are called “Reaction” and “Interaction.” This means that those games violate the assumption of strategic independence.

Imagine any one-shot game between Ann and Bob. Suppose that Ann can choose only between the actions Up and Down, whereas the action set of Bob consists of Left and Right. In order to understand the meaning of strategic independence, we should consider each player a decision maker, who is faced with his *own* state space. Let Ω_A be the state space of Ann and Ω_B be Bob’s state space. Correspondingly, each player has his own scenario matrix. The scenarios on the left-hand side of Table 4.1 represent the strategies of Bob and the columns that can be found in the scenario matrix on the right-hand side of Table 4.1 are the strategies of Ann. To sum up, $\Omega_A = s_{A1} \cup s_{A2} \cup s_{A3} \cup s_{A4}$ is the state space of Ann and $\Omega_B = s_{B1} \cup s_{B2} \cup s_{B3} \cup s_{B4}$ is the state space of Bob.

Table 4.1: Scenario matrices of Ann (left) and Bob (right).

Action	Scenario			
	s_{A1}	s_{A2}	s_{A3}	s_{A4}
Up	Left	Right	Left	Right
Down	Left	Left	Right	Right

Action	Scenario			
	s_{B1}	s_{B2}	s_{B3}	s_{B4}
Left	Up	Down	Up	Down
Right	Up	Up	Down	Down

Strategic independence requires that Ann’s true state of the world belongs to Scenario s_{A1} or Scenario s_{A4} , whereas Bob’s true state of the world must belong to Scenario s_{B1} or Scenario s_{B4} . This means that Bob’s action cannot depend on Ann’s action and vice versa. Since Ann cannot make her own choice dependent on Bob’s, the scenarios s_{B2} and s_{B3} , which are contained in *Bob’s* scenario matrix (see the right-hand side of Table 4.1), are impossible from the perspective of an outside observer. The same holds true, mutatis mutandis, for *Ann’s* scenario matrix, which means that the scenarios s_{A2} and s_{A3} on the left-hand side of Table 4.1 are impossible, too. We conclude that the strategic-independence assumption is satisfied.

However, this does not mean that Ann or Bob actually know, a priori, that the heterogeneous scenarios are impossible. For this reason, these scenarios can very well belong to their individual state spaces Ω_A and Ω_B . Further, strategic independence does

not imply, e.g., that $P(s_{A2} \cup s_{A3}) = 0$, where P is Ann's prior. Under specific circumstances, she might very well think that Bob's action can depend on hers and the same holds true, *mutatis mutandis*, for Bob. I will come back to this point in Section 4.2.

Strategic independence shall not be confounded with stochastic independence. In the subjectivistic framework, the state space is not common to the players. More precisely, the action of Ann is not a random variable from her own perspective and the same holds true, correspondingly, for Bob. For this reason, their actions cannot have any joint probability distribution and so, in this context, it makes no sense at all to apply the notion of stochastic independence.

In Section 1.5, I already mentioned that each decision maker is equipped with a private information set. Strategic independence results from the fact that both players possess the trivial information, i.e., Ω_A or Ω_B , respectively. This means that the players have no information about the strategy of their opponent. More precisely, Ann is not able to localize the true state of the world ω_0 in Ω_A and so she is not able to say anything about the strategy of Bob. By contrast, if she were equipped with the information set $\mathcal{I}_A = \{s_{A1}, s_{A2}, s_{A3}, s_{A4}\}$, she could very well say, a posteriori, which scenario happens, i.e., she would be informed about Bob's strategy. Of course, the same holds true, *mutatis mutandis*, for Bob. This leads us to the following basic definition:

Definition 11 (Perfect information). A player in a 2-person game is said to have perfect information if and only if he knows, a posteriori, the strategy of the other player.

Thus, a player has perfect information if and only if his private information set, \mathcal{I} , is a subset of $\{s_1, s_2, s_3, s_4\}$.⁶⁴ In this case, the player knows, *ex ante*, the response of his opponent—irrespective of which action he decides to choose.

For example, consider Ann's state space and assume that $\omega_0 \in s_{A2}$. Then Ann knows that Bob decides to go right if she decides to go up and she knows that he goes left if she goes down. By contrast, in the case of $\omega_0 \in s_{A1}$, Ann knows that Bob goes left, irrespective of whatever she does, i.e., his action does not depend on hers, etc. Put another way, Bob decides to go left *categorically*.

In our subjectivistic framework, a rational subject treats each adversary like Nature and Nature like an adversary. Hence, all possible actions of his opponents and all possible events belong to the chance nodes of his decision tree, whereas the decision nodes contain only his own actions. Thus, perfect information means that each chance node in the player's decision tree has only one outgoing branch. Simply put, the player knows the consequence of each available strategy.⁶⁵

However, in this chapter, we assume that the private information sets of Ann and Bob are trivial and so their information is *imperfect*, which is the reason for strategic

⁶⁴ In the case in which it is a *proper* subset of $\{s_1, s_2, s_3, s_4\}$, the player knows, a priori, that some scenarios cannot happen. In traditional game theory it is typically assumed that he ignores s_2 and s_3 .

⁶⁵ In the context of decision theory, we say that the subject is faced with a decision under certainty.

independence. In Chapter 5 we will analyze situations in which one player has perfect information, whereas the other has imperfect information, and Chapter 6 treats situations in which both players have perfect information.

Since Ann's information set is trivial, i.e., $\mathcal{I}_A = \{\Omega_A\}$, her action cannot depend on the true state of the world. More precisely, she cannot make her own action dependent on Bob's strategy. For example, suppose that Ann chooses Up if Scenario s_{A1} happens and Down otherwise. Well, in this case, her choice would determine the location of $\omega_0 \in \Omega_A$: If Ann chooses Up, we have that $\omega_0 \in s_{A1}$, and if she chooses Down, we obtain $\omega_0 \in \Omega_A \setminus s_{A1}$. This situation is illustrated in Table 4.2. However, deliberation crowds out prediction. This means that Ann cannot control the location of ω_0 in Ω_A and so her action must be constant on Ω_A .

Table 4.2: If Ann could make her action dependent on Bob's strategy.

Action	ω_0			
	s_{A1}	s_{A2}	s_{A3}	s_{A4}
Up	Left	Right	Left	Right
Down	Left	Left	Right	Right

Action	ω_0			
	s_{A1}	s_{A2}	s_{A3}	s_{A4}
Up	Left	Right	Left	Right
Down	Left	Left	Right	Right

Action	ω_0			
	s_{A1}	s_{A2}	s_{A3}	s_{A4}
Up	Left	Right	Left	Right
Down	Left	Left	Right	Right

Action	ω_0			
	s_{A1}	s_{A2}	s_{A3}	s_{A4}
Up	Left	Right	Left	Right
Down	Left	Left	Right	Right

By contrast, if Ann's private information set were $\mathcal{I}_A = \{s_{A1}, \Omega \setminus s_{A1}\}$, or even finer, then indeed she could choose Up if $\omega_0 \in s_{A1}$ and Down otherwise. The reason for knowing that the true state of the world, ω_0 , belongs to s_{A1} or $\Omega \setminus s_{A1}$ is not because she can deliberately choose Up if Scenario s_{A1} happens and Down otherwise. It is the other way around: Ann can deliberately choose Up if $\omega_0 \in s_{A1}$ and Down otherwise because she knows that the true state of the world, ω_0 , belongs to s_{A1} or $\Omega \setminus s_{A1}$, respectively! Simply put, decision is based on evidence, but evidence is not based on decision.

4.2 The Subjectivistic Solution Concept

Strategic independence just means that no player can react to the other, and if we assume that the players believe in strategic independence, the subjectivistic solution concept turns out to be particularly simple. This shall be illustrated by some basic examples. Some examples that are discussed in this book can be found in Rieck (2016), who provides a very nice introduction to traditional game theory. However, to a large extent, it seems that they can be considered folklore. Indeed, many of them appear, in different variants, also in other monographs or textbooks on the subject matter.

Throughout this book, I will use the given examples in order to explain the subjectivistic approach but also to demonstrate the differences between subjectivistic and traditional, i.e., objectivistic, game theory.

4.2.1 The Entrepreneurship

Suppose that there is an entrepreneur who aims at using some new technology in an atomistic market. He fears that his competitors will jump on the same bandwagon and so he is thinking about choosing between three levels of engagement: low, moderate, and high.⁶⁶ The number of competitors is large and so we may assume that his action has no influence on the others' actions, which means that the strategic-independence assumption is satisfied.

We may consider the competitors one entity and assume that they can decide between a low, moderate, or a high engagement, too. This leads us to the (full) scenario matrix in Table 4.3. In principle, we could imagine that the competitors decide to engage to a moderate degree if the entrepreneur chooses a low engagement, whereas they choose a low engagement otherwise. This is expressed by Scenario s_2 in Table 4.3. However, the assumption of strategic independence rules out such scenarios. That is, the entrepreneur has no influence on the action of his competitors and thus only the scenarios s_1 , s_{14} , and s_{27} are possible.

Table 4.3: Full scenario matrix of the entrepreneur.

Action	Scenario								
	s_1	s_2	s_3	...	s_{14}	...	s_{25}	s_{26}	s_{27}
Low	Low	Moderate	High	...	Moderate	...	Low	Moderate	High
Moderate	Low	Low	Low	...	Moderate	...	High	High	High
High	Low	Low	Low	...	Moderate	...	High	High	High

Hence, let us assume that the entrepreneur *thinks* that he has no influence on the engagement of his opponents. Thus, we can eliminate all heterogeneous scenarios, i.e., the scenarios that contain different events, and obtain the reduced scenario matrix in Table 4.4. Remember that we have already applied the same principle in the homeowner's problem, which has been discussed in Section 2.3.4.1.

The entrepreneur will have some profit or loss, which depends on his own engagement and the engagement of the competitors. The decision matrix in Table 4.5 contains his potential profits and losses. It refers only to the 3 homogeneous scenarios that are

⁶⁶ A similar game can be found in Rieck (2016, p. 20).

Table 4.4: Reduced scenario matrix of the entrepreneur.

Action	Scenario		
	s_1	s_{14}	s_{27}
Low	Low	Moderate	High
Moderate	Low	Moderate	High
High	Low	Moderate	High

Table 4.5: Decision matrix of the entrepreneur with profits and losses.

Action	Competitors		
	Low	Moderate	High
Low	\$50,000	\$30,000	\$10,000
Moderate	\$90,000	\$50,000	\$20,000
High	\$150,000	\$20,000	–\$20,000

given by the reduced scenario matrix in Table 4.4 and not to all $3^3 = 27$ homogeneous and heterogeneous scenarios that can be found in the full scenario matrix in Table 4.3.

The dominance principle immediately reveals that the entrepreneur prefers a moderate to a low engagement and so we are able to eliminate the first row in the decision matrix. Note that the given dominance is *strict*, which means that we need not know the zero probabilities of the entrepreneur (see Section 3.6). However, the dominance principle is not very effective. In most practical applications it does not lead to a unique solution and so we must dig further. Usually, we cannot avoid to specify the utilities and subjective probabilities of the decision maker, which is done in Table 4.6. We conclude that the entrepreneur considers a moderate engagement optimal.

Table 4.6: Decision matrix of the entrepreneur with (expected) utilities.

Action	Competitors			Expected utility
	20 %	50 %	30 %	
	Low	Moderate	High	
Moderate	(8)	(6)	(4)	(5.8)
High	(10)	(4)	(0)	(4.0)

Now, we could argue that the competitors act in a similar way and conclude that they prefer to engage in the moderate degree, too. Of course, this argument might have influenced also the opinion of the entrepreneur about his competitors. To be more precise, since he chooses a moderate engagement, he could believe that most of the others will do the same, etc. However, his personal considerations have already culminated in the subjective probabilities that are given in Table 4.6. Hence, these

probabilities reflect his state of mind when he comes to his final conclusion. Given that also the others think like that, this sort of behavior might lead to a self-fulfilling prophecy, which means that all competitors choose a moderate level of engagement because they believe that each other prefers to engage in the moderate degree, too.

4.2.2 Rebel Without a Cause

The following example is based on the famous movie “Rebel Without a Cause” with James Dean: Jim and Buzz are driving their (stolen) cars side by side. They quickly approach a seaside cliff and can choose to jump out of the car or keep driving. The first driver who jumps out of the car loses the game. This means that the other one wins the game, but nonetheless he loses his life if he falls into the abyss. If both drivers jump out simultaneously, the game ends in a draw.

This is an example of a continuous-time game, since the actions of the players take place in continuous time $t \in [0, T]$, where $T > 0$ is the time point at which both drivers reach the cliff. Thus, we have an uncountable set of *simultaneous* actions until the first driver jumps out of the car. For choosing his action at time t , a player can use only the information that is available to him *before* t because he needs some (arbitrarily) short period of time in order to react. Hence, the decisions of Jim and Buzz at each time t are always independent. Put another way, they have imperfect information and so the strategic-independence assumption is satisfied. Similar situations are frequently encountered later on in this book.

Consider the payoff matrix in Table 4.7, which represents the situation of Jim and Buzz at any time point $t < T$, i.e., while they are approaching the cliff. We may assume that both players *believe* that nobody can react to the other at the same point in time. Put another way, the players believe that the strategic-independence assumption is satisfied. Obviously, Drive *strictly* dominates Jump. Hence, we may conclude that Jim and Buzz will keep driving while approaching the cliff.

Table 4.7: Payoff matrix while Jim and Buzz are approaching the cliff ($t < T$).

Jim	Buzz	
	Jump	Drive
Jump	(0, 0)	(-1, 1)
Drive	(1, -1)	(0, 0)

Their situation changes dramatically *at* the cliff, i.e., at $t = T$, where a driver inevitably dies if he decides to drive further but stays alive if he jumps out of the car. Table 4.8 contains the given payoffs of Jim and Buzz. Now, Jump (strictly) dominates Drive. Hence, if the players are rational, they will jump out of the car, simultaneously. Well, in the

movie, however, Buzz falls to his death, which could mean that he was tired of life or that he had no good reflexes.

Table 4.8: Payoff matrix when Jim and Buzz are at the cliff ($t = T$).

Jim	Buzz	
	Jump	Drive
Jump	(0, 0)	(-1, -10)
Drive	(-10, -1)	(-10, -10)

Although we have solved the game without knowing the subjective probabilities of Jim and Buzz, we must still assume that both players believe to have no influence on the other's choice. This means that they must believe in strategic independence. The next section illustrates what can go wrong otherwise.

4.2.3 The Game Show

We have seen that the dominance principle must be applied with care in the context of decision theory. Now, we must even be *more* careful because game theory typically deals with payoff rather than decision matrices. A payoff matrix combines the *reduced* decision matrices of two players, for which reason it is sometimes called bimatrix in the literature. By using a payoff matrix, we implicitly presume that the players believe that the strategic-independence assumption is satisfied.

The following example is inspired by Rieck (2016, p. 26): Ann and Bob are candidates in a game show, each one sitting on a table with two buttons. Ann can choose between Up and Down, whereas Bob can choose between Left and Right. Nobody can observe or has any influence on the choice of the other. Depending on the other's choice, Ann and Bob can win the amounts of money that are given in the (monetary!) payoff matrix in Table 4.9. The left component of each tuple represents the payoff of Ann, whereas the right component is the payoff of Bob. We assume that both players know this payoff matrix.

Table 4.9: Prizes of Ann and Bob in the game show.

Ann	Bob	
	Left	Right
Up	\$3,000; \$5,000	\$10,000; \$3,000
Down	\$0; \$0	\$10,000; \$10,000

Ann and Bob have no evidence about the choice of the other and so nobody is able to respond to his opponent. Put another way, the strategic-independence assumption is

satisfied. Thus, we may assume also that Ann *thinks* that her choice has no influence on Bob's choice and the same argument applies to Bob. The reduced decision matrices of Ann and Bob can be found in Table 4.10.

Table 4.10: Reduced decision matrices of Ann and Bob.

Ann	Bob	
	Left	Right
Up	\$3,000	\$10,000
Down	\$0	\$10,000

Bob	Ann	
	Up	Down
Left	\$5,000	\$0
Right	\$3,000	\$10,000

Note that we can use the reduced decision matrices only because we presume that the heterogeneous scenarios, which are contained in the *full* decision matrices but not in the reduced ones, are null. Moreover, the reader should keep in mind that every player has his *own* state space. Ann's strategy is part of Bob's scenario set, whereas Bob's strategy belongs to Ann's scenario set. Ann and Bob can choose their own strategies at their discretion. If the state space, Ω , were common to the players, Ann could make some event happen by choosing a certain strategy. Of course, the same argument holds true for Bob. However, this contradicts Levi's postulate: Deliberation crowds out prediction. Thus, Ann and Bob cannot share the same state space.

We can easily solve this game by means of subjectivistic decision theory. A rational solution of the game is just a tuple of optimal actions. First of all, we could try to apply the dominance principle. We can see that Up dominates Down in the *weak* sense. Does this imply that Ann should choose Up? The answer is "No"! The problem is that we did not yet specify the zero probabilities of Ann, but these probabilities are essential (see Section 3.6):

- If Ann is not convinced that Bob chooses Right, i.e., if the probability of Left is positive, then Up indeed dominates Down, but
- if she is sure that he goes right, which means that the probability of Left is zero, then she is *indifferent* among Up and Down!

Due to this reason, Bob need not assume that Ann goes up, which means that Ann need not assume that Bob goes left, etc. Thus, it could very well happen that Ann chooses Down and Bob chooses Right. However, this is possible only if Ann is *convinced* that Bob goes right. We must keep in mind that the given probabilities are not objective and each player has his own (subjective) probability measure.

Hence, in order to solve the game, we need to apply another method. In the first step, we have to transform the (monetary) payoffs of Ann and Bob into their utilities. The corresponding result is exemplified by the payoff matrix that is given in Ta-

ble 4.11.⁶⁷ Of course, we could have chosen any other utilities, but our choice should somehow reflect Ann's and Bob's individual preferences. In any case, it is worth pointing out that Ann need not know Bob's utilities and vice versa!

Table 4.11: Payoff matrix of the game show.

Ann	Bob	
	Left	Right
Up	(2, 1)	(9, 0)
Down	(1, -1)	(9, 8)

Let p be Bob's subjective probability of Up. That is, he believes that Ann goes up with probability p . Correspondingly, let q be Ann's subjective probability of Left, i.e., she thinks that Bob goes left with probability q . Note that the probabilities p and q are not objective. These probabilities—as well as the associated utilities—just represent the individual preferences of Ann and Bob.

Ann prefers Up to Down if and only if her expected utility of Up is greater than her expected utility of Down:

$$2q + 9(1 - q) > 1q + 9(1 - q),$$

which is equivalent to $q > 0$. This result is simple and intuitive: Ann prefers to go up if she fears that Bob goes left. Indeed, in this case Up dominates Down.

On the contrary, Bob prefers Left to Right if and only if his expected utility of Left is greater than his expected utility of Right:

$$1p + (-1)(1 - p) > 0p + 8(1 - p).$$

This is equivalent to $p > \frac{9}{10}$. Hence, if Bob is pretty sure that Ann goes up, he goes left. Subjectivistic game theory enables us to *quantify* the meaning of “pretty sure.”

We conclude that the rational solution of the game show depends essentially on the subjective probabilities of the players and, of course, on their personal utilities. The question of how and why they come to their own conclusions is not of primary importance. The number of possible reasons is infinite and we could imagine quite different arguments. This is precisely the topic of epistemic game theory (Brandenburger, 2014; de Bruin, 2010; Perea, 2012). In any way, Ann and Bob end up with their subjective probability measures and utility functions. This is sufficient for us in order to solve the game.

⁶⁷ As already discussed in Section 3.3.1, this payoff matrix contains utilities, not monetary payoffs.

The possible solutions of the game show can be illustrated by the best-response diagram in Figure 4.1. Note that its point of origin is on the lower *right*, not on the lower left. By contrast, in most textbooks on game theory, the point of origin can be found on the lower left (see, among many others, Gibbons, 1992, p. 35), which seems more familiar in the first instance. However, placing the point of origin on the lower right has a big advantage: We can immediately associate each corner of the best-response diagram with the corresponding position in the payoff matrix.⁶⁸




Figure 4.1: Best-response curve of Ann (red) and best-response curve of Bob (black). The green point, i.e., Up/Left, is the rational solution if $p > \frac{9}{10}$ and $q > 0$.

Ann's best-response curve is marked red in Figure 4.1. It shows the optimal action of Ann given her probability q that Bob goes left. As we can see, she always prefers to go up unless she is convinced that Bob goes right ($q = 0$). In the latter case, she is indifferent among Up and Down, which is indicated by the red vertical part of her best-response curve at $q = 0$.

Bob's best-response curve is marked black. The optimal action of Bob depends on whether p is greater or lower than $\frac{9}{10}$. In the former case, he prefers to go left, in the latter case he prefers to go right, and in the case of $p = \frac{9}{10}$ he is indifferent among Left and Right. However, at least in my opinion, it seems quite restrictive to assume that p is *exactly* $\frac{9}{10}$.

Hence, except for $p = \frac{9}{10}$ or $q = 0$, we obtain a *unique* rational solution. This is a nice feature of subjectivistic game theory: In most cases, the rational solution of a game turns out to be unique. Note that the rational solution specifies the optimal *actions* of the players. By contrast, the traditional approach to (noncooperative) game theory is based on the idea that the players either choose a random variable (Aumann, 1987) or some probability distribution on their action set (Nash, 1951). Their particular choice is referred to as a (mixed) strategy and a solution represents an n -tuple of strategies, where n is the number of players. The traditional approach leads us to a unique

⁶⁸ I would like to thank an anonymous student, who brought me to that splendid idea.

rational solution only in exceptional cases. That is, it suffers from a multiple-solutions problem (Colman, 2004). I will discuss the traditional approach in much more detail in Chapter 7.

Now, suppose that Ann does not know whether or not Bob can see what she is doing before he makes his own decision. Well, then she might take the possibility into account that Bob *has* evidence about her choice. In this case, he probably goes left if she goes up, whereas he goes right if she goes down. That is, we can no longer use Ann's reduced decision matrix and, due to the same reason, the best-response diagram makes no sense either.

Ann's composite scenario and decision matrix can be found in Table 4.12. We conclude that the payoff matrix of the game show fails entirely to explain rational behavior if Ann is optimistic, i.e., if she hopes that Bob chooses his action to her advantage. In particular, the dominance principle is no longer applicable because $P(s_3) > 0$. Interestingly, even the smallest amount of optimism is sufficient for making the dominance principle inapplicable in this game.

Table 4.12: Ann's composite scenario and decision matrix of the game show.

Action	Scenario			
	5 %	0 %	5 %	90 %
	s_1	s_2	s_3	s_4
Up	Left (2)	Right (9)	Left (2)	Right (9)
Down	Left (1)	Left (1)	Right (9)	Right (9)

Remember that decision makers can also be pessimistic. For example, consider the situations discussed in Section 2.4.2 and Section 2.4.3, in which the decision maker is extremely pessimistic, as well as Ellsberg's paradox in Section 2.7, where we referred to ambiguity aversion. Hope and fear, i.e., optimism and pessimism, play an important role in the subjectivistic approach to rational choice.

The game can still be solved in a simple way by means of subjectivistic decision theory. We can just calculate Ann's expected utilities of Up and Down, i.e.,

$$EU(\text{Up}) = 0.1 \cdot 2 + 0.9 \cdot 9 = 8.3$$

and

$$EU(\text{Down}) = 0.05 \cdot 1 + 0.95 \cdot 9 = 8.6.$$

Hence, Ann's optimal action is to go down. We can use also the decision tree in Figure 4.2 in order to find her optimal choice. Note that the action of Bob is considered an *event* from Ann's perspective.




Figure 4.2: Decision tree of the game show from the perspective of Ann.

Of course, we may drop also our assumption that Bob presumes that Ann is unable to see what he is doing. In this case, we could obtain the decision tree in Figure 4.3. As we can see, Bob prefers to go left. Both players act in an optimal way. Nonetheless, they realize their worst outcome, i.e., Ann and Bob win nothing. This situation might seem annoying. We could say that Ann is too optimistic or that her subjective probabilities are “incorrect,” but this argument is *prescriptive*. The given solution is rational. It just reflects the individual preferences of the players and is immanently *descriptive*.




Figure 4.3: Decision tree of the game show from the perspective of Bob.

The host of the game show can have a substantial impact on the individual preferences and thus on the subjective probabilities of Ann and Bob. For example, suppose that the host tells Ann that Bob can see, beforehand, what she is doing, whereas Bob is told that his action is hidden from Ann. It does not matter at all that both players, in fact, have imperfect information. In this case, the solution derived above seems to be even more plausible, given that Ann lets herself be influenced by the host. These arguments play a crucial role in strategic conflicts, in which the protagonists typically try to affect the state of mind of their opponent.

4.3 Typical 2×2 Games

In this section, the subjectivistic solution concept shall be demonstrated by means of typical 2×2 games, namely anti-coordination, discoordination, and coordination games. These games are very convenient from a didactic point of view. They can always be represented by a payoff matrix of the form

$$P = \begin{bmatrix} (u_{11}, v_{11}) & (u_{12}, v_{12}) \\ (u_{21}, v_{21}) & (u_{22}, v_{22}) \end{bmatrix}. \quad (4.1)$$

For example, consider a game with the payoff matrix

$$\begin{bmatrix} (3, 1) & (-1, 1) \\ (2, 0) & (0, 2) \end{bmatrix}.$$

The particular names of the players and of their available actions do not matter at all for identifying a 2-person game. Hence, transposing the payoff matrix, i.e., making each row a column and each column a row, and interchanging the rows and columns, does not change the game. More generally, two finite games in normal form that can be mapped onto one another by renaming the players and their actions, and by applying some positive affine transformation to their payoffs, are said to be isomorphic (Harsanyi and Selten, 1988, p. 72).

Isomorphic games are always considered identical. For example, a game with the payoff matrix

$$\begin{bmatrix} (-3, 3) & (5, 3) \\ (-1, 6) & (3, 0) \end{bmatrix}$$

is identical to the previous one, since I have just interchanged the columns, i.e., the available actions of Player 2, transformed the payoffs of Player 1 by $u \mapsto 2u - 1$, and transformed the payoffs of Player 2 by $v \mapsto 3v$.

Throughout this section, p denotes the subjective probability of Player 2 that Player 1 chooses Action 1, i.e., the first row of the payoff matrix, and q is the subjective probability of Player 1 that Player 2 chooses Action 1, i.e., the first column. If Player 1 is indifferent among Action 1 and Action 2, we have that

$$q^* u_{11} + (1 - q^*) u_{12} = q^* u_{21} + (1 - q^*) u_{22}.$$

The solution to this equation is

$$q^* = \frac{u_{22} - u_{12}}{u_{22} - u_{12} + u_{11} - u_{21}}, \quad (4.2)$$

provided that $u_{22} - u_{12} + u_{11} - u_{21} \neq 0$. Note that $u_{22} - u_{12} + u_{11} - u_{21} = 0$ is equivalent to $u_{21} - u_{11} = u_{22} - u_{12}$. In this case, either Action 1 strictly dominates Action 2 or vice

versa and so Player 1 cannot be indifferent among those actions. Similarly, for Player 2 we obtain

$$p^* = \frac{v_{22} - v_{21}}{v_{22} - v_{21} + v_{11} - v_{12}}, \quad (4.3)$$

given that $v_{22} - v_{21} + v_{11} - v_{12} \neq 0$.

The simple formulas given by Equation 4.2 and Equation 4.3 will prove to be very helpful when solving 2×2 games. The numbers p^* and q^* are referred to as critical thresholds.⁶⁹ In the case of $p \neq p^*$, Player 2 will prefer either Action 1 or Action 2. Similarly, Player 1 will prefer either Action 1 or Action 2 if $q \neq q^*$. Only if $p = p^*$ or $q = q^*$ some player is indecisive. The critical thresholds enable us to draw a best-response diagram, which can be used to find the rational solution of the game. I am going to demonstrate the procedure in the following sections. However, in the subjectivistic framework, I do not use best-response diagrams to calculate Nash equilibria (Nash, 1951). Although Nash equilibrium is a cornerstone in traditional game theory, it plays no specific role in subjectivistic game theory. Nonetheless, it will be explained and discussed in Section 7.2. In Section 7.2.2, I will show how to deduce the Nash equilibria of a 2-person game by using best-response diagrams.

4.3.1 Anti-Coordination Games

Anti-coordination games are characterized by the fact that

- $u_{21} > u_{11}$ and $u_{12} > u_{22}$ as well as
- $v_{12} > v_{11}$ and $v_{21} > v_{22}$

in the payoff matrix P that can be found in Equation 4.1. For example, consider the payoff matrix

$$\begin{bmatrix} (1, 1) & (0, 2) \\ (2, 0) & (-1, -1) \end{bmatrix}.$$

4.3.1.1 The Chicken Game

Many conflicts in real life represent an anti-coordination game. A well-known example is the chicken game, which is also referred to as the hawk-dove game: Andy and Bob are driving their cars on a collision course. If they go straight, both drivers die, whereas they survive if at least one of them swerves. A driver is called “chicken” if he swerves alone. Table 4.13 contains the payoff matrix of this game.

How will the players act? They must make their choices simultaneously and so we may assume that (they believe that) their actions have no influence on the other’s

⁶⁹ Note that the critical thresholds are invariant to positive affine transformations of the payoffs.

Table 4.13: Payoff matrix of the chicken game.

Andy	Bob	
	Swerve	Straight
Swerve	(1, 1)	(0, 2)
Straight	(2, 0)	(-9, -9)

action. The critical thresholds are given by

$$p^* = q^* = \frac{-9 - 0}{-9 - 0 + 1 - 2} = \frac{9}{10}. \quad (4.4)$$

The solution of that game can be illustrated by using the best-response diagram in Figure 4.4. The best-response curve of Andy is marked red, whereas the best-response curve of Bob is marked black. The break point of Bob's best-response curve equals p^* , i.e., his critical threshold, whereas the break point of Andy's best-response is q^* . We already know that the critical thresholds equal 90 % in the present case. Hence, if both drivers fear that the other goes straight, they swerve. To “fear” that the other goes straight means to possess a subjective probability of Straight greater than 10 %.




Figure 4.4: Best-response curves in the chicken game. Both drivers swerve if they fear that the other goes straight, i.e., if $p, q < 0.9$.

This solution is perfectly understandable and can be observed in everyday life. However, it does not represent a Nash equilibrium (see Section 7.2): In fact, Andy could do better by going straight if Bob swerves and vice versa, but the problem is that no player knows the decision of his opponent when making his own choice! Hence, in real life, the particular decision of each player is based on his personal *belief* about his opponent's behavior.

By contrast, if Andy is confident enough that Bob is going to swerve, he will go straight. More precisely, Andy must believe that Bob swerves with probability greater than 90 %. For example, he could have made his own impression about Bob's willpower before or during the game. In any case, he has some subjective probabilities regarding Bob's potential actions and if he thinks that Bob is a coward, Andy will risk his life and go straight.

This explains very well why players in an anti-coordination game should try, as best as they can, to convince their opponent that they are prepared to risk anything in order to prevail. Any sort of verbal or nonverbal communication that serves this purpose can influence the other's subjective probabilities into the desired direction. It is absolutely plausible (and by far not irrational) that the players show their claws, before the showdown takes place, but get cold feet at the last moment.

The overall situation becomes hazardous only if both drivers are quite sure that their opponent is a coward. In this case, they will die in the crash. Nonetheless, they are still rational. Note that the critical thresholds depend essentially on the consequences that occur if both drivers go straight. For example, if dying leads to a payoff of -99 , not only -9 , whereas all other payoffs remain constant, the critical thresholds are $p^* = q^* = 0.99$. Then the drivers must be *very* confident that the other is a coward in order to end in the disaster.

Should a driver commit himself to go straight, for example, by screwing off his steering wheel? Well, this is a good idea only if the other driver does not commit himself to go straight, too. If it is evident to Andy that Bob has screwed off his wheel, he will certainly not do the same but swerve. By contrast, if nobody can see whether or not the other has screwed off his wheel, this instrument loses its deterrent effect. In this case, the entire game can be readily explained like before.

Self-commitment represents a decision in its own. Suppose that Andy thinks about screwing off and defenestrate his wheel on the collision course. Of course, Bob can do the same, either concomitantly or subsequently. In any case, Andy's decision is independent of Bob's. Now, each driver has to make two decisions. The decision tree of this sequential problem is depicted in Figure 4.5.

If Andy screws off his wheel, he cannot decide to swerve later on, which is the reason why Andy's decision nodes on the upper part of the decision tree have only one outgoing branch, i.e., "Straight." Moreover, when he sees that Bob screwed off his wheel, too, he knows that also Bob must go straight. Hence, we can find only the outgoing branch "Straight" behind the chance node on the upper right of the decision tree as well. If Bob did not screw off his wheel, Andy can be sure that Bob will swerve in order to save his life. By contrast, in the case in which Andy does not screw off his wheel, but observes that Bob did, he knows that Bob cannot swerve anymore, and so we can find only the outgoing branch "Straight" behind the two corresponding chance nodes. The rest of the decision tree just reflects the chicken game without self-commitment.

The decision tree in Figure 4.5 reveals that Andy decides to screw off his wheel. This is because his subjective probability that Bob screws off his wheel, too, is only 10%. Andy considers this probability sufficiently small and so he takes the risk. Of course, this is dangerous. Our final conclusion is not essentially different from the previous one: A smart player should make the other believe that he is prepared to risk anything in order to prevail, i.e., he should try to do his very best to intimidate his opponent. He could also try to make the negative impact of his potential attack as

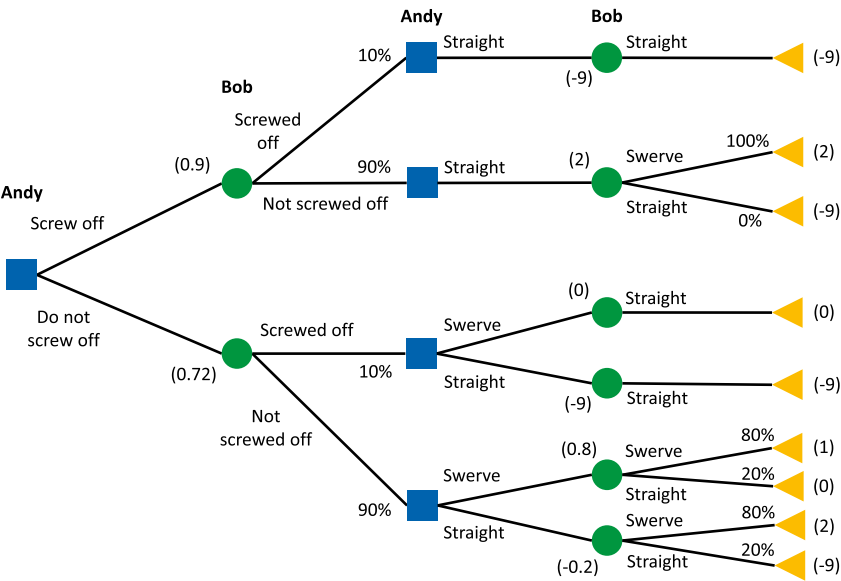


Figure 4.5: Andy's decision tree of the chicken game if the drivers can decide to screw off their steering wheels. He makes use of this instrument of self-commitment.

unpleasant as possible. The better the players succeed, the more we can be sure that they will *not* follow their own words. That is, at the end, both players will swerve. This kind of conclusion is typical for the subjectivistic approach to game theory.

4.3.1.2 Invade or Retreat

The game of chicken can be represented also by the following example: Two countries are in conflict over some region that is claimed by both of them. Each country can decide to invade the region or to withdraw his troops. If both countries decide to withdraw, they will live in peace and harmony. If one country marches forward and the other retreats, the former wins the battle, but if both countries invade the region, their will suffer serious losses. This can be illustrated by the decision matrix in Table 4.14.

Table 4.14: Payoff matrix of Invade or Retreat.

Country A	Country B	
	Retreat	Invade
Retreat	(0, 0)	(−1, 1)
Invade	(1, −1)	(−10, −10)

The overall situation is not essentially different from the chicken game. It is clear that both countries will try to intimidate the hostile country. As long as both countries expect to suffer from a sufficiently high loss if they decide to invade, they will retreat.

However, the situation changes if, e.g., Country A considers the consequence of a confrontation negligible. This is reflected by the payoff matrix in Table 4.15. Suppose that Country A's probability that Country B retreats is positive. In this case, Invade dominates Retreat and so Country A will march in.

Table 4.15: Payoff matrix if the potential loss of Country A is negligible.

Country A	Country B	
	Retreat	Invade
Retreat	(0, 0)	(-1, 1)
Invade	(1, -1)	(-1, -10)

More generally, let $x \in \mathbb{R}$ be the consequence of Country A in case of a confrontation. If we have that

$$q1 + (1 - q)x > q0 + (1 - q)(-1),$$

then Country A prefers to invade the region. More precisely, it marches in if

$$q > 1 + \frac{1}{x}$$

for any $x \leq -1$. In the case of $x > -1$, Country A prefers to invade the region irrespective of q because then Invade strictly dominates Retreat.

This result demonstrates how the critical threshold, q^* , depends on x , i.e., the potential loss of Country A. Thus, in order to create a situation in which each country decides to retreat, i.e., if we want to keep the peace, we must guarantee that the potential impact of a confrontation is sufficiently high. Put another way, the potential threat must not be negligible from the perspective of both countries.

4.3.2 Discoordination Games

Discoordination games are characterized by the fact that

- $u_{11} > u_{21}$ and $u_{22} > u_{12}$ as well as
- $v_{12} > v_{11}$ and $v_{21} > v_{22}$

in the payoff matrix P that is given by Equation 4.1. A typical example is

$$\begin{bmatrix} (1, -1) & (-1, 1) \\ (-1, 1) & (1, -1) \end{bmatrix}.$$

4.3.2.1 Coin Games

A well-known example is Matching Pennies: Ann and Bob each have one penny. They must secretly turn their penny to heads or tails. After revealing their choice, each player either wins the other's penny or loses his own penny. More precisely, if the pennies match, then Ann wins Bob's penny and otherwise Bob wins Ann's penny.

Hence, the set of consequences consists only of two elements, i.e., $-\$0.01$ and $\$0.01$. Due to the representation theorem, we may choose any arbitrary utilities for Ann and Bob in order to express their individual preferences. Of course, it is apparent that the utility of $-\$0.01$ must be lower than the utility of $\$0.01$. Thus, without loss of generality, let us suppose that $u(-\$0.01) = -1$ and $u(\$0.01) = 1$, which leads us to the payoff matrix in Table 4.16.⁷⁰

Table 4.16: Payoff matrix of Matching Pennies.

Ann	Bob	
	Heads	Tails
Heads	(1, -1)	(-1, 1)
Tails	(-1, 1)	(1, -1)

Ann's expected utilities are

$$EU(\text{Heads}) = q1 + (1 - q)(-1) = -1 + 2q$$

and

$$EU(\text{Tails}) = q(-1) + (1 - q)1 = 1 - 2q.$$

Thus, she will choose Heads if $q > 0.5$, i.e., if she believes that Bob chooses Heads with probability greater than 50 %, but if she believes that Bob chooses Tails with probability greater than 50 %, then she prefers Tails. Only if Bob chooses Heads and Tails each with probability 50 %, she is indifferent among Heads and Tails. Similar arguments apply, mutatis mutandis, to Bob.

A possible solution of Matching Pennies can be found by using the best-response diagram in Figure 4.6, whose break points are given by $p^* = q^* = 0.5$. As we can see, Ann believes that Bob tends to Heads ($q > 0.5$) and Bob thinks that Ann tends to Tails ($p < 0.5$). Hence, the rational solution is Heads/Heads, which is marked green in Figure 4.6. We could imagine also any other solution, depending on whether the subjective probabilities of Ann and Bob are either greater or lower than 0.5.

70 Matching Pennies and Odds and Evens (see Section 3.3.1) are isomorphic and thus identical.




Figure 4.6: Best-response diagram of Matching Pennies.

Except for the quite unlikely case in which Ann or Bob consider Heads and Tails equally probable, there exists one and only one rational solution, which consists of the optimal *actions* of both players. Even if a player is indecisive, he does not randomize his action. When he is indifferent among Heads and Tails, he still chooses either Heads or Tails. As an outside observer, we are not able to say which of the two actions he is going to perform. However, this does not mean that anybody makes his choice by throwing a dice or applying any other random generator. This means that the player's decision is still deterministic. I will come back to this crucial point in Section 7.2.4.

Suppose that this game is played repeatedly. It is often argued that, on the long run, the relative number of Heads (and of Tails) must converge to 50 %. For example, if Ann plays Heads more often than Tails, then Bob could improve his chance of success by choosing Tails more often than Heads, but if Bob plays Tails more often than Heads, then Ann could improve her chance of success by playing Tails more often than Heads, etc. Only when both players choose Heads and Tails equally often, i.e., with *objective* probability 50 %, we have reached a (Nash) equilibrium. That is, both players have no reason to change their behavior.

The latter implies that the players use empirical observations in order to manage their decisions. This line of argument is frequentistic and assumes that Ann's and Bob's expected utilities are calculated on the basis of the objective probability measure. The hypothesis that the relative number of Heads (and of Tails) converges to 50 % could be justified by the assumption that the subjective probabilities of Ann and Bob correspond to the objective ones. However, this assumption seems to be quite artificial. The subjectivistic approach is based on the idea that subjectivistic and objectivistic probabilities need not coincide. Moreover, it does not require the law of large numbers. It just tries to explain the rational solution of a *particular* (round of the) game, irrespective of whether it is a one-shot or a repeated game.

Here is another example, which is named after the French mathematician Pierre Rémond de Montmort: A father asks his son to guess in which hand behind his back he hides a gold coin. If the son says "Left" and the coin is, in fact, in the left hand of his father, he wins the coin. If the son says "Right" and this is correct, he wins even two coins. Otherwise, the son goes away empty-handed. The payoff matrix of this game can be found in Table 4.17.


Table 4.17: Payoff matrix of the de-Montmort game.

Father	Son	
	Left	Right
Left	$(-1, 1)$	$(0, 0)$
Right	$(0, 0)$	$(-2, 2)$

As we can see, this is a discoordination game, whose critical thresholds are given by

$$p^* = q^* = \frac{2 - 0}{2 - 0 + 1 - 0} = \frac{2}{3}.$$

This leads us to the best-response diagram in Figure 4.7.

**Figure 4.7:** Best-response diagram of the de-Montmort game.

If the father believes that his son chooses the left hand with a probability greater than $\frac{2}{3}$, then he chooses the right hand. This probability is significantly greater than 50 %, which can be explained by the fact that the father's potential loss of Right is twice the potential loss of Left. If the son thinks that his father chooses Left with a probability greater than $\frac{2}{3}$, then he chooses the left hand. Due to the same reasons, also this probability is significantly greater than 50 %. The rational solution is marked green in the best-response diagram: The father decides to hide the coin in his left hand, whereas the son guesses Right, which means that he goes away empty-handed. Of course, we could have found also any other result of this game—depending on the subjective probabilities of father and son.

4.3.2.2 The Cat-And-Mouse Game

Now, consider the following game of cat and mouse: Andy has a crush on Betty, but Betty does not requite Andy's love. Moreover, she already has a lover: Bob. Betty is going to be either in Joey's Sports Bar or in the Crazy Club this evening. Bob is working as a bartender in Joey's Sports Bar and Betty would like to meet him, but she fears that Andy follows her. If this happens, the overall situation becomes quite inconvenient for

the threesome. Thus, in order to avoid any discussion with Bob, Andy prefers to meet Betty in the club, not in the bar.

This is a discoordination game because Andy (the “cat”) tries to catch Betty (the “mouse”), whereas she tries to escape from him. It can be explained by the payoff matrix in Table 4.18. The payoffs are motivated like this: If Andy goes to the bar and Betty to the club, nobody has any benefit nor damage. By contrast, if Andy meets Betty in the bar, she feels unpleasant because of meeting him. This costs her 1 util, which is compensated by being together with her boyfriend, i.e., Bob.⁷¹ Nonetheless, she loses 2 additional utils because of the inconvenience that arises if Andy meets Bob. By contrast, Andy wins 1 util because of meeting Betty in the bar.

Table 4.18: Payoff matrix of the cat-and-mouse game.

Andy	Betty	
	Bar	Club
Bar	(1, -2)	(0, 0)
Club	(0, 1)	(2, -1)

Further, if Andy goes to the club and Betty to the bar, he goes away empty-handed. By contrast, Betty is lucky because of meeting her boyfriend alone, which means that she wins 1 util. Finally, if Andy meets Betty in the club, he wins 2 utils instead of 1 because in this case he need not confront with Bob. However, Betty loses 1 util, since she feels unpleasant in the presence of Andy.

The critical threshold of Betty amounts to

$$p^* = \frac{-1 - 1}{-1 - 1 + (-2) - 0} = \frac{1}{2},$$

whereas Andy’s critical threshold is

$$q^* = \frac{2 - 0}{2 - 0 + 1 - 0} = \frac{2}{3}.$$

The best-response diagram of this game can be found in Figure 4.8.

Thus, Andy will go to the bar if he expects that Betty goes to the same place with probability greater than $\frac{2}{3}$. By contrast, if Andy thinks that Betty will be in the club with probability greater than $\frac{1}{3}$, he goes to the club. He is indifferent among Joey’s Sports Bar and the Crazy Club only in the singular case of $q = \frac{2}{3}$. Further, Betty will go to the bar if she thinks that Andy will go to the club with probability greater than $\frac{1}{2}$, whereas she goes to the club if she believes that he goes to the bar with probability greater than $\frac{1}{2}$. Finally, only in the singular case of $p = \frac{1}{2}$, Betty is indifferent among

⁷¹ A “util” is a unit of utility.




Figure 4.8: Best-response diagram of the cat-and-mouse game.

the bar and the club. Figure 4.8 depicts a situation in which we have that $p > \frac{1}{2}$ and $q > \frac{2}{3}$. That is, Andy goes to the bar and Betty to the club. This rational solution is represented by the green point in Figure 4.8.

4.3.2.3 Penalty Shoot-Out

Now, consider a penalty shoot-out. The goalkeeper tries to catch the ball, whereas the scorer aims at scoring the goal. For the sake of simplicity, but without loss of generality, we may assume that there are only two possible actions: Shooting the ball into the right or into the left corner, where “Right” and “Left” shall be understood from the perspective of the scorer. Further, let us assume that both players must act instantly, i.e., as soon as the referee gives the whistle signal. This assumption is quite artificial and it will be dropped in Section 5.3. However, for the time being, we maintain this highly simplifying assumption in order to guarantee that the strategic-independence assumption is satisfied. The payoff matrix of this game is given in Table 4.19.

Table 4.19: Payoff matrix of the penalty shoot-out.

Keeper	Scorer	
	Left	Right
Left	(1, -1)	(-1, 1)
Right	(-1, 1)	(1, -1)

In a discoordination game, the players typically try to deceive their opponent. For example, the goalkeeper could try to make the scorer believe that he jumps to the left corner. In reality, however, he plans to jump to the right one. However, if the goalkeeper fears that the scorer sees straight through his game, he should choose the left corner instead, etc. Similar arguments apply to the scorer. This kind of manipulation takes place before the referee starts the signal.

There are plenty of reasons why one player should tend to believe that the other chooses Right, but we can find sufficient arguments for the opposite opinion, too. Such

arguments are explored in great detail in epistemic game theory.⁷² For my part, I fear that thinking about what the goalkeeper thinks about the scorer, what he thinks that the other thinks about himself, and so on, is a road to nowhere. The problem is that we have no idea about the state of mind of any player and the particular way how a rational subject comes to his conclusions can be quite simple, very complicated, or something in between. The players themselves suffer from the same ambiguity and they have not much time to make their decisions.

In terms of behavioral decision theory, we could say that the players must think fast (Kahneman, 2011). This means that they must decide in an emotional, unconscious, or instinctive way. On the contrary, if they would have (much) more time, they could think slowly, i.e., in a logical, conscious, or intuitive way. In the end, however, they *must* come to some final conclusion. Subjectivistic game theory solves this problem by representing their conclusions in the form of (subjective) probabilities, which express their individual degrees of belief, without scrutinizing their particular arguments. It can even happen that they have no arguments at all.

In any case, a rational goalkeeper will jump into the corner that he assigns the higher probability. Correspondingly, a rational scorer will kick the ball into the corner that he assigns the lower probability. It is just as simple as that.

4.3.3 Coordination Games

Coordination games are characterized by the fact that

- $u_{11} > u_{21}$ and $u_{22} > u_{12}$ as well as
- $v_{11} > v_{12}$ and $v_{22} > v_{21}$

in the payoff matrix P of Equation 4.1. They have gained great popularity by the contributions of Thomas Schelling (1980). In its most simple form, a coordination game can be represented by the stylized payoff matrix

$$\begin{bmatrix} (1, 1) & (0, 0) \\ (0, 0) & (1, 1) \end{bmatrix}.$$

4.3.3.1 The Reunion Game

Suppose that Ann and Bob are visiting a city and have lost sight of one another in a crowd. Now, they have no possibility to communicate. Let us assume for the sake of simplicity that there are only two points at which they can meet again: The meeting point at the main station and the foyer of the city museum. The payoff matrix of this game is given in Table 4.20. Its subjectivistic solution is astonishingly simple: Ann will go to the main station if she believes that Bob chooses the same spot with a probability

⁷² This relatively new branch of game theory will be discussed in Chapter 8.

greater than 50 %, i.e., $q > 0.5$, whereas she prefers the museum if $q < 0.5$. In the case of $q = 0.5$ she is indecisive and we cannot say whether she chooses the main station or the museum. The same holds true for Bob.

Table 4.20: Payoff matrix of the reunion game.

Ann	Bob	
	Station	Museum
Station	(1, 1)	(0, 0)
Museum	(0, 0)	(1, 1)

Meeting at an official meeting point, provided that we are strangers somewhere, is firmly established in our culture. Hence, we can expect that Ann and Bob tend to believe that the other decides to go to the main station. The main station represents a focal point (Schelling, 1980). On page 57, Schelling writes: “Finding the key, or rather finding a key—any key that is mutually recognized as the key becomes *the* key—may depend on imagination more than on logic.” Well, this statement is precisely in the spirit of the subjectivistic approach to game theory.

As Schelling points out, convention leads to coordination and vice versa. In the most extreme case, the players are even *convinced* that the other makes the obvious choice and then it happens that they are right. This phenomenon represents a nice case study in order to distinguish between belief and knowledge.

Suppose, for example, that Ann is convinced that Bob goes to the main station, i.e., we have that $q = 1$, and Bob *de facto* decides to go there. Did she know what Bob is going to do? Of course, the answer is “No”! Her conviction does not depend on his actual choice: If Bob had gone to the museum, then she would have still believed that he is approaching the main station. That is, she has no evidence.

Alternatively, consider the game of numbers: Ann and Bob are asked to choose, independently, a number from $\{1, 2, \dots, 9, 42\}$. If their choices match, they win \$100, but otherwise they do not earn anything. There are only two consequences, i.e., \$0 and \$100. Thus, we have that $u(\$0) = 0$ and $u(\$100) = 1$, which means that the expected utility of any number, which represents an action in this game, corresponds to the probability that the other player chooses the same number. The number 42 can be considered a focal point. This holds true because 42 stands out from the given set of numbers, not only because “42” is also The Answer to the Ultimate Question of Life, the Universe, and Everything (Adams, 1979).

Hence, we can expect that Ann’s and Bob’s subjective probabilities on the support $\{1, 2, \dots, 9, 42\}$ are maximal for the element 42. Put another way, 42 is the mode of their subjective probability distributions. Thus, Ann and Bob decide to choose 42. Note that we need not require common belief. This means that Ann and Bob need not believe that the other chooses 42, that the other believes that the other chooses 42 and so

on, ad infinitum. It suffices that both players think that the other tends to choose the same number, which means only that the modes of their probability distributions must coincide. Hence, coordination can simply be characterized in probabilistic terms and the underlying epistemic conditions are, in fact, minimal.

4.3.3.2 The Win-Win Situation

Another form of a coordination game is the win-win situation, which can be explained like this: Some time ago, Ann and Bob planned to visit a restaurant this evening. They had their favorite restaurant, namely Trattoria Palermo, and another somewhat more mediocre place, called The Butcher, at their disposal. Both players find the trattoria better than the butcher, but Ann likes the trattoria even more than Bob does. This game is reflected by the payoff matrix in Table 4.21.

Table 4.21: Payoff matrix of the win-win situation.

Ann	Bob	
	Palermo	Butcher
Palermo	(3, 2)	(0, 0)
Butcher	(0, 0)	(1, 1)

The problem is that Ann and Bob did not agree on some restaurant and now there is absolutely no possibility to tell one another where to go. The trattoria represents a focal point and so the players believe that the other tends to visit their favorite restaurant. How large must their degrees of belief, i.e., their subjective probabilities of “Palermo,” be in order to get to their favorite place? The critical thresholds are

$$p^* = \frac{1 - 0}{1 - 0 + 2 - 0} = \frac{1}{3} \quad \text{and} \quad q^* = \frac{1 - 0}{1 - 0 + 3 - 0} = \frac{1}{4},$$

which can be seen also in the best-response diagram in Figure 4.9.

Hence, if Ann believes that Bob decides to visit their favorite place with a probability greater than $\frac{1}{4}$ and Bob thinks that Ann goes to the trattoria with a probability greater than $\frac{1}{3}$, they will realize the first-best solution. Note that the critical thresholds of Ann and Bob depend on their preferences: If $x > 1$ is the payoff of a player when meeting the other at their favorite place, the individual threshold just amounts to $(x + 1)^{-1}$.⁷³ According to Figure 4.9, Ann and Bob are sufficiently confident that the other goes to the trattoria and so they realize the first-best solution, i.e., Palermo/Palermo.

⁷³ For Ann we have that $x = 3$ and in the case of Bob it holds that $x = 2$.




Figure 4.9: Best-response diagram of the win-win situation.

4.3.3.3 Battle of the Sexes

The next coordination game is typically referred to as the battle of the sexes. The overall situation is similar to the win-win situation: At breakfast, Ann and Bob discussed about visiting either a concert or a football match after work. Bob was in a hurry and, unfortunately, they have forgotten to fix their appointment. Now, there is no way to take up contact to the other. Ann would like to go to the concert, whereas Bob wants to see the football match. They enjoy to be together and so the worst case we can imagine is that Ann goes to the football match, whereas Bob visits the concert. This situation is represented by the payoff matrix in Table 4.22.

Table 4.22: Payoff matrix of the battle of the sexes.

Ann	Bob	
	Concert	Football
Concert	(3, 2)	(1, 1)
Football	(0, 0)	(2, 3)

This game is a little bit different from the win-win situation. In the win-win situation both players share the same opinion regarding the favorable place, but here they have different opinions about the place to be. Hence, although this is a coordination game, the players pursue opposing interests and so the battle of the sexes represents a strategic conflict. Its critical thresholds are

$$p^* = \frac{3 - 0}{3 - 0 + 2 - 1} = \frac{3}{4} \quad \text{and} \quad q^* = \frac{2 - 1}{2 - 1 + 3 - 0} = \frac{1}{4},$$

which lead to the best-response diagram in Figure 4.10.

Ann is quite sure that Bob goes to the football match. Her subjective probability that Bob goes to the concert, q , is lower than 25 % and so she prefers to go to the match. Bob thinks that Ann goes either to the concert or to the match with equal probability, i.e., $p = 0.5$. Hence, he decides to go to the match. There can be many reasons why Ann and Bob think so. For example, Ann knows that Bob is a football maniac and Bob




Figure 4.10: Best-response diagram of the battle of the sexes.

is unsure about how much Ann really prefers to be at the concert. In any case, their choices are rational.

Suppose that Ann and Bob each have a cell phone and so, in principle, they are able to communicate. However, Ann knows that Bob can be very persuasive and fears that she could be overpowered. She is a smart person and, to avoid some unpleasant conversation, Ann writes Bob a short message: “I will be at the concert!” Immediately after seeing that he received her message, she shuts down her cell phone.⁷⁴ She expects that he will be angry and try to contact her by phone. Since she has shut down her phone, he has no chance to do so and even if he answers her short message, he will not get any electronic receipt, which confirms that she even did not receive, let alone read, his message.

Now, due to her short message, Ann expects that Bob decides to come to the concert with probability 80 %. In fact, her message has made its impact and thus Bob is quite sure that Ann is going to the concert. Let us say that his probability that this happens is 95 %. This leads us to the best-response diagram in Figure 4.11. We conclude that both Ann and Bob go the concert.




Figure 4.11: Best-response diagram of the battle of the sexes with a short message.

Hence, although the strategic-independence assumption is still satisfied, the players may very well be able to control the subjective probabilities of their opponent and thus

⁷⁴ At least in Germany, it is possible to request a confirmation of receipt when writing a short message.

to influence the solution of the game. This is exactly what we can observe in everyday life.

4.3.3.4 The Stag Hunt

Ann and Bob go out on a hunt. Each one can choose to hunt either a stag or a hare and their choices are independent. If one player decides to hunt the stag, the other player must rush to help him or her in order to be successful. In this case they share their quarry. By contrast, in order to hunt a hare, a player need not the help of the other. The stag hunt is called also Assurance Game or Trust Dilemma. It is a prototypical example of social behavior and cooperation.

Table 4.23 contains the payoffs of the game. The question is whether a player will risk to go for the stag, in which case he can earn all or nothing, or to hunt a hare just to be safe. The critical thresholds are given by

$$p^* = q^* = \frac{1 - 0}{1 - 0 + 2 - 1} = \frac{1}{2}.$$

Table 4.23: Payoff matrix of the stag hunt.

Ann	Bob	
	Stag	Hare
Stag	(2, 2)	(0, 1)
Hare	(1, 0)	(1, 1)

In Figure 4.12 we can see the best-response curves of that game. We conclude that both players decide to hunt the stag if they believe that the other will follow him or her with probability greater than 50 %.




Figure 4.12: Best-response diagram of the stag hunt.

The critical thresholds essentially depend on the payoffs of the players and it can happen that one player trusts in the other, whereas the latter is sceptical about the former.

If one player is (too) sceptical, he will defect. Interestingly, a player who trusts in a sceptical adversary will lose the game because he was trustful, whereas the other was sceptical, which might explain the alternative name of this game, i.e., Trust Dilemma. Hence, it does not suffice to trust in people. We should rather guarantee that they trust in us! This underpins the important role of trust in social societies. To be more precise, it is important to guarantee that most members of a society are confident with the system, since it can break down if a critical number of people lose their trust.

4.4 Zero-Sum Games

A zero-sum game is a game in which the payoffs of the players sum up to zero in each cell of the payoff matrix.

Many parlor games are zero sum, e.g., Chess, Poker, Monopoly, etc. However, we can find many other examples in real life. Indeed, most sports bets and competitions can be considered zero sum. Another typical example is trading on the stock market. However, in the following I concentrate on 2-person zero-sum games.

4.4.1 Rock-Paper-Scissors

Ann and Bob play Rock-Paper-Scissors: Each one is asked to form either Rock, i.e., a closed fist, Paper, i.e., a flat hand, or Scissors, i.e., a “V” with his fingers. The players must make their choices simultaneously, which means that the strategic-independence assumption is satisfied. There are three possible outcomes of the game: Either Ann wins or Bob wins or the game ends in a draw. This can be illustrated by the payoff matrix in Table 4.24.

Table 4.24: Payoff matrix of Rock-Paper-Scissors.

Ann	Bob		
	Rock	Paper	Scissors
Rock	(0, 0)	(−1, 1)	(1, −1)
Paper	(1, −1)	(0, 0)	(−1, 1)
Scissors	(−1, 1)	(1, −1)	(0, 0)

This is a typical zero-sum game because the utilities of the players sum up to zero at each position in the payoff matrix. We have already seen many other zero-sum games in which the players had only two possible choices.⁷⁵ If the action set of a player con-

⁷⁵ For example, the coin games in Section 4.3.2.1 fall into this category.

tains more than two elements, the game can no longer be solved by using a best-response diagram. However, the subjectivistic solution of Rock-Paper-Scissors is totally unspectacular: If Ann thinks that Bob tends to Rock, she will choose Paper, if she believes that he tends to Paper, she will choose Scissors, and if she thinks that he tends to Scissors, she prefers Rock. Simply put, the mode of her subjective probability distribution on {Rock, Paper, Scissors} uniquely determines her optimal choice.

Even if there are two modes, her optimal choice is unique. For example, consider the decision tree in Figure 4.13. Ann believes that Bob chooses Rock and Paper each with probability 40 %. Hence, choosing Rock or Scissors is suboptimal. By choosing Rock, she would win the game only with probability 20 %, but she would lose it with probability 40 %. By choosing Scissors, she would win or lose it each with probability 40 %. Finally, by choosing Paper, she wins the game with probability 40 %, but she loses it only with probability 20 %. Thus, her optimal choice is Paper.

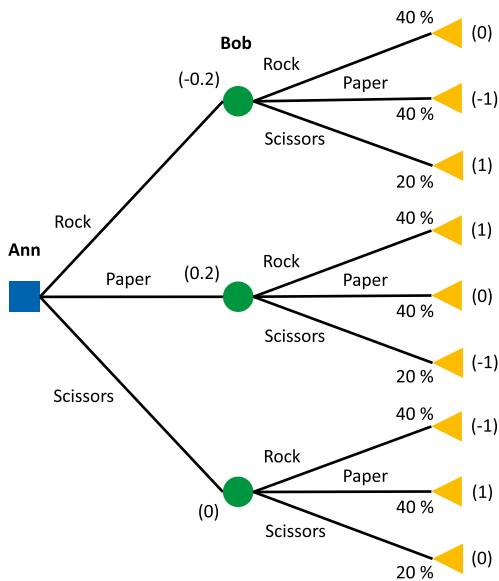


Figure 4.13: Ann's decision tree of Rock-Paper-Scissors.

Ann is indecisive only if her subjective probability distribution has three modes, i.e., if she believes that Bob chooses Rock, Paper, and Scissors each with probability $\frac{1}{3}$. In this singular case, we cannot predict her action. In fact, each one is optimal. The same arguments hold true for Bob. That is, we can neither predict Bob's action only if his subjective probabilities of Rock, Paper, and Scissors are equal.

Is it plausible to assume that Ann's and Bob's subjective probabilities are *precisely* $\frac{1}{3}$? Well, this question cannot be answered without prejudice and I guess that most objectivists would say "Yes." However, we have seen many examples in which the subjective probabilities of the players depend very much on behavioural aspects

and environmental factors. In my opinion, assuming that all subjective probabilities in Rock-Paper-Scissors are precisely $\frac{1}{3}$ is just the same as propagating the hypothesis that Earth is a perfect sphere.

Thus, at least to me, it seems quite artificial to argue that Ann's and Bob's subjective probability distributions are, or even *must be*, uniform. Nonetheless, I admit that the uniform-distribution hypothesis is appealing from a mathematical point of view. In Rock-Paper-Scissors it is typically justified by von Neumann's celebrated minimax theorem, which marks the starting point of traditional game theory (von Neumann, 1928). The minimax theorem will be discussed in great detail in Section 7.5.

Before finishing this section, I would like to emphasize an important issue that has already been discussed in Section 2.4.3: Indeed, the uniform-distribution hypothesis tells us that the players are indecisive. However, irrespective of whether the players are indecisive or not, i.e., whether their subjective probability distributions have three modes or less, they do not randomize their actions! Applying a mixed strategy makes no sense at all when we are indifferent among two or more alternatives. We can just realize the same (expected) utility by choosing some optimal action, deliberately. I will come back to this point in Section 7.2.4.

4.4.2 The Mallorca Game

I found the following game in Perea (2012, Section 2.1) and present it here in a slightly modified version: Ann and Bob each plan to open a pub on a street on Mallorca, which is about 300 meters long and heavily populated by tourists. The connoisseur might have already recognized that I am alluding to the so-called "Schinkenstraße," a very popular place of German tourists on Mallorca.⁷⁶

Suppose that there are only seven places on the street that are worth taking into consideration, and each one is at equal distance to the next. Hence, we have 50 meters between all neighboring locations. Further, we may assume for the sake of simplicity that Ann and Bob are the first ones who open a pub on that street and that the tourists are lazy. That is, when they become thirsty, they choose the nearest pub. Moreover, the distribution of tourists on the street shall be uniform.

For example, if Ann places her pub at the second location and Bob decides to take the position in the center of the street, she accumulates only $\frac{1}{3}$ of all customers, whereas he attracts $\frac{2}{3}$ of them. This situation is illustrated in Figure 4.14. In the case in which Ann and Bob choose the same place for their pub, we may assume that each gets $\frac{1}{2}$ of the customers. Finally, let us suppose that the utility functions of Ann and Bob are linear, so that we have a zero-sum game.

⁷⁶ Interestingly, Perea (2012, p. xvii) mentions that the idea for writing his book came to him during his Christmas holidays on Mallorca, but he has chosen this example for quite different reasons. In fact, it is based on Hotelling's industrial model for spatial competition (Hotelling, 1929).




Figure 4.14: Where to locate the pub on Mallorca?

Ann and Bob make their choices independent of one another. Ann’s reduced decision matrix is given in Table 4.25, which contains her market shares—depending on Bob’s place. What is Ann’s optimal choice? It depends on her subjective probability distribution on $\{a, b, \dots, g\}$. Nonetheless, we can see that b strictly dominates a and f strictly dominates g . Hence, in any case, she will not choose a or g .

Table 4.25: Ann’s reduced decision matrix of the Mallorca game.

Ann	Bob						
	a	b	c	d	e	f	g
a	$\frac{6}{12}$	$\frac{1}{12}$	$\frac{2}{12}$	$\frac{3}{12}$	$\frac{4}{12}$	$\frac{5}{12}$	$\frac{6}{12}$
b	$\frac{11}{12}$	$\frac{6}{12}$	$\frac{3}{12}$	$\frac{4}{12}$	$\frac{5}{12}$	$\frac{6}{12}$	$\frac{7}{12}$
c	$\frac{10}{12}$	$\frac{9}{12}$	$\frac{6}{12}$	$\frac{5}{12}$	$\frac{6}{12}$	$\frac{7}{12}$	$\frac{8}{12}$
d	$\frac{9}{12}$	$\frac{8}{12}$	$\frac{7}{12}$	$\frac{6}{12}$	$\frac{7}{12}$	$\frac{8}{12}$	$\frac{9}{12}$
e	$\frac{8}{12}$	$\frac{7}{12}$	$\frac{6}{12}$	$\frac{5}{12}$	$\frac{6}{12}$	$\frac{9}{12}$	$\frac{10}{12}$
f	$\frac{7}{12}$	$\frac{6}{12}$	$\frac{5}{12}$	$\frac{4}{12}$	$\frac{3}{12}$	$\frac{6}{12}$	$\frac{11}{12}$
g	$\frac{6}{12}$	$\frac{5}{12}$	$\frac{4}{12}$	$\frac{3}{12}$	$\frac{2}{12}$	$\frac{1}{12}$	$\frac{6}{12}$

Now, let us assume that $q(d) = 1$, i.e., that she believes that Bob chooses the place in the center. Then she will choose Place d , too. By contrast, if she believes that Bob decides to place his pub at b , she will choose Place c , etc. Of course, we could also assume that Ann’s personal probabilities are mixed. In this case, we must calculate her expected utility in order to specify her optimal choice. The same arguments apply, mutatis mutandis, to Bob.

In epistemic game theory, it is typically assumed that the players have common belief in rationality and the suggested solution of this game is (d, d) (Perea, 2012, p. 71). I will explain that solution—which occurs after an iterated elimination of strictly dominated actions—in Section 8.2.1.1. However, subjectivistic game theory is not based on the common-belief assumption and so it allows us to accept a much wider range of possible solutions. We could even assume that Ann doubts that Bob is rational at all

and vice versa. In any case, the optimal choice of each player is always a result of his subjective probability distribution on the set of all possible actions.

4.5 The Prisoners' Dilemma

4.5.1 The Classic Dilemma

In order to understand the aforementioned arguments, we can use also the well-known prisoners' dilemma: Two gangsters have committed a crime and have been arrested. Apart from illegal possession of arms, there is no evidence against them. Now, each prisoner can either deny or confess. If both prisoners deny, they serve only one year in prison. If one player denies and the other confesses, the latter is set free as a principal witness, whereas the former is sentenced to five years. Otherwise, if both players confess, each of them serves four years in prison.

The number of years that each prisoner must spend in jail is given by the penalty matrix in Table 4.26. Mary is *not* aware of the action of Joe, and the same holds true for Joe. This means that nobody knows what the other is going to do. We may assume also that each prisoner believes that his action has no influence on the other's action. Thus, we can apply the dominance principle: Confess dominates Deny—even in the strict sense. Hence, if the prisoners are rational, they should confess and so we obtain the standard result of the prisoners' dilemma, i.e., the noncooperative solution.

Table 4.26: Penalty matrix of the prisoners' dilemma.

Mary	Joe	
	Deny	Confess
Deny	(1, 1)	(5, 0)
Confess	(0, 5)	(4, 4)

Now, suppose that Mary and Joe meet in the courtyard before they make their decisions. Joe tells Mary that he will certainly deny and he tries to convince her that she should do the same. He pulls out all the stops to make her believe that he is trustworthy. Does Mary *know* that he will deny? Put another way, is it possible that he will confess although he assures that he is going to deny? Of course, it is possible! This means that she does not know his action and the same arguments apply, *mutatis mutandis*, to Joe. We conclude that communication does not solve the basic issue, namely that Mary and Joe are not equipped with evidence (see Section 2.5.3).

The main problem of the prisoners' dilemma is that the players believe that their own action has no influence on the other's action. That is, they believe that the

strategic-independence assumption is satisfied. In this case, it cannot be optimal to deny—irrespective of whether the other player denies or confesses.

Now, assume that Mary tells Joe that she has a very good friend at court who will tell her Joe's action after he has made his choice. Hence, allegedly, Mary can make her own choice dependent on the given information about Joe. Mary assures that she will deny if Joe denies but confess if he confesses. Hence, she tells him that she will play tit for tat.⁷⁷ Her story is so convincing that Joe comes to the conclusion that the strategic-independence assumption is violated. Joe is not completely sure about Mary's intentions. However, her statement has taken its effect on his subjective probabilities, which are given by Joe's scenario matrix in Table 4.27.

Table 4.27: Joe's scenario matrix.

Action	Scenario			
	0 %	0 %	40 %	60 %
	s_1	s_2	s_3	s_4
Deny	Mary denies	Mary confesses	Mary denies	Mary confesses
Confess	Mary denies	Mary denies	Mary confesses	Mary confesses

Table 4.28 contains both the potential consequences, i.e., penalties, of Joe's action and his associated utilities in parentheses. Joe's expected utility of Deny amounts to

$$EU(\text{Deny}) = 0.4 \cdot 1 + 0.6 \cdot (-4) = -2,$$

whereas the expected utility of Confess is

$$EU(\text{Confess}) = 0.4 \cdot (-3) + 0.6 \cdot (-3) = -3.$$

Table 4.28: Joe's decision matrix with utilities in parentheses.

Action	Scenario			
	0 %	0 %	40 %	60 %
	s_1	s_2	s_3	s_4
Deny	1 year (1)	5 years (-4)	1 year (1)	5 years (-4)
Confess	0 years (5)	0 years (5)	4 years (-3)	4 years (-3)

Thus, although he is rational, Joe will deny! This is optimal because, from his own point of view, the strategic-independence assumption is violated and the probability

⁷⁷ I will come back to this important strategy in Section 5.5.

that Mary denies if (and only if) he denies, too, is sufficiently high in order to be convinced that denying is better than confessing.

This can be illustrated even better by the decision tree in Figure 4.15. It does not matter at all whether or not Joe seems to be silly from our own perspectives. Subjectivistic decision theory is descriptive rather than prescriptive!




Figure 4.15: Joe's decision tree.

Joe's behavior seems to make sense only if he doubts that Mary is rational. Otherwise, he should expect that she will break her word because, according to her own tale, she apparently *knows* that his decision is irreversible and he has no means to punish her afterwards. At first glance, the argument that Joe must doubt that Mary is rational seems obvious, but actually it is somewhat problematic: Joe need not know Mary's utility function and we could very well imagine that she feels guilty after her betrayal.

Thus, let us suppose that Mary is not only interested in the level of penalty, i.e., number of years that she must spend in jail, but also in keeping a clean record. That is, moral and loyalty play a major role in Mary's world. As we can see from her decision tree in Figure 4.16, it is optimal for Mary to deny if Joe denies but to confess if he confesses. This means that she will play tit for tat although she knows Joe's decision when making her own choice.

Now, let us come back to the original game, in which the players (believe that they) act independently. We could assume that also Joe lays considerable emphasis on moral and loyalty, which can be understood as a code of honor. In order to create a situation in which both prisoners might consider denying optimal, we must guarantee that being exempted from one year in prison is overcompensated by the guilt that the prisoners feel because of betraying the other. This is shown in the payoff matrix that is given in Table 4.29.

Well, the overall story is the same, but the game is no longer a "prisoners' dilemma" in the classical sense. Interestingly, we obtain a *coordination game* whose




Figure 4.16: Mary's decision tree.

Table 4.29: Payoff matrix of the prisoners' dilemma with moral and loyalty.

Mary	Joe	
	Deny	Confess
Deny	(10, 10)	(0, 9)
Confess	(9, 0)	(4, 4)

critical thresholds are given by

$$p^* = q^* = \frac{4 - 0}{4 - 0 + 10 - 9} = 0.8.$$

It is even a win-win-situation (see Section 4.3.3.2). Hence, Joe will deny if he thinks that Mary denies, too, with probability greater than 80 % and the same holds true, *mutatis mutandis*, for Mary. By contrast, Joe will confess if he believes that she is going to confess with probability greater than 20 %. Only in the (quite unrealistic) case in which Joe's probability that Mary denies is *precisely* 80 %, he turns out to be indifferent among Deny and Confess. We conclude that denying in the prisoners' dilemma can very well be optimal, provided moral and loyalty play a major role, and that the prisoners' level of trust (in their opponent) is sufficiently high.

For a couple of years, I play this game with my students at the University of the Federal Armed Forces and at the Bundeswehr Command and Staff College in Hamburg. The students are asked to choose their actions anonymously, i.e., nobody knows his opponent, and the game is not repeated. Most of the students are officers. Some of them belong to Germany's top-level military staff. The reader might expect that, due to the special audience, the result of the prisoners' dilemma typically turns out to be cooperative rather noncooperative. In fact, it regularly happens that more than 50 % of the students deny. Thus, most of them cooperate! In the light of the aforementioned arguments, it is hard for me to consider these people irrational. I guess that this result is exactly what everybody expects in a branch in which loyalty is a key factor.

A meta-study about the prisoners' dilemma conducted by Mengel (2018) takes civil participants into account. It reveals that the rate of cooperation is roughly 37 %, provided that nobody knows his adversary and that the game is one shot. Although the rate of cooperation is smaller for civil participants compared with military participants,⁷⁸ it might seem surprising to the reader that it is still quite large. Lave (1962) reports that the percentage number of denying players even increases with the number of trials after playing the prisoners' dilemma repeatedly. Hence, by becoming more and more acquainted with the game, the players increasingly opt for cooperation!

In my opinion, this finding underpins the social aspect of the prisoners' dilemma and, at least from a descriptive point of view, we should not be too rash when predicting people's behavior by (mis-)using the dominance principle. In any case, the given observations do not violate the basic principles of rational choice and we (still) need not assume that people are altruistic in order to explain the cooperative solution of the prisoners' dilemma. They tend to cooperate just because they are *egoistic*. Indeed, we must not mistake guilt for compassion.

4.5.2 Split or Steal

Split or Steal was part of a British game show called Golden Balls: Ann and Bob are sitting in front of one another. Each one has a ball labeled "Split" and a ball labeled "Steal." There is some amount of money in the jackpot. If both Ann and Bob choose Split, each player wins half the jackpot. If one player chooses Steal and the other chooses Split, the former wins the entire jackpot and the latter leaves with nothing. Finally, if both players choose Steal, they go away empty-handed. The players must make their choices independently, but they may speak to one another and to the host before making up their minds.

Suppose that the players are interested only in the monetary consequences of their decisions. Then we could analyze the game by the payoff matrix in Table 4.30. This is an interesting variant of the prisoners' dilemma: Steal dominates Split only in the *weak* sense. Thus, we must know the zero probabilities of the players in order to predict their behavior (see Section 3.6).

Table 4.30: Payoff matrix of Split or Steal.

Ann	Bob	
	Split	Steal
Split	(1, 1)	(0, 2)
Steal	(2, 0)	(0, 0)

⁷⁸ It is even *significantly* smaller, but I will not go into the details here.

Split is an optimal choice only if the player is convinced that the other decides to steal. By contrast, if a player assigns Split the slightest amount of probability, he will try to steal the entire jackpot. Since the players are able to communicate, they should try to convince one another to split, and if a player is foolish enough, the other one can easily catch the entire jackpot. Nonetheless, how should a rational player be convinced to split the jackpot if his only optimal choice is Steal just *because* he thinks that his adversary might decide to split? Since Ann and Bob put much effort into making the other believe that they should split the jackpot, the game typically ends up in a situation in which both players assign Split a positive probability, which can be arbitrarily small. However, the only rational choice of Ann and Bob is *Steal* and this is precisely what the dilemma is all about!

In real life, most players are not driven only by monetary considerations. They are influenced also by social norms like moral, loyalty, and honor. Further, they try to avoid a bad reputation. For example, if Ann promises Bob in front of rolling cameras that she will split, whereas she actually decides to steal the entire jackpot, she might lose her public reputation. Golden Balls had about 2 million viewers and thus each player was certainly aware of the potential effect on his image. The question is whether or not winning the entire jackpot overcompensates the loss of reputation and the guilt that the cheating player feels.⁷⁹ Hence, we may expect that many players in that game were willing to split the jackpot. In fact, van den Assem et al. (2012) report that 53 % of all players of Golden Balls chose Split.

Consider Ann's (reduced) decision matrix in Table 4.31: If Bob cooperates, Ann wins half the jackpot and feels good, which provides a utility of 1. By contrast, if Bob decides to steal, Ann wins nothing and feels bad because of being cheated, which costs her one util. If she decides to steal, whereas Bob wants to split, she wins the entire jackpot, but she loses her reputation and feels guilty because of cheating Bob. This leads to a utility of 0, which is lower than the utility of winning half the jackpot. Finally, if also Bob chooses Steal, she goes away empty-handed. In this case, we may assume that she does not lose her reputation and she does not feel guilty because Bob has cheated, too. Thus, the utility of Ann is 0.

Table 4.31: Ann's reduced decision matrix of Split or Steal.

Ann	Bob	
	Split	Steal
Split	Win half the jackpot and feel good because Bob cooperates (1)	Win nothing and feel bad because of being cheated (−1)
Steal	Win the entire jackpot but lose reputation and feel guilty (0)	Go away empty-handed without any loss of reputation (0)

⁷⁹ Of course, this could depend also on the amount in the jackpot.

Now, the critical threshold of Ann is

$$q^* = \frac{0 - (-1)}{0 - (-1) + 1 - 0} = \frac{1}{2}.$$

This means that she will split if she believes that Bob splits with some probability greater than $\frac{1}{2}$. By contrast, if she thinks that Bob chooses to steal the jackpot with a probability greater than $\frac{1}{2}$, she will decide to steal, too. Ann is indifferent among Split and Steal only in the singular (and thus quite unlikely) case of $q = \frac{1}{2}$.

We conclude that Ann's decision depends essentially on her personal impression about Bob's character. Trust plays the most important role in this game, which is driven by verbal *and* nonverbal communication. In my opinion, it does not suffice for Bob just to promise that he will split. In order to convince Ann to split, Bob's promise must be credible, which might depend on psychological factors that are hard to observe. If Ann thinks that Bob is unscrupulous and has no reputation at all, Bob's verbal attempts to persuade Ann will probably miss the mark. A more detailed analysis of Golden Balls can be found in van den Assem et al. (2012).

4.5.3 Nuclear Threat

The next conflict is prototypical in game theory: Two nuclear powers are facing one another. Each one could defeat the other at one stroke. No country is able to react, and the country that strikes first will destroy its enemy. This is a continuous-time game that takes place at each time $t \in [0, \infty)$. It stops only when a country strikes first and it is possible that both countries strike first at the same time.

It is common practice in game theory to consider each country a rational subject. Hence, what is a country going to do in such a situation,⁸⁰ which is described by the payoff matrix in Table 4.32? Destroying the enemy is rewarded by one util, whereas being destroyed is penalized by 10 negative utils. We can see that Strike *strictly* dominates Restrain. In fact, even if the other country strikes too, striking is better than holding still. This is because one can be sure that the bomb will drag the enemy into the same vortex, which alleviates the own downfall a little bit. The reader can see that this game is nothing other than some prisoners' dilemma.

Table 4.32: Payoff matrix of the countries.

Country A	Country B	
	Restrain	Strike
Restrain	(0, 0)	(-10, 1)
Strike	(1, -10)	(-9, -9)

⁸⁰ I do not ask what a country *should* do. That question would be prescriptive rather than descriptive.

The countries make their decisions simultaneously. This means that the standard assumption of strategic independence is satisfied. Hence, we must conclude that each (rational) nation will destroy its enemy, immediately, provided that the first strike is final! Both from a political and from an ethical point of view, this situation is clearly devastating and unacceptable. In the next chapter, I will discuss the reason why nuclear powers typically do *not* act in this manner.

4.6 Conclusion

Strategic independence is a result of imperfect information. It represents a standard assumption of noncooperative game theory. Conflicts in which the players believe in strategic independence can easily be solved by means of subjectivistic game theory. In most cases, we are able to provide a unique rational solution. This holds true even if some player doubts that his action has no influence on the action of his opponent. In that slightly more complicated case we can make use of decision trees in order to solve the strategic conflict.

Many strategic conflicts in real life can be expressed by a 2×2 payoff matrix. We can distinguish between anti-coordination, discoordination, and coordination games. In an anti-coordination game, the players can typically choose to be either defensive or offensive. These games are characterized by the fact that the players usually decide to be defensive in order to avoid a confrontation. However, the players typically try by all means to convince one another of the contrary. If both players are quite sure that their adversary is defensive, they decide to be offensive and the game ends in a disaster.

The solution of a discoordination game depends essentially on how much the players are able to deceive one another. More precisely, each player should make the other believe that he is going to choose Action *a*, whereas he actually chooses Action *b*. The player who succeeds wins the game. In any case, the given solution is rational because each player makes an optimal choice, based on his subjective probability distribution on the action set of the other player.

In a coordination game, the players try to coordinate their actions. Real-life behavior is often driven by convention. To be more precise, in order to facilitate coordination, societies create focal points, i.e., actions that are performed out of habit. Nonetheless, even coordination games can represent genuine conflicts if the protagonists do not pursue the same interests. Then we can often observe that people try to manipulate one another in order to force the desired solution. In any case, trust and belief play an essential role in coordination games.

In a 2-person zero-sum game the gain of one player corresponds to the loss of the other. Such games can often be observed in real life. Their subjectivistic solution is quite easy. In many applications it suffices to consider the modes of the subjective probability distributions of the players in order to find a unique rational solution. The

players turn out to be indecisive only in singular cases, which are unlikely to occur in reality. To sum up, the subjectivistic approach does not treat zero-sum games differently from any other strategic conflict.

The prisoners' dilemma represents a touchstone in game theory. The traditional approach holds that each rational subject should be noncooperative. However, many empirical studies show that people often decide to cooperate. Their behavior can be explained by subjectivistic game theory. We could assume that some player does not believe in strategic independence. Another argument is that human beings are not only interested in monetary consequences or material incentives. Social norms like moral, loyalty, and honor can play a major role, too. Other factors like public reputation and guilt might be important as well. If we take all these factors seriously into account, the prisoners' dilemma can turn into a coordination game, in which cooperation is a natural phenomenon.

5 Reaction

In this chapter, I suppose that the action of a player can have an influence on the action of the other, whereas the action of the former cannot depend on the action of the latter. Put another way, this chapter deals with dynamic games and it is clear that the strategic-independence assumption does no longer hold true in such a situation.

For example, consider once again a game between Ann and Bob in which Ann's action set is $A = \{\text{Up, Down}\}$, whereas Bob has the action set $B = \{\text{Left, Right}\}$. Suppose that Ann starts making her decision, which is then revealed to Bob. This means that Bob knows her choice when making his own decision. For this reason, Bob is able to *react*, whereas Ann can only *act*. Action and reaction can be explained by the private information sets of Ann and Bob. More precisely,

- Ann has only the trivial information set $\mathcal{I}_A = \{\Omega_A\}$, whereas
- Bob possesses the information set $\mathcal{I}_B = \{s_{B1}, s_{B2}, s_{B3}, s_{B4}\}$.

Hence, Bob has perfect information, whereas Ann's information is imperfect.

In traditional game theory (von Neumann and Morgenstern, 1953, Chapter 15), dynamic games in which each player knows the previous moves of his adversaries (and of Nature) are called games with perfect information. Typical examples are Chess and Backgammon (von Neumann and Morgenstern, 1953, pp. 51–52). In this book, I refrain from using that terminology. Here, a game with perfect information refers to a game in which all players know the strategy of each other and thus are able to interact. Put another way, a game with perfect information is a coherent game. Coherent games will be described in Chapter 6.

Ann's action cannot depend on Bob's action. For this reason, the scenarios s_{B2} and s_{B3} in Bob's scenario matrix are actually impossible, whereas the scenarios s_{A2} and s_{A3} are very well possible. However, Ann does not know which scenario happens, i.e., she is unaware of Bob's strategy. Moreover, although the heterogeneous scenarios s_{B2} and s_{B3} are impossible from an *outside observer's* point of view, we need not assume that Bob knows that Ann possesses the trivial information. Hence, the scenarios s_{B2} and s_{B3} might very well be part of his state space Ω_B . By contrast, if we assume that Bob knows that Ann's private information is trivial, and if we rule out pessimism and optimism, he considers s_{B2} and s_{B3} a priori impossible. In this case, his state space corresponds to $\Omega_B = s_{B1} \cup s_{B4}$ and his private information set equals $\mathcal{I}_B = \{s_{B1}, s_{B4}\}$. However, a posteriori, Bob is always aware that the true state of the world, ω_0 , either belongs to s_{B1} or to s_{B4} , respectively, which means that he knows Ann's action.

5.1 The Ultimatum Game

The following dynamic one-shot game was first described by Güth et al. (1982) and is called ultimatum game: Ann is endowed with \$100 and must split this amount of

<https://doi.org/10.1515/9783110596106-005>

money with Bob, who can either accept or reject Ann's proposal. If he accepts, the money is split accordingly. Otherwise, both Ann and Bob go away empty-handed. The potential outcomes of the game are known to both players.

Suppose, for the sake of simplicity, that Ann can make either a fair split (50/50) or an unfair split (90/10). Ann's decision tree of the ultimatum game can be found in Figure 5.1. Ann has imperfect information and so she assigns Bob's possible actions a subjective probability. She believes that Bob accepts her proposal with probability 90 % if her split is fair, but he will reject it with probability 80 % if it is unfair. According to her own utilities, which are given in parentheses behind the end nodes in Figure 5.1, she decides to make a fair split.




Figure 5.1: Ann's decision tree of the ultimatum game.

Subjectivistic game theory considers each player a decision maker. Hence, in order to understand the solution of the ultimatum game, we must draw *two* decision trees, i.e., one for each player. In some cases it suffices to consider only one decision tree or no decision tree at all. This can be done, e.g., if the game that we are taking into consideration is symmetric, which means that each player is equipped with the same action set, information set, utility function, etc. However, here we presume that one player has more information than the other, in which case it is usually recommended to analyze the game by drawing two decision trees.

Why should Bob decide to reject any positive amount of money at all if he is rational? Well, this can be seen in Figure 5.2, which contains Bob's decision tree. The problem is that he might feel betrayed if Ann splits the money in an unfair way. The bad feeling wipes out the good feeling of winning \$10. If Bob rejects, he gets nothing, but at least he feels satisfied because he has punished Ann for being unfair. This explains the utilities that are depicted on the lower right of Figure 5.2. Hence, Bob will accept if Ann makes a fair split, but he will reject if she makes an unfair one. Due to his own utility function, this behavior is completely rational.

Since Ann does not know Bob's utility function, she is uncertain about his behavior. Thus, she does not know whether he will accept or reject if she makes an unfair




Figure 5.2: Bob's decision tree of the ultimatum game.

split, which explains why she assigns some probabilities to Bob's actions on the lower right of Figure 5.1. She even does not know whether Bob agrees to a fair split, which explains why she assigns some probabilities to Bob's actions on the upper right of Figure 5.1, too. The subjectivistic approach to game theory does not require complete information, i.e., that the players know all relevant aspects of the game. I will come back to this point in Section 7.1.1.

Experimental results over different countries and cultures show that the average rejection rate is 16 % and, on average, people choose to split the money by a ratio of 60/40 (Oosterbeek et al., 2004). The offered share is usually smaller if the total amount of money is larger and it becomes larger if the players are inexperienced. Further, the rejection rate is typically lower the larger the total amount of money and the larger the offered share. On average, for each percentage point that is offered by the proposer (i.e., "Ann") to the responder (i.e., "Bob"), the rejection rate decreases by half a percentage point.

The experiments clearly reveal that people are, in general, not unfair and that they are punished if they are unfair. In my opinion, this result is not surprising at all. It explains why and how societies work. There should be some people who are willing to punish or, at least, blame evildoers even if this is unfavorable (in a pure monetary sense) or dangerous for themselves. This keeps people without any remorse from being unfair. As in the prisoners' dilemma with moral and loyalty, a rational subject is not fair because he is altruistic. It is because he is *egoistic*! I think that we can find many examples that demonstrate that this is a typical phenomenon in human culture.

What happens if the responder has no possibility to punish the proposer? That is, whether the responder accepts or rejects the offered share does not have any impact at all. For example, let us suppose that the proposer is a dictator who is asked to spend some money for his citizens. This is the so-called dictator game. Actually, the dictator game is an ordinary decision problem, not a strategic conflict. However, since it represents a straightforward simplification of the ultimatum game, it seems reasonable to me to discuss the dictator game here and not in Part I of this book.

Is it true that the dictator must offer nothing if he is rational? Experiments show that most people offer a positive amount of money although there seems to be no incentive for them to do so. Obviously, this result contradicts the homo-economicus model, which is prevalent in neoclassical economics. I guess that most neoclassical economists would agree that some people are fair, but they would say that fairness is negligible from a macro-economical point of view. Nonetheless, fairness *plays* an important role in human societies.

To some up, most of us have a bad feeling when refusing to share some given amount of money with others. Proposing \$0 is worse than proposing \$1 or \$2. I observed this kind of feeling many times myself when being asked for some money by homeless people. At least from a subjectivistic point of view, giving money for charity is a rational behavior, provided we want to be fair. My mentor, Karl Mosler, once told me: “Do good and make it known!” So do it if you feel better and enjoy that feeling.

5.2 Nuclear Threat

In real life, nuclear powers are usually not able to destroy an enemy on the first strike. The nuclear arms race during the cold war and the military alliances that have been forged in the second half of the 20th century have created a situation in which a country that strikes first must fear a retaliatory strike by its enemy or some ally. This situation is depicted by the decision tree in Figure 5.3.

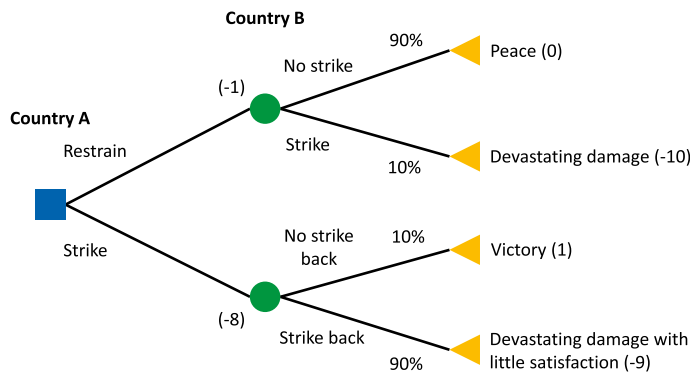


Figure 5.3: Decision tree of Country A, which fears a retaliatory strike of Country B.

Note that the strategic-independence assumption is violated because the country that has been attacked knows that it has been attacked and now it can decide whether to strike back or to hold still. Country A assumes that Country B strikes back with probability 90 % if it strikes first, whereas Country B holds still with the same probability if Country A does not strike first. This situation cannot be explained by the payoff matrix

in Table 4.32, which can be found in Section 4.5.3. A payoff matrix is no longer appropriate if the strategic-independence assumption is violated. Instead, we should use a decision tree, which enables us to illustrate that one player can react to the other.

In contrast to our findings in Section 4.5.3, Strike no longer dominates Restrain and Country A will decide to hold still. The same applies, *mutatis mutandis*, for Country B and thus no country dares to strike first. However, the whole line of argument makes sense only if Country A can really *expect* that Country B possesses the necessary means to strike back and vice versa. For this reason, Country B should try to make use of deterrence. That is, it should guarantee that its enemy's subjective probability that it will strike back is sufficiently high.

Figure 5.3 reveals that Country A will certainly restrain if

$$p_0 \cdot 0 + (1 - p_0)(-10) > p_1 \cdot 1 + (1 - p_1)(-9),$$

where p_0 is the probability of Country A that Country B holds still if the former restrains and p_1 is the probability that Country B holds still otherwise. Hence, we come to the conclusion that Country A holds still if $p_0 - p_1 > 0.1$. That is, in order to avoid a first strike, Country B must guarantee that

$$1 - p_1 > (1 - p_0) + 0.1.$$

This means that the probability that Country B strikes *back* should exceed the probability that it strikes *first* by more than 0.1. This result is crucial: It says that Country B should try to make Country A believe that it does not strike first, but *if* Country A dares to attack Country B, it will retaliate without hesitation. Simply put, it should try to convince its enemy that it plays tit for tat.

Hence, the typical reason why countries have nuclear weapons is because they try to *avoid* a violent confrontation, not because they plan to use them. Nowadays, military intelligence is so far-reaching that a country that wants to deter its enemy must expose his nuclear arsenal. Hence, it should *not* hide but show off (some, not all of) its nuclear weapons. Otherwise, the enemy could doubt its military capabilities, in which case the weapons miss their actual purpose, i.e., they fail to be a credible threat. Thus, possessing nuclear weapons, and showing them off to the enemy, can be considered both a necessary and a sufficient condition for peace.

Are the previous arguments still valid if a country is able to destroy an enemy at one stroke? It depends on the military intelligence of the enemy. For example, if the countries are able to monitor one another by satellite, it does not make any sense to strike first, since the enemy can immediately react and strike back before it is too late. Hence, each country should try to make the other believe that its military intelligence is far-reaching enough in order to be able to strike back immediately.

Deterrence is a necessary means only if a country is faced with a nuclear power or, at least, with an enemy that could become a nuclear power in the future. However, given the international political situation and the growing flashpoints worldwide, I fear that this military instrument will continue to be important for a long time.

To sum up, we are very well able to explain the global proliferation of nuclear weapons by subjectivistic game theory. For more details on that topic see, e.g., Schelling (1980).

5.3 Penalty Shoot-Out

Reconsider the penalty shoot-out that has already been described in Section 4.3.2.3. The keeper stands in his goal and waits for the scorer to act. In reality, the scorer has some time to kick the ball after the referee gives the whistle signal. During this time, the goalkeeper can decide whether to jump before or after the scorer kicks the ball. If he decides to jump into one corner beforehand, he should do this as late as possible, i.e., just when the scorer is kicking the ball,⁸¹ in order to keep him from choosing the other corner. However, the problem is that the goalkeeper does not know when the scorer is going to kick the ball, which means that he could be too rash when jumping into some corner. In any case, if the goalkeeper does not wait for the kick, he cannot react to the scorer. This means that he is risking to jump into the wrong corner or to be already on the floor when the ball approaches the goal. By contrast, if he decides to wait for the kick, he is able to react to the scorer, but then he could be too late to catch the ball. Hence, what will the goalkeeper do?

For a better illustration of this game, I have split the goalkeeper's decision tree into two parts. Figure 5.4 contains the decision tree given that the goalkeeper jumps before the scorer kicks the ball. By contrast, Figure 5.5 contains the decision tree given that the goalkeeper jumps after the scorer kicks the ball. First of all, let us discuss the decision tree in Figure 5.4.

The goalkeeper can jump either into the left or into the right corner. Then he has to accept the scorer's choice. After the scorer sees the goalkeeper jumping, he is able to react accordingly. Thus, if the scorer observes that the goalkeeper jumps into one corner, he will kick the ball into the other. The goalkeeper thinks that the scorer is able to choose the other corner with probability 70 %. Moreover, he believes to catch the ball with probability 90 % if the scorer was not able to react. By contrast, if he was able to choose the other corner, it is impossible to catch the ball. It is obvious that the goalkeeper is indifferent among Left or Right. In any case, the expected utility amounts to -0.46 .

Now, we come to the decision tree in Figure 5.5. The goalkeeper is able to react to the scorer, who chooses Left or Right each with equal probability. If the goalkeeper decides to jump into the right corner, he catches the ball with probability 80 %. By contrast, if he decides to jump into the wrong corner, it is impossible to catch the ball. Of course, he will decide to jump into the right corner, in which case he realizes an

81 Indeed, from a decision-theoretic (but not a physical) point of view, jumping at the same time as the scorer kicks the ball means to act *beforehand*. I will come back to this point in Section 6.2.




Figure 5.4: Goalkeeper's decision tree if he jumps before the scorer kicks the ball.




Figure 5.5: Goalkeeper's decision tree if he jumps after the scorer kicks the ball.

expected utility of 0.6. This exceeds the expected utility of jumping beforehand. We conclude that the goalkeeper decides to wait for the kick.

It is intuitively clear that the scorer will try to wait as long as possible in order to be able to react to the goalkeeper's choice. However, since the goalkeeper decides to wait, too, the scorer must kick before the goalkeeper jumps. That is, the scorer cannot react to the goalkeeper and so he has imperfect information. That problem can be solved by a decision tree, too. We can easily imagine that the scorer will choose precisely that corner with the lower probability of being chosen by the goalkeeper. In the case that the corners have equal probability, he will be indifferent among Left and Right. Of

course, these arguments implicitly presume that the probability that the goalkeeper catches the ball if he jumps into the right corner does not depend on the corner itself.

In my opinion, the most interesting part of the solution of this game is not the result itself but the way we come to our conclusion. Subjectivistic game theory requires us to think about the specific events that may have a substantial impact on the potential outcomes of the game, i.e., the players' consequences, and thus on the individual preferences of the players. This gives us an impression about how people decide even in situations of conflict in which the choice must be made in split seconds. As already mentioned in Section 4.3.2.3, we do not have to assume that they make their decisions in a conscious way. Indeed, the cognitive process can be very fast.

5.4 The Game of Chess

Chess is *the* prototypical 2-person game in which chance plays no role, i.e., the outcome of the game depends only on the moves and countermoves of Black and White. It is played on a square board of eight rows and eight columns. At the beginning of the game, each player has 16 pieces. The players move their pieces alternately. Usually, the player whose turn it is can move only one piece from one place to another.⁸² There is no hidden information, i.e., each player has the same information regarding the moves and countermoves that have been made at any point in time. Hence, everybody knows the available strategies of the other. The goal is to checkmate the other's king by placing it under an inescapable threat of capture. The player who achieves this goal wins and the other one loses the game. However, Chess can end also in a draw. The Chess rules contain several criteria for a draw, which guarantee that the game always comes to an end. Thus, both players have a finite number of strategies, which means that Chess represents a finite game. Let us suppose that each player assigns winning utility 1, losing utility -1 , and a draw utility 0. That is, Chess is a zero-sum game.

The game starts with the initial position on the chessboard and White moves first (see Figure 5.6). The initial position is the first node in the decision tree of each player. From the perspective of White, it represents a decision node, but from the viewpoint of Black, it is a chance node, etc. Making a move, in general, means to change the position of the pieces on the chessboard. However, we can imagine also moves that are not made on the chessboard. For example, a player could offer a draw to his opponent or he might agree to any such offer that has already been made in the past by the other player. Further, the game can end also by resignation of a player, which can be considered a move, too. Although these moves lead to another position in the decision tree, they do not change the picture on the chessboard. Hence, the term "position" should be understood as a node in the decision tree, which contains the complete history of

⁸² The only exception is castling, which involves two pieces, namely the king and a rook.




Figure 5.6: Initial Chess position.

the game up to that point—including not only all preceding moves and counter moves but also all draw offerings and promotions of pawns that have been done in the past, etc. Each decision and chance node in the decision tree specifies whether it is Black’s or White’s turn, whereas an end node indicates that the game is over.

5.4.1 Zermelo’s Chess Theorem

The first formal theorem in game theory is introduced by Zermelo (1913)⁸³:

Theorem 4 (Zermelo’s Chess theorem). *Either White can force a win, or Black can force a win, or both Black and White can force at least a draw.*

This fundamental theorem of game theory shall be explained in more detail. Let P_0 be an arbitrary node in the decision tree. It is not relevant whether we consider the decision tree of Black or White. Suppose, without loss of generality, that it is White’s turn. White can move to some succeeding node P_1 , then Black can move to another succeeding node P_2 , etc. This leads to a so-called endgame $P = (P_0, P_1, \dots, P_n)$.⁸⁴

White is in a winning position if and only if

1. there exists some endgame P such that White wins and
2. if Black alternates any moves, White can find some appropriate counter moves such that he wins also in the new endgame P' .

Hence, White must be able to beat Black irrespective of how Black changes his moves. In this case, we say that White can force a win.

The same arguments apply, mutatis mutandis, to Black. In fact, although Black moves from P_1 to P_2 , etc., we can still consider P_0 a starting point. That is, Black can

⁸³ A nice overview of Zermelo’s Chess theorem can be found in Schwalbe and Walker (2001).

⁸⁴ The term “endgame” is somewhat misleading, since P_0 could be also the initial position on the chessboard. However, it is used by Zermelo throughout his original work.

force a win if and only if there exists some endgame in which Black wins and Black is able to beat White irrespective of how White changes his moves. It is clear that Black and White cannot be in a winning position together.

Further, we say that White is in an unbeatable position if and only if

1. there exists some endgame P such that White does not lose and
2. if Black alternates any moves, White can find some appropriate counter moves such that he does not lose also in the new endgame P' .

Hence, White must be able to avoid a loss irrespective of how Black changes his moves. Put another way, White can force at least a draw. This could mean that White wins or that the game ends in a draw. Of course, if White is in a winning position, he is also in an unbeatable position, but the converse is not true.

White is *not* in an unbeatable position if and only if Black is in a winning position. This can be seen as follows: If White is not in an unbeatable position, Black has some (alternating) moves that cannot be countered by White in order to avoid a loss, which means that Black can force a win. Conversely, if Black is in a winning position it is obvious that White cannot be in an unbeatable position. Once again, the same arguments apply also to Black. Thus, if one player is not in an unbeatable position, we can say that he is in a losing position.

Moreover, if White is in an unbeatable position, but not in a winning position, the same must hold true for Black: If Black would not be in an unbeatable position, White would be in a winning position, and Black cannot be in a winning position because otherwise White cannot be in an unbeatable position. This completes our (verbal) proof and understanding of Zermelo's Chess theorem.

5.4.2 The Subjectivistic Explanation of Chess

The Chess theorem is a beautiful contribution to game theory. It demonstrates how far we can go with logical reasoning. At the beginning of the game we can expect that both players are in an unbeatable position, which means that no one can force a win. So why do (rational) players lose any game at all?

Chess statistician Jeff Sonas reports that only about 50 % of all Chess games end in a draw. The relative number of draws depends on the player's capability, which can be measured by his Elo points. Interestingly, the better the player the higher the rate of draws that one can usually observe in his track record. However, on average, even the best players in the world have draws in less than 60 % of their games against other elite players or more mediocre ones. If we match only games of players who exceed 2700 Elo points, the relative number of draws equals 65 %.⁸⁵

⁸⁵ I found this information on February 8, 2019, under the following URL: <https://en.chessbase.com/post/sonas-what-exactly-is-the-problem->

In order to understand Chess from a subjectivistic point of view, first we may start considering the decision matrices of Black and White. Usually, both players have a tremendous number of possible strategies at hand when starting at Position P_0 . Nonetheless, the number of strategies is always finite. Hence, we may assume, for the sake of simplicity but without loss of generality, that both Black and White can choose only between 3 strategies. On the left-hand side of Table 5.1 we can see White's decision matrix, whereas Black's decision matrix is given on the right-hand side. The columns s_1, s_2, s_3 in White's decision matrix represent Black's available strategies and the rows r_1, r_2, r_3 are the available strategies of White. Since P_0 can be any position that has been obtained during the game, we could speak also about partial strategies, since they determine only the moves and countermoves of the corresponding endgame.

Table 5.1: Decision matrices of White (left) and Black (right).

Strategy	Scenario		
	s_1	s_2	s_3
r_1	(1)	(1)	(1)
r_2	(-1)	(0)	(1)
r_3	(1)	(-1)	(1)

Strategy	Scenario		
	r_1	r_2	r_3
s_1	(-1)	(1)	(-1)
s_2	(-1)	(0)	(1)
s_3	(-1)	(-1)	(-1)

White is in a winning position because the first row on the left-hand side of Table 5.1 contains only ones. Equivalently, Black is in a losing position because he has no unbeatable strategy. It is clear that r_1 dominates each other strategy in the weak sense. However, is it really true that we must consider White irrational if he does *not* prefer Strategy r_1 , in which case it may happen that he does not force a win? Well, the answer is “No”! Suppose, for example, that White assigns Strategy s_2 probability 0. Then r_3 is just as good as r_1 . Indeed, a rational player must not apply a dominated strategy, but we already know from Section 2.6 and Section 3.6 that dominance depends essentially on the null scenarios of the decision maker. Simply put, if White believes, a priori, that Black does not perform Strategy s_2 , he does not take this strategy into consideration. In this case, we are left with ambiguity and cannot say whether White will perform Strategy r_1 or Strategy r_3 .

Now, consider the situation in Table 5.2. The given decision matrices reveal that both Black and White are in an unbeatable position. If White applies Strategy r_1 , he can force at least a draw. The same holds true for Black if he applies Strategy s_2 . However, does this mean that White must prefer Strategy r_1 if he is rational or, equivalently, that r_1 is the only “correct” strategy of White, as is claimed by Zermelo (1913)?

Once again, from a subjectivistic point of view, the answer is “No”! For example, suppose that White's subjective probabilities are $P(s_1) = 0.4$, $P(s_2) = 0$, and $P(s_3) = 0.6$. In this case, his expected utilities are $EU(r_1) = 0.4$, $EU(r_2) = 0.2$, and $EU(r_3) = 0.6$, which means that White prefers Strategy r_3 , not the unbeatable strategy r_1 !

Table 5.2: Decision matrices of White (left) and Black (right).

Strategy	Scenario		
	s_1	s_2	s_3
r_1	(1)	(0)	(0)
r_2	(-1)	(0)	(1)
r_3	(0)	(-1)	(1)

Strategy	Scenario		
	r_1	r_2	r_3
s_1	(-1)	(1)	(0)
s_2	(0)	(0)	(1)
s_3	(0)	(-1)	(-1)

Strategy s_2 is unbeatable, whereas Strategy s_3 is quite unfavorable from Black's viewpoint. Nonetheless, White does not believe that Black chooses Strategy s_2 . By contrast, he assigns Strategy s_3 a relatively high probability. Thus, we could say that White considers Black somewhat silly. For example, given Black's previous moves, White might have come to the conclusion that he is not a good player. For this reason, he may also play badly in order to maximize his expected utility. Although we can expect that this is not the typical case in a professional Chess match, this situation can very well happen if the players' capabilities are far apart.

The more White is uncertain about the strategy of Black, the less concentrated are his subjective probabilities on a small number of Black's strategies. In the case in which White does not neglect any strategy at all, except for those strategies that are unavailable to Black, he assigns each strategy of Black a positive probability. Nonetheless, even then it is not necessarily true that White prefers an unbeatable strategy. For example, let White's subjective probabilities be $P(s_1) = 0.3$, $P(s_2) = 0.1$, and $P(s_3) = 0.6$, in which case his expected utilities are $EU(r_1) = EU(r_2) = 0.3$ and $EU(r_3) = 0.5$. That is, White still prefers Strategy r_3 to Strategy r_1 .

In the special case in which White's subjective probabilities are $P(s_1) = P(s_2) = \frac{1}{2}$ and $P(s_3) = \frac{1}{2}$, he is indifferent among all of his available strategies. Then White could choose any strategy and so the game need not end in a draw at all. Hence, the choice of a rational Chess player depends essentially on his subjective probabilities, but he can still fail and lose the game. For example, if White chooses r_3 and Black chooses the unbeatable strategy s_2 , the latter wins the game. Nonetheless, White is still rational in the sense of subjectivistic game theory.

Now, suppose that White believes that Black chooses the unbeatable strategy s_2 , whereas Black believes that White chooses the unbeatable strategy r_1 . Well, in this case, r_1 and s_2 are, in fact, optimal strategies and the solution (r_1, s_2) , in which the game ends in a draw, represents a Nash equilibrium.⁸⁶ However, there is no reason why White should be *convinced* that Black chooses Strategy s_2 and the same argument applies, mutatis mutandis, to Black. Moreover, even if both players are convinced that the other chooses his only unbeatable strategy, they could be wrong.

I conclude this section by making the following assertions:

⁸⁶ As already mentioned before, Nash equilibrium will be explained and discussed in Section 7.2.

- If a player is in a winning position and believes that the other player acts in a reasonable way, he will force a win.
- If a player is in an unbeatable position and believes that the other player acts in a reasonable way, he will force at least a draw.

Acting in a “reasonable way” means to force a win when being in a winning position and to force at least a draw when being in an unbeatable position.⁸⁷ Hence, if the players are in an unbeatable position and believe that their opponent acts in a reasonable way, they will force at least a draw and so the game must end in a draw. This is illustrated in the next section by backward induction.

5.4.3 Backward Induction

Chess is a finite game and so it can be solved by backward induction. To me it makes no sense at all trying to solve Chess by using decision matrices or a payoff matrix. The game of chess can be a (highly) irregular decision problem and thus decision matrices can suffer from the usual shortcomings, which I have already discussed in Section 2.8.2. Simply put, we are hardly able to say how a rational Chess player acts during the game, i.e., a posteriori, just by analyzing his situation a priori. Thus, in the subsequent analysis I will concentrate on decision trees in order to apply the method of backward induction.

Every Chess match eventually comes to an end, which is symbolized by the end nodes in the decision tree. It is either Black’s or White’s move that leads to an end. Let us consider the game, without loss of generality, from the perspective of White. His own moves in the decision tree represent actions, whereas Black’s moves are considered events. That is, White’s moves are the outgoing branches of each decision node, whereas Black’s moves can be found after each chance node in White’s decision tree. For a better illustration, we may assume that each node has only two outgoing branches. It is clear that this is a massive understatement. In reality, the players usually have an overwhelming number of possible choices, but this has no influence at all on our principal arguments.

A prototypical example of a decision tree is given in Figure 5.7, which reveals that White is in a winning position at Decision Node 1. The bold branches indicate the optimal actions of the players, given that they believe that their opponent acts in a reasonable way. Since the players have complete information, they know the optimal moves and countermoves of the other, which can be derived by backward induction.

For example, at Decision Node 4 White prefers to checkmate Black’s king and so the upper branch is bold.⁸⁸ Further, White believes that Black will not resign but

⁸⁷ Note that a player who acts in a reasonable way cannot act in an unreasonable way, too.

⁸⁸ Such a position is illustrated in Figure 5.8.

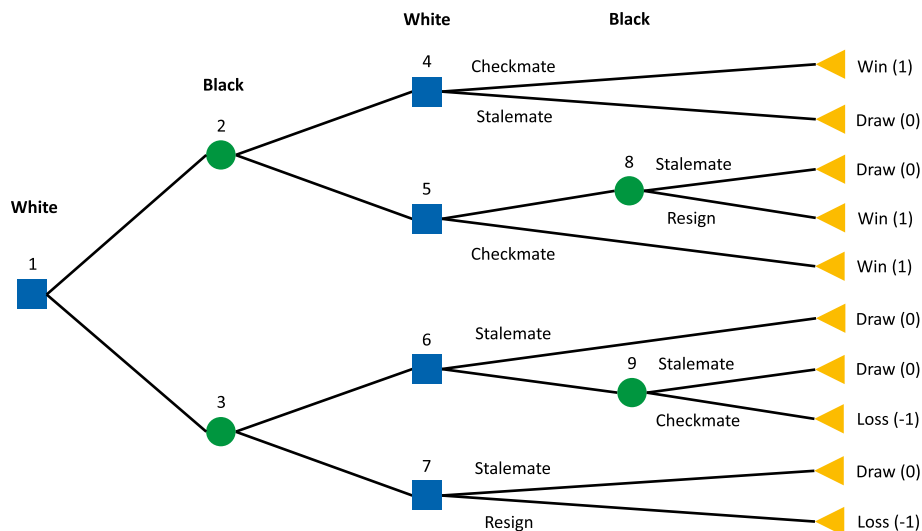


Figure 5.7: White's decision tree, who is in a winning position.




Figure 5.8: White can win by moving his queen to h8.

choose to stalemate at Chance Node 8. Hence, White assigns the upper, i.e., bold, branch probability 1. Thus, when being at Decision Node 5, White prefers to checkmate Black, i.e., to move along the lower branch. Black anticipates White's optimal moves, which means that he is indifferent among the upper and lower branch at Chance Node 2, since White wins anyway. The same arguments can be applied to the lower part of Figure 5.7. However, if White chooses the lower branch at Decision Node 1, he can at best achieve a draw, and so he would abandon his winning position. The same holds true if White moves up at Decision Node 5. Thus, White prefers to walk along the upper branch of Decision Node 1 in order to force a win.

Now, consider the decision tree in Figure 5.9. Here, White is in an unbeatable position. Let us concentrate on the lower part of the decision tree. At Decision Node 6, White will choose to stalemate because otherwise he expects that Black will checkmate his king at Chance Node 9. Moreover, at Decision Node 7 White will clearly checkmate Black's king. Hence, he cannot be beaten by Black at Chance Node 3, where Black will walk along the upper branch in order to avoid a loss, etc. At Decision Node 1, White is indifferent among the upper and lower branch because he expects that the game will anyway end in a draw.

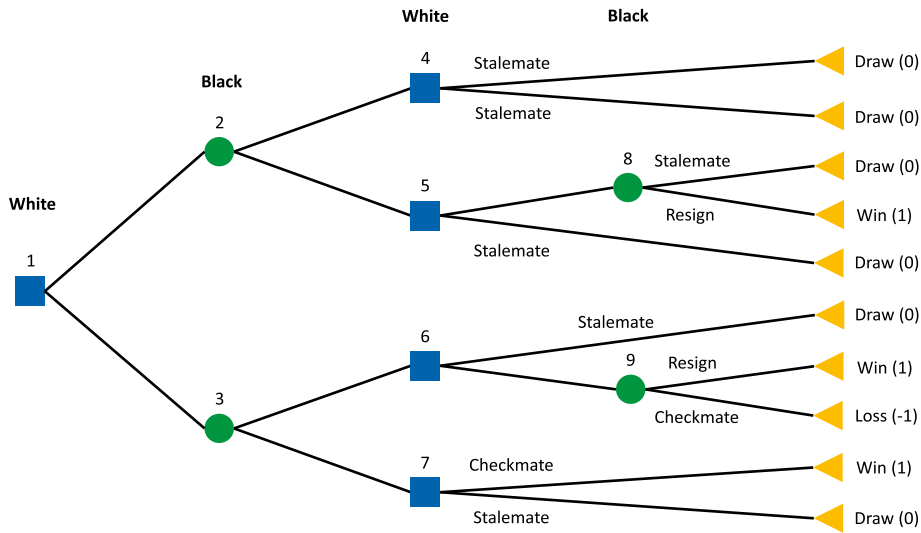


Figure 5.9: White's decision tree, who is in an unbeatable position.

We conclude that the players can always find an unbeatable strategy, i.e., force at least a draw if they are in an unbeatable position, by backward induction. In the same way, they can always find a winning strategy, i.e., force a win, provided they are in a winning position. A player in a winning position who does not force a win will lose his winning position if the other player acts in a reasonable way. This means that the game will either end in a draw or the other player wins. Thus, it makes no sense for a player in a winning position, who believes that the other acts in a reasonable way, to abandon a winning strategy. To sum up, if a player is in a winning position and believes that the other player acts in a reasonable way, he will force a win.

Moreover, a player in an unbeatable position who does not force at least a draw will lose his unbeatable position if the other player acts in a reasonable way. This means that he is going to lose the game. Hence, the player will never perform a beatable strategy, provided that he believes that the other acts in a reasonable way. There-

fore, if a player is in an unbeatable position and believes that the other player acts in a reasonable way, he will force at least a draw.

The preceding arguments lead us to the following theorem:

Theorem 5 (Subjectivistic Chess theorem). *If a Bellman-rational Chess player believes that his opponent acts in a reasonable way, he prefers to act in a reasonable way, too.*

I consider a formal proof of this theorem unnecessary. It should already be clear from backward induction. Note that the theorem does not require a Bayes-rational player: The preceding arguments do not refer to any conditional probability at all!

In the modus tollens, Theorem 5 reads thus: If a Bellman-rational Chess player does not prefer to act in a reasonable way, he must doubt that his opponent acts in a reasonable way. For example, consider White's strategy in Figure 5.10. Obviously, he applies a beatable strategy, since Black is able to checkmate White's king at Chance Node 9. Hence, if White is rational (and his strategy is optimal), we must conclude that he does not believe that Black acts in a reasonable way.

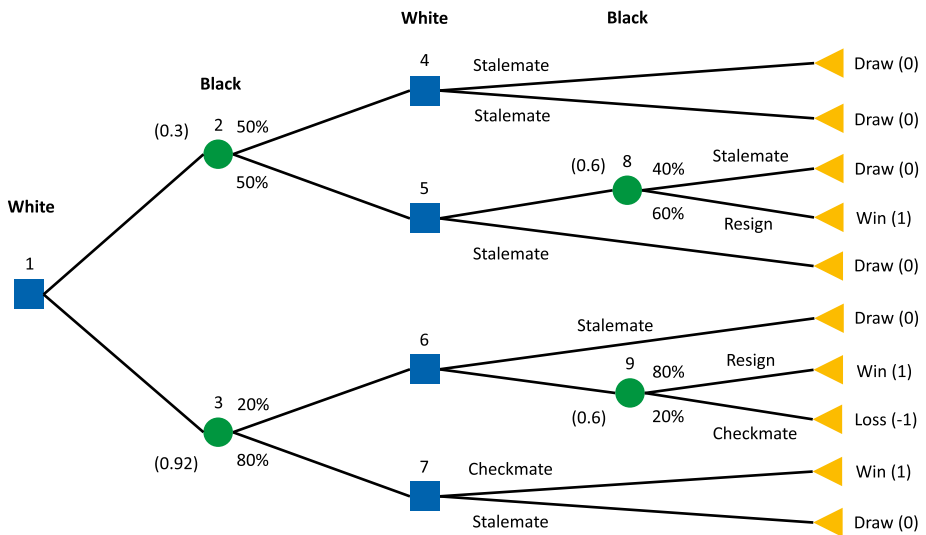


Figure 5.10: White is in an unbeatable position but performs a beatable strategy.

Consider White's subjective probabilities that are given in Figure 5.10. White believes that Black resigns with probability 60 % at Chance Node 8. For this reason, White prefers to walk along the upper branch at Decision Node 5 and thus, in contrast to Figure 5.9, he is no longer indifferent among the upper and lower branch. However, at Decision Node 4, White is still indifferent. Moreover, he assigns the branches at Chance Node 2 a fifty-fifty chance. The reader can easily verify that his expected utility at this chance node amounts to 0.3.

Now, let us turn to the lower part of Figure 5.10. White thinks that Black resigns with probability 80 % at Decision Node 9. For this reason, he decides to walk along the lower branch after Decision Node 6. This is in direct contrast to Figure 5.9, in which White decides to stalemate. The reason might be that White believes that Black is a bad player and thus fails to recognize that he can checkmate White's king by choosing the lower branch. Of course, at Decision Node 7, White will still decide to checkmate Black's king.

Further, White believes that Black walks along the lower branch at Decision Node 3 with probability 80 %. Once again, this is in contrast to Figure 5.9. However, White expects that Black plays badly, which explains his subjective probabilities. This leads to White's expected utility of 0.92 at Chance Node 3. Hence, White walks along the lower branch after Decision Node 1. Afterwards, if Black goes up, White goes down, and if Black goes down, White goes up. From White's perspective, this strategy is optimal. However, it is clearly beatable. White believes that Black is a bad player and thus he dares to apply a beatable strategy in order to maximize his expected utility. However, Black can easily beat White by making the right moves in order to come into a winning position. In fact, Chance Node 9 represents a winning position for Black. That is, White drops his unbeatable position by going down at Decision Node 6, after which Black can checkmate White's king.

It is worth emphasizing that the players do not know which strategy the other performs. In contrast to von Neumann and Morgenstern (1953), I refrain from calling Chess a game with perfect information because each player is unaware of the future moves of his opponent.⁸⁹ Indeed, it is true that both players are able to deduce the *optimal* strategies of one another, given that everyone believes that his opponent acts in a reasonable way. However, there could be many optimal strategies, and even if there is only one optimal strategy, the corresponding player could still perform a suboptimal one. As already mentioned in Section 2.1, Savage's postulates of rational choice do not make any procedural or behavioral requirements. Hence, even though we assume that the players are rational, they can still behave in a suboptimal way.

We may assume that the game starts with both players being in an unbeatable position. Hence, if the players believe that their opponent acts in a reasonable way, they will force at least a draw, which means that the game must end in a draw. Conversely, if the game does *not* end in a draw, then at least one player did not believe that his opponent acts in a reasonable way. Of course, our conclusions hold true only if we presume that no player makes any unfavorable move by mistake.

To sum up, there are plenty of reasons why Chess matches often do not end in a draw, even on a very high, i.e., professional, level. For example, some player might be irrational, in which case we cannot apply our normative theory of rational choice at

89 As already mentioned at the beginning of this chapter, in this book the term “perfect information” is reserved for situations in which some player knows the (entire) strategies of his adversaries.

all. However, even if both players are rational, we can imagine that some player fails to perform a favorable strategy. For example, he could be disturbed by noise or any other environmental influence and thus make an unfavorable move just by accident. Another possibility is that he simply overlooks some moves that are available to him in the endgame.⁹⁰ This could be the most common reason why a Chess match does not end in a draw. In fact, Chess is a very complex game.

Moreover, as we have seen above, it could even be *favorable* to perform a beatable strategy. However, this happens only if a (rational) player is not convinced that the other acts in a reasonable way. Then he might renounce acting in a reasonable way, too, in order to maximize his expected utility. Further, some player might not be interested in winning a game at all because he does not want to offend or discourage his opponent. This typically happens in private situations when a superior player wants to spare or motivate an inferior one.

We conclude that subjectivistic game theory is able to explain why players in a deterministic zero-sum game do not behave in a manner that is typically suggested by traditional game theory. Nonetheless, we must not hide the fact that a rational player should act, at least, *as if* he were making use of backward induction. In a very complex game like Chess, which contains an abundant number of strategies, this requirement might be too ambitious. However, we need not assume that a rational player is aware of his (and the other's) strategies or that he makes his choices in a conscious way—as long as his individual preferences do not violate Savage's axioms of rational choice.

5.5 The Iterated Prisoners' Dilemma

Let us come back to the prisoners' dilemma, which is usually considered a one-shot game (see Section 4.5). Now, we assume that the game is played repeatedly.

5.5.1 The Axelrod-Hamilton Experiment

The iterated prisoners' dilemma has attracted many scholars. Its fascination goes back to the seminal work of Axelrod and Hamilton (1981).⁹¹ They refer to cooperation and defection as possible choices of human beings, animals, bacteria, etc. Hence, their point of view is rather biological, but we can find an overwhelming number of examples in which individuals must either cooperate or defect in many different circumstances. Astonishingly enough, subjectivistic game theory is not restricted to human interaction. Our notion of rationality is purely descriptive (see Section 2.2) and has no

⁹⁰ Note that this is a procedural mistake. Nonetheless, the player is still rational.

⁹¹ It is further elaborated and popularized by Axelrod (1984).

procedural aspects (see Section 2.1). In particular, the subjectivistic approach does not require a homo economicus or a rational subject who is *aware* of his own decisions.

Consider the payoff matrix in Table 5.3. The game is played repeatedly. That is, it starts in Round 0 and is repeated in Round $r = 1, 2, \dots, N$. Axelrod and Hamilton (1981) suppose that the players act simultaneously. This means that in each round of the game, Ann and Bob must make their choices independently and their decisions are revealed to the other *before* the next round. The payoffs are accumulated after each round and handed out to the players at the end of the game, i.e., after Round N .

Table 5.3: Payoff matrix of the iterated prisoners' dilemma.

Ann	Bob	
	Cooperate	Defect
Cooperate	(3, 3)	(0, 5)
Defect	(5, 0)	(1, 1)

Suppose that Ann and Bob know N , i.e., the number of consecutive rounds. Ann has nothing to lose in Round N . That is, she has already accumulated some amount of utils and she knows that her choice in Round N has no influence on Bob's decision in the same round. Since Defect strictly dominates Cooperate, she will defect and the same argument holds true for Bob. If Ann believes that Bob behaves in the same manner, it makes no sense for her to cooperate in Round $N - 1$, etc.⁹² Hence, the players will defect right from the start.

What happens if Ann and Bob know that the game will end *not later* than after Round M ? This means that M is an upper bound for N . In this case, Ann will ponder like this: "If we are in Round M , I will defect and the same holds true for Bob. If we are in Round $M - 1$, it makes no sense to cooperate either because Bob will defect anyway in Round M , provided that the game will not already stop after Round $M - 1$. Thus, Bob will not cooperate in Round $M - 1$, too, etc."

The problem is that cooperation makes sense for a player only if he thinks that this has a positive impact on the behavior of his opponent in one of the *next* rounds. However, this condition is not satisfied if $M \geq N$ is common knowledge. Actually, we even need not assume that the players *know* that M is an upper bound for N . If it is common *belief* that M is a maximum number of (consecutive) rounds, nobody will ever cooperate. This means that everybody will defect from start to finish.

Axelrod and Hamilton (1981) conducted a computer tournament between game theorists in order to find out which strategy works best. The authors had revealed the

⁹² This argument requires at least bounded knowledge. I will come back to this point in Section 7.1.1.

number of rounds (200) to the players before the game started. Nonetheless, categorical defection was *not* the best strategy! The actual recommendation based on the results of this tournament is unambiguous: You should play tit for tat. More precisely, start to cooperate and then cooperate if and only if the other player has cooperated in the previous round. On the long run, this simple strategy beats even the most sophisticated ones.

This statement is purely prescriptive. In fact, the overall approach of Axelrod and Hamilton (1981) is evolutionary and thus it refers to natural selection. Their aim is to point out what a species *should* do in order to survive. Hence, rational individuals should not, and they also *do* not, behave as suggested by traditional game theory. We can readily explain their behavioral pattern from a subjectivistic point of view. The subjectivistic approach does not presume common knowledge and not even common belief.⁹³ For this reason, it does not imply that the players defect, categorically, even if they know the end of the game. The subjectivistic explanation of the iterated prisoners' dilemma is provided in the following section.

5.5.2 General Form of the Dilemma

From now on, I consider the prisoners' dilemma in its general form. This is represented by the payoff matrix in Table 5.4. The payoff of mutual cooperation amounts to C , whereas the payoff of mutual defection is given by D . It is assumed that $0 \leq D < C < 1$. Savage's representation theorem guarantees that the utility functions of the players are unique up to any positive affine transformation and so the payoff of single cooperation, 0, and the payoff of single defection, 1, are chosen in this way only for the sake of simplicity but without loss of generality.

Table 5.4: Payoff matrix of the prisoners' dilemma with $0 \leq D < C < 1$.

Ann	Bob	
	Cooperate	Defect
Cooperate	(C, C)	(0, 1)
Defect	(1, 0)	(D, D)

In the Axelrod-Hamilton experiment we have that $C = \frac{3}{5}$ and $D = \frac{1}{5} > 0$. Nonetheless, our general setting allows us to analyze the prisoners' dilemma also with $D = 0$, in which case Defect dominates Cooperate only in the weak sense. For example, in Split or Steal (see Section 4.5.2) we have that $C = \frac{1}{2}$ and $D = 0$.

⁹³ Common knowledge and belief will be explained in Section 7.1.1 and in Section 8.2.1, respectively.

Suppose that the payoffs are paid out after each round and accumulated until the end of the game. In this case, it is often assumed that the players' reward for mutual cooperation (C) is greater than the average of the payoff that a player receives if he defects alone (1), which is referred to as the "temptation payoff," and the payoff in the case in which he cooperates alone (0), i.e., the so-called "sucker's payoff." Put another way, it must hold that $C > \frac{1}{2}$, since otherwise it could be better for the players to cooperate and defect, alternatingly.

We may readily assume that the number of rounds is not known to the players in advance. However, from a mathematical point of view, this makes the usual construction of the iterated prisoners' dilemma somewhat problematic: We cannot accumulate payoffs until infinity. More precisely, let u_r be Ann's payoff after Round $r \in \mathbb{N}$. She does not know the final number of consecutive rounds, N , which means that she considers the end of the game an *event*. Now, if she does not expect that the game will end after any finite number of rounds, there would be some cases in which she receives the accumulated utility $\sum_{r=0}^{\infty} u_r$. In general, this series does not converge (to a finite number) and so it is not well-defined.

In order to guarantee that the overall construction makes sense, we should allow the players to become tired. Let $0 < \phi < 1$ be a parameter that measures Ann's stamina. The larger ϕ , the more she can enjoy later payoffs, but the smaller ϕ , the more Ann prefers to receive her payoffs at an early stage of the game. After Round r she receives only $\phi^r < 1$ times the corresponding payoff that is given in Table 5.4. Hence, ϕ^r is a discount factor, which assigns each payoff a "time value."

At the end of the game, which can be anytime, Ann has a payoff of

$$\sum_{r=0}^{\infty} \phi^r u_r \leq \frac{1}{1-\phi} < \infty$$

and so her (final) utility is bounded.⁹⁴ This formula presumes that Ann's payoff, u_r , vanishes for each $r > N$.⁹⁵ Bob has his own stamina, φ , which may differ from Ann's stamina ϕ , and thus his (final) utility is bounded by $1/(1-\varphi)$.

In order to make the dilemma somewhat more tangible, from now on I assume that the payoffs are *not* accumulated but paid out only at the end of the game. More precisely, the payoffs of Ann and Bob are based on their *final* decisions.⁹⁶ Thus, we may suppose that $0 < \phi, \varphi \leq 1$, i.e., her or his stamina may equal one, in which case the player does not become tired at all. Further, we will no longer assume, explicitly, that $C > \frac{1}{2}$ in order to analyze the game without prejudice.

⁹⁴ This is also a necessary condition for the representation theorem to hold true (see Section 1.4.2).

⁹⁵ For example, if Ann believes that $N = n$, we have that $\sum_{r=0}^{\infty} \phi^r u_r = \sum_{r=0}^n \phi^r u_r$.

⁹⁶ However, this does not mean that Ann and Bob make their choices independent of the evolution of the game, i.e., the previous actions of their adversary.

How will Ann act and react in this game? Put another way, which strategy will she perform? Well, this depends essentially on *Bob's* strategy. His behavior, which can be observed by Ann during the game, has an essential impact on her subjective probabilities regarding Bob's *forthcoming* actions. Thus, Bob's strategy has an influence on Ann's strategy and vice versa. Since the decision trees of Ann and Bob contain paths that do not lead to any end node, we must solve the iterated prisoners' dilemma by forward deduction.

5.5.3 Forward Deduction

I will compare only the following strategies with each other⁹⁷:

1. **Cooperate** from now on and forever.
2. **Defect** from now on and forever.
3. **Tit for tat**: Cooperate now and repeat the previous action of the opponent in each forthcoming round.
4. **Tat for tit**: Defect now and repeat the previous action of the opponent in each forthcoming round.

The last two strategies differ from one another only in that Tit for tat proposes cooperation, whereas Tat for tit proposes defection in the first round.

The players may change their strategies whenever they want and it is clear that they take only the forthcoming actions into account when making up their minds. Figure 5.11 shows the possible evolutions of the game, depending on the strategies of Ann and Bob. Our investigation may start from any round $r \in \mathbb{N}$.

Suppose that Ann thinks that the game continues with probability $0 \leq \pi < 1$ after each round. In the special case of $\pi = 0$, she believes that the game is one shot. If she decides to cooperate, categorically, and believes that Bob does the same, her expected utility amounts to

$$\begin{aligned} EU_{\text{Ann}}(\text{Cooperate}) &= (1 - \pi)C + \pi(1 - \pi)\phi C + \pi^2(1 - \pi)\phi^2 C + \dots \\ &= (1 - \pi)C \sum_{r=0}^{\infty} (\pi\phi)^r = \frac{1 - \pi}{1 - \pi\phi} C. \end{aligned}$$

By contrast, if she defects, we have that

$$\begin{aligned} EU_{\text{Ann}}(\text{Defect}) &= (1 - \pi)1 + \pi(1 - \pi)\phi 1 + \pi^2(1 - \pi)\phi^2 1 + \dots \\ &= (1 - \pi) \sum_{r=0}^{\infty} (\pi\phi)^r = \frac{1 - \pi}{1 - \pi\phi}. \end{aligned}$$

⁹⁷ Since we can imagine an infinite number of strategies, this analysis is incomplete, but I think that our main conclusion does not change after taking any additional strategy into account.

		Cooperate			Defect			Tit for tat			Tat for tit		
Cooperate	Ann	c	c	c	c	c	c	c	c	c	c	c	c
	Bob	c	c	c	d	d	d	c	c	c	d	c	c
Defect	Ann	d	d	d	d	d	d	d	d	d	d	d	d
	Bob	c	c	c	d	d	d	c	d	d	d	d	d
Tit for tat	Ann	c	c	c	c	d	d	c	c	c	c	d	c
	Bob	c	c	c	d	d	d	c	c	c	d	c	d
Tat for tit	Ann	d	c	c	d	d	d	d	c	d	d	d	d
	Bob	c	c	c	d	d	d	c	d	d	d	d	d

Figure 5.11: Evolutions of the game, where “c” stands for “cooperate” and “d” for “defect.” The rows refer to Ann’s strategies and the columns represent Bob’s strategies.

Further, if she plays tit for tat, her expected utility is the same as for Cooperate, i.e.,

$$EU_{\text{Ann}}(\text{Tit for tat}) = \frac{1-\pi}{1-\pi\phi}C.$$

Finally, if Ann defects in the current round and then moves on playing tit for that, i.e., if she chooses Tat for tit, we have that

$$\begin{aligned}
 EU_{\text{Ann}}(\text{Tat for tit}) &= (1-\pi)1 + \pi(1-\pi)\phi C + \pi^2(1-\pi)\phi^2 C + \cdots \\
 &= (1-\pi)(1-C) + (1-\pi)C \sum_{r=0}^{\infty} (\pi\phi)^r \\
 &= (1-\pi)(1-C) + \frac{1-\pi}{1-\pi\phi}C.
 \end{aligned}$$

Note that

$$\frac{1-\pi}{1-\pi\phi} > (1-\pi)(1-C) + \frac{1-\pi}{1-\pi\phi}C > \frac{1-\pi}{1-\pi\phi}C$$

whenever $\pi > 0$ and so we obtain the following preference order

$$\text{Defect} > \text{Tat for tit} > \text{Tit for tat} \sim \text{Cooperate},$$

given that Ann believes that Bob cooperates. Thus, Ann will defect in this case.

It can very well happen that Bob will *not* cooperate in Round 0 or in any forthcoming round. This means that Ann was wrong in her assessment at the beginning of the

game, i.e., she is surprised about Bob's behavior, in which case she will simply change her mind about Bob. The same principle applies if Ann initially expected a one-shot game but discovers after Round 0 that the game is not over. In any case, Ann will update her subjective probabilities about all subsequent events, and it is clear that the same arguments hold true, *mutatis mutandis*, for Bob. Hence, the iterated prisoners' dilemma might be an irregular decision problem both from Ann's perspective and from Bob's perspective. However, as already discussed in Section 2.8.2, this poses no real challenge for us.

Similar calculations lead us to Table 5.5, which contains Ann's expected utilities, depending on the supposed strategy of Bob. This table reveals that, in general, Ann prefers to defect if she believes that Bob defects, too.⁹⁸ This is because we have that

$$\frac{1-\pi}{1-\pi\phi}D > \frac{(1-\pi)\pi\phi}{1-\pi\phi}D > 0$$

for $\pi > 0$ and $D > 0$, which implies that

$$\text{Defect} \sim \text{Tat for tit} > \text{Tit for tat} > \text{Cooperate}.$$

We conclude that Ann will decide to defect—not only if she thinks that Bob is going to cooperate but also if she believes that he plans to defect, categorically.

Table 5.5: Ann's expected utilities depending on her and Bob's strategy.

Ann's strategy	Bob's strategy			
	Cooperate	Defect	Tit for tat	Tat for tit
Cooperate	$\frac{1-\pi}{1-\pi\phi}C$	0	$\frac{1-\pi}{1-\pi\phi}C$	$\frac{(1-\pi)\pi\phi}{1-\pi\phi}C$
Defect	$\frac{1-\pi}{1-\pi\phi}$	$\frac{1-\pi}{1-\pi\phi}D$	$B + \frac{1-\pi}{1-\pi\phi}D$	$\frac{1-\pi}{1-\pi\phi}D$
Tit for tat	$\frac{1-\pi}{1-\pi\phi}C$	$\frac{(1-\pi)\pi\phi}{1-\pi\phi}D$	$\frac{1-\pi}{1-\pi\phi}C$	$\frac{(1-\pi)\pi\phi}{1-\pi^2\phi^2}$
Tat for tit	$A + \frac{1-\pi}{1-\pi\phi}C$	$\frac{1-\pi}{1-\pi\phi}D$	$\frac{1-\pi}{1-\pi^2\phi^2}$	$\frac{1-\pi}{1-\pi\phi}D$

$$A = (1-\pi)(1-C), B = (1-\pi)(1-D), 0 \leq D < C < 1, 0 \leq \pi < 1, 0 < \phi \leq 1$$

Now, suppose that Ann thinks that Bob plays tit for tat, in which case Cooperate and Tit for Tat are equivalent. She prefers Tit for tat (or, equivalently, Cooperate) to Defect if and only if

$$\frac{1-\pi}{1-\pi\phi}C > (1-\pi)(1-D) + \frac{1-\pi}{1-\pi\phi}D,$$

⁹⁸ In this case, Tat for tit is just equivalent to Defect (see Figure 5.11).

i.e.,

$$\pi\phi > \frac{1-C}{1-D}.$$

Ann could think also about playing tat for tit, in which case both players are actually playing tit for tat, but their actions are out of phase (see Figure 5.11). It turns out that Tit for Tat is better than Tat for tit if and only if

$$\frac{1-\pi}{1-\pi\phi}C > \frac{1-\pi}{1-\pi^2\phi^2} = \frac{1-\pi}{(1-\pi\phi)(1+\pi\phi)},$$

i.e.,

$$\pi\phi > \frac{1-C}{C}.$$

Note that $\pi\phi < 1$ and so the latter condition can be satisfied only if $C > \frac{1}{2}$.

Further, Tat for tit is better than Defect if and only if

$$\frac{1-\pi}{1-\pi^2\phi^2} > (1-\pi)(1-D) + \frac{1-\pi}{1-\pi\phi}D,$$

which is equivalent to

$$\pi\phi > \frac{D}{1-D}.$$

This condition implies that $D < \frac{1}{2}$. We conclude that

$$\text{Tit for tat} \sim \text{Cooperate} > \text{Tat for tit} > \text{Defect}.$$

The product $\pi\phi \in [0, 1)$ can be interpreted as Ann's planning horizon. Hence, if her planning horizon is big enough and she thinks that Bob plays tit for tat, she will play tit for tat, too. This means that Defect is no longer optimal in this case.

Finally, what happens if Ann believes that Bob performs Tat for tit, which means that he is going to defect in the first round and then to play tit for tat afterwards? In this case, Defect and Tat for tit turn out to be equivalent because if Ann defects in the first round, Bob will defect one round later and so on (see Figure 5.11).

Ann prefers Tit for tat to Defect (or, equivalently, Tat for tit) if and only if

$$\frac{(1-\pi)\pi\phi}{1-\pi^2\phi^2} > \frac{1-\pi}{1-\pi\phi}D,$$

i.e.,

$$\pi\phi > \frac{D}{1-D}.$$

Further, she prefers Cooperate to Tit for tat if and only if

$$\frac{(1-\pi)\pi\phi}{1-\pi\phi}C > \frac{(1-\pi)\pi\phi}{1-\pi^2\phi^2},$$

which is equivalent to

$$\pi\phi > \frac{1-C}{C}.$$

Finally, she prefers Cooperate to Defect (and thus to Tat for tit) if and only if

$$\frac{(1-\pi)\pi\phi}{1-\pi\phi}C > \frac{1-\pi}{1-\pi\phi}D,$$

i.e.,

$$\pi\phi > \frac{D}{C}.$$

We conclude that

$$\text{Cooperate} > \text{Tit for tat} > \text{Defect} \sim \text{Tat for tit}.$$

Hence, if Ann's planning horizon is sufficiently big and she believes that Bob plays tat for tit, she will *not* play tit for tat but cooperate! Given that Ann believes that Bob performs Tat for tit, this is just equivalent to cooperating now, forgiving in the next round, and moving on playing tit for tat.

In order to understand why this behavior makes sense for Ann, we should take a look at Figure 5.11. Cooperation enables her to avoid being trapped in a vicious circle of alternating cooperation and defection or even of mutual defection. This can be done by starting to play tit for tat in the *next* round. Indeed, playing tit for tat is important for Ann because otherwise Bob will tend to defect, categorically, as soon as he observes that Ann cooperates forever. This statement is implicitly based on the fact that we can apply the same arguments to Bob, which can be justified by assuming that also Bob's planning horizon, $\rho\phi$, is big enough, where ρ quantifies his probability that the game continues after each round and ϕ represents his stamina.

The reader can find Ann's preferences and their conditions in Table 5.6. We conclude that the players will play tit for tat, and temporarily forgive whenever they believe that this avoids being trapped in a vicious circle, if

$$\pi\phi > \max\left\{\frac{D}{C}, \frac{D}{1-D}, \frac{1-C}{C}, \frac{1-C}{1-D}\right\} < \rho\phi, \quad (5.1)$$

which implies that $C > \frac{1}{2} > D$.

Equation 5.1 tells us that the payoff of mutual defection, D , must be sufficiently small, whereas the payoff of mutual cooperation, C , must be sufficiently large if we

Table 5.6: Ann's preferences and their conditions.

Preference	Condition	Preference	Condition
Ann believes that Bob cooperates		Ann believes that Bob defects	
Defect > Tat for tit	$\pi > 0$	Defect ~ Tat for tit	–
Tat for tit > Tit for tat	–	Tat for tit > Tit for tat	$\pi, D > 0$
Tit for tat ~ Cooperate	–	Tit for tat ~ Cooperate	$D > 0$
Ann believes that Bob plays tit for tat		Ann believes that Bob plays tat for tit	
Tit for tat ~ Cooperate	–	Cooperate > Tit for tat	$\pi\phi > \frac{1-C}{C}$
Cooperate > Tat for tit	$\pi\phi > \frac{1-C}{C}$	Cooperate > Defect	$\pi\phi > \frac{D}{C}$
Cooperate > Defect	$\pi\phi > \frac{1-C}{1-D}$	Tit for tat > Defect	$\pi\phi > \frac{D}{1-D}$
Tat for tit > Defect	$\pi\phi > \frac{D}{1-D}$	Defect ~ Tat for tit	–

want to guarantee that the players act as described above. Here, it is implicitly assumed that their planning horizons $\pi\phi$ and $\rho\phi$ are fixed. Conversely, if we fix C and D , their planning horizons must be sufficiently big. For example, in the Axelrod-Hamilton experiment the payoffs are $C = \frac{3}{5}$ and $D = \frac{1}{5}$. The reader can easily verify that the critical threshold for the planning horizons is $\frac{2}{3}$ in this case.

If some player tries to defect, once and forever, his opponent will quickly see straight through this manoeuvre and start to defect, too, since in this situation defection is optimal for him. Hence, each player has a credible threat. Thus, he will cooperate if and only if the other one cooperates, too, which means that the players *de facto* cooperate on the long run. Our arguments are based on the assumption that the number of consecutive rounds, which is unknown to the players, is large enough. Moreover, their planning horizons should exceed the critical threshold given by Equation 5.1.

Suppose, for the sake of simplicity, that $\phi = 1$, which means that Ann does not become tired at all. Further, Ann's expectation about the playing time, i.e., the number of additional rounds,⁹⁹ amounts to

$$E_{\text{Ann}}(N+1) = (1-\pi)1 + \pi(1-\pi)2 + \pi^2(1-\pi)3 + \dots = \frac{1}{1-\pi}.$$

In the Axelrod-Hamilton experiment we must guarantee that $\pi > \frac{2}{3}$, which means that Ann will act according to our previous findings if she expects that the game will last more than $1/(1-2/3) = 3$ (additional) rounds.

Hence, in order to guarantee that the players mutually cooperate in the iterated prisoners' dilemma, at least on the long run, we should construct a situation in which they believe that the game takes a long time, i.e., “many rounds.” Additionally, later payoffs must not lose too much attraction compared with earlier payoffs. This goal

⁹⁹ Note that the players already *are* in Round $r \in \mathbb{N}$.

could be achieved by making the time interval between each consecutive round as short as possible. If the time interval is too long, the players will tend to defect.

For example, consider a peace-war game, i.e., an iterated prisoners' dilemma in which we substitute "Cooperate" with "Peace" and "Defect" with "War." Between each consecutive round there is a certain period of time. The players mutually decide to cooperate, i.e., to be peaceful, only if their planning horizons are long enough. To be more precise, they should believe that the game will not come to an end soon and they should also have enough interest in the payoffs that occur on the long run. This implies that the players believe that their enemy has enough time and capacity to strike back. Moreover, they should consider late damages sufficiently harmful. These substantial conditions could be violated if the time interval between one instance to another, after which a player can earliest execute a retaliatory strike, is too long compared with the typical lifetime of the subjects.

I think that these conclusions might appear somewhat frightening to the reader. In fact, strategic conflicts can be very daunting. Thus, it is even *more* important to find a way, i.e., to construct some situation, in which human beings (or other kinds of players) prefer peace to war. In my opinion, this is one of the most important applications of game theory. For more details on the subject matter see, e.g., Schelling (1980).

5.6 Conclusion

Subjectivistic game theory is a straightforward extension of subjectivistic decision theory to situations of conflict. We are already used to solve sequential decision problems by applying decision trees. Thus, solving dynamic games, indeed, turns out to be quite natural and intuitive in the subjectivistic framework. We can solve a dynamic 2-person game in a relatively simple way by treating the game as two separate decision problems and using decision trees at discretion. Of course, we can apply the same principle to dynamic games with more than two players. In most cases we are able to calculate the rational solutions by backward induction.

As we have already seen in the previous chapter, subjectivistic game theory does not presume that a rational subject is interested only in monetary consequences or material incentives. Of course, this holds true for dynamic games, which explains why many players in the ultimatum game refuse to accept some amount of money if they have the impression that they were treated unfairly. The same principle applies to the dictator game, which explains why people give money for charity although, apparently, this makes no sense from an objectivistic point of view.

It is crucial to abandon the static model if the strategic-independence assumption is violated. This can be observed, for example, in the case in which two nuclear powers are facing one another. If we neglect the fact that a country is usually able to strike back after its enemy strikes first, we come to erroneous conclusions. The static model leads us to a solution in which both countries decide to destroy one another, whereas

the dynamic model predicts that they decide to hold still, i.e., it comes to the opposite result, which seems to be much more realistic.

We are able to solve also more complicated dynamic games, in which the players can make their decisions either unconsciously or consciously. For example, in a penalty shoot-out the players think fast, whereas in a normal Chess match they usually think slowly. The subjectivistic Chess theorem explains why it need not be optimal for a rational player to act in a reasonable way. In order to maximize his expected utility, the player might abandon an unbeatable position if he thinks that his adversary plays badly. By contrast, traditional game theory claims that a Chess player should always force at least a draw if he is in an unbeatable position.

Some dynamic games cannot be solved by backward induction. For example, if the players in the iterated prisoners' dilemma think that the game might continue after each round, their decision trees contain paths without any end point. In this case we can make use of forward deduction and come to the conclusion that, on the long run, rational players prefer to play tit for tat if their planning horizons exceed a critical threshold, which depends on the payoffs of the game.

6 Interaction

In this chapter, I assume that the action of a player can have an influence on the action of the other and vice versa. More precisely, the players can act and react to one another *at the same time* and so they are able to *interact*. What distinguishes “interaction” from “action” and “reaction”? Both action and reaction take place unilaterally. This means that the players either act independently or one after another. By contrast, interaction is a *multilateral* phenomenon, i.e., the players coordinate their actions. It is clear that the standard assumption of strategic independence is violated in such a situation. I call games in which the players are able to interact coherent. We cannot use a payoff matrix or a game tree in order to solve a coherent 2-person game. Thus, we will come to conclusions that go far beyond the standard results of traditional game theory.

Suppose that Ann and Bob are facing one another in a one-shot game. In order to keep things simple we may assume, once again, that Ann can choose only between Up and Down, whereas Bob can decide only between Left and Right. Imagine that Ann and Bob can see the action of one another. Moreover, assume that they are able to react at the same time. This means that Ann and Bob can make their decisions *instantaneously*. Since Ann is able to react to the action of Bob at the same time at which Bob performs his action, she *knows* his action when making her own decision. Hence, Bob’s action is part of her private information and so her action can depend on Bob’s. The same argument holds true for Bob. This means that his action represents a reaction to Ann’s action and thus it may depend on her action, too. This situation can be understood only by counterfactual reasoning (see Section 2.3).

Interaction can be explained by the private information sets of Ann and Bob. Consider the scenario matrices of Ann and Bob in Table 4.1. It is assumed that

- Ann has the information set $\mathcal{I}_A = \{s_{A1}, s_{A2}, s_{A3}, s_{A4}\}$ and
- Bob has the information set $\mathcal{I}_B = \{s_{B1}, s_{B2}, s_{B3}, s_{B4}\}$.

That is, both players have perfect information.

Strategic independence was explained in Chapter 4. It results from the fact that both players possess only the trivial information. Hence, no player knows the action of the other. By contrast, in Chapter 5 we analyzed the situation in which one player has the trivial information, whereas the information of the other player is perfect. For example, Ann does not know the action of Bob, whereas Bob knows the action of Ann. In this chapter, we assume that both players have perfect information, i.e., Ann and Bob know the action of one another.

6.1 Coherent vs. Quasicoherent Games

Imagine the following ideal model of a coherent 2-person game: Ann and Bob must make their choices at Time 0 while they can see one another. Hence, Ann can take

<https://doi.org/10.1515/9783110596106-006>

Bob's action into account when making her own choice, whereas Bob is able to react to Ann's action, simultaneously. Since both players act and react *at the same time*, each action can be considered a reaction to the opponent's action.

Coherent games cannot be often encountered in real life because human beings are usually not able to make a choice immediately, i.e., when receiving all (relevant) information. The players typically need some time before they are able to react to their opponent. However, we can frequently observe situations in which the players can act and react to one another an infinite number of times before they come to their *final* conclusions. By making his preliminary decisions, each player reveals a specific behavior, which leads to a situation in which the players *think* that the other performs a certain strategy. I call such games quasicohherent.

For example, suppose that Ann and Bob must make their final decisions at any time $t \geq 0$ before $T > 0$. However, they may choose an interim action, i.e., pick out an element from their own action set, at each time before t . Let $\{X_t\}_{0 \leq t < T}$ be the process of Ann's preliminary decisions. Bob can see Ann's interim action, X_t , at each time $0 \leq t < T$. However, he can react only at any later time point $s \in (t, T)$ and then his decision is only preliminary, too. We may assume that Ann's decision process, $\{X_t\}_{0 \leq t < T}$, is càdlàg, i.e., right-continuous with left limits,¹⁰⁰ and her final decision is the (left) limit of $\{X_t\}_{0 \leq t < T}$ at time T (see Figure 6.1). The same principle applies, mutatis mutandis, to Bob's decision process $\{Y_t\}_{0 \leq t < T}$.




Figure 6.1: Ann's decision process.

To be more precise, Ann's action set is finite and so there exists some point in time $t < T$ such that her decision process $\{X_t\}_{0 \leq t < T}$ is constant on (t, T) . The corresponding value represents Ann's last action, whereas Bob's limiting value marks his last action. The crucial point is that the players can revise their decisions at each time before T . Thus, Ann's last action can be considered a reaction to Bob's last action and vice versa. The game comes to an end at time T , i.e., when the final decisions are fixed. This is illustrated in Figure 6.2.

¹⁰⁰ Correspondingly, the abbreviation càdlàg stands for “continue à droite, limite à gauche.”




Figure 6.2: Decision processes of Ann and Bob.

Alternatively, we can assume that Ann and Bob can choose their (interim) actions in every round $r \in \mathbb{N}$. Each player starts making a choice in Round 0. Then his (preliminary) decision is revealed to the other. In Round 1 the players have the possibility to revise their decisions. Once again, their choices are then revealed to one another, irrespective of whether or not any player has changed his mind. This procedure is repeated in each forthcoming round. It is assumed that there exists some round $s \in \mathbb{N}$ after which the players do not change their minds anymore, i.e., maintain their previous actions. The decisions of the players in Round s are considered final and constitute the payoffs of the game.

In principle, the two models presented above differ only in that the first one takes place in continuous time and the game comes to an end at some predefined time point $0 < T < \infty$, whereas the second one does not refer to time at all and the end of the game remains unspecified. Nonetheless, we could associate each $r \in \mathbb{N}$ with some point in time, e.g.,

$$r \mapsto T\left(1 - \frac{1}{2^r}\right),$$

in which case the game inevitably ends before Time T .

What is the difference between a dynamic game and a quasicoherent game? A quasicoherent game can be considered a hybrid form between a static and a dynamic game. In fact, a quasicoherent game has both a static and a dynamic component:

- The players act simultaneously and thus independently, but
- they also react to one another after some (infinitesimal) delay.

Moreover, the payoffs of the game depend only on the *final* decisions of the players but not on their history of actions. This is precisely the reason why all actions that take place before their final decisions are considered interim.

In the continuous-time model the game ends at some fixed time point $T > 0$, but it does not harm to assume that some instance stops the game at any *unspecified* point in time. In this case, T represents a positive random variable, i.e., a stopping time.¹⁰¹ Similarly, in the iterative model, in which the players make their choices in each round

101 Once again, the last actions of the players *before* the stopping time constitute their final choices.

$r \in \mathbb{N}$, there could be some referee who stops the game after Round $s \in \mathbb{N}$. In any case, the end of a quasicohherent game is not optional to the players.¹⁰²

The reader might have already observed that the iterated prisoners' dilemma in its version described by Axelrod and Hamilton (1981) (see Section 5.5.1) can easily be transformed into a quasicohherent game after some slight modifications:

1. The payoffs are paid out only at the end of the game, i.e., after Round N , and
2. the decisions of Ann and Bob in each round before N are only preliminary, i.e., they do not have any influence on their (final) payoffs.

Precisely these modifications have been applied in Section 5.5.2 in order to solve the iterated prisoners' dilemma by forward deduction.

6.2 Private Information and Interaction

In a quasicohherent game in continuous time, a player can see the interim action of his adversary at each time $t \in [0, T)$. However, he is not able to react to his opponent instantaneously, i.e., at the same time. His reaction can take place only at any later time point $s \in (t, T)$. Hence, does a player *know* the (interim) action of the other at Time t ? The answer is “No”! In decision theory, “knowing” means to be able to make a decision based on the given information.

Since the players in a quasicohherent game can react only after an infinitesimal delay, their information flow is not càdlàg but càglàd, i.e., left-continuous with right limits. Hence, a player does not know what the other player is doing now, but he knows what his opponent has done in the past. To sum up, both decision processes are in fact càdlàg, but Bob's decision process appears to be càglàd from Ann's perspective, whereas Ann's decision process appears to be càglàd from Bob's perspective (see Figure 6.3). Thus, each player has imperfect information at every time $t \in [0, T)$.¹⁰³




Figure 6.3: Ann's information flow (red) regarding Bob's decision process (black).

102 Rebel Without a Cause (see Section 4.2.2) is a continuous-time game in which a player can decide to stop the game at any time by jumping out of the car. Hence, this game is not quasicohherent.

103 This is also the reason why we often assume that the players in a static game act “simultaneously.” We just want to express that they have imperfect information when they make their choices.

Why does it make sense to allow each player in a quasicohherent game to make any preliminary decision at all? Every player has enough time to observe the behavior of his opponent and thus to form his own opinion about the other's strategy. Independent at which time a player performs some action, the other one has always the possibility to respond. This creates a paradoxical situation: No player knows the final decision of the other, i.e., nobody has perfect information, but every player knows that they can respond to one another. Thus, Ann may very well assume that Bob makes his choice dependent on hers and the same holds true for Bob.

Although Ann and Bob can react to one another, their private information sets are still *trivial* because when making their own choices, they do not know whether or not the other changes his mind later on. Thus, no player knows the other's strategy. Nonetheless, Ann and Bob might very well *believe* that the other performs a certain strategy. Moreover, since each player can make a preliminary decision at any occasion, he can influence the other's belief. More precisely, he can make the other believe that he performs some strategy just by *performing* that strategy. Hence, players in a quasicohherent game usually behave *as if* they had perfect information, i.e., they interact—although, in fact, their information is imperfect! The simplifying assumption that the players have perfect information, which distinguishes coherent from quasicohherent games, just serves its purpose in order to solve quasicohherent games, which often appear in real life, in the same way as coherent games.

Before I explain and demonstrate the solution concept, I would like to clarify the meaning of private information in more detail. In order to understand the following arguments, the reader should take a look at Table 6.1, which has already been discussed in Section 4.1. Recall that this table contains Ann's and Bob's scenario matrices of a 2×2 game in which Ann can choose between Up and Down, whereas Bob can choose between Left and Right. Ann can always choose either to go up, categorically, or to go down, categorically, where “always” means “irrespective of her private information set \mathcal{I}_A ” and “categorically” means “independent of whatever Bob does.” The same holds true for Bob, who can always decide to go left or right, categorically. Nonetheless, whether a player can make his own action dependent on the action of the other player depends essentially on his private information.

Table 6.1: Scenario matrices of Ann (left) and Bob (right).

Action	Scenario			
	s_{A1}	s_{A2}	s_{A3}	s_{A4}
Up	Left	Right	Left	Right
Down	Left	Left	Right	Right

Action	Scenario			
	s_{B1}	s_{B2}	s_{B3}	s_{B4}
Left	Up	Down	Up	Down
Right	Up	Up	Down	Down

Suppose that Ann wants to choose Down if and only if Bob chooses Left. It is clear that the scenarios s_{A1} and s_{A4} must be distinguishable for Ann in order to be able to

perform this strategy. This means that they must not belong to the same element of her private information set \mathcal{I}_A . By contrast, it is not relevant to which element of \mathcal{I}_A the scenarios s_{A2} and s_{A3} belong. However, note that Ann cannot perform the desired strategy if $\omega_0 \in s_{A3}$, i.e., if Bob chooses Left if and only if Ann chooses Up: If Ann chooses Down, Bob chooses Right, but in this case Ann chooses Up, not Down, etc. Hence, we always obtain a contradiction.¹⁰⁴

I would like to demonstrate the overall concept by Figure 6.4, which shows two response diagrams. Each one contains two blue and two green response curves. The response curves represent the potential strategies of Bob. Each strategy of Bob is a scenario from the perspective of Ann. More precisely, s_{A1} corresponds to Strategy 1 of Bob, s_{A2} is Bob's second strategy, etc.




Figure 6.4: Private information sets of Ann.

Let us focus on the left-hand side of Figure 6.4. The blue vertical response curve represents the scenario s_{A1} , i.e., Bob's first strategy, and the blue diagonal response curve is associated with the scenario s_{A2} , i.e., Bob's second strategy. Further, the green diagonal response curve represents the scenario s_{A3} and the green vertical response curve is associated with the scenario s_{A4} , which stand for Bob's third and fourth strategy. The given colors on the left-hand side of Figure 6.4 shall indicate that the scenarios s_{A1} and s_{A2} belong to the same element of \mathcal{I}_A . Correspondingly, also the scenarios s_{A3} and s_{A4} are indistinguishable for Ann. That is, Ann's private information set is $\mathcal{I}_A = \{s_{A1} \cup s_{A2}, s_{A3} \cup s_{A4}\}$. By contrast, on the right-hand side of Figure 6.4 Ann's private information set is $\mathcal{I}_A = \{s_{A1} \cup s_{A3}, s_{A2} \cup s_{A4}\}$.

Assume that Bob performs either Strategy 1 or Strategy 2, i.e., one of the two strategies that are marked blue on the left-hand side of Figure 6.4. Then Ann is informed only about the fact that $\omega_0 \in s_{A1} \cup s_{A2}$, in which case she may choose Down. By choosing Down, Bob decides to choose Left and so we come full circle. What happens if Bob performs Strategy 4? Ann is informed only about the fact that $\omega_0 \in s_{A3} \cup s_{A4}$. Now, based on her private information, she can choose Up, in which case Bob chooses Right. To sum up, if Bob chooses Strategy 1 or Strategy 2, Ann chooses Down, which means that

¹⁰⁴ I have already mentioned in Section 3.4 that such a behavior is impossible.

Bob chooses left, and if Bob chooses Strategy 4, Ann chooses Up, which means that Bob chooses Right. Well, this is precisely Ann's desired strategy: "Choose Down if and only if Bob chooses Left."

Now, suppose that Ann's private information set is $\mathcal{I}_A = \{s_{A1} \cup s_{A3}, s_{A2} \cup s_{A4}\}$. This situation is depicted on the right-hand side of Figure 6.4. If Bob performs Strategy 1, Ann can decide to choose Down, and if he performs Strategy 2 or Strategy 4, she may choose Up. Thus, once again, Ann applies Strategy 2, i.e., the diagonal strategy "Choose Down if and only if Bob chooses Left." We could have made the blue diagonal line on the left-hand side of Figure 6.4 green or the green diagonal line on the right-hand side blue. In any case, Ann would still be able to apply her desired strategy. Of course, she could perform the same strategy also if her private information set were even *finer* than $\{s_{A1} \cup s_{A2}, s_{A3} \cup s_{A4}\}$ or $\{s_{A1} \cup s_{A3}, s_{A2} \cup s_{A4}\}$. We have to guarantee only that s_{A1} and s_{A4} do not belong to the same element of \mathcal{I}_A .

Let us turn to the strategy "Choose Up if and only if Bob chooses Left." We can apply the same arguments as before in order to conclude that Ann can perform this strategy if and only if the scenarios s_{A1} and s_{A4} are distinguishable for her. Hence, Ann can make her own action dependent on Bob's action if and only if the scenarios s_{A1} and s_{A4} do not belong to the same element of her private information set \mathcal{I}_A . For example, we could have that $\mathcal{I}_A = \{s_{A1}, s_{A2} \cup s_{A3} \cup s_{A4}\}$. However, we should keep in mind that even if Ann's information set is fine enough, she can perform Strategy 2 if and only if Bob does not perform Strategy 3, whereas she can perform Strategy 3 if and only if he does not perform Strategy 2.

Of course, the same conclusions hold true, *mutatis mutandis*, for Bob. This means that the players can react to one another if and only if they are able to distinguish between the first and fourth scenario in their private information sets. However, this is not sufficient for interaction. Indeed, we have to assume also that Ann always knows the reaction of Bob and vice versa. Suppose that Ann performs Strategy 3, i.e., she goes up if and only if Bob goes left. Consider the left-hand side of Figure 6.4 and assume that Bob chooses either Strategy 1 or Strategy 2, i.e., one of the strategies marked blue. Then Ann goes up and so she is uncertain about Bob's reaction. The problem is that she cannot distinguish between s_{A1} and s_{A2} , which means that Bob could go either left or right. The same problem occurs on the right-hand side of Figure 6.4: If Bob chooses Strategy 2 or Strategy 4, i.e., one of the strategies marked green, then Ann goes down, in which case she does not know whether Bob goes left or right.

The second strategy of Bob, which is marked blue on the left-hand side and green on the right-hand side of Figure 6.4, always prevents Ann from knowing the reaction of Bob if she applies Strategy 3—unless the scenario s_{A2} stands alone in her private information set. Hence, given that Ann applies Strategy 3, she always knows Bob's reaction if and only if the scenario s_{A2} is not combined with any other scenario in her private information set. Correspondingly, if Ann applies Strategy 2, she always knows Bob's reaction if and only if the scenario s_{A3} stands alone in \mathcal{I}_A . Table 6.2 summarizes all possible information sets and their particular meanings from the perspective of Ann.

Table 6.2: Private information sets of Ann and their meanings.

Information set	Meaning
$\{s_{A1} \cup s_{A2} \cup s_{A3} \cup s_{A4}\}$	Ann cannot react to Bob and she never knows his action.
$\{s_{A1}, s_{A2} \cup s_{A3} \cup s_{A4}\}$	Ann can react to Bob, but she knows his action only if he chooses categorically Left.
$\{s_{A2}, s_{A1} \cup s_{A3} \cup s_{A4}\}$	Ann cannot react to Bob and she knows his action only if he performs Strategy 2, in which case she cannot perform Strategy 3.
$\{s_{A3}, s_{A1} \cup s_{A2} \cup s_{A4}\}$	Ann cannot react to Bob and she knows his action only if he performs Strategy 3, in which case she cannot perform Strategy 2.
$\{s_{A4}, s_{A1} \cup s_{A2} \cup s_{A3}\}$	Ann can react to Bob, but she knows his action only if he chooses categorically Right.
$\{s_{A1} \cup s_{A2}, s_{A3} \cup s_{A4}\}$	Ann can react to Bob, but she knows Bob's action only if she chooses Down.
$\{s_{A1} \cup s_{A3}, s_{A2} \cup s_{A4}\}$	Ann can react to Bob, but she knows Bob's action only if she chooses Up.
$\{s_{A1} \cup s_{A4}, s_{A2} \cup s_{A3}\}$	Ann cannot react to Bob and she never knows his action.
$\{s_{A1}, s_{A2}, s_{A3}, s_{A4}\}$	Ann can react to Bob and she always knows his action.

To sum up, (i) the first and fourth scenario must be distinguishable by a player in order to be able to react to the other player and (ii) the second and the third scenario must stand alone in order to be always aware of the (re-)action of the opponent. Hence, both players interact if and only if they have *perfect* information!

At this point it is worth noting that *complete* information, which plays an essential role in traditional game theory and will be described in Section 7.1.1, is not required at all in a coherent game.

I would like to finish this section with Newcomb's paradox, which has been discussed in Section 2.5.2. Although this thought experiment describes a decision problem, not a strategic conflict, it can be analyzed by means of game theory. Hence, suppose that the predictor is a player with perfect information. He knows that the subject's response cannot depend on whether he puts nothing or \$1 million into the box. This means that the predictor's information set is $\mathcal{I}_p = \{s_{p1}, s_{p4}\}$. The basic formulation of the game implies that the predictor performs Strategy 3. Thus, he responds with \$0 if the subject chooses A+B, i.e., if $\omega_0 \in s_{p1}$, whereas he puts \$1 million into Box B if $\omega_0 \in s_{p4}$ (see Figure 6.5). Whether or not this strategy is optimal for the predictor is out of the question.

The subject is not an omniscient being and so he has imperfect information. We may assume that he has the information set $\mathcal{I}_s = \{s_{s1} \cup s_{s2} \cup s_{s4}, s_{s3}\}$. This means that he knows the predictor's reaction only if the latter performs Strategy 3, which is de facto true. Since s_{s1} and s_{s4} are indistinguishable, the subject is not able to react to the predictor. Hence, Newcomb's paradox can be considered a game in which the players are *not* able to interact. However, the subject already knows, a posteriori, that




Figure 6.5: Response curve of the predictor in Newcomb's paradox.

the predictor performs Strategy 3, i.e., he knows that $\omega_0 \in s_{S3}$. Thus, based on his private information, it is clearly optimal for the subject to choose B. In this way, he can force the predictor to put \$1 million into Box B.

Why should the predictor behave in this way if we assume that he is a rational subject? Well, the basic formulation of Newcomb's paradox does not require the predictor to be rational at all. However, this does no longer hold true in strategic games, in which we assume that all players are rational.

6.3 Solving Coherent Games

6.3.1 Response Diagrams

In a static one-shot game, every action represents a (simple) strategy and thus we need not distinguish between “action” and “strategy.” This does no longer hold true in a coherent one-shot game, in which the players have perfect information. That is, they are able both to act and to react at the same time. The action of a player describes how he acts, whereas his strategy tells us how he reacts. The interaction scheme of the players in a coherent game can be analyzed by using response diagrams. They have been shortly described in Section 3.4 and used in the previous section in order to motivate that interaction requires perfect information. In this section, I will show how to apply this important tool in order to solve coherent games.

Each response curve in Figure 6.6 represents an available strategy of Ann. She has four available strategies, which have already been described in Section 3.4:

1. Go up, categorically.
2. Go down if Bob goes left and go up otherwise.
3. Go up if Bob goes left and go down otherwise.
4. Go down, categorically.

Accordingly, Bob's available strategies are as follows:

1. Go left, categorically.
2. Go right if Ann goes up and go left otherwise.
3. Go left if Ann goes up and go right otherwise.
4. Go right, categorically.




Figure 6.6: Response curves of Ann, i.e., Strategy 1 (upper left), Strategy 2 (upper right), Strategy 3 (lower left), and Strategy 4 (lower right).

These strategies are represented by the response curves in Figure 6.7. Remember that each strategy of Ann corresponds to a scenario from the perspective of Bob, whereas each strategy of Bob is a scenario from the perspective of Ann. Note also that a response curve has nothing to do with a *best-response* curve:

- Actually, a response curve is only a set of two points.
- Further, the indicated response of a player need not be a *best* response to the action of the other.

By contrast, a best-response curve indicates the best response of a player, given his subjective probability regarding the action of the other (see, e.g., Section 4.2.3).

By combining the response curves of Ann and Bob, we are able to find the possible solutions of a coherent game. I call each combination of response curves of Ann and Bob a play. The solution of a coherent (2-person) game must always be an intersection point of the two response curves and it can very well happen that some play leads us to two possible solutions or to no solution at all.

For example, if Ann goes categorically down and Bob goes categorically left, we obtain the solution on the lower left of the response diagram that is depicted on the upper left of Figure 6.8. Alternatively, if Ann goes down, categorically, but Bob goes left if and only if Ann goes up, we obtain the solution on the lower right of the response diagram that can be found on the upper right of Figure 6.8. We can easily check why it is impossible to obtain, e.g., the solution on the lower left of the same response diagram: If Ann chooses Down, then Bob chooses Right, not Left. Similar arguments can be applied in order to show that the solutions on the upper left and upper right of that response diagram are impossible, too, given that Ann and Bob perform the given strategies.




Figure 6.7: Response curves of Bob, i.e., Strategy 1 (upper left), Strategy 2 (upper right), Strategy 3 (lower left), and Strategy 4 (lower right).




Figure 6.8: Plays and the corresponding solutions of the game (green points).

We already know that the play that is shown on the lower left of Figure 6.8 leads to no solution at all (see Section 3.4). It is simply impossible for Ann and Bob to interact in such a way, i.e., to perform the diagonal strategies that are depicted on the lower left of Figure 6.8 *together*. Finally, on the lower right of Figure 6.8 we can find two possible solutions, which are marked green. In fact, if Ann goes up, Bob goes right and if he goes right, Ann goes up. Alternatively, if Ann goes down, Bob goes left and if he goes left, Ann goes down. Hence, in this case, there is no contradiction at all. However, in order to find the *rational* solution of the game, we have to specify the payoffs of the

players, which must be unique only up to a strictly increasing transformation. This will be explained in the next section.

First of all, we can make the following general conclusions:

- There are $4^2 = 16$ plays, but two of them are impossible;
 1. The play in which Ann chooses Down if and only if Bob chooses Left and Bob chooses Left if and only if Ann chooses Up as well as
 2. the play in which Ann chooses Up if and only if Bob chooses Left and Bob chooses Right if and only if Ann chooses Up.
- Hence, there remain 14 possible plays. Two of them lead to two possible solutions:
 1. Ann chooses Down if and only if Bob chooses Left and Bob chooses Left if and only if Ann chooses Down as well as
 2. Ann chooses Up if and only if Bob chooses Left and Bob chooses Left if and only if Ann chooses Up.
- The other 12 plays lead to one and only one possible solution.

6.3.2 Rational Solutions

Let us reconsider the game show, which has been discussed in Section 4.2.3. Its payoff matrix is recapitulated in Table 6.3. Now, in contrast to Section 4.2.3, assume that Ann and Bob *can* observe the choice of the other, and that they are able to act and react at the same time. This leads us to a coherent one-shot game, in which the strategic-independence assumption is clearly violated.

Table 6.3: Payoffs of Ann and Bob in the game show.

Ann	Bob	
	Left	Right
Up	(2, 1)	(9, 0)
Down	(1, -1)	(9, 8)

The solution of a coherent game is said to be rational if and only if no player can increase his payoff by choosing another action. Put another way, it is rational if and only if the action of each player is optimal (see Definition 10). However, whether the action of some player is optimal or not depends on the response curve, i.e., the strategy, of his adversary. Hence, it depends on the *play*, and a play is said to be rational if and only if it leads to a rational solution.

Suppose that Ann goes categorically up, whereas Bob goes categorically left. This is depicted by the response diagram on the upper left of Figure 6.9. The solution of this play is indicated by the green point on the upper left of the response diagram, i.e., Ann goes up and Bob goes left. This play is rational because Ann cannot increase

her payoff by moving down and Bob cannot do better by moving right.¹⁰⁵ This means that the corresponding solution is rational.




Figure 6.9: Possible strategies of Ann if Bob chooses categorically Left. Only the play on the upper left is rational.

Now, consider the solution on the lower left of the upper right response diagram in Figure 6.9. This solution is *not* rational:

- Ann knows that she can receive the payoff 2 instead of 1 by moving up and
- Bob knows that he can get 0 instead of -1 by going from left to right.

Since the only possible solution of the play that is depicted in that response diagram is irrational, the given play is irrational, too. Hence, if we assume that Ann and Bob are rational, they will not interact in this way.

The solution on the upper left of the lower left response diagram in Figure 6.9 is not rational either because Bob can increase his payoff by going from left to right. However, Ann has made her optimal choice in this situation. Note that the *same* solution appears on the upper left response diagram in Figure 6.9, but in that case it turns out to be rational! In fact, two different plays can lead to the same solution of the game and whether the solution is rational or not depends essentially on the given response curves. In our case, the play on the upper left of Figure 6.9 is rational, whereas the play on the lower left is irrational.

Finally, the solution on the lower left of the lower right response diagram in Figure 6.9 is irrational, too. In this situation, Ann can do better by moving up and Bob can

105 Note that Ann *knows* that Bob holds still if she changes her mind and the same holds true for Bob.

improve his situation by going right. To sum up, the only rational solution, provided that Bob chooses categorically Left, is that Ann chooses categorically Up.

Now, consider the response diagrams in Figure 6.10, in which Bob chooses categorically Right. The reader can easily verify that the solution in the upper left response diagram is irrational because Bob can do better by moving left. By contrast, the *same* solution on the upper right response diagram turns out to be rational. Indeed, in this case, no player can increase his payoff by moving from Up to Down or from Right to Left. Moreover, it is clear that the solutions on the lower left and lower right response diagrams of Figure 6.10 are rational, too.




Figure 6.10: Possible strategies of Ann if Bob chooses categorically Right. All solutions are rational, except for the solution on the upper left.

In the same way, we can proceed further in order to search for all rational solutions of the game show. Figure 6.11 contains all plays that lead us to a rational solution. All solutions but Down/Left can be rational. Whether a given solution is rational or not depends on how the players interact, i.e., on the particular play, but we cannot find any play of the game show that makes Down/Left rational. The set of rational solutions is indicated by the blue points on the lower right of Figure 6.11.

Now, it should be clear to the reader that we can find a large number of rational solutions of a coherent game. That is, in general, the set of rational solutions of a coherent game is not a singleton, and the number of plays that lead us to a rational solution often exceeds the number of rational solutions of the game. Can it happen that some coherent game has no rational solution at all? In Section 6.6, I will prove that we can always find a rational solution in a finite coherent game—even if the number of players, n , exceeds 2. Further, in most other practical applications, infinite coherent games with a finite number of players possess a rational solution, too.




Figure 6.11: Rational plays of the game show and their rational solutions (lower right). The irrational solutions are marked green.

6.3.3 Refinement

Game theory usually tries to predict the outcome of a strategic conflict. As far as we are concerned with a coherent game, in most cases this goal can be accomplished only by refinement. This means that we have to eliminate all solutions that seem to be implausible—given that the players have perfect information and are rational.

Refinement procedures play a big role not only here but also in traditional game theory, which does not take interaction into consideration. I will come back to those refinement procedures in Section 7.6.

Searching for the solution of a coherent game reminds us a little bit of trying to predict the solution of a cooperative (2-person) game in which the players are able to make an enforceable agreement after a process of bargaining (see, e.g., Nash, 1950a, 1953). Nash's bargaining solution is based on the idea that the players can increase their payoffs by cooperation. Bargaining theory is essentially based on the assumption that the players are able to achieve convex combinations of payoffs. This is done by applying a mixed strategy. I will come back to this point in Section 7.2.

In a coherent game, the players cannot generate any convex combination of payoffs because each player performs one and only one *action*. Moreover, it is hard to imagine that players in a coherent game, which starts and stops at Time 0, are able to bargain. Nonetheless, we will see that cooperation is a typical result of coherent games, too. Further, a coherent game is just a simplification of a quasicoherent game, which can very well be understood as a bargaining process.

However, I will not try to develop a refinement procedure that propagates a *unique* solution for every coherent game, which is typically done when solving bargaining problems. I think that we can solve coherent games more appropriately by a case-by-case analysis and it can very well happen that a coherent game has more than one reasonable solution.¹⁰⁶ The underlying assumption of a coherent game is perfect information, not the possibility to make an enforceable agreement. Indeed, such an agreement can lead to cooperation, but there exist many other situations in which the players in a game that is usually called “noncooperative” are willing to cooperate just because they have perfect information. In these cases, cooperation takes place by self-enforcement. A surprising result of our analysis will be that self-enforcement generally does *not* lead to a Nash equilibrium.

Let \mathcal{R} be the set of solutions of a coherent game that are considered reasonable. It is clear that this set should depend only on the payoffs of the game. Moreover, isomorphic games are identical (see Section 4.3). Thus, after renaming the players, their actions, or applying some positive affine transformation to their payoffs, we must not come to a different conclusion regarding \mathcal{R} . Finally, each reasonable solution should be rational. Hence, our refinement procedure shall satisfy the following criteria:

1. \mathcal{R} should be a function of the payoff matrix of the game.
2. \mathcal{R} should be invariant to any isomorphism, i.e., renaming the players, their actions, or applying some positive affine transformation to their payoffs.
3. Each element of \mathcal{R} should be rational and payoff efficient.

The concept of payoff efficiency goes back to Harsanyi and Selten (1988).¹⁰⁷

Definition 12 (Payoff efficiency, I). A rational solution is payoff efficient if and only if there is no other rational solution in which all players have a higher payoff.

¹⁰⁶ The term “reasonable” often occurs in this book. Its meaning depends on the particular context.

¹⁰⁷ Payoff efficiency will be discussed in more detail in Section 7.6.1.

The rational solution on the upper left of Figure 6.11 is not payoff efficient and thus it can be eliminated. By contrast, all other (rational) solutions that can be found in that figure are payoff efficient. Let us consider the play on the upper right of Figure 6.11. In this situation, Ann receives her best payoff, 9, whereas Bob gets only 0, which is still better than -1 . Is this play reasonable? The answer is “No”! Bob could simply perform Strategy 3 in order to force Ann to choose Down, in which case he receives his best payoff, i.e., 8. That is, Strategy 3 is Bob’s favorable strategy.

Conversely, Ann’s favorable strategies are Strategy 2, Strategy 3, and Strategy 4. Hence, she can always force Bob to choose Right unless she performs Strategy 1. If Bob chooses his favorable strategy, i.e., Strategy 3, Ann cannot choose Strategy 2. Thus, we may assume that she performs either Strategy 3 or Strategy 4. In fact, we have just eliminated 5 out of the 7 rational plays that are depicted in Figure 6.11. Finally, why should Ann perform Strategy 4, i.e., choose Down, categorically? She could simply go up if Bob goes left in order to get 2 instead of 1.

To sum up, both Ann and Bob will perform Strategy 3. This means that the only reasonable play, i.e., play that leads us to a reasonable solution, is that Ann goes up if and only if Bob goes left and vice versa. To be more precise, the (only) reasonable solution of the game show is that Ann goes down and Bob goes right. This result is quite intuitive, since the game show represents a coordination game.

The reason why we can ignore all unreasonable solutions of the game show is that both players have perfect information. Simply put, it is evident to Ann and Bob what the other is going to do. This means that they can coordinate their actions and thus act in a coherent way. Hence, the results of coherent games are obtained by self-enforcement, which leads us to interesting solutions and many of them are not considered in traditional game theory. I will illustrate the solution concept in the following sections by a number of examples.

The strategy of a player is called best-response if and only if it assigns the actions of his opponents a best response. The best-response strategies of Ann are Strategy 1 and Strategy 3, whereas Strategy 3 is the (only) best-response strategy of Bob. Whenever the players perform a best-response strategy, the result of the game is a Nash equilibrium. This can typically be observed in coordination games, like the game show. However, Nash equilibria will no longer prevail in other coherent games, whose reasonable solutions can be legitimated by self-enforcement, too.

6.4 Typical 2×2 Games

6.4.1 Anti-Coordination Games

6.4.1.1 Trapped in a Cave

Andy and Bob are two speleologists who are trapped in a cave. A big rock blocks their way out. They can decide either to sit down and wait for things to happen or to stand

up and get rid of the stone. If a player remains trapped, he earns nothing, but if he gets free, he wins 3 utils. Although the rock is quite heavy, one player alone is able to remove it when making an effort. Removing the rock alone, however, costs a player 2 utils. If the other player comes to his aid, they can share their workload and lose only 1 util. Both players act under the same conditions and nobody is in a dependent relationship with the other. The payoff matrix of this game is given in Table 6.4. It shows that this is an anti-coordination game.

Table 6.4: Payoff matrix of the speleologists.

Andy	Bob	
	Stand up	Sit down
Stand up	(2, 2)	(1, 3)
Sit down	(3, 1)	(0, 0)

Traditional game theory presumes that the players make their choices independently. It suggests that there are three possible solutions:

1. Andy sits down and Bob stands up.
2. Bob sits down and Andy stands up.
3. Andy and Bob each flip a coin and decide either to sit down or to stand up, depending on the outcome.¹⁰⁸

Why should Andy and Bob behave like this in such a situation? Andy can see what Bob does if he decides to stand up and it is evident to him what he does also if he chooses to sit down. At some point in time the players must have come to a conclusion and so this game is quasicohherent. The only reasonable solution of this game is that both players help one another, i.e., decide to remove the rock together. This means that they will cooperate, but not because they are good-hearted—the reason is that each player can force the other to cooperate!

Suppose that Andy decides to sit down, categorically. Then the best what Bob can do is to stand up, alone, and we obtain the solution on the lower left of the payoff matrix, which is preferable for Andy. I do not think that this solution is reasonable. The problem is that both players act under equal conditions and thus we could invert our argument in order to conclude that Bob should decide to sit down, categorically, and wait for Andy to stand up. What is wrong with this argument? If both players decide to sit down, categorically, i.e., to perform their favorable strategies, we obtain the solution on the upper left of Figure 6.12. This solution is irrational because Andy and Bob can do better by standing up.

108 If some player has no coin at hand, he can use a small stone.




Figure 6.12: Some plays of the speleologists. The green point on the upper left represents an irrational solution, whereas the blue points indicate rational solutions.

Andy and Bob will choose the strategies on the lower right of Figure 6.12. More precisely, Andy stands up if and only if Bob stands up and Bob stands up if and only if Andy stands up. Hence, they will play tit for tat, in which case the only rational solution is that both players stand up. This solution is not a Nash equilibrium and thus it is widely ignored by traditional game theory. Nonetheless, the players will coordinate their actions, although the payoff matrix in Table 6.4 suggests that this is an anti-coordination game. Interaction can lead to surprising results—at least from the perspective of traditional game theory.

I would like to explain my arguments more precisely. We can ignore each play that leads to the solution in which both players sit down because this solution is payoff inefficient. Moreover, we have seen that no player will sit down, categorically, since then the other player could force him to stand up by sitting down, too. Put another way, sitting down, categorically, represents a favorable strategy both for Andy and for Bob. However, we already know that the solution of this game is irrational if the players perform their favorable strategies *together*.

Interestingly, in this game, the favorable strategies of the players are not their best-response strategies. For example, consider the solution on the upper right of Figure 6.12, where Bob performs his best-response strategy, i.e., Strategy 2, but not his favorable strategy, i.e., Strategy 4. Indeed, the given solution is rational, but it makes not much sense to assume that Bob stands up if and only if Andy sits down, since then Andy could simply achieve his goal just by sitting down, categorically. Thus, performing his best-response strategy is clearly self-defeating for Bob, and the same holds true for Andy, which can be seen on the lower left of Figure 6.12. The only reasonable solution of this game is that the players help one another.

It is worth emphasizing that the preceding arguments are not based on the assumption of common knowledge, which plays a fundamental role in traditional game theory. I even do not presume that the players have complete information.¹⁰⁹ The given arguments are based only on the assumption of perfect information.

6.4.1.2 The Chicken Game

Chicken Game is *the* anti-coordination game per se. It has already been elaborated in Section 4.3.1.1 and the reader can find its payoff matrix in Table 6.5. There we assumed that Andy and Bob are not able to interact. Is this assumption really tenable? Let us suppose just for the sake of simplicity that Andy and Bob are driving towards one another with constant speed. Thus, at Time $T > 0$ they will collide unless somebody swerves at any time before T . Since each player is able to act and to react to the other player an arbitrary number of times before T , this game is quasicohherent. This means that the players are, in fact, able to interact.

Table 6.5: Payoff matrix of the chicken game.

Andy	Bob	
	Swerve	Straight
Swerve	(1, 1)	(0, 2)
Straight	(2, 0)	(-9, -9)

The game starts at Time 0 and, without loss of generality, we may suppose that the final time is $T = 1$. Assume that we are at $t = 0.95$ while both players are on the collision course. Andy can see that Bob goes straight and vice versa. Going straight before Time 1 means to apply the *interim* action Straight. The players can react to one another at any arbitrary time $s \in (0.95, 1)$. Suppose that Andy decides to swerve at time $t = 0.96$, which means that he makes the *preliminary* decision Swerve. Since this decision is preliminary, he can always come back on the track at any time before 1.

Bob appears to be unimpressed by Andy's peace offer. This means that he is still going straight at $t = 0.97$. Andy sees that Bob is stubborn and so he decides at $t = 0.98$ to come back on the track. At $t = 0.99$ Bob becomes nervous and decides to swerve. Now, Andy is in a huff and chooses to go straight. At $t = 0.995$ Andy is still going straight and so Bob comes back on the track at $t = 0.999$. Andy begins to understand that going straight could be dangerous and thus at $t = 0.9995$ he swerves. At $t = 0.9997$ Bob is still going straight and so Andy decides at $t = 0.9998$ to come back on the track, etc. Eventually, each player recognizes that, whenever he decides to go straight, categorically, his adversary comes back on the track. Finally, they start to play tit for tat.

¹⁰⁹ Complete information and common knowledge will be treated in Section 7.1.1.

This means that Andy makes a new peace offer at $t = 0.9999$, which is then accepted by Bob at $t = 0.99995$. Put another way, Andy and Bob swerve and no player comes back on the track after $t = 0.99995$. Hence, the game comes to an end at Time 1 without any collision.

Although the chicken game is quasicohherent, we can explain the behavior of the players *as if* they were playing a coherent game. We have already applied the same principle when solving Trapped in a Cave and we could take many other examples of anti-coordination games into consideration. I think that the quintessence is clear to the reader: In a coherent or quasicohherent anti-coordination game, the players coordinate their actions by playing tit for tat. This means that they finally cooperate and do not harm one another. This result does not require us to specify the players' subjective probability measures at all and also the specific payoffs of the game are not essential.

Our findings are in direct contrast to the results of Chapter 4, where it is shown that cooperation in anti-coordination games *essentially* depends on the priors and utilities of the players. If the players choose their actions independently, as it is assumed in Chapter 4, each one should try to make the other believe that he is prepared to risk anything. Thus, in this case, verbal and nonverbal communication play an important role. By contrast, if the players are able to interact, which is the basic assumption of this chapter, a player just needs to perform his favorable strategy in order to force his opponent to behave in the desired way. In principle, this makes communication superfluous unless we consider interaction some form of communication or signaling (see Section 2.5.3).

6.4.2 Discoordination Games

6.4.2.1 Matching Pennies

Suppose that Ann and Bob are playing Matching Pennies, but now they reveal their actions to one another. Hence, their choices are no longer hidden. Moreover, assume that each player is able to react to the other instantaneously. In this case, the players have perfect information and we obtain a coherent discoordination game, whose simple payoff matrix can be found in Table 6.6.

Table 6.6: Payoff matrix of Matching Pennies.

Ann	Bob	
	Heads	Tails
Heads	(1, -1)	(-1, 1)
Tails	(-1, 1)	(1, -1)

Matching Pennies cannot have a rational solution if both players perform a nondiagonal strategy. For example, if Ann chooses categorically Heads and Bob chooses

categorically Tails, then Ann can do better by moving to Tails. We can always find a better choice for some player if Ann performs a horizontal and Bob performs a vertical strategy. Moreover, Ann cannot perform the strategy “Choose Tails if and only if Bob chooses Heads” if Bob performs the strategy “Choose Heads if and only if Ann chooses Heads” and the same holds true, *mutatis mutandis*, for Bob. Thus, we can eliminate 6 out of 16 plays right from the start.

Now, consider the two response diagrams in Figure 6.13, in which the diagonal strategies of Ann and Bob overlap. In both cases we have two rational solutions, in which both players are indifferent among Heads and Tails: On the left-hand side of Figure 6.13, Bob wins the game, irrespective of whatever he does, whereas Ann cannot lose the game on the right-hand side of Figure 6.13. However, why should any rational player in a zero-sum game with perfect information perform a strategy that enables his opponent to win the game, irrespective of whatever he does? These plays are not reasonable.




Figure 6.13: Overlapping strategies in Matching Pennies.

There are 8 plays left, in which precisely one player performs a nondiagonal strategy. Figure 6.14 contains two of them. On the left-hand side of this figure we can see a rational solution, whereas the right-hand side depicts an irrational one. The reader can verify that only half of the remaining 8 plays are rational. Ann prevails in two of those rational plays and the same holds true for Bob (see Figure 6.15).




Figure 6.14: Rational (left) and irrational solution (right) of Matching Pennies.

Ann's favorable strategy is Strategy 3 and Bob's favorable strategy is Strategy 2. The upper part of Figure 6.15 contains the two plays in which Ann performs her favorable




Figure 6.15: Rational plays of Matching Pennies. Ann prevails in the upper part, whereas Bob prevails in the lower part.

strategy, i.e., “Choose Heads if and only if Bob chooses Heads,” and so she wins the game. In both cases, Bob capitulates and chooses either Left or Right, categorically. The lower part of Figure 6.15 shows the two plays in which Bob performs his favorable strategy, i.e., “Choose Heads if and only if Ann chooses Tails,” whereas Ann gives up and chooses either Up or Down, categorically.

Matching Pennies is a typical example in which we cannot find a unique reasonable solution. In fact, all solutions of this coherent game are reasonable. However, some player must win the game. It is precisely that player who is more assertive, but without any further information, we cannot say whether this is Ann or Bob. In its coherent version, Matching Pennies is a power play: By performing the favorable strategy, a player can keep his adversary from performing his own favorable strategy. In contrast to an anti-coordination game, it is impossible that both players perform their favorable strategies together, since this would produce a contradiction.

6.4.2.2 The Cat-And-Mouse Game

Let us come back to the cat-and-mouse game, which has already been considered in Section 4.3.2.2. In that section, we assumed that Andy and Betty cannot interact. For example, we could have supposed that Andy and Betty choose their locations in advance and do not change their mind when they arrive at the corresponding place.

In this section, we assume that Andy and Betty are able to act and to react at any time before the game comes to an end. For example, when Betty sees Andy in the bar, she can immediately go back to the club, etc.¹¹⁰ This means that Betty has not made her

¹¹⁰ We may suppose that Bob does not recognize that Andy meets Betty unless they finally decide to stay in the bar.

final choice as long as she does not keep staying in the bar or in the club, respectively. Of course, the same holds true for Bob. For the sake of simplicity, we may suppose also that, at each time before the game comes to an end, Andy and Betty must either be in the bar or in the club. Thus, if Andy does not see Betty in the bar, then he knows that she must be in the club, etc. This quasicohherent game ends at some time $T > 0$, when both players have made their final decisions.

The payoff matrix of this game can be found in Table 6.7. Once again, there is no rational solution in which no player performs a diagonal strategy. Further, why should Andy try to escape from Betty or Betty try to meet Andy? Hence, the plays in which Andy performs Strategy 2 or Betty performs Strategy 3 are not reasonable. Only 3 out of the 4 remaining plays, in which precisely one player performs a nondiagonal strategy, are rational (see Figure 6.16).

Table 6.7: Payoff matrix of the cat-and-mouse game.

Andy	Betty	
	Bar	Club
Bar	(1, -2)	(0, 0)
Club	(0, 1)	(2, -1)




Figure 6.16: Plays of the cat-and-mouse game. Betty prevails on the top, whereas Andy prevails on the bottom. However, the play on the lower left is irrational.

On the top of Figure 6.16, Betty wins the battle. She just keeps running away from Andy, who gives up at the end either by staying in the bar or in the club. In fact, he is indifferent among the two locations because he loses anyway. By contrast, Andy

prevails on the bottom of Figure 6.16, but the play on the lower left is irrational and thus it will not manifest. If Andy prevails, Betty gives up and stays in the club. She will *not* stay in the bar in order to avoid a confrontation with Bob. The more persistent player wins the cat-and-mouse game.

We have solved the game as if it were coherent. Alternatively, we could have presented the overall situation in its quasicohherent form, in which the players can act and react at any time before T : Betty goes to the bar and observes that Andy is waiting for her. She immediately escapes to the club. Andy observes that Betty has escaped to the club and so he follows her. Betty sees that Andy comes to the club and so she runs to the bar. After arriving at the bar, Betty observes that Andy is back again, etc.

Who wins the cat-and-mouse game? It is the player who is more assertive. Recall that the decision processes of the players in a quasicohherent game are càdlàg. Hence, the cat-and-mouse game cannot continue, on and on, until Time T . This means that the players must come to a final conclusion at some time before T .¹¹¹ However, without making any further assumption, it is impossible to judge whether Andy or Betty eventually gives up. For example, by modifying the payoffs of the players, we could take fatigue and surfeit into account. Then it can happen that some player prefers to stop playing cat and mouse at some point in time just because running away produces too much frustration and pain.¹¹²

Thus, suppose that Ann is in the club after running away a large number of times and observes that Bob, once again, has followed her. If she gives up, she receives the payoff -1 . Bob is following her right from the start and so Ann expects that the game will continue in the same way. Now, she is thinking about giving up or not. She expects that she will come back to the same situation if she runs away to the bar again. Assume that she loses any amount of utils each time she is back to square one. In this case, she will stop at the club.

The quintessence of this section is that, in coherent discoordination games, we typically cannot find just one reasonable play that leads to a rational solution. The problem is that each player can try to adapt his action to the action of his adversary in order to obtain a better outcome. Finally, some player backs down. However, without any further arguments, we are not able to predict which player eventually becomes tired and concedes victory to the other.

The reader might think that a quasicohherent model is more appropriate than a coherent model to describe the overall situation. I do not really disagree. However, I think that the principal result will not change at all. A coherent model just enables us to apply response diagrams in order to solve the game. Using response diagrams can be considered a shortcut whenever it is clear to us that the players are able to act and

¹¹¹ In fact, if we drop the càdlàg assumption, there is no solution at all.

¹¹² We already used this technique in order to solve the iterated prisoners' dilemma (see Section 5.5.3).

react to one another, which means that they reveal their behavioral patterns, until the game finally comes to an end and the payoffs are handed out to the players.

6.4.3 Coordination Games

We already know that players in a coherent anti-coordination game with perfect information coordinate their actions. Hence, it should not be surprising at all that players coordinate their actions, a fortiori, in a coordination game with perfect information. However, “coordination” in a coherent anti-coordination game means that the players do *not* perform their best-response strategies, whereas the opposite is true in a (coherent) coordination game, which shall be demonstrated in this section.

6.4.3.1 The Win-Win Situation

In Section 4.3.3.2 we have discussed a win-win situation in which Ann and Bob try to meet one another in a restaurant. The problem is that they cannot see where the other is going and so it can happen that they miss one another. This situation typically occurs if Ann and Bob have incomplete information, in which case a player might not know whether his partner prefers the one or the other restaurant.

Now, suppose that Ann and Bob each have a cell phone that is equipped with GPS tracking and both are able to track the GPS signal of the other.¹¹³ This means that they can see the direction in which the other is running. It is intuitively clear that Ann and Bob will meet in the trattoria, even if their information is incomplete. In fact, a payoff-efficient (rational) solution can be found only on the upper left of the payoff matrix in Table 6.8. Figure 6.17 contains all plays that make that solution, i.e., Palermo/Palermo, rational.

Table 6.8: Payoff matrix of the win-win situation.

Ann	Bob	
	Palermo	Butcher
Palermo	(3, 2)	(0, 0)
Butcher	(0, 0)	(1, 1)

In this game, it makes no sense to assume that some player eschews his best-response strategy. Hence, the only reasonable play is on the lower right of Figure 6.17. That is, Ann will choose the trattoria if and only if Bob goes there and vice versa. Since both

¹¹³ This example might seem artificial, but it is not too farfetched. Nowadays, GPS tracking can be done with appropriate software applications. However, the technical details shall not bother us here.




Figure 6.17: Rational plays of the win-win situation.

players prefer the same restaurant, they will meet in the trattoria. We can easily imagine that this rational solution occurs in real life without much fuss.

6.4.3.2 Battle of the Sexes

Now, reconsider the battle of the sexes, whose payoff matrix is given in Table 6.9. Assume that the concert hall is in the North, whereas the football stadium can be found in the South of the city. Ann and Bob are just in the middle of the city and they can track the other's GPS signal.

Table 6.9: Payoff matrix of the battle of the sexes.

Ann	Bob	
	Concert	Football
Concert	(3, 2)	(1, 1)
Football	(0, 0)	(2, 3)

We can immediately discard all plays that are irrational and that lead to a payoff-inefficient solution, i.e., Football/Concert and Concert/Football. There remain only 4 rational plays (see Figure 6.18), but the plays on the upper left and lower right of Figure 6.18 are not reasonable. Why should Bob decide to go to the concert, categorically, if Ann goes categorically to the concert? In this case, they will anyway meet at the concert hall and so Bob can perform his best-response strategy, i.e., Strategy 3. The same conclusion holds true, *mutatis mutandis*, for Ann. The only reasonable solutions of the battle of the sexes can be found on the upper right and lower left of Figure 6.18.




Figure 6.18: Rational plays of the battle of the sexes.

Let us discuss in more detail why it will not happen that both Ann and Bob adhere to their favorable strategies. More precisely, why will Ann not categorically go to the concert hall and Bob not categorically go to the football stadium? The problem is that then they would end up with the solution given in Figure 6.19, i.e., Concert/Football, which is irrational because both Ann and Bob could do better by changing their actions. Thus, if both players perform their favorable strategies, the solution of this game is irrational.¹¹⁴




Figure 6.19: Irrational solution of the battle of the sexes.

Suppose that Bob notices that Ann is approaching the concert hall although he is still walking to the football stadium. Does it help for Bob to continue in the current direction? Well, he could try to force Ann to follow him. However, if Ann maintains her strategy, i.e., if she keeps staying on her track, Bob will eventually turn around and

¹¹⁴ Note that we already observed the same phenomenon in Section 6.4.1, i.e., when analyzing coherent anti-coordination games.

go back to the concert hall because it is not preferable for him to go to the football match alone. Of course, the same argument holds true, *mutatis mutandis*, for Ann. We conclude that the player with the strongest nerves wins the game unless there are no other factors that determine the outcome.

Hence, we cannot say which place Ann and Bob will choose without making any further assumption. The result of this game might depend on many factors that go beyond the payoff matrix. For example, it could be a long tradition in the relationship of Ann and Bob to go to the football match. In this case, the football stadium represents a focal point (Schelling, 1980), in which Ann and Bob decide to choose Football out of habit. Then we obtain the solution on the lower left of Figure 6.18. Alternatively, Ann could have adopted a dominant position in her relationship with Bob, in which case Bob caves in and follows her to the concert hall. This leads us to the solution on the upper right of Figure 6.18.

In any case, the player who performs his favorable strategy succeeds.¹¹⁵ Persistence pays off in this coordination game and there can be only one player who prevails. Interestingly, although this is a coordination game, i.e., the players have similar interests, it is still a strategic conflict. Their interests are not identical and so each player tries to force his adversary to renounce his favorable strategy.

6.5 The Prisoners' Dilemma

6.5.1 The Classic Dilemma

Here we come back to the prisoners' dilemma, which has already been discussed as a static game in Section 4.5.1, whereas in Section 5.5 we treated iterated versions of the prisoners' dilemma. Here, we revive the classic dilemma, in which Mary and Joe must decide either to deny or to confess. Hence, I refer to the penalty matrix in Table 6.10 and the game is one shot.

Table 6.10: Penalty matrix of the prisoners' dilemma.

Mary	Joe	
	Deny	Confess
Deny	(1, 1)	(5, 0)
Confess	(0, 5)	(4, 4)

Now, Mary and Joe are sitting in a room together with their lawyer, who is instructed to announce their testimonies in court. He asks the prisoners to make their choices.

¹¹⁵ It is implicitly assumed that the players make their decisions in a timely manner, so that they do not miss the corresponding event.

In this situation, each player is able to verify the action of the other, i.e., the other's action is *evident* to him. Suppose that Mary proposes to deny, but Joe rejects her offer and wants to confess. Why should Mary accept Joe's reply? It is clearly better for her to change her mind, i.e., to confess, too.¹¹⁶ Joe observes that Mary has changed her mind because he had chosen to confess. Thus, he could try now to deny. Mary could still reject his offer, but then he can simply change his mind, too, etc. In fact, he has nothing to lose and the same holds true for Mary.

We could imagine that the discussion between Mary and Joe goes back and forth a couple of times, but eventually they will understand that the other prisoner is playing tit for tat: Mary confesses if and only if Joe confesses and vice versa. The reason why they act in such a way is that both Mary and Joe have a credible threat. Hence, the players interact and the dominance principle fails entirely to explain the rational solution of this quasicohherent game: Both players will finally deny!

For a couple of years I play this game in my lectures on game theory. I have observed that the students *always* decide to deny after revealing myself as a lawyer, and they usually make their decisions almost instantly: They *immediately* agree on the cooperative solution! As is indicated by Axelrod and Hamilton (1981), mutual cooperation seems to be an instinctive behavior of human beings and other species.

The reader might wonder why traditional game theory has so much difficulties to explain the cooperative solution of the prisoners' dilemma. The reason is simple: It considers the prisoners' dilemma static but not coherent or, at least, quasicohherent. However, the static model makes no sense in a situation in which the players know that they are able to act and react to one another at any time before the game comes to an end. Put another way, they are able to interact.

In fact, Mary and Joe cooperate if there is a lawyer in the prisoners' dilemma, but not primarily because they have the possibility of a binding agreement. This is just one of many possible instruments that guarantee that players cooperate. Mary and Joe cooperate because each one anticipates that the other defects if and only if he defects. That is, cooperation evolves only because Mary and Joe *believe* that the other plays tit for tat. It does not matter at all whether or not their decisions are binding. Hence, when the game is over, the lawyer might be recognized as a fake. Nonetheless, Mary and Joe would still have cooperated. Once again, it is worth emphasizing that the players cooperate not because they form any coalition in the sense of cooperative game theory but because they recognize that defection is self-destructing.

This shall be clarified by another example: Suppose that Ann is Mary's best friend in jail, whereas Joe's best mate in prison is Bob. Joe has bribed Ann in order to find out Mary's intentions and he is convinced that she tells him the truth. Well, Mary has engaged her own informer: Bob. Joe decides to confess and speaks about his choice with Bob while they are strolling through the prison yard. Afterwards Bob goes to Mary

116 Note that in the given situation no player has made any final decision, yet.

and tells her that Joe is going to confess. For this reason, she decides to confess, too, which is later on revealed by Ann to Joe, etc.

The players recognize that the other is able to react, but they do not know why and so they have no clue how to betray one another. Otherwise, it would be easy for a player to conceal his own intentions. Indeed, a rational prisoner always prefers to defect if he believes that his own action has no impact on the other's action.¹¹⁷ However, the results of Axelrod and Hamilton (1981), which are discussed in Section 5.5.1, and also our general analysis carried out in Section 5.5.3 suggest that Mary and Joe will play tit for tat and so they will, finally, cooperate.

Now, let us solve the prisoners' dilemma with lawyer by means of response curves. In this game, there exist only two rational plays, which are depicted in Figure 6.20. The right-hand side of the figure provides the noncooperative solution, which is typically advocated in traditional game theory, whereas the left-hand side contains the cooperative solution. In fact, both solutions are rational, but only the cooperative one is payoff efficient. Hence, the only reasonable play can be found on the left-hand side of Figure 6.20. This means that Mary and Joe play tit for tat and so they deny!




Figure 6.20: Rational plays of the prisoners' dilemma with lawyer. Only tit for tat (left) is a reasonable play. The numbers in parentheses are levels of penalties.

On the right-hand side of Figure 6.20, Mary and Joe perform their best-response strategies and we have seen that this is unreasonable. By contrast, on the left-hand side of that figure, they perform their favorable strategies. In fact, by performing Strategy 2 or Strategy 4, Mary cannot force Joe to deny in order to receive her best payoff, i.e., to be released from prison. However, at least she can force him to deny by performing Strategy 3, in which case she will deny, too. Hence, her favorable strategy is Strategy 3 and the same holds true for Joe.

We conclude that cooperation is not (only) a matter of the payoffs. Indeed, the penalty matrix of the prisoners' dilemma is always the same, irrespective of whether there is a lawyer or not. This demonstrates that analyzing a game by its payoff matrix can be misleading if the players do not choose their actions independently. Hence,

¹¹⁷ For this reason, communication alone does not solve the basic issue (see Section 4.5.1).

the payoff matrix is not an appropriate tool for solving the prisoners' dilemma with lawyer and, due to the same reason, the dominance principle fails to explain the cooperative solution of this game. Further, if the players are able to interact, the prisoners' dilemma cannot be analyzed by a game tree either. The problem is that they can respond to one another at any time before the game comes to an end, which cannot be described in an appropriate way by a game tree. This becomes even more obvious if we treat the prisoners' dilemma as a coherent game. Nonetheless, solving a coherent, or a quasicohherent, game is relatively simple by means of response diagrams.

6.5.2 Split or Steal

Once again, Ann and Bob are sitting in front of one another playing Split or Steal (see Section 4.5.2), but now we assume that they must reveal their choices before the game comes to an end. Hence, this variant of Split or Steal is quasicohherent. The payoff matrix of this game can be found in Table 6.11.

Table 6.11: Payoff matrix of Split or Steal.

Ann	Bob	
	Split	Steal
Split	(1, 1)	(0, 2)
Steal	(2, 0)	(0, 0)

In Section 4.5.2 it was assumed that the players must make their choices independently and we have come to the conclusion that a player tries to steal the jackpot whenever he thinks that the other decides to split with positive probability. The game typically ends up in a situation in which both players decide to steal. Of course, this requires us to assume that each player is driven only by monetary considerations and, in order to keep things simple, we shall maintain the homo-economicus assumption in this section. Nonetheless, now it turns out that the players decide to split the jackpot because their choices are evident!

Figure 6.21 contains the rational plays of this game. The rational play on the lower right of the figure, in which Ann and Bob decide to steal the jackpot, is payoff-inefficient and thus it can be discarded. Further, the rational play on the upper left of the figure is not reasonable, too: Why should Bob decide to split the jackpot, categorically, if he knows that Ann steals it, categorically? It is better for him to move to Steal if she decides to split the jackpot, which is depicted on the upper right of Figure 6.21. It is clear that Ann accepts this solution because then she receives her best payoff. The same arguments hold true, *mutatis mutandis*, for Ann. This means that it makes no sense for her to perform Strategy 1 if she knows that Bob performs Strategy 4. Then

she would rather prefer Strategy 2, in which case Bob receives his best payoff. This is illustrated on the lower left of Figure 6.21.




Figure 6.21: Rational plays of Split or Steal.

However, is it reasonable at all for any player to perform Strategy 2, which is a losing strategy? The answer is “No”! If Ann performs Strategy 4, i.e., if she decides to steal the jackpot, categorically, Bob has nothing to lose and so he can decide to make the same. Thus, Bob has a credible threat and thus he can simply keep Ann from winning the jackpot by performing Strategy 4. The same principle applies, *mutatis mutandis*, to Ann, which means that she can keep Bob from performing the same strategy, too. The only reasonable solution of this game can be found on the left-hand side of the middle of Figure 6.21. We conclude that Ann and Bob will play tit for tat and so they finally decide to split the jackpot. Once again, the given strategies are not the best-response strategies of Ann and Bob and the resulting solution is not a Nash equilibrium, although it is generated by self-enforcement.

6.6 Coherent n -Person Games

Before we finish this part of the book, and turn towards alternative concepts of game theory, I would like to make an excursion to coherent n -person games with $n > 2$. In

the previous chapters, I have already mentioned that the transition from $n = 2$ to $n > 2$ is straightforward for static and dynamic games, but in the case of coherent games, things become a little bit more intricate. Hence, it is worth elaborating coherent multiperson games in more detail. However, I have to mention that this section turns out to be somewhat more technical than the previous ones. Readers who are not so much interested in a general treatment of coherent games may skip the following analysis.

6.6.1 Preliminary Remarks

6.6.1.1 Response vs. Information

In a coherent 2-person game, the response curve of Ann performs two functions:

1. It represents Ann's potential responses to Bob *and*
2. it describes Bob's private information about Ann.

Of course, the same holds true, *mutatis mutandis*, for Bob's response curve.

For example, consider the play depicted in Figure 6.22, where Ann's response curve is red and Bob's response curve is black. The red curve indicates that Ann goes categorically down, which is just a description of Ann's potential responses to Bob, i.e., it describes her strategy. Further, Bob has perfect information, which means that he knows Ann's potential responses and so the red curve also describes Bob's private information about Ann. Analogously, the black curve tells us that Bob performs the strategy "Go left if and only if Ann goes up." Moreover, Ann knows that Bob goes left if and only if she goes up and so the black curve also describes Ann's private information about Bob.




Figure 6.22: A play of a coherent 2-person game.

Now, consider a coherent 3-person game with Ann, Bob, and Claire. Let

- $A = \{\text{Up}, \text{Down}\};$
- $B = \{\text{Left}, \text{Right}\},$ and
- $C = \{\text{Front}, \text{Back}\}$

be the corresponding action sets of the players. The empty response diagram of this game, which contains no graph, is given in Figure 6.23.




Figure 6.23: Empty response diagram of a coherent 3-person game.

Ann's response depends on Bob's *and* Claire's actions. This means that her response is a function of *two* actions, i.e., $\text{Res}_A : B \times C \rightarrow A$. Similarly, the response function of Bob is $\text{Res}_B : A \times C \rightarrow B$ and the response function of Claire is $\text{Res}_C : A \times B \rightarrow C$. Hence, Bob's response depends on Ann's and Claire's actions, whereas the response of Claire depends on Ann's and Bob's actions.

The response functions of the players represent their strategies. For example, if we assume that Ann chooses categorically Up, Bob chooses categorically Right, and Claire chooses categorically Back, we obtain the response graphs that are depicted in Figure 6.24. As we can see, each graph represents a 2-dimensional plane and the solution of the game must be an element of the intersection of the response graphs of Ann, Bob, and Claire. In our example, there is only one possible solution, namely the green point in Figure 6.24.




Figure 6.24: Ann's response graph on the upper side (red), Bob's response graph on the right side (black), and Claire's response graph on the back side (blue). The given solution is marked green.

Ann has perfect information. This means that she knows Bob's and Claire's responses to each action $a \in A$. Thus, her private information is a function of *one* action only, i.e., $\text{Inf}_A : A \rightarrow B \times C$. Correspondingly, the information function of Bob is $\text{Inf}_B : B \rightarrow A \times C$ and the information function of Claire is $\text{Inf}_C : C \rightarrow A \times B$. Our conclusion is that, in a coherent n -person game with $n > 2$, we must clearly distinguish between response and information, whereas in a coherent 2-person game this distinction is not necessary.

Let us consider Figure 6.25 and step into Ann's shoes. She knows that Bob chooses Right and Claire chooses Back if she chooses Down because, according to Bob's and Claire's strategies, Right and Back are the only possible responses to Down. More precisely, there is only one intersection point of the response graph of Bob (black) and the response graph of Claire (blue) that comes into question if Ann chooses Down. This intersection point is (Down, Right, Back) (see the green point in Figure 6.25).




Figure 6.25: If Ann chooses Down, categorically, she knows that Bob chooses Right and Claire chooses Back (green point), i.e., $\text{Inf}_A(\text{Down}) = (\text{Right}, \text{Back})$.

However, is it true that the responses of Bob and Claire always follow from Ann's action? The answer is “No”! The problem is that Bob's response is a function of Ann's and Claire's action, whereas Claire's response is a function of Ann's and Bob's action. Hence, on the basis of Bob's and Claire's response functions, in general we cannot predict Bob's and Claire's action by specifying Ann's action alone.

For example, consider Figure 6.26 in which both Bob and Claire perform a diagonal strategy. If Ann chooses Down, it can happen that Bob chooses Left and Claire chooses Back or that Bob chooses Right and Claire chooses Front. However, we assume that Ann *has* perfect information and so she knows Bob's and Claire's responses if she chooses Down. Of course, the same holds true if she chooses Up.




Figure 6.26: Bob and Claire perform a diagonal strategy. Now, there are two possible solutions if Ann chooses Down (green points).

Hence, having perfect information about Bob and Claire is *more* than knowing their strategies. Of course, the same arguments apply, *mutatis mutandis*, to Bob and Claire. By contrast, in a coherent 2-person game, having perfect information is just the *same* as knowing the other's strategy, i.e., his response curve. In any case, the information function of a player reflects perfect information and thus assigns each element of his own action set one and *only one* response vector, which contains the associated responses of his opponents.

We conclude that the graph of the information function of Player $i = 1, 2, \dots, n$ must be a subset of the intersection of the response graphs of all other players. Further, it is clear that the solution of the game must belong to the intersection of all information graphs, i.e., the graphs of the information functions of all players, which implies that it is an element of the intersection of all response graphs.

6.6.1.2 Handling Response Diagrams

Altering the action of a player in a response diagram might be paradoxical. Consider the solution on the lower right of Figure 6.22 and assume that Ann decides to go up. Bob's (black) response curve tells us that then he goes left, but according to Ann's (red) response curve, she goes down in this case, not up, which is a contradiction.

Does this mean that Ann cannot go up, which implies that she has no freedom of choice? Indeed, this would thwart our overall approach to game and decision theory, which is based on the assumption that every subject has a free will.¹¹⁸ So, where is the misconception? Here comes the answer: By moving from Down to Up, Ann changes her strategy. More precisely, she does no longer perform the strategy "Go down, categorically," i.e., she changes her own response curve. This is illustrated in Figure 6.27.¹¹⁹




Figure 6.27: If Ann moves up, she changes her own response curve.

Hence, whenever we consider some solution in the response diagram of a coherent 2-person game and change the action of any player, we must take into account that

¹¹⁸ Whether or not free will is an illusion does not matter at all in the subjectivistic framework.

¹¹⁹ Here, Ann's new strategy is Strategy 3. It could also be Strategy 1, i.e., "Go up, categorically."

this change, in general, leads to another response curve of the corresponding player, whereas the response curve of the other player is always fixed.¹²⁰

The same principle applies to all coherent n -person games with $n > 2$, in which case altering the action of Player $i \in \{1, 2, \dots, n\}$ generally means that this player changes his response function, i.e., Res_i , whereas the response function of each other player $j \neq i$, i.e., Res_j , is always fixed. Conversely, if Player i changes his action, his own information function Inf_i always remains fixed, but the information function Inf_j of each player $j \neq i$ usually changes, since everybody has perfect information and thus Player j knows if Player i changes his action.

6.6.2 Main Theorems

Consider any n -person game with $n \geq 2$ in normal form: Player $i = 1, 2, \dots, n$ has some nonempty action set $A_i \subseteq \mathbb{R}^{r_i}$, so that $A := \times_{i=1}^n A_i$ is the action space of the game. Each element $x = (x_1, x_2, \dots, x_n) \in A$ of the action space is said to be a solution. Every player has a payoff function $f_i : A \rightarrow \mathbb{R}$ that assigns each solution a real number. The game can either be static, dynamic, or coherent.

The reader may substitute “payoff” with “expected utility.” However, in the case in which the players have perfect information, i.e., the game is coherent, probabilities and expectations play no role. Further, it is worth emphasizing that in that case we need not specify the response and information functions of the players because these objects are just a *result*, not a prerequisite, of a coherent game.

Definition 13 (Equilibrium). The solution $x = (x_1, x_2, \dots, x_n) \in A$ is said to be an equilibrium if and only if

$$f_i(x_1, x_2, \dots, x_i, \dots, x_n) \geq f_i(x_1, x_2, \dots, y_i, \dots, x_n)$$

for every action $y_i \in A_i$ of each player $i \in \{1, 2, \dots, n\}$.

Equilibrium is a pivotal notion in game theory and I will discuss it in much more detail in Section 7.2. Figure 6.28 shows an equilibrium of the coherent 3-person game that has been illustrated in Section 6.6.1.1.

According to Definition 10, a solution of a game is said to be rational if and only if all players maximize their expected utilities. In a *coherent* n -person game the players have perfect information. Hence, a solution is said to be rational if and only if no player can increase his payoff by choosing another action. More specifically, let Inf_i be the information function of Player $i = 1, 2, \dots, n$, i.e., $\text{Inf}_i(x_i) \in \times_{j \neq i} A_j$ is the vector of responses of the other players if Player i chooses Action $x_i \in A_i$. The solution $x \in A$ is said to be rational if and only if

120 There exist situations in which a player need not change his response curve when changing his action. In a coherent 2-person game this happens if both players perform a diagonal strategy.




Figure 6.28: Equilibrium of a coherent 3-person game (blue point).

$$f_i(x) \geq f_i(y_i, \text{Inf}_i(y_i))$$

for every action $y_i \in A_i$ of each player $i \in \{1, 2, \dots, n\}$. Note that, in general, Inf_i is not invariant to the action of any player i . In the special case in which Inf_i is invariant to the action of Player $i = 1, 2, \dots, n$, the rational solution represents an equilibrium. However, we have seen that many rational solutions are not an equilibrium and many equilibria cannot be considered a reasonable solution.

We could call a rational solution a “coherent equilibrium” because no player can increase his payoff by changing his action or, more generally, his strategy. The fact that all other players hold their strategies fixed if one player changes his need not be stated explicitly. The other players can very well change their *actions*, i.e., they can react to the player, but whatever responses he gets—they just reflect the *given* strategies of his adversaries. Hence, our understanding of rational solution fits quite well with the concept of equilibrium. However, calling a rational solution “coherent equilibrium” could be somewhat misleading and, to be honest, I do not wish to participate in the inflationary proliferation of the notion of “equilibrium,” which can be observed in traditional game theory during the last decades.¹²¹

Each equilibrium of a coherent game is a rational solution, provided that all players perform a categorical strategy. However, the problem is that not all *finite* coherent games, i.e., coherent games in which the number of players, n , and their action sets A_1, A_2, \dots, A_n are finite, have an equilibrium. In fact, we have already discussed some finite coherent 2-person games without any equilibrium. Typical examples are the discoordination games elaborated in Section 6.4.2.

The following theorem ensures that a finite coherent game always has at least one rational solution. It represents the main result for finite coherent games.

Theorem 6 (Finite coherent game). *Every finite coherent game has a rational solution.*

The proof of that theorem is constructive and shall be provided, without loss of generality, by means of the coherent 3-person game depicted in Figure 6.28. In that

¹²¹ The equilibrium doctrine will be discussed in Chapter 7.

case, the action sets of the players can be set to $\{0, 1\}$, where 0 stands for Up, Left, or Front, and 1 stands for Down, Right, or Back, depending on the player who is taken into consideration. Applying the proof to any case with $n \neq 3$ players and other (finite) action spaces as well as payoff functions is straightforward.

First of all, we ignore Claire, i.e., we focus on Ann and Bob by assuming that Claire chooses Front, categorically. Suppose that also Bob performs a categorical strategy. If he chooses Left, Ann's best response is Down, whereas it is Up if he chooses Right. This constitutes Ann's strategy, given that Claire chooses Front. Let us step into Bob's shoes. He knows that Ann goes down if he goes left but up if he goes right. Given her response curve, Bob maximizes his own payoff by choosing Left, categorically.¹²²

Now, we assume that Claire chooses Back, categorically, and repeat our procedure. To be more precise, Ann gives a best response to each action of Bob, whereas Bob maximizes his own payoff given Ann's response curve. It turns out that Ann goes up if Bob goes left, whereas she goes down if he goes right. In this case, it is better for Bob to choose Right, categorically. Finally, we turn our attention to Claire. She knows the strategies of Ann and Bob. More precisely, she knows their responses, depending on whether she chooses Front or Back. Claire maximizes her own payoff by choosing Back, categorically.

The solution of this game can be found in Figure 6.29. We have constructed a play where Ann performs a strategy whose response graph, i.e., the graph of her response function, is bent in two dimensions, Bob performs a strategy whose response graph is bent in one dimension, and Claire performs a strategy whose response graph is not bent at all. The solution is rational by construction, i.e., due to the response graphs that are depicted in Figure 6.29, no player can increase his payoff by changing his action: If Ann moves up, she receives only the payoff 0 instead of 2, if Bob moves left, he receives only the payoff 1 instead of 3, and if Claire moves to the front, she receives only the payoff -1 instead of 1. However, the given solution is not an equilibrium!




Figure 6.29: Rational solution of a coherent 3-person game (green point).

¹²² Even if we would consider a game in which Ann's best response is not unique, Bob would still know Ann's actual choice because he has perfect information.

Now, consider any coherent n -person game. The iterative procedure goes like this: Player 1 creates his strategy by giving a best response to Player 2, 3, ..., n , Player 2 creates his strategy by maximizing his payoff given the strategy of Player 1 and giving a best response to Player 3, 4, ..., n , etc., and Player n creates a categorical strategy by maximizing his payoff given the strategy of Player 1, 2, ..., $n - 1$.

The resulting play has a specific geometrical structure, which can be observed in Figure 6.29: The rational solution is constructed in such a way that the response graph of Player $i \in \{1, 2, \dots, n\}$ can be bent, at most, in $n - i$ dimensions. To be more precise, the response graph of Player 1 can be bent along the dimension of Player 2, 3, ..., n , the response graph of Player 2 can be bent along the dimension of Player 3, 4, ..., n , etc., whereas the response graph of Player n is not bent at all. Of course, we may consider any permutation of 1, 2, ..., n in order to construct another rational play, but then we usually obtain also another (rational) solution.

Consider any player $i \in \{1, 2, \dots, n\}$ and let

$$x_{-i} := (x_1, x_2, \dots, x_{i-1}, x_{i+1}, \dots, x_n) \in \bigtimes_{j \neq i} A_j$$

be the vector of actions of his opponents. This enables us to use the shorthand notation $x = (x_i, x_{-i})$. Hence, a solution $x \in A$ represents an equilibrium if and only if $f_i(x) \geq f_i(y_i, x_{-i})$ for every $y_i \in A_i$ and $i = 1, 2, \dots, n$. Further, a payoff function f_i is said to be quasiconcave in x_i , i.e., in its own argument, if and only if $f_i(\cdot, x_{-i})$ is quasiconcave for all $x_{-i} \in \bigtimes_{j \neq i} A_j$, i.e.,

$$f_i(\pi x_i + (1 - \pi)y_i, x_{-i}) \geq \min \{f_i(x), f_i(y_i, x_{-i})\}$$

for all $0 \leq \pi \leq 1$, $x_i, y_i \in A_i$, and $x_{-i} \in \bigtimes_{j \neq i} A_j$. Here, it is implicitly assumed that A_i is a convex subset of \mathbb{R}^{r_i} .

The next theorem generalizes Nash's existence theorem, which will be discussed in much more detail in Section 7.2. It has already been developed by Debreu (1952), Fan (1952), and Glicksberg (1952).

Theorem 7 (Equilibrium). *Consider any n -person game and suppose that the action sets are compact and convex. Further, assume that the payoff functions are continuous and quasiconcave in their own arguments. The game has an equilibrium.*

Proof. Fix any solution $x \in A = \bigtimes_{i=1}^n A_i$. Let $\text{Res}_i(x_{-i}) \subseteq A_i$ be the set of all best responses of Player $i = 1, 2, \dots, n$ to his opponents, given that they perform the actions in $x_{-i} \in \bigtimes_{j \neq i} A_j$. Since A_i is compact and f_i is continuous, it follows from the extreme-value theorem that $\text{Res}_i(x_{-i})$ is nonempty. Further, let $\text{Res}(x) \subseteq A$ be the joint set of all best responses of the players to each other, where the action space A is nonempty, compact, and convex. Each payoff function f_i is quasiconcave in x_i and thus every convex combination of two best responses is a best response, too. Hence, $\text{Res}(x)$ is

nonempty and convex for all $x \in A$. Further, since the payoff functions of the players are continuous, Res is a (set-valued) function on A with closed graph. Now, Kakutani's fixed-point theorem guarantees that there exists a solution $x \in A$ such that $x \in \text{Res}(x)$, which represents an equilibrium. \square

Each equilibrium of a coherent game is a rational solution if we assume that all players perform a categorical strategy. Thus, every coherent game that satisfies the conditions of Theorem 7 has a rational solution. However, we can imagine many coherent games that have a rational solution although they *violate* the conditions of Theorem 7. Indeed, following the proof of Theorem 6, we are able to provide a very simple sufficient condition for the existence of a rational solution, without making any requirements about the action sets and payoff functions of the players.

Sufficient condition. A coherent n -person game has a rational solution if

- there exists a best-response function for Player 1 with respect to the actions of Player 2, 3, \dots , n such that
- there exists a best-response function for Player 2 with respect to the actions of Player 3, 4, \dots , n , given the best-response function of Player 1, such that
- there exists a best-response function for Player 3 with respect to the actions of Player 4, 5, \dots , n , given the best-response functions of Player 1, 2, etc.

It is clear that every finite coherent game satisfies this sufficient condition. The proof of the following proposition, which shows that many infinite coherent 2-person games have a rational solution, too, is based on the same condition.

Proposition 2. *Every coherent 2-person game with compact action sets and continuous payoff functions has a rational solution.*

Proof. The action set A_1 is nonempty and compact. Further, the payoff function f_1 is continuous. Thus, for each $x_2 \in A_2$ we are able to maximize $f_1(\cdot, x_2)$. Let $M_1(x_2) \subseteq A_1$ be the set of all maximum points at x_2 . It is clear that $M_1(x_2) \neq \emptyset$ for all $x_2 \in A_2$ and since f_1 is continuous, the graph of M_1 is a closed subset of $A = A_1 \times A_2$. Further, the payoff function f_2 is continuous, too, and thus for each $x_2 \in A_2$ we can maximize $f_2(\cdot, x_2)$ on $M_1(x_2)$. Let us choose any maximum point of $f_2(\cdot, x_2)$ in $M_1(x_2)$ as a response $\text{Res}_1(x_2)$ of Player 1 to Player 2 given that the latter chooses Action x_2 . Due to our construction of Res_1 , the function $f_2^*(\cdot) := f_2(\text{Res}_1(\cdot), \cdot)$ is upper semicontinuous on the nonempty and compact action set A_2 . Thus, we are able to maximize f_2^* . Let $M_2 \subseteq A_2$ be the (nonempty) set of all maximum points. Now, we can choose any element of M_2 as a categorical response of Player 2 to Player 1, which means that $\text{Res}_2(x_1) \in M_2$ is independent of $x_1 \in A_1$. The intersection of the graph of Res_1 and the graph of Res_2 is nonempty and contains a rational solution of the game. \square

For example, consider the following coherent game, which contains a rational solution but no equilibrium: Ann and Bob must decide to sit down somewhere on a

bench. Each player can choose between any location in $A_1 = A_2 = [0, 1]$, where 0 means “far left” and 1 means “far right.” Ann does not like Bob and so she wants to sit away as far as possible from Bob, whereas Bob likes Ann very much and would like to be as close as possible to Ann. Let Ann’s payoff be $|x_1 - x_2|$, where $x_1 \in A_1$ is Ann’s action and $x_2 \in A_2$ is Bob’s action. Correspondingly, $1 - |x_1 - x_2|$ is Bob’s payoff if they choose Action x_1 and Action x_2 , respectively.

This is a coherent discoordination game with infinite actions sets. Moreover, it is a constant-sum game and thus a zero-sum game.¹²³ The best-response curves of Ann and Bob can be found on the left-hand side of Figure 6.30. As we can see, the best-response curves do not overlap and so this game cannot have any equilibrium, although the action sets are compact and convex, and the payoff functions are continuous. The reason is that Ann’s payoff function, f_1 , is not quasiconcave in x_1 .




Figure 6.30: Best-response curves (left) and rational play (right) of the bench game.

The right-hand side of Figure 6.30 illustrates a rational play of the bench game:

- Ann chooses Location 1 if Bob decides to sit down at Location $x_2 \in [0, 0.5]$. Otherwise, she chooses Location 0.
- Bob decides to sit down at Location 0.5 categorically, i.e., independent of Ann’s choice $x_1 \in [0, 1]$.

In fact, Ann performs a best-response strategy and Bob cannot do better as sitting down in the middle of the bench given Ann’s strategy. Hence, a rational solution of this game is $(1, 0.5)$, which means that Ann goes to the right of the bench and Bob sits down in the middle. In this case, both players receive a payoff of 0.5.

The bench game has an infinite number of rational solutions. For example, Bob could decide to perform his best-response strategy, i.e., to follow Ann wherever she goes, in which case Ann were indifferent among all locations between 0 and 1. This means that each solution (x, x) with $x \in [0, 1]$ is rational, too, if we assume that Ann chooses Location x , categorically, and Bob performs his best-response strategy. In that case, Ann receives only the payoff 0, whereas Bob has a payoff of 1. This solution seems

¹²³ Due to Savage’s representation theorem (see Section 1.4), we can always transform the payoffs of the players in a constant-sum game into an isomorphic zero-sum game.

unreasonable compared with the solution above. The reason is that Ann can simply prevent being persecuted by Bob just by performing her best-response strategy. In fact, as we already know, it is impossible that both players perform their best-response strategies, simultaneously.

The last theorem of this chapter represents the main result for finite and infinite coherent n -person games. Its proof is based on Wald's maximin rule (see Section 2.4.2).

Theorem 8 (Coherent n -person game). *Every coherent n -person game with compact action sets and continuous payoff functions has a rational solution.*

Proof. In order to find a rational solution, we must specify only the information function of each player in a proper way. Once we have the desired information functions, we can imagine some appropriate response functions.¹²⁴ However, as we will see below, this is not necessary at all. Since the action sets are compact and the payoff functions are continuous, we can fix any action

$$x_1^* \in \arg \max_{x_1} \min_{x_{i>1}} f_1(x_1, x_{i>1})$$

for Player 1.¹²⁵ Next, we fix any action

$$x_2^* \in \arg \max_{x_2} \min_{x_{i>2}} f_2(x_1^*, x_2, x_{i>2})$$

for Player 2, etc. Finally, for Player n we just fix any action

$$x_n^* \in \arg \max_{x_n} f_n(x_1^*, x_2^*, \dots, x_n).$$

The information function of Player 1 is

$$\text{Inf}_1 : x_1 \mapsto \begin{cases} x_{i>1}^*, & x_1 = x_1^*, \\ y_{i>1} \in \arg \min_{x_{i>1}} f_1(x_1, x_{i>1}), & x_1 \neq x_1^*. \end{cases}$$

Correspondingly, the information function of Player 2 is

$$\text{Inf}_2 : x_2 \mapsto \begin{cases} (x_1^*, x_{i>2}^*), & x_2 = x_2^*, \\ (x_1^*, y_{i>2}) \text{ with } y_{i>2} \in \arg \min_{x_{i>2}} f_2(x_1^*, x_2, x_{i>2}), & x_2 \neq x_2^*, \end{cases}$$

etc. By construction, we can always find some response graphs that lead to the graphs of the given information functions.¹²⁶ In particular, it holds that

124 Remember that the graph of the information function of some player is a subset of the intersection of the graphs of the response functions of all other players. Further, the solution of the game must belong to the intersection of the graphs of all information functions (see Section 6.6.1.1).

125 It is implicitly assumed that $x_i \in A_i$ and $x_{j>i} \in \prod_{j>i} A_j$ for $i = 1, 2, \dots, n-1$.

126 More precisely, we can find a response graph for Player 1 that is not bent at all, a response graph for Player 2 that is bent, at most, along the dimension of Player 1, etc.

$$\begin{aligned}
 f_1(x_1, \text{Inf}_1(x_1)) &= \min_{x_{i>1}} f_1(x_1, x_{i>1}) \leq \min_{x_{i>1}} f_1(x_1^*, x_{i>1}) \\
 &\leq f_1(x_1^*, x_{i>1}^*) = f_1(x_1^*, \text{Inf}_1(x_1^*))
 \end{aligned}$$

for all $x_1 \neq x_1^*$. This means that Player 1 cannot improve by moving away from x_1^* and it is easy to see that the same holds true regarding x_i^* for Player $i = 2, 3, \dots, n$. Hence, the solution $x^* = (x_1^*, x_2^*, \dots, x_n^*)$ is rational. \square

Every finite coherent game satisfies the conditions of Theorem 8: The action sets of a finite coherent game are finite, which means that they are compact and the payoff functions are continuous. Hence, Theorem 8 can be considered a generalization of Theorem 6. Nonetheless, finite games play a major role in game theory and so we can let Theorem 6 stand on its own for practical reasons.

To sum up, all finite coherent games have a rational solution and we have seen that the same holds true for many infinite coherent n -person games. Thus, we are able to find a rational solution of a coherent game in most practical applications. Finite coherent 2-person games have already been treated throughout this chapter. Further, infinite coherent 2-person games that satisfy the conditions of Theorem 7 will be discussed in more detail in Section 7.5.3.

6.7 Conclusion

Interaction can often be observed in real life. In order to understand the phenomenon of interaction, we may treat a strategic conflict in which the players are able to interact as a coherent game. In such a game all players have perfect information, i.e., know the strategies of each other. This simplifying assumption enables us to solve quasicoherent games, i.e., games in which the players act *as if* they had perfect information, in a relatively simple way by means of response curves.

Thus, we can derive the rational solutions of both coherent games and quasicoherent games by making use of response diagrams. However, in general, the set of rational solutions is not a singleton and thus we must apply a refinement procedure in order to eliminate all unreasonable solutions. This was demonstrated by solving anti-coordination, discoordination, and coordination games, as well as the prisoners' dilemma. The corresponding solutions typically differ essentially from those that are obtained under the strategic-independence assumption.

For example, it turns out that players in an anti-coordination game cooperate, since each player can force his adversary to act in the desired way. This means that cooperation takes place by self-enforcement and it does not represent a Nash equilibrium. The solution of a discoordination game depends essentially on which player is able to carry through his favorable strategy. However, without making any further assumption, we are not able to predict which player prevails. In a coordination game,

the players cooperate, but if they pursue opposing interests, the solution might still depend on the power of each player or on convention.

The most striking example that demonstrates the difference between static and coherent games is the prisoners' dilemma. The prisoners confess if they are not able to interact, given that they are rational and interested only in minimizing their penalties. If the prisoners are able to interact, they deny. The reason is that each prisoner has a credible threat and can force the other to cooperate by playing tit for tat. Hence, once again, it turns out that cooperation takes place by self-enforcement. This explains why cooperation is a phenomenon that is often observed in real life—although the players should apparently defect from the perspective of traditional game theory.

The theory of coherent games can be extended to the case of $n > 2$. Whenever a coherent game has an equilibrium, it has a rational solution, but many coherent games possess a rational solution although they do not have any equilibrium. In fact, every coherent n -person game with compact action sets and continuous payoff functions has a rational solution. This enables us to solve coherent games in most practical applications, irrespective of whether they have an equilibrium or not. Despite these circumstances we will see in the next part of this book that the notion of equilibrium is a central element of traditional game theory.

Part III: Alternative Concepts of Game Theory

7 The Traditional Approach

The basic assumption of noncooperative game theory is strategic independence. It claims that the players make their choices independently, which means that the decision of a player has no influence on the decision of any other player. This situation has been thoroughly analyzed in Chapter 4. The notion of equilibrium is implicitly based on the strategic-independence assumption. Hence, in most parts of this chapter, it is assumed that the game is static and we will restrict ourselves to one-shot games without loss of generality. An exception is Section 7.5, in which we discuss von Neumann's minimax solution. This solution can be interpreted also from a coherent perspective.

A main motivation for writing this book was to present a simple but general solution concept that goes beyond the equilibrium doctrine of game theory. In the subjectivistic framework, the rational solution of a (static) game need not be an equilibrium. To be more precise, it need not be an ordinary Nash equilibrium (Nash, 1951), a Bayesian Nash equilibrium (Harsanyi, 1967–1968), or a correlated equilibrium (Aumann, 1987), which are the predominant solution concepts in traditional game theory. In this chapter, I will explain these solution concepts in more detail and compare them with the subjectivistic approach.

7.1 Transparency of Reason

The transparency of reason is a key principle in traditional game theory and can be explained like this (Bacharach, 1987)¹²⁷:

“Broadly, the principle of the Transparency of Reason is the claim that if reason suffices for a player who has certain data to come to a given conclusion—say, as to what he should do—then, if a second player believes the first to have these data and to be rational, reason suffices for him to come to the conclusion that the first will come to his.”

Simply put, if the players can step into the shoes of each other, then they are able to deduce the strategies of their opponents. In order to justify the transparency of reason, it is usually assumed that the players have

1. complete information and
2. common knowledge.

These terms play a major role in traditional game theory and so they shall be explained in this section, before I discuss the traditional solution concepts.

¹²⁷ Bacharach notes that the term “transparency of reasoning” goes back to Bjerring (1978).

7.1.1 Complete Information and Common Knowledge

Roughly speaking, a player is said to have complete information if and only if he knows all relevant aspects of the game and thus can act like an outside observer. The following definition of complete information applies also to dynamic games.

Definition 14 (Complete information). A player is said to have complete information if he knows the rules of the game, the number of players, their action sets, their private information sets, their payoffs, and thus that all players are rational.

Actually, this definition represents a sufficient condition because which aspect of the game is “relevant” in order to be able to step into the shoes of each other depends on the specific situation at hand.

Classical arguments of noncooperative game theory are further based on the assumption of common knowledge. The notion of common knowledge goes back to Lewis (1969) and is first formalized by Aumann (1976). Let f be any fact and $\alpha(f)$ the fact that f is known to all people. The latter fact is referred to as mutual knowledge. Common knowledge of f means that $\alpha^k(f)$ for $k = 1, 2, \dots$, i.e., everybody knows f and everybody knows that everybody knows f , etc. For example, if two persons have witnessed a murder and recognized the murderer, he is mutually known. In addition, if the witnesses have seen one another while recognizing the murderer, he is commonly known. Common knowledge is called also unbounded knowledge. Bounded knowledge simply means that the people have no common knowledge about f , but at least it holds that $\alpha^k(f)$ for some positive integer k .

The key idea of traditional game theory is that the players have complete information and that their information is common knowledge. Put another way, they know all relevant aspects of the game and know also that each other knows all relevant aspects of the game, etc., where “relevant” means that the players are able to step into each other’s shoes and deduce his strategy, given that each other steps into each other’s shoes, too, etc. To be more precise, Γ shall be an entire description of the game, i.e., it contains the rules of the game, the number of players, their action sets, their private information sets, and their payoffs, which means that all players are rational. Thus, Γ contains all relevant aspects of the game. A game with complete information is a strategic conflict in which $\alpha(\Gamma)$ holds true, i.e., all players have complete information about the game and the rationality of all players.

Common knowledge goes even further: It says that every player has complete information and that every player knows that everybody has complete information and so on, ad infinitum. This provides us with a beautiful solution concept—at least from an academic point of view. Nonetheless, many behavioral studies show that the common-knowledge assumption is violated in most real-life strategic conflicts. For a nice discussion about this issue see, e.g., the study by Nagel (1995) and the references contained therein. I will come back to this study in Section 7.1.3.2.

7.1.2 What Does Common Knowledge Tell Us?

Apart from the problem of human reasoning and the question of common knowledge, we may think about a very simple philosophical question: Can a decision maker know which choice he will make? In other words, can he know his own action? I think that the answer is “No”! In our subjectivistic framework, “knowledge” always refers to some event, but deliberation crowds out prediction. Hence, the action of a decision maker cannot be an event from his *own* perspective. A decision maker may very well decide which action he is going to perform, but it makes no sense to assume that he knows that action, i.e., that his own choice is evident to himself.

In Section 2.4.1, I argued that the decision maker knows that he chooses Heads *if* he chooses Heads and he knows also that he chooses Tails *if* he chooses Tails. These trivial tautologies are true because substantive conditionals are reflexive (see Section 2.3.1). More generally, let A be the action set of the decision maker. He knows that he will perform Action $a \in A$ if he decides to perform Action a and, of course, the same substantive conditional holds true for any other element of A . However, the decision maker cannot know, e.g., that he performs Action a because this would mean that he knows that he performs Action a also in the case in which he decides to perform Action $b \neq a$ ($b \in A$) instead of Action a . This is a contradiction.

The same principle applies to the players in a strategic conflict. Each player is a decision maker. However, a player *can* know the action (and even the strategy) of some opponent because he considers the action (or strategy) of the other an event. Knowing the other’s action means to know that he will choose, or already has chosen, this action—irrespective of whatever action I choose.¹²⁸ The reason why this is possible is just because all players live in their *own* world, i.e., they do not share the same state space. Thus, we can apply Levi’s postulate: Deliberation crowds out prediction or, equivalently, prediction crowds out deliberation. If somebody *predicts* an action, he cannot *choose* this action.

What I am trying to say is that the notion of common knowledge should be used with care. Some authors suggest that common knowledge implies that all players *know* their own actions and the actions of each other. As we have already seen, this is impossible. If a player knows his own action, he has no choice at all. For example, consider two trains on a collision course. The locomotives are very fast and not far enough away to prevent the catastrophe. Each train driver is looking into the eyes of the other and so it is common knowledge that they are going to collide. I guess that most of us would not consider this hopeless situation any “game” or “decision problem” because nobody can do anything against the inevitable disaster.

¹²⁸ Nonetheless, the other player still has a freedom of choice, which has already been demonstrated in Section 2.5.2 by Newcomb’s paradox.

In traditional game theory, the common-knowledge assumption is typically used for static games, i.e., for games in which the strategic-independence assumption is satisfied. Of course, in a static game it is very well possible that every player knows which action each other *prefers* and the preferred actions might even be common knowledge. However, this does not mean that any player knows that his adversaries will actually *perform* their optimal actions. That is, under the common-knowledge assumption, the players still suffer from imperfect information!

Subjectivistic game theory does not presume that any player in a static game knows the actions of his adversaries. Indeed, a player can very well be able to deduce the optimal actions of his opponents. However, since the subjectivistic approach does not pose any behavioral requirements on rationality, it may very well happen that a rational player performs a suboptimal action. Consequently, in a static game, complete information and common knowledge do not lead to situations in which the players know the action of each other, but they may know their *optimal* actions.

7.1.3 Some Examples

7.1.3.1 The Cournot Duopoly

The Cournot duopoly (Rieck, 2016, p. 30) is similar to the entrepreneurship, which has been discussed in Section 4.2.1.¹²⁹ The main difference is that in the Cournot duopoly the entrepreneur is no longer confronted with an unspecified number of competitors. Now, there are only two firms acting against one another. However, each firm can still choose between a low, moderate, or high engagement. We may assume that the firms know (or, at least, that they believe) that their choices are independent. The payoff matrix of the Cournot duopoly is given in Table 7.1.

Table 7.1: The Cournot duopoly.

Firm A	Firm B		
	Low	Moderate	High
Low	(5, 5)	(3, 6)	(0, 7)
Moderate	(6, 3)	(4, 4)	(1, 3)
High	(7, 0)	(3, 1)	(-1, -1)

The transparency of reason goes like this: Firm A and Firm B are rational, they know the payoff matrix, and they know that their choices are independent. Low is strictly

¹²⁹ It is named after Antoine Augustin Cournot, a French philosopher, who is well-known for his seminal contributions to mathematical economics (Cournot, 1897).

dominated by Moderate. Firm A is rational and so it will not prefer to engage to a low degree. Moreover, Firm A knows that Firm B is rational, too, which means that Firm A knows that also Firm B will abandon Low. Now, it turns out that also High is strictly dominated by Moderate and so Firm A eliminates High. The same arguments apply, *mutatis mutandis*, to Firm B. This means that both firms, finally, prefer to engage to a moderate degree.

Common knowledge implies even more¹³⁰: Firm A knows that Firm B knows that Firm A is rational. Thus, it knows that Firm B knows that Firm A abandons Low, too. Put another way, Firm A knows that Firm B eliminates also High and so Firm A knows that Firm B prefers Moderate, too. To sum up, if the firms have complete information and common knowledge, they prefer to engage to a moderate degree and they also know that their opponent prefers the same action.

The subjectivistic approach does not require common knowledge and so the firms may have any arbitrary subjective probability distribution on the action set {Low, Moderate, High}. For example, Firm A might doubt that Firm B is rational and so it comes to the conclusion that Firm B chooses Low with probability 5 %, Moderate with probability 85 %, and High with probability 10 %. Then we have the expected utilities

- $EU_A(\text{Low}) = 0.05 \cdot 5 + 0.85 \cdot 3 + 0.1 \cdot 0 = 2.8$;
- $EU_A(\text{Moderate}) = 0.05 \cdot 6 + 0.85 \cdot 4 + 0.1 \cdot 1 = 3.8$; and
- $EU_A(\text{High}) = 0.05 \cdot 7 + 0.85 \cdot 3 + 0.1 \cdot (-1) = 2.8$.

Hence, Firm A still prefers Moderate. Nonetheless, if Firm B thinks, for whatever reason, that Firm A is lazy and thus assigns Low probability 1, it will choose High. In this case, we obtain the rational solution Moderate/High, which is not predicted by traditional game theory. In order to justify such a solution, we must assume that the common-knowledge assumption is violated. For example, this can happen if Firm B doubts that Firm A is rational or if it does not know Firm A's payoffs, etc. Alternatively, Firm B could simply assume that Firm A chooses Low by mistake.

7.1.3.2 The Guessing Game

The guessing game is described by Nagel (1995). It nicely demonstrates that the common-knowledge assumption is violated in real-life situations of strategic conflict. Here, I present the game in a slightly modified version: A number $n > 1$ of persons are asked to choose some element from the closed interval $[0, 100]$. No player knows the choice of any other. Let $0 \leq \pi < 1$ be some predetermined parameter and $\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$ the average of all chosen numbers. The closer the choice of Player i , i.e., x_i , is to the

¹³⁰ In this game, it would suffice to assume that the players have bounded knowledge. However, with common knowledge we are always on the safe side.

target number $\tau := \pi \bar{x}$, the more money he wins.¹³¹ For example, we could assume that the payoff is $\$(100 - |x_i - \tau|)$.

Suppose that the players have complete information and common knowledge. Player 1 thinks about his choice x_1 . Choosing a number $x_1 > \pi 100$ makes no sense at all. We have that $\bar{x} \leq 100$ and thus $\tau \leq \pi 100$. Thus, Player 1 can always increase his payoff by choosing $\pi 100$ instead of any number greater than $\pi 100$, where “always” means “for all choices $x_2, x_3, \dots, x_n \in [0, 100]$ of his adversaries.”¹³² This means that $x_1 = \pi 100$ strictly dominates each other choice of Player 1 greater than $\pi 100$.

Each other player comes to the same conclusion, i.e., no player will prefer any number $x_i > \pi 100$. Hence, Player 1 believes that $x_2, x_3, \dots, x_n \leq \pi 100$ and so it is not reasonable for him to choose any number greater than $\pi^2 100$. The same holds true for each other player, etc. Common knowledge enables us to repeat our arguments an infinite number of times. This means that no player will prefer any number greater than $\lim_{k \rightarrow \infty} \pi^k 100 = 0$. Hence,

$$x_1 = x_2 = \dots = x_n = 0$$

is the only solution of the guessing game that is consistent with the assumption of common knowledge.

The guessing game was played with thousands of readers of science and economics journals with $\pi = \frac{2}{3}$ (Selten and Nagel, 1998). Although the rules of the game were slightly different between the audiences,¹³³ the results are quite similar:

- The distribution of numbers between 0 and 100 is right-skewed.
- Less than 20 % of the players choose 0. Usually, the relative number of people choosing 0 is about 10 %. Further, about 1–2 % of the people choose 100.
- There are always two peaks at 22 and 33.¹³⁴
- The average number, \bar{x} , is between 18 and 26, i.e., the target number, τ , is between 12 and 17.

Hence, the given results clearly violate the assumption of common knowledge. Are the readers of the journals irrational? Well, at least they do not behave in a way that is suggested by traditional game theory.

In fact, a player does even *better* not to believe that the others are rational and that the others believe that each other is rational, too, etc. We could say that “rationality does not pay if the others are not rational, too.” However, this statement must

¹³¹ In the original game described by Nagel (1995), there is a fixed amount of money that is fairly divided among all players who are closest to τ . All other players go away empty-handed.

¹³² Note that τ depends also on his own choice x_1 .

¹³³ Some journals asked their readers to choose an integer from $[0, 100]$ or a real number from $[1, 100]$.

¹³⁴ This result can be explained as follows: Many readers thought that $\bar{x} = 50$ and thus they have chosen $\frac{2}{3} 50 \approx 33$. Some readers anticipated this result and have chosen $\frac{1}{2} 33 = 22$.

be made with care. First of all, we must define the meaning of rationality. In the traditional sense of game theory, rationality incorporates common knowledge. However, subjectivistic game theory does not require common knowledge at all.

The subjectivistic solution of the guessing game is quite simple and intuitive: We may assume that Player 1 thinks that his own choice has no influence on the choice of any other player. For each combination of $x_2, x_3, \dots, x_n \in [0, 100]$, Player 1 calculates the payoff given his own choice $x_1 \in [0, 100]$. He has some subjective probability distribution on $[0, 100]^{n-1}$ regarding the choices of the other players. Thus, Player 1 chooses some number from $[0, 100]$ that maximizes his expected payoff. The same arguments hold true for each other player.¹³⁵

More precisely, Player 1 considers the choice of each other player a random variable X_i ($i > 1$) but his own choice, x_1 , a control variable. The average number of all players (including himself) amounts to

$$\bar{X} = \frac{x_1 + \sum_{i=2}^n X_i}{n} = \frac{1}{n}x_1 + \frac{n-1}{n} \left(\frac{1}{n-1} \sum_{i=2}^n X_i \right).$$

Hence, if Player 1 chooses the number x_1 , he expects the average number

$$E(\bar{X}) = \frac{1}{n}x_1 + \frac{n-1}{n}\mu,$$

where

$$\mu := E \left(\frac{1}{n-1} \sum_{i=2}^n X_i \right)$$

denotes the mean of the subjective probability distribution of Player 1 regarding the choices X_2, X_3, \dots, X_n of the other players.

We conclude that Player 1 will solve the fixed-point equation

$$x = \pi \left(\frac{1}{n}x + \frac{n-1}{n}\mu \right),$$

which leads to the optimal choice

$$x = \left(\frac{n-1}{n-\pi} \right) \pi \mu.$$

Since the quotient $(n-1)/(n-\pi)$ is lower than 1, we always have that $x < \pi\mu$. For example, in the case that $n = 2$, $\pi = \frac{2}{3}$, and $\mu = 50$ we obtain $x = 25$. If the number of players, n , is large, it holds that $(n-1)/(n-\pi) \approx 1$ and thus $x \approx \pi\mu$. This leads us to $x \approx 33$, which can often be observed in the real-life experiments mentioned above.

¹³⁵ Here, we assume without loss of generality that the utility functions of the players are linear.

7.2 Nash Equilibrium

Nash's (1951) seminal work on noncooperative game theory has had a tremendous impact on our present understanding of strategic conflicts. It is the basic solution concept of traditional game theory and has found its way from mathematics into many other disciplines, such as economics, biology, psychology, philosophy, and politics. In this section, I am going to compare Nash's solution concept with the subjectivistic approach. Since Nash equilibrium is used in so many different areas, the arguments presented in this section might be interesting to a broad audience.

7.2.1 Basic Model

Nash considers a finite game in normal form. This means that there are n players and each player possesses a finite action set. He calls each action a "pure strategy." In a (finite) 2-person game, a pure strategy is nothing other than a row or a column of the payoff matrix. The payoff of a player is the utility that he receives if he performs some pure strategy, i.e., chooses some element from his action set, and the other player chooses an element from his own action set.

Further, Nash defines a "mixed strategy," which is a convex combination of pure strategies. More precisely, a mixed strategy s_i assigns a nonnegative number to each element of the set of pure strategies of Player i in such a way that the numbers sum up to 1. Hence, we have that $s_i \in S_i$, where S_i denotes the standard simplex in \mathbb{R}^{r_i} and r_i is the number of pure strategies of Player i . Nash (1951) does not explain how to interpret mixed strategies, but in Nash (1950b) he mentions that s_i represents a (discrete) probability distribution on the set of pure strategies of Player i , i.e., on his action set.

Thus, applying a mixed strategy, s_i , means to select an element from the own action set at random, where the components of s_i specify the probability of each action.¹³⁶ To be more precise, applying a mixed strategy means to assign positive probability to more than one action. By contrast, a player is said to perform a pure strategy if and only if he assigns probability 1 to one action only, in which case it is assumed that he chooses the corresponding action. Actually, the terminology is somewhat misleading because, if the player performs a pure strategy, he could still choose another action with probability 0. For the same reason, it is misleading to treat the terms "pure strategy" and "action" synonymously. However, this is common practice in traditional game theory and I will come back to this issue in Section 7.2.4.

Nash presumes that the actions of the players in a noncooperative game are stochastically independent.¹³⁷ In fact, Nash (1951) writes that

¹³⁶ Throughout this work, s_i is considered a column vector.

¹³⁷ Alternatively, we might consider the mixed strategies random variables and say that the random variables are mutually independent. However, this is not the path that is taken by Nash (1951).

“Our theory [...] is based on the absence of coalitions in that it is assumed that each participant acts independently, without collaboration or communication with any of the others.”

The stochastic-independence assumption is in direct contrast to his approach to cooperative games (Nash, 1950a, 1953), i.e., the bargaining solutions, in which the players are able and willing to coordinate their actions in order to achieve correlation and thus to increase their expected payoffs by cooperation. I will come back to correlated strategies in Section 7.4.

Hence, in a 2-person game, the expected utility of Player 1 amounts to

$$EU_1(s_1, s_2) = s_1' U s_2,$$

where $U = [u_{ij}] \in \mathbb{R}^{r \times s}$ is the decision matrix of Player 1. Here, u_{ij} denotes the utility of Player 1 given that he chooses the i th element and Player 2 chooses the j th element from his action set. Further, $V = [v_{ji}] \in \mathbb{R}^{s \times r}$ shall be the decision matrix of Player 2. Thus,

$$EU_2(s_1, s_2) = s_2' V s_1$$

represents his expected utility. Correspondingly, the expected utility of Player $i = 1, 2, \dots, n$ in a (finite) n -person game is $EU_i(s)$, where $s = (s_1, s_2, \dots, s_n) \in S := \prod_{i=1}^n S_i$ represents the solution of the game and S is the set of solutions.

The following definition goes back to Nash (1951) and can be considered a “mixed-strategy counterpart” to our (more general) definition of equilibrium for coherent n -person games (see Definition 13).

Definition 15 (Nash equilibrium). The solution $s \in S$ is said to be a Nash equilibrium if and only if

$$EU_i(s_1, s_2, \dots, s_i, \dots, s_n) \geq EU_i(s_1, s_2, \dots, t_i, \dots, s_n)$$

for every strategy $t_i \in S_i$ of each player $i \in \{1, 2, \dots, n\}$.¹³⁸

Hence, a Nash equilibrium is a solution of the game in which no player can increase his expected utility by changing his strategy, provided that no other player changes his strategy, too. Obviously, Nash presumes that the players choose their strategies independently, since each opponent holds still while the player himself might change his own strategy in order to improve his position. We conclude that Nash makes use of, or even *introduces*, a standard assumption of noncooperative game theory: strategic independence.

138 A Nash equilibrium is called strict if and only if the inequalities are strict for all players.

Note that strategic independence is distinct from stochastic independence. The former refers to the choice of strategy, whereas the latter refers to the joint probability distribution of actions. Nash presumes that both the strategic-independence assumption and the stochastic-independence assumption are satisfied.

Now, Nash's famous existence theorem reads:

Theorem 9 (Nash equilibrium). *Every finite game has a Nash equilibrium.*

In a short communication (Nash, 1950b), this is proved by applying Kakutani's fixed-point theorem, whereas in the original work (Nash, 1951), i.e., in his PhD thesis, Nash uses Brouwer's fixed-point theorem. Debreu (1952), Fan (1952), and Glicksberg (1952) generalize Nash's existence theorem. Their result has already been used in Section 6.6.2 (see Theorem 7) and it will be useful also in Section 7.5.2.

My personal impression is that Nash's stroke of genius was to use the notion of mixed strategy and to observe that every (Nash) equilibrium is a fixed point. However, it should be noted that mixed strategies had already been introduced by von Neumann (1928) in order to analyze (finite) zero-sum games, and von Neumann (1937) even noticed the importance of Brouwer's fixed-point theorem in that context. Further, a little bit later, Kakutani (1941) himself demonstrated his generalization of Brouwer's fixed-point theorem just by showing that all (proper) 2-person zero-sum games have at least one minimax solution, which is nothing other than a Nash equilibrium. I will come back to this point in Section 7.5.

7.2.2 Some Examples

7.2.2.1 The Game Show

For example, let us consider the game show, which has already been discussed in Section 4.2.3 and whose payoff matrix can be found in Table 7.2. It has infinitely many Nash equilibria, i.e., the points of intersection of the best-response curves that are depicted in Figure 7.1:

- Ann goes up and Bob goes left.
- Bob goes right and Ann goes up with probability lower or equal to $\frac{9}{10}$.

Table 7.2: Payoff matrix of the game show.

Ann	Bob	
	Left	Right
Up	(2, 1)	(9, 0)
Down	(1, -1)	(9, 8)




Figure 7.1: Nash equilibria of the game show (blue intersection points).

Which Nash equilibrium will be manifested? Unfortunately, traditional game theory fails to generate any clear answer to this important question. This is the so-called multiple-solutions problem of traditional game theory (see, e.g., Colman, 2004). Several attempts are made in order to avoid, or at least to mitigate, this problem. However, in my opinion, none of them appear to be entirely convincing. I will explain my reasons in Section 7.6.

By contrast, the subjectivistic approach provides a clear answer, which can be derived from the best-response diagram, too (see Figure 4.1):

- Ann goes up if she doubts that Bob goes right.
- Bob goes left if he thinks that Ann goes up with probability greater than $\frac{9}{10}$. By contrast, he goes right if he expects that Ann goes down with probability greater than $\frac{1}{10}$.

There are only two (nonexcluding) cases in which subjectivistic game theory fails to deliver an unambiguous solution:

1. Ann is convinced that Bob goes right or
2. Bob thinks that Ann goes up with probability $\frac{9}{10}$.

I think that it is very unlikely that Bob's subjective probability is *exactly* $\frac{9}{10}$. Hence, from a practical point of view, the second case is negligible. It remains to think about the first case. Is it possible that Ann is *convinced* that Bob goes right? Let us assume that the monetary payoffs of Ann and Bob in the game show are common knowledge. This means that Ann and Bob know the payoff matrix, they know that the other knows the payoff matrix, they know that the other knows that the other knows the payoff matrix, etc. Now, if Ann believes that Bob believes that she will go up if she is not convinced that he goes right, she could fear that he is quite sure that she goes up, which means that he goes left. In this case, in fact, she will go left, which can be viewed as a self-fulfilling prophecy.

There are plenty of reasons why Ann might not be convinced that Bob goes right but also some reasons that support the opposite argument. For example, Ann and Bob could have agreed to choose Down/Right backstage, i.e., before the game show started, and Ann might be truly convinced that she can rely on Bob. In any case, the key finding of our subjectivistic approach is that Ann must be *convinced* that Bob goes

right, whereas Bob should only *tend to believe* that Ann goes down, in order to guarantee the first-best solution Down/Right. This means that we should concentrate on Ann's sentiments if our goal is to provoke the first-best solution. This can be done by communication and, if possible, by signaling (see Section 2.5.3). The latter might be impossible in the game show, but nonetheless it can very well be possible in other situations that occur in real life.

7.2.2.2 The Chicken Game

Now, let us consider the chicken game (see Section 4.3.1.1), whose payoff matrix is depicted in Table 7.3. In a 2×2 game we can easily characterize any solution $s = (s_1, s_2)$ by (p, q) , where p is the probability of Player 1 to choose his first action and q is the probability of Player 2 to choose his first action.

Table 7.3: Payoff matrix of the chicken game.

Andy	Bob	
	Swerve	Straight
Swerve	(1, 1)	(0, 2)
Straight	(2, 0)	(-9, -9)

The chicken game possesses three Nash equilibria, i.e., the pure-strategy Nash equilibria (0, 1) and (1, 0) plus the mixed-strategy Nash equilibrium (0.9, 0.9). The probability of 90 %, which constitutes the mixed-strategy Nash equilibrium, has already been calculated by Equation 4.4. It represents the change points of the best-response curves in Figure 4.4. The Nash equilibria correspond to the points of intersection, which are marked blue in Figure 7.2.




Figure 7.2: Nash equilibria of the chicken game (blue intersection points).

The Nash equilibria can be interpreted thus:

1. Andy goes straight and Bob swerves.
2. Bob goes straight and Andy swerves.
3. Both drivers swerve with probability 90 %.

What does this solution tell us? It says that one driver swerves, whereas the other driver goes straight.¹³⁹ Another possibility is that both drivers randomize their actions in order to swerve with probability 90 %.

Interestingly, in order to calculate the strategy of a player in a mixed-strategy Nash equilibrium, we must take the utilities of the *other* player into account. Hence, according to traditional game theory, the optimal strategy of a player depends on the individual preferences of his opponent! This is in direct contrast to the subjectivistic solution, which has been presented in Section 4.3.1.1. In subjectivistic game theory, 90 % is not a frequentistic probability. It represents a critical threshold, and for calculating the critical threshold of a player, we need to know only his *own* utilities. That is, the individual preferences of the opponent are irrelevant.

The subjectivistic solution tells us that both drivers swerve if they fear that the other goes straight, i.e., if their subjective probabilities of Straight exceed 10 %. This solution is rational. However, it is not a Nash equilibrium. Nonetheless, the subjectivistic solution can frequently be observed in everyday life and it is hard to doubt that swerving is a wise decision if a driver fears that the other goes straight. Moreover, in the subjectivistic framework, each player commits himself to a definite *action*. That is, the players do not choose their actions at random.

We conclude that, although the strategic-independence assumption is satisfied in the chicken game, Nash equilibrium is not a necessary condition for rationality. This can be stated thus: $R \not\Rightarrow N$. Here, “*R*” stands for the rationality of the players and “*N*” means that the solution of the game is a Nash equilibrium.

7.2.2.3 The Ultimatum Game

We just noticed that $R \not\Rightarrow N$. The following example demonstrates that $N \not\Rightarrow R$, which means that Nash equilibrium is neither a sufficient condition for rationality. More precisely, Nash equilibria can be subgame imperfect. This observation goes back to Selten (1965, 1973).¹⁴⁰

For example, consider the ultimatum game, which has already been discussed in Section 7.2.2.3. Figure 7.3 contains the corresponding game tree on the upper left. Remember that a game tree is not a decision tree, since it combines the decision nodes of Ann and Bob (see Section 3.3.2). A similar argument holds true for the payoff matrix, which combines the (reduced) decision matrices of both players (see Section 3.3.1).

The given payoffs shall indicate that now, in contrast to Section 7.2.2.3, both players focus on the monetary consequences of the ultimatum game. This means that fair-

¹³⁹ Actually, it can still happen that a driver chooses his suboptimal action just by accident, but this possibility is typically ignored in traditional game theory.

¹⁴⁰ The term “subgame perfect” is first used in Selten (1975). In his earlier works, Selten just refers to “perfect” Nash equilibria. I will come back to this point in Section 7.6.3.

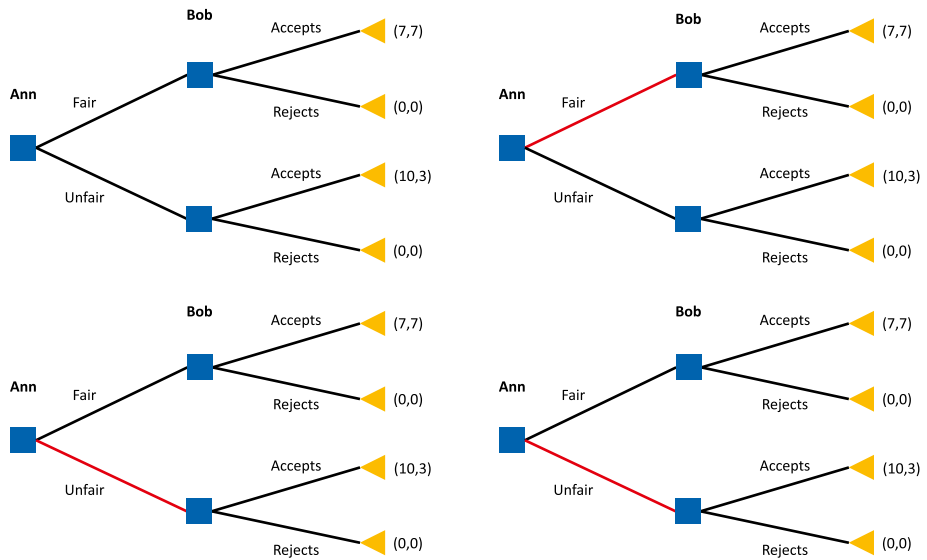


Figure 7.3: Game tree (upper left) and Nash equilibria of the ultimatum game.

ness plays no role in the traditional context. Ann can decide only to make either a fair or an unfair split, whereas Bob has four strategies:

1. Accept irrespective of whether or not Ann is fair.
2. Reject if Ann is fair and accept if she is unfair.
3. Accept if Ann is fair and reject if she is unfair.
4. Reject irrespective of whether or not Ann is fair.

We can find three (pure-strategy) Nash equilibria, which are depicted on the upper right and the lower part of Figure 7.3. The Nash equilibria on the upper and lower right violate the Bellman postulate, i.e., they are subgame imperfect (Selten, 1975): Bob will not reject Ann's offer if she is unfair (upper right) and he will not reject it either if she is fair (lower right). Put another way, in each subgame, he will act according to the basic principles of rational choice.

In order to discover subgame-imperfect Nash equilibria, we must focus on Bob's decision nodes that are *not* reached after Ann has performed her optimal action. For example, on the upper right of Figure 7.3, Bob could have either accepted or rejected Ann's offer if she *were* unfair, even though she actually decided to be fair. Put another way, subgame-imperfect Nash equilibria can be discovered only by counterfactual reasoning. The reader can find the unique subgame-perfect Nash equilibrium of this game on the lower left of Figure 7.3.

Traditional game theory claims that it is irrelevant whether we solve a game in the normal or extensive form. According to von Neumann and Morgenstern (1953, p. 85), these two forms are strictly equivalent. I think that this assertion is seriously misleading and it is Selten's contribution to demonstrate that solving a dynamic game in the

normal form can lead to wrong conclusions. Of course, the sequential structure of a dynamic game matters, but the problem is that it cannot be formulated in an appropriate way by using a payoff matrix. In the subjectivistic context, I always refer to a static or a coherent one-shot game when using the normal form. By contrast, I usually refer to a dynamic game when using the extensive form. In some cases it is convenient to use the extensive form not only for dynamic but also for static games.¹⁴¹

The subjectivistic solutions of the ultimatum game, which we already discussed in Section 7.2.2.3, do not suffer from subgame imperfectness. They are simply derived by backward induction and so the Bellman postulate is always satisfied.

Subgame imperfectness reminds us of the usual deficiencies that occur if we try to solve an irregular decision problem by forward deduction, i.e., by using a decision matrix (see Section 2.8.2). For example, suppose that Bob is convinced, *a priori*, that Ann makes an unfair split (see Figure 7.4). However, when he is surprised, *a posteriori*, by Ann's charity, he immediately revises his opinion and accepts her offer. This kind of behavior cannot be explained when trying to solve Bob's decision problem by forward deduction. The problem is that, *a priori*, it makes no difference to him whether or not he accepts Ann's offer if she is fair because he is sure that she will be unfair.




Figure 7.4: Bob's decision problem in the ultimatum game can be irregular.

This example shows that solving a dynamic game in the normal form can, at best, be ineffective. That is, it can lead to multiple solutions, which can be usually avoided by using a decision tree and applying the method of backward induction.

7.2.3 Bacharach's Transparency of Reason

Nash (1951) leaves us with the fundamental question of whether or not the players know the strategy of each other. We could be inclined to think that Nash's approach is

¹⁴¹ This has already been demonstrated, e.g., in Section 4.2.3.

objectivistic. Then the players know not only their own strategies but also the strategy of each other. In fact, a strategy, let it be mixed or pure, is just a probability distribution and the objectivistic approach holds that all subjects agree on the same, i.e., objective, probability measure. However, if the strategies are known to each player, why should their strategy choices be independent? We argued in Chapter 6 that the players interact if they know the strategy of each other and the same principle can be applied in this context: Performing a strategy means to choose some random generator, but if some opponent knows that I choose Random Generator A, he could respond with Random Generator B, etc. Hence, let us suppose that the players do *not* know the strategy of any other in order to legitimate the strategic-independence assumption. More precisely, let c_{ij} be the conjecture of Player i about Player $j \neq i$. This means that Player i believes that Player j performs Strategy c_{ij} . Further, let s_i be the *true* strategy of Player i .¹⁴²

Recall that in the subjectivistic framework a solution is said to be rational if and only if all players choose an optimal action or strategy (see Definition 10). Here, the term “strategy” is understood as a pure or mixed strategy in the sense of Nash (1951). Thus, a solution of the game is considered rational if and only if all players maximize their expected utilities given their conjectures about each other.

The following theorem provides a sufficient (but not a necessary) condition for Nash equilibrium and can be attributed to Bacharach (1987).

Theorem 10 (Transparency of reason). *If the conjectures of the players about each other are correct and the solution of the game is rational, it must be a Nash equilibrium.*

Proof. If the solution $s = (s_1, s_2, \dots, s_n)$ of the game is rational, we must have that

$$EU_i(c_{i1}, c_{i2}, \dots, s_i, \dots, c_{in}) \geq EU_i(c_{i1}, c_{i2}, \dots, t_i, \dots, c_{in})$$

for each strategy t_i . Further, if the conjectures of the players about each other are correct, we have that $c_{ij} = s_j$ for $i, j = 1, 2, \dots, n$. Hence, after substituting c_{ij} with s_j for all $j \neq i$, the strategy tuple s constitutes a Nash equilibrium. \square

Theorem 10 is labeled “transparency of reason” because it presumes that the players can step into the shoes of each other and deduce his strategy. Then the resulting solution must be a Nash equilibrium—provided that each player performs an optimal strategy.

A typical argument in favor of Nash equilibrium, which is related to the transparency of reason, goes like this: Suppose that all players in a static one-shot game have complete information and common knowledge. Further, assume that there exists one and only one Nash equilibrium. Each player knows that no player has an incentive to change his strategy if and only if the strategies are in equilibrium. In other

¹⁴² It is implicitly assumed that $c_{ii} = s_i$ for $i = 1, 2, \dots, n$, i.e., the players’ conjectures about themselves are always correct.

words, the Nash equilibrium is the only stable solution. He knows also that everybody else knows this fact. Further, he knows that everybody else knows that everybody else knows this fact, etc. This means that the players will perform their equilibrium strategies by self-enforcement.

This argument sounds great—and I must admit that I was very taken by it at the beginning of my studies in game theory. However, now I think that this is a misuse of the concept of common knowledge. In fact, it is true that everybody knows that everybody knows, etc., that there exists one and only one Nash equilibrium in that game, but this does not mean at all that anybody will actually perform his equilibrium strategy. The reason is simple: It is a static one-shot game. Thus, each player chooses his strategy independently and nobody must fear anything if he deviates from the Nash equilibrium. Apart from that, most games possess more than one Nash equilibrium, in which case the overall argument breaks down, since nobody knows which Nash equilibrium is the “right one.” I will come back to this important issue in Section 7.6.

7.2.4 Do We Make Our Choices at Random?

Why should the conjectures of the players about their adversaries be correct? More specifically, why should the transparency of reason make any sense? This question has already been tackled on a quite general level in Section 7.1.1: The typical argument of traditional game theory is common knowledge. In Section 7.1.3 we discussed some examples that aim at supporting the transparency of reason. However, in Section 8.1 we will see that common knowledge does *not* imply the transparency of reason, but this shall not bother us here. There is still one open question: Should the players apply a mixed strategy at all?

A player who applies a mixed strategy, i.e., who chooses his action at random, sacrifices a control variable. Applying a mixed strategy means that the player substitutes a decision node with a chance node in his decision tree. Of course, since deliberation crowds out prediction, we cannot have both a decision node and a chance node at the same point. Is it meaningful to make our choices at random, even if we strictly follow the equilibrium doctrine? I think that the answer is “No”!

The overall concept of Nash equilibrium tells us that a player should combine only actions that are *optimal*. For example, consider the mixed-strategy Nash equilibria of the game show, which are depicted in Figure 7.1. Ann is indifferent among Up and Down if Bob goes right with probability 1. For this reason, she can simply choose to go up or to go down. Hence, why should Ann apply any mixed strategy at all? The same argument applies to Bob: If Ann goes up with probability 90 %, he can simply choose to go left or to go right. There is no reason for him to make his action at random. Note that our arguments are based on the standard assumption of noncooperative game theory, i.e., strategic independence.

Here is another example, which I have found in Rubinstein (1991): An employer can monitor only one of two workers, say *A* and *B*, at the same time. The workers can choose between two levels of effort, i.e., either High or Low. They cannot see whether or not they are being monitored, but they know that if they are caught being lazy, the employer will cut their salary. Suppose that each worker's utility function is such that he prefers to choose a high level of effort if the employer decides to monitor him with probability greater or equal to 50 % (see Figure 7.5). Here, we assume that the probabilities are *objective*.




Figure 7.5: Worker's decision problem in the monitoring game ($0 < \varepsilon < 5$).

The employer wants to guarantee that *both* workers choose a high level of effort because cutting their salary is only a drop in the bucket, compared with the harm a worker causes to the company if he decides to be lazy. Let $p \in [0, 1]$ be the probability that he monitors Worker *A*. It is assumed that the employer knows that his workers prefer to choose a high effort if $p \geq 0.5$ and that p is common knowledge. If the employer chooses $p < 0.5$, Worker *A* could decide to be lazy and in the case of $p > 0.5$, Worker *B* could lose his motivation. In particular, if the employer chooses $p \in \{0, 1\}$, it is sure that either worker decides to be lazy. Hence, the best choice of the employer is $p = 0.5$, in which case it is always optimal for *both* workers to choose a high effort.

Does this mean that it is better for the employer to perform a mixed rather than a pure strategy, and that performing a mixed strategy is even better than making a *definite* choice? Well, a closer inspection reveals that the employer does not perform any mixed strategy at all! As is pointed out by Rubinstein (1991), his action set is $[0, 1]$, not $\{\text{Monitor A, Monitor B}\}$. Indeed, it would make no (essential) difference whether the employer picks out Worker *A* or Worker *B*. Either way, nobody knows whether or not he is being monitored. It is much more important to select the parameter p in an optimal way, since this parameter is common knowledge and, de facto, determines the individual behavior of the employees.

Taking a subjectivistic perspective, we conclude that the employer should only guarantee that each worker *believes* that he is being monitored with probability

greater or equal to 50 %. This can be achieved by communication or, which might be more expensive, by signaling. We could imagine, for example, that the firm installs an unequivocal surveillance system, in which case the workers *know* that they are being monitored. However, it remains questionable if anybody would really enjoy working for Big Brother, but this is a different kettle of fish.

Choosing a number p from the interval $[0, 1]$ does not mean to apply a mixed strategy in the sense of Nash.¹⁴³ When the employer decides to choose the number 0.5 from his action set $[0, 1]$, he makes this choice deliberately. Applying a random generator in order to choose the parameter p makes no sense at all because, anyhow, its choice should be restricted to the optimal subset of $[0, 1]$, i.e., to $\{0.5\}$. However, even if this condition is satisfied, it could still happen that the random device produces a wrong parameter just by accident, i.e., with probability 0.

The same question arises in every other decision problem: Why should a rational subject in a game against Nature randomize his action? Usually, this leads to even more uncertainty. If he mixes up optimal and suboptimal actions, it can happen that he chooses a suboptimal one. Otherwise, randomization makes no sense at all, since the decision maker could simply choose *any* optimal action, deliberately, in order to receive the same expected utility. The same observations have already been made in Section 2.4.3.

7.2.5 Subjectivistic Nash Equilibrium

The transparency of reason claims that the players are able to deduce the strategies of their opponents by stepping into their shoes. Complete information and common knowledge are typical arguments of traditional game theory in order to justify the transparency of reason. However, the specific reasons why it should be possible for a player to step into the shoes of each other are of secondary importance. In the subjectivistic context, we could use any argument that is sufficient to guarantee that the players are able to deduce the actions of their opponents.

In order to understand the content of the following theorem, recall that an action in a finite 2-person game in normal form corresponds to a row or a column in the payoff matrix. Here, we still assume that the game is static, but the action set of each player can be infinite and even uncountable. However, we continue making the standard assumption that each player believes that his action has no influence on the action of his adversary. Moreover, in the subjectivistic framework, a solution of the game always refers to the actions of the players—not to pure or mixed strategies—and a rational solution is just a combination of optimal actions.

143 Moreover, since the action set of the employer, $[0, 1]$, is uncountable, the game is not finite.

Theorem 11 (Subjectivistic Nash equilibrium). *Consider any static 2-person game in normal form and fix the optimal actions. If both players believe that the other chooses an optimal action, their subjective probabilities constitute a Nash equilibrium.*

Proof. If a player believes that the other chooses an optimal action, the support of his probability distribution is a subset of the set of optimal actions of his opponent. Hence, the subjective probabilities constitute a Nash equilibrium. \square

Thus, if the players in a finite 2-person game are able to deduce the optimal actions of the other, and if they are convinced that the adversary does not choose a suboptimal action, their conjectures must form a Nash equilibrium.

This holds true if the players know the conjectures and payoffs of one another. Note that this is *more* than complete information, since complete information does not require that any player knows the conjecture of another. However, this is precisely what we are talking about, namely the transparency of reason. To sum up, whenever we assume that both players know which actions their adversary prefers, and that they believe that the other does not fail to perform an optimal action, we obtain a (subjectivistic) Nash equilibrium.

Theorem 11 translates the transparency of reason, i.e., Theorem 10, into our subjectivistic framework and demonstrates the role of Nash equilibrium in this context. However, it is worth emphasizing that our understanding of “Nash equilibrium” differs substantially from traditional game theory. Here, we do not speak about mixed strategies but rather about the subjective probability distributions of the players.

Note that Theorem 11 does not presume that each player knows the optimal actions of his adversary. It only presumes that both players believe that the other chooses an optimal action. However, we must be very careful when interpreting the sentence “If a player believes that the other chooses an optimal action.” When we read this sentence separately, it almost becomes void, since then it just suggests that the player believes that his adversary makes an optimal choice *in any case*. However, this is not meant here. The theorem presumes that the optimal actions of both players are *fixed* in advance. Hence, if a player believes that the other chooses an optimal action, then he must assign the actions that are considered optimal by his opponent *in the particular case* a positive probability.

In real life, it can very well happen that a player considers every action that is suboptimal for the other player null and the same can be true, *mutatis mutandis*, for the other player.¹⁴⁴ This shall be demonstrated by the following example: Suppose that Ann and Bob are playing some finite game repeatedly. After a while, Ann observes that Bob regularly avoids some actions that belong to a subset $A_{\text{sub}} \subset A$ of his action set A . She comes to the conclusion that the actions in A_{sub} are suboptimal for Bob

¹⁴⁴ In particular, if each player has only *one* optimal action, and decides to choose this action, the solution of the game represents a pure-strategy Nash equilibrium.

and so she considers the corresponding columns in the payoff matrix null, whereas she assigns all other columns a positive probability. Let us assume that Ann is right, which means that Bob avoids the actions in A_{sub} just because they are, in fact, not optimal for him. By contrast, he considers all actions in $A \setminus A_{\text{sub}}$ optimal, which is the reason why he does not avoid those actions. If we apply the same arguments to Bob, we obtain a situation in which the conditions of Theorem 11 are satisfied, in the course of time, even without complete information or common knowledge.

Since Ann and Bob make their choices repeatedly, their actions appear to be stochastic. Put another way, it seems to us that the players apply a mixed strategy, but in fact they are just indifferent among their optimal actions and thus alternate their choices. Suppose that the players act in a way such that the (strong) law of large numbers is satisfied.¹⁴⁵ Then the relative number of each action converges to some limit, which is called “objective probability.” Moreover, if we assume that the subjective probabilities of the players get, more and more, closer to the objective ones, we end up with a Nash equilibrium in the traditional sense. Here subjectivistic game theory meets traditional game theory. Similar results are frequently obtained in evolutionary game theory (Smith, 1982).

Let us consider Matching Pennies, which has been discussed in Section 4.3.2.1. Its payoff matrix is given in Table 7.4. Suppose that Ann thinks that Bob chooses Heads and Tails each with probability 50 %, in which case both Heads and Tails are optimal for Ann. The same holds true for Bob if he thinks that Ann chooses Heads and Tails each with probability 50 %. Since all actions are optimal for Ann and for Bob, the conditions of Theorem 11 are clearly satisfied. Hence, the subjective probabilities of Ann and Bob must constitute a Nash equilibrium. In fact, this is true, as we can see in the best-response diagram of Matching Pennies, which is given in Figure 7.6.

Table 7.4: Payoff matrix of Matching Pennies.

Ann	Bob	
	Heads	Tails
Heads	(1, -1)	(-1, 1)
Tails	(-1, 1)	(1, -1)




Figure 7.6: Best-response diagram of Matching Pennies.

¹⁴⁵ More precisely, their actions form an ergodic stationary sequence.

Now, assume that Ann is convinced that Bob chooses Heads, which means that she prefers Heads. Conversely, Bob is convinced that Ann chooses Tails and thus, de facto, he prefers Heads, too. However, Bob is wrong and so the conditions of Theorem 11 are violated. For this reason, we must not apply Theorem 11 in the *modus ponens*. Indeed, the response diagram in Figure 7.6 reveals that Ann's and Bob's subjective probabilities do not constitute a Nash equilibrium. However, we can apply Theorem 11 in the *modus tollens*: Either Ann or Bob must *not* be convinced that the other applies an optimal action. In fact, it is Bob who doubts that Ann chooses Heads, although Heads is actually her (only) optimal action.

Thus, it is impossible for both players to have only *one* optimal action and to deduce the (optimal) action of the other player by the transparency of reason—given that no player doubts that his adversary performs his optimal action! If each player is convinced that the other one chooses his optimal action, then their subjective probabilities must constitute a Nash equilibrium, but the only Nash equilibrium of Matching Pennies consists of mixed strategies, which means that no player can be convinced that the other one chooses his (unique) optimal action. This quintessence holds true for all finite 2-person games without a pure-strategy Nash equilibrium.

7.3 Bayesian Nash Equilibrium

Although the notion of complete information plays no role in Nash (1951), we already know that this fundamental assumption is often used in order to justify the equilibrium doctrine. It is assumed also that the players have common knowledge about all relevant aspects of the game. Then the players can step into the shoes of their opponents and are able to deduce their strategies by the transparency of reason. Hence, the rational solution of the game must be a Nash equilibrium.

Well, at least this is the basic storyline that I followed for many years in my lectures on game theory, but we will see in Section 8.1 that this was a fallacy! It will be shown that the rational solution of a game need not be a Nash equilibrium—even if the players have complete information and the assumption of common knowledge is satisfied. However, here I follow the usual arguments of traditional game theory, which suggest that the rational solution of the game must be a Nash equilibrium if the players have complete information and common knowledge.

The problem is that in most real-life situations we cannot see through the persons that are opposed to us. Of course, this holds true for each other player in the game. This means that the players have incomplete information, in which case the common-knowledge assumption is violated, too. Can we still expect that the solution of the game is a Nash equilibrium? This question is answered by Harsanyi (1967–1968). His nifty idea is to transform incomplete into imperfect information and then to proceed in the usual way. How does this work?

7.3.1 Basic Model

Harsanyi makes the simple assumption that, in a preliminary step, Nature has chosen some type of person. More precisely, before the game starts, each player has been chosen at random from a specific set of types. Every type contains a complete specification of all things that the players must know in order to deduce the strategy of each other by the transparency of reason. Here, the term “strategy” is understood in the sense of Nash (1951).

We may focus on a 2-person game for the sake of simplicity. Let $T = \{1, 2, \dots, m\}$ be some (index) set of possible types. Every player has been chosen at random from the set T . Let p be the corresponding probability distribution on $T \times T$. It is not necessary to allow both players to belong to each type in T . For example, if we wish to assume that Player 1 is of some type in $T_1 \subseteq T$, whereas Player 2 is of some type in $T_2 \subseteq T$ (with $T_1 \neq T_2$), we can simply set $p(T_1 \times T_2) = 1$. However, in the following I assume that the players have a common type set, T , for the sake of simplicity but without loss of generality.

Each player knows his own type but not the type of the other. That is, Player 1 has a *conditional* probability distribution, i.e., a posterior, on T , which expresses his individual uncertainty about the type of Player 2. Let p_{ij} be the (conditional) probability that Player 2 belongs to Type $j \in T$ given that Player 1 has Type $i \in T$. Further, let U_{ij} be the decision matrix of Player 1 given that he belongs to Type i and Player 2 is of Type j . Hence, the payoffs of Player 1 may depend both on his own type and on the type of Player 2.

Each type performs some strategy. The strategy of Player 1, who is of Type i , is denoted by s_{1i} . Correspondingly, s_{2j} is the strategy of Player 2, who belongs to Type j . The expected utility of Player 1 reads

$$EU_{1i}(s_{1i}, s_{21}, s_{22}, \dots, s_{2m}) = s'_{1i} \sum_{j=1}^m p_{ij} U_{ij} s_{2j}$$

and the expected utility of Player 2 is

$$EU_{2j}(s_{2j}, s_{11}, s_{12}, \dots, s_{1m}) = s'_{2j} \sum_{i=1}^m p_{ji} V_{ji} s_{1i},$$

where V_{ji} denotes his decision matrix.

Now, Player 1 maximizes his expected utility, given the strategies $s_{21}, s_{22}, \dots, s_{2m}$ of Player 2 and the same holds true for Player 2, given the strategies $s_{11}, s_{12}, \dots, s_{1m}$ of Player 1. If all strategies are a best response to each other, i.e., if

$$EU_{1i}(s_{1i}, s_{21}, s_{22}, \dots, s_{2m}) \geq EU_{1i}(t_{1i}, s_{21}, s_{22}, \dots, s_{2m})$$

and

$$EU_{2j}(s_{2j}, s_{11}, s_{12}, \dots, s_{1m}) \geq EU_{2j}(t_{2j}, s_{11}, s_{12}, \dots, s_{1m})$$

for every t_{1i}, t_{2j} and all $i, j \in T$, we have a Bayesian Nash equilibrium.

Games that are constructed in such a way are called “Bayesian” because the players act conditionally on their own, i.e., private, information. We could take much more complicated Bayesian games into consideration, in which the number of players exceeds 2 and each player has his own type set, etc., but this would not alter Harsanyi’s principal idea. To be more precise, the players in a Bayesian game have complete information because they (supposedly) know all things that are necessary in order to deduce the strategy for each type of the other player, but in general they do not know to which type the other player belongs and so they suffer (only) from imperfect information.

Harsanyi (1967–1968) assumes that the players share the same prior P . This means that the joint probability distribution of types is common to all players. Nonetheless, each player has his own private information and so the players may come to different conclusions. More precisely, their posteriors differ from one another because they know their own types but not the type of the other player. Hence, their individual assessments are different, but not because they have different subjective probability measures. The reason is that the players are equipped with different information.

The common-prior assumption plays a major role in traditional game theory and is sometimes called Harsanyi doctrine (Aumann, 1987). In my opinion, it violates the principal idea of subjectivistic decision theory. If the probability measure is common to all players, is it then still subjective? I would clearly say “No!” If everybody agrees upon the prior, it becomes intersubjectively valid and thus objective (Kant, 1787). Thus, in my opinion, Harsanyi’s model is indeed Bayesian, but it is far away from being subjectivistic. I will come back to this issue in Section 7.4.4.1.

7.3.2 Some Examples

7.3.2.1 The Sheriff’s Dilemma

For example, consider the following dilemma: A sheriff sees an armed suspect, but he does not know whether he is a criminal or a civilian. Both the sheriff and the suspect must decide independently whether or not to shoot the other. Hence, the suspect can be either of the type Criminal or of the type Civilian. Let p be the probability that the suspect is a criminal. The payoff matrices of the sheriff’s dilemma are given in Table 7.5.

The payoffs can be explained like this: If nobody shoots, the players do not harm one another and so their payoffs are zero. Whenever the sheriff shoots, he has a bad conscience and loses 1 util. The same holds true for the civilian, whereas the criminal has no conscience at all. He feels even happy if he shoots the sheriff and wins 2 utils. If any player is shot, he loses 2 utils, but the sheriff wins 2 utils if he shoots a shooter because he has just made his job—irrespective of whether it is a criminal or a civilian.

Table 7.5: Criminal's (left) and civilian's (right) payoffs.

Sheriff	Criminal	
	Shoot	Hesitate
Shoot	$(-1, 0)$	$(-1, -2)$
Hesitate	$(-2, 2)$	$(0, 0)$

Sheriff	Civilian	
	Shoot	Hesitate
Shoot	$(-1, -3)$	$(-1, -2)$
Hesitate	$(-2, -1)$	$(0, 0)$

The sheriff knows that the suspect is a criminal with probability p and a civilian with probability $1 - p$. Moreover, it is clear to him that a criminal considers Shoot a strictly dominant action, whereas Hesitate is a strictly dominant action from a civilian's perspective. Suppose that the sheriff decides to shoot. Then his expected utility is

$$p(-1) + (1 - p)(-1) = -1.$$

By contrast, if he decides to hesitate, his expected utility is

$$p(-2) + (1 - p)0 = -2p.$$

It follows that the sheriff will shoot if $p > \frac{1}{2}$, he will hesitate if $p < \frac{1}{2}$, and in the case of $p = \frac{1}{2}$ he is indifferent, which means that he can choose any mixed strategy. However, a criminal will always shoot, whereas a civilian will always hesitate. The Bayesian Nash equilibria can be found in Table 7.6

Table 7.6: Bayesian Nash equilibria of the sheriff's dilemma.

	Sheriff	Criminal	Civilian
$p > \frac{1}{2}$	Shoot	Shoot	Hesitate
$p < \frac{1}{2}$	Hesitate	Shoot	Hesitate
$p = \frac{1}{2}$	Indifferent	Shoot	Hesitate

7.3.2.2 Friend or Foe

The following game is similar to the sheriff's dilemma: Two soldiers face one another on a battlefield and cannot see whether the other soldier is a friend or a foe.¹⁴⁶ They must immediately decide whether to shoot or not. The type set of this game is $T = \{\text{Country A, Country B}\}$. Suppose that p is the probability that a soldier from Country A is faced with a friend, whereas q is the probability that a soldier from Country B is faced with a friend. The payoffs of this game are given in Table 7.7. If a player shoots an enemy, he wins 1 util, but if he shoots a friend, he loses 1 util. Further, if any player is shot, he loses 2 utils.

¹⁴⁶ Hence, I call this game Friend or Foe. However, the same name is sometimes used for Split or Steal, which is a variant of the prisoners' dilemma (see Section 4.5.2).

Table 7.7: Payoff matrix of friends (left) and foes (right).

Friend	Friend		Foe	Foe	
	Shoot	Hesitate		Shoot	Hesitate
Shoot	(-3, -3)	(-1, -2)	Shoot	(-1, -1)	(1, -2)
Hesitate	(-2, -1)	(0, 0)	Hesitate	(-2, 1)	(0, 0)

Player 1 performs the strategy $s_{11} = (a, 1 - a)$ if he is from Country A and $s_{12} = (b, 1 - b)$ if he is from Country B. Correspondingly, Player 2 performs the strategy $s_{21} = (a, 1 - a)$ if he is from Country A and $s_{22} = (b, 1 - b)$ if he is from Country B. Here, a and b are the probabilities to shoot, whereas $1 - a$ and $1 - b$ are the probabilities to hesitate. Suppose that Player 1 is from Country A and decides to shoot. Then his expected utility amounts to

$$\begin{aligned} EU_A(\text{Shoot}) &= p[a(-3) + (1 - a)(-1)] + (1 - p)[b(-1) + (1 - b)1] \\ &= (1 - p)(1 - 2b) - p(1 + 2a). \end{aligned}$$

By contrast, if he decides to hesitate, he obtains the expected utility

$$EU_A(\text{Hesitate}) = pa(-2) + (1 - p)b(-2) = -2[pa + (1 - p)b].$$

Now, let us suppose that $a = 0$ and $b = 1$. This means that a player hesitates if he is from Country A, whereas he shoots if he is from Country B. Since Player 1 is from Country A, he prefers to hesitate. In fact, this happens if and only if $p > \frac{1}{2}$: The expected utility of Hesitate amounts to $-2 + 2p$ and the expected utility of Shoot is -1 . Thus, Hesitate is better than Shoot if and only if $p > \frac{1}{2}$. Hence, if it is more probable that Player 2 comes from Country A, too, rather than from Country B, Player 1 will prefer to hesitate, whereas he prefers to shoot if and only if it is more probable that Player 2 comes from the enemy-country B.

Now, suppose that Player 1 stems from Country B. Then we have the expected utility

$$\begin{aligned} EU_B(\text{Shoot}) &= q[b(-3) + (1 - b)(-1)] + (1 - q)[a(-1) + (1 - a)1] \\ &= (1 - q)(1 - 2a) - q(1 + 2b) \end{aligned}$$

in the case in which he decides to shoot and

$$EU_B(\text{Hesitate}) = qb(-2) + (1 - q)a(-2) = -2[qb + (1 - q)a]$$

if he decides to hesitate. Given that our initial assumption $a = 0$ and $b = 1$ holds true, the former amounts to $1 - 4q$ and the latter is $-2q$. Hence, Player 1 prefers to shoot if and only if $q < \frac{1}{2}$. Remember that q is the probability that Player 2 comes from Country B, too, whereas $1 - q$ is the probability that he comes from the enemy-country A. Thus,

we have proved that $a = 0$ and $b = 1$ constitute a Bayesian Nash equilibrium, given that $p > \frac{1}{2}$ and $q < \frac{1}{2}$.

How can we interpret the probabilities p and q ? Suppose that Friend or Foe is a so-called population game in which Nature chooses the type of the players at random with replacement from the set T . Then the probability of Country A is p and the probability of country B is $q = 1 - p$. Thus, $p > \frac{1}{2}$ is just equivalent to $q < \frac{1}{2}$. Whenever we have that $p > \frac{1}{2}$, the population of Country B represents a minority. In this case, a soldier coming from Country B prefers to shoot because the probability to kill a friend is relatively low. By contrast, a soldier coming from Country A will hesitate, since the probability to kill a friend is relatively high. It is clear that we obtain the opposite result if $p < \frac{1}{2}$, in which case the population of Country A represents a minority and so we have that $a = 1$ and $b = 0$.

Thus, we may distinguish three cases, i.e., $p > \frac{1}{2}$, $p < \frac{1}{2}$, and $p = \frac{1}{2}$. It remains to analyze the case in which the populations are even. If we assume that Player 1 stems from Country A, his expected utility turns out to be $-(a + b)$, irrespective of whether he shoots or not. The reader can easily verify that the same holds true if Player 1 belongs to Country B, and precisely the same arguments can be applied to Player 2. Hence, a soldier can choose any pure or mixed strategy, i.e., all combinations of a and b are optimal in the case of $p = \frac{1}{2}$. The Bayesian Nash equilibria of Friend or Foe can be found in Table 7.8.

Table 7.8: Bayesian Nash equilibria of Friend or Foe.

	Country A	Country B
$p > \frac{1}{2}$	$a = 0$	$b = 1$
$p < \frac{1}{2}$	$a = 1$	$b = 0$
$p = \frac{1}{2}$	$0 \leq a \leq 1$	$0 \leq b \leq 1$

The previous arguments make sense only if both players have incomplete information or, equivalently, if they do not know the type of the other player. Otherwise, a soldier will always hesitate to shoot a friend, whereas he will always shoot a foe. The reason is that Hesitate strictly dominates Shoot on the left-hand side and Shoot strictly dominates Hesitate on the right-hand side of Table 7.7. However, the problem is that the soldiers do not know the type of one another and so they will perform the wrong action with probability $\min\{p, 1 - p\}$.

7.3.2.3 The Subjectivistic Solutions

In our subjectivistic framework, Bayesian thinking is a matter of course. Since we need not stick to Nash equilibrium, the overall solution concept is less complicated and more flexible. First of all, this shall be demonstrated by the sheriff's decision tree in

Figure 7.7. It is very easy to take incomplete information into account when working with decision trees and a special theory is not required at all. The type of the player is nothing other than a possible event. This example nicely demonstrates that the location of nodes in a decision tree need not have any chronological meaning. Indeed, in the sheriff's decision tree, Nature's choice occurs *after* the sheriff's action. This is because the sheriff does not know the type of the suspect. Nonetheless, it is clear that Nature actually makes its choice *before* the sheriff decides whether to shoot or not.




Figure 7.7: Sheriff's decision tree.

The decision tree reveals that the sheriff shoots if he believes that the suspect is a criminal with probability 80 %. Moreover, the sheriff is convinced that a criminal shoots, whereas he thinks that a civilian hesitates.¹⁴⁷ Of course, we could change the (conditional) probabilities in an arbitrary way. For example, the sheriff might not be convinced about the behavior of criminals and civilians. He even could make the unconditional probabilities of Nature dependent on his own action: If he shoots, Nature chooses a criminal with probability 60 %, but if he hesitates, Nature delivers a criminal with probability 100 %. In this case, the sheriff would be pessimistic. It is worth emphasizing that the probabilities in the decision tree are *subjective* and so we are not bound to any physical laws. Moreover, the common-prior assumption plays no role in our framework because each player has his own state space.

¹⁴⁷ Nonetheless, even though his belief is perfectly understandable, the sheriff does not *know* that the suspect behaves in such a way, which means that we must not ignore the corresponding branches.

Now, let us turn to Friend or Foe. The decision tree of a soldier coming from Country A is depicted in Figure 7.8. It reveals that the soldier will hesitate. The principal arguments are the same as in the sheriff's dilemma: He hesitates because he believes that the other is a friend with high probability. Note that the soldier is not convinced that a friend will hesitate, but he is quite sure that a foe will shoot, which is expressed by the conditional probabilities. In the subjectivistic framework, we need not treat any player of the same type identical. Hence, if a soldier from Country A comes to the conclusion that it is better to hesitate, he need not think that a friend comes to the same conclusion, etc. This makes the subjectivistic approach much more flexible and realistic than the traditional approach. In particular, it is less cumbersome because we need not guarantee that the players give a best response to one another.




Figure 7.8: Decision tree of a soldier from Country A.

7.4 Correlated Equilibrium

The idea that players use a random generator in order to perform a mixed strategy leads to an obscure and nonsensical description of real-life situations. Rubinstein (1991) puts it like this: "Outside of Las Vegas we do not spin roulettes."

To the best of my knowledge, Aumann (1974, 1987) is the first who provides another, more profound, understanding of mixed strategy, which leads to his celebrated correlated equilibrium. Here, I follow his presentation in Aumann (1987). His model differs from that of Nash (1951) in three essential aspects:

1. The players do not randomize their actions;
2. a strategy is a random variable, not a probability distribution, and
3. the strategies need not be stochastically independent.

7.4.1 Basic Model

Aumann's model describes a strategic conflict between n players in normal form. It consists of the following primitives:

- A state space Ω ;
- a general σ -algebra \mathcal{F} ;
- the private information sets $\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_n$ of the players;
- their subjective probability measures, i.e., priors, P_1, P_2, \dots, P_n ;
- their action sets A_1, A_2, \dots, A_n , and
- their utility functions u_1, u_2, \dots, u_n .

Aumann (1987) presumes that the state space Ω is finite.¹⁴⁸ This means that \mathcal{F} is a Boolean algebra. The private information set of Player i , \mathcal{I}_i , is a partition of Ω and its elements belong to \mathcal{F} . It is clear that the private information sets $\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_n$ are finite, too. Further, Aumann explicitly assumes that the players share the same prior, i.e., $P_1 = P_2 = \dots = P_n =: P$. Hence, he retains Harsanyi's (1967–1968) common-prior assumption, which has already been mentioned in Section 7.3.1. Since all players share the same prior, P , general σ -algebra, \mathcal{F} , and state space, Ω , Aumann's approach is inherently objectivistic. More precisely, P represents the *objective* probability measure.

Nature chooses some true state of the world $\omega_0 \in \Omega$. Player i is informed, a posteriori, about the event $I_i \in \mathcal{I}_i$ that is such that $\omega_0 \in I_i$. Based on his private information, I_i , he can choose any element, a_i , from his action set A_i . Hence, whenever he receives some private information $I_i \in \mathcal{I}_i$, he performs a specific action. In this way, Player i implicitly creates a function s_i from Ω to A_i that is constant on each element of \mathcal{I}_i . This function represents his strategy. In fact, s_i is nothing other than an \mathcal{I}_i -measurable random variable.¹⁴⁹ Further, Player i possesses a utility function u_i , which assigns each n -tuple (a_1, a_2, \dots, a_n) of actions a real number. That is, $u_i(a_1, a_2, \dots, a_n)$ represents the utility of Player i given that he performs Action a_i and the others perform the actions $a_1, a_2, \dots, a_{i-1}, a_{i+1}, \dots, a_n$. Thus, $EU_i(s) := E(u_i(s))$ with $s = (s_1, s_2, \dots, s_n)$ is the expected utility of Player i .

Aumann assumes that each player is Bayes rational. More precisely, given his particular information $I_i \ni \omega_0$, Player i aims at maximizing his conditional expected utility $EU_i(a_i, s_{-i} | I_i)$ by choosing an appropriate action $a_i \in A_i$.¹⁵⁰ This is done for all elements of \mathcal{I}_i , which leads to a situation in which we have that $EU_i(s) \geq EU_i(t_i, s_{-i})$ for every \mathcal{I}_i -measurable random variable, i.e., strategy, t_i . Since all players are rational,

¹⁴⁸ However, as is pointed out in Section 1.1, this contradicts Savage's postulates of rational choice.

¹⁴⁹ Note that there is no contradiction between deliberation and prediction, since the player does not decide upon any state of the world, $\omega \in \Omega$, but he chooses some action $a_i \in A_i$.

¹⁵⁰ Here, " s_{-i} " denotes the vector of strategies of the opponents of Player i (see Section 6.6.2).

Aumann concludes that the solution of the game, $s = (s_1, s_2, \dots, s_n)$, must be such that

$$EU_i(s_1, s_2, \dots, s_i, \dots, s_n) \geq EU_i(s_1, s_2, \dots, t_i, \dots, s_n)$$

for every strategy t_i of each player $i \in \{1, 2, \dots, n\}$. He calls the rational solution, s , a correlated equilibrium. Hence, a correlated equilibrium is a solution in which no player can increase his expected utility by performing another strategy, given that the other players hold still if he decides to alter his own strategy.

7.4.2 Aumann's Notion of Strategy

Consider a 2-person game in which the players perform the strategies that are depicted in Figure 7.9. The red lines indicate the strategy of Player 1, whereas the black lines represent the strategy of Player 2. Player 1 possesses the private information set $\mathcal{I}_1 = \{I_{11}, I_{12}, I_{13}\}$, whereas $\mathcal{I}_2 = \{I_{21}, I_{22}, I_{23}\}$ is the private information set of Player 2. The strategy of Player 1 is constant on each element of his private information set and the same holds true for the strategy of Player 2.




Figure 7.9: Aumann strategies in a 2-person game.

How can we interpret the given strategies in Aumann's model? Suppose that the true state of the world, ω_0 , belongs to I_{11} . Then Player 1 knows that the event I_{11} happens and, based on this private information, he chooses some element from his action set A_1 . He knows only that ω_0 belongs to I_{11} but not more than that. More precisely, the particular location of ω_0 in I_{11} is unknown to him, which means that he always picks the same action out of A_1 throughout I_{11} . In the case that $\omega_0 \in I_{11}$, we have also that $\omega_0 \in I_{21}$. Hence, Player 2 has the information I_{21} and chooses some action from A_2 . His action cannot depend on the particular location of ω_0 in I_{21} either, which explains why his strategy is constant on I_{21} . The same arguments hold true, analogously, for each other elements of \mathcal{I}_1 and \mathcal{I}_2 .

In the case that $\omega_0 \in I_{11}$, Player 1 knows Player 2's action, but Player 2 does not know Player 1's action. This is because we have that $\omega_0 \in I_{11} \subseteq I_{21}$ but not $\omega_0 \in I_{21} \subseteq I_{11}$.

(see Section 2.5). Further, if $\omega_0 \in I_{12}$, then Player 1 does not know Player 2's action, whereas Player 2 knows Player 1's action if $\omega_0 \in I_{22}$, etc. Of course, we can create other situations in which one player never knows the action of the other, whereas the latter always knows the action of the former. Then the latter player has perfect information, whereas the former one has imperfect information. If both players have the same private information set, we have a game with perfect information. However, in Aumann's model it is typically assumed that all players suffer from imperfect information.

Aumann's model of strategic conflict reveals that the players do *not* randomize their actions. Each player chooses his own action, deliberately, on the basis of his private information. By definition, his own action is known to himself, but since the players have different information, no one knows the action of any other. Hence, from the viewpoint of each single player, the others' actions appear to be stochastic, whereas his own action is, in fact, deterministic. I think that this is a beautiful explanation of "mixed strategy" and it provides a nice framework for asymmetric information, which is the normal case in real life.

7.4.3 A Simple Example

I would like to demonstrate Aumann's model of strategic conflict by a well-known example (Aumann, 1974). Consider the payoff matrix of a 2-person game in Table 7.9. Partition Ω into the upper left (F_1), upper right (F_2), lower left (F_3), and lower right (F_4) part of the payoff matrix. Moreover, suppose that the lower right part obtains with probability 0, whereas the other elements of the partition each have equal probability, i.e., $\frac{1}{3}$. Assume that Ann's private information set is $\{F_1 \cup F_2, F_3 \cup F_4\}$, whereas Bob's private information set corresponds to $\{F_1 \cup F_3, F_2 \cup F_4\}$. Hence, Ann and Bob suffer from imperfect information. Ann knows whether the true state of the world belongs to the upper or lower part, whereas Bob knows whether it belongs to the left or right part of the payoff matrix. For example, we could assume that Nature "tells" Ann whether she should go up or down, whereas Bob is told whether he should go left or right. However, the players know only their own proposals.

Table 7.9: Aumann's example of correlated equilibrium.

Ann	Bob	
	Left	Right
Up	(6, 6)	(2, 7)
Down	(7, 2)	(0, 0)

Suppose that Ann applies the following strategy: If the upper part obtains, she chooses Up and otherwise she chooses Down. Correspondingly, Bob goes left if the left part

obtains, but otherwise he goes right. This is a correlated equilibrium: If Ann decides to go always up, she receives the expected utility $\frac{2}{3} \cdot 6 + \frac{1}{3} \cdot 2 = \frac{14}{3}$, and if she decides to go always down, her expected utility is $\frac{2}{3} \cdot 7 + \frac{1}{3} \cdot 0 = \frac{14}{3}$, too. Further, if she decides to go down if Nature proposes to go up and vice versa, her expected utility equals only $\frac{1}{3}(7 + 0 + 6) = \frac{13}{3}$. By contrast, if Ann applies the equilibrium strategy, she yields an expected utility of $\frac{1}{3}(6 + 2 + 7) = \frac{15}{3}$.¹⁵¹ The same arguments apply, mutatis mutandis, to Bob. Note that the equilibrium strategies are indeed correlated, since $P(\text{Left} \mid \text{Up}) = 0.5$ but $P(\text{Left} \mid \text{Down}) = 1$.

How can we interpret the correlated equilibrium in the subjectivistic context? First of all, each player has his own prior. Further, Nature's proposals and the adversary's actions are considered events. Let us step into the shoes of Ann and suppose that Nature proposes Up with (subjective) probability $\frac{2}{3}$. She can either obey Nature's proposal or do the opposite. Of course, the same holds true if Nature proposes Down. Since Ann suffers from imperfect information, Bob's actions are part of the chance nodes that can be found behind Ann's decision nodes in Figure 7.10. The subjective probabilities follow immediately from Aumann's example and shall indicate that Ann believes that Bob follows Nature's proposal. Hence, she decides to act in the same way. Now, we just have to translate those arguments to Bob in order to explain why he follows Nature's proposal, too. Note that this does not require us to assume that Ann and Bob have the same probability measure.

Although Aumann's example represents an anti-coordination game, the players coordinate their actions in order to achieve a higher (expected) payoff. However, this holds true just because we assume that Ann and Bob *trust* one another. Since each player has imperfect information, nobody knows whether or not the other player obeys Nature's proposal and so their actions are *strategically* independent. This means that Ann cannot make her own action dependent on Bob's action and vice versa. Nonetheless, Ann's action might very well dependent in the *stochastic* sense on Bob's action, provided that we argue from the viewpoint of the very first chance node in Figure 7.10 and fix Ann's strategy. In this case, both Ann's actions and Bob's actions can be considered events from an outsider's perspective.

What happens if Ann believes that Bob does *not* make his own choice dependent on Nature's signal? More precisely, suppose that Ann believes that Bob goes left with probability 50 %, irrespective of whether Nature proposes Left or Right. Then we obtain the decision tree in Figure 7.11. It turns out that Ann will always go up and so she will ignore Nature's signal, too. The same arguments hold true, mutatis mutandis, for Bob. It follows that the players always decide to be defensive, which is a typical result of anti-coordination games (see Section 4.3.1 and Section 6.4.1).

151 If Ann's information were perfect, the equilibrium strategy would no longer be optimal for her. In that case, she could do better by switching to Down if the true state of the world belongs to the upper left part of the payoff matrix and switching to Up if it belongs to the lower right part.

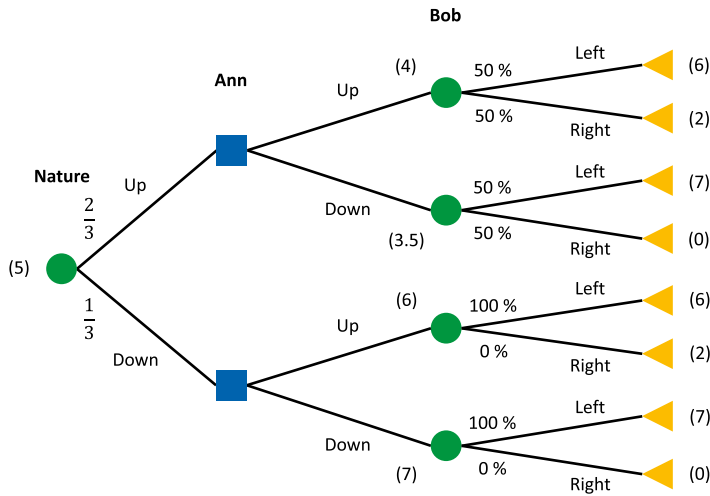


Figure 7.10: Ann's decision tree if she trusts Bob.




Figure 7.11: Ann's decision tree if she mistrusts Bob.

Hence, the rational solution of this game depends essentially on whether the players trust one another or not. The problem is that no player knows whether his adversary behaves in any specific way. Nature's signal might be evident to the players. However, the main point is that they cannot see whether or not that signal has any impact on the adversary's behavior. For this reason, their conclusions are based on belief rather than knowledge. We conclude that a rational solution need not be a correlated equilibrium.

7.4.4 Controversy

7.4.4.1 Common Prior

Aumann calls his approach subjectivistic because each player has his own private information set and so the players may come to different conclusions. Nonetheless, he assumes that the players share the same prior. His arguments go like this (Aumann, 1987, p. 7):

“The subjective probability of a player is his posterior given his information; these may well be different. Roughly speaking, the common prior assumption says that differences in probability estimates of distinct individuals should be explained by differences in information and experience.”

Hence, the posteriors of the players are not different because their subjective probability measures differ from each other, but because they have different information. Aumann vividly defends the common-prior assumption in his original work on correlated equilibrium and also in Aumann (1998). Aumann (1987, p. 12) points out that “Common priors are explicit or implicit in the vast majority of the differential information literature in economics and game theory.” For a broad overview of the common-prior assumption in economics see Morris (1995).

Statements like “Strategy s_1 is stochastically independent of Strategy s_2 ,” “Player 1 chooses Action a_1 with probability $\frac{1}{2}$,” etc., remain vague unless we specify the underlying probability measure. Hence, the common-prior assumption leads to a substantial simplification of Aumann’s model of strategic conflict. Nonetheless, despite its wide acceptance in traditional game theory, it was (Aumann, 1998; Gul, 1998; Morris, 1995) and is still the subject of controversy. Aumann generally supports the common-prior assumption, but in Aumann (1987, p. 12) he mentions that it “is not a tautological consequence of the Bayesian approach.” Gul (1998) goes even further and states that “the assumption of common priors in this context is antithetical to the Savage-established foundations of statistics (i.e., the ‘Bayesian view’), since it amounts to asserting that at some moment in time everyone must have identical beliefs.”

In fact, Savage (1954, p. 3) points out that his personalistic interpretation of probability does “not deny the possibility that two reasonable individuals faced with the same evidence may have different degrees of confidence in the truth of the same proposition.” Morris (1995) concludes that it makes little sense, on the one hand, to allow for subjective probabilities and, on the other hand, to impose the common-prior assumption, which postulates that the players start with the same probability beliefs before calculating their posteriors. In my opinion, that hits the nail on the head. Simply put, the common-prior assumption just makes P an *objective* probability measure.

7.4.4.2 Common State Space

The reader might have already recognized another inconsistency of correlated equilibrium: Since the state space is common to all players, they know all random variables and thus all strategies. To be more precise, they know their own strategies *and* the

strategies of their opponents. Of course, this does not mean that they have perfect information because the players might still be unaware of each other's *action*. However, a strategy is a function on Ω , i.e., a random variable, and so the players know the *strategies* of their opponents. Well, if they know the strategy of each other, why should they make their choices independently?

Interestingly, Aumann (1995) himself points out that:

“Making a decision means choosing among alternatives. Thus one must consider hypothetical situations—what would happen if one did something different from what one actually does.”

This is counterfactual reasoning. However, the entire concept of equilibrium is based on the implicit assumption of strategic independence, which guarantees that the decision of one player has no influence on the decision of any other player. In fact, the equilibrium doctrine tells us that each player can maximize his own expected utility, i.e., choose an optimal strategy, while the strategies of the other players stay fixed. However, in general, the others will not be passive when they observe that some opponent alternates his strategy, which inevitably happens, at least *hypothetically*, when someone is searching for an optimal strategy.

Probability theory deals with random variables that are considered fixed. For this reason, the probabilistic viewpoint reflects a situation in which the players have already made up their minds and so it is *ex post*. However, in game theory we want to analyze the situation of rational choice, in which each player evaluates his possible alternatives before he comes to a final conclusion. Hence, the decision-theoretic viewpoint is *ex ante*. By distinguishing between these different points of view, we can avoid serious misunderstandings that can arise when discussing strategic conflicts.¹⁵²

If we step into the shoes of some player, we may consider the strategies of his opponents fixed *after* they have made their decisions, but then the corresponding player has already made up his mind, too. This means that the whole procedure of rational choice is already completed. This would not serve our purpose. Since we want to analyze the strategic behavior of the players, we must adopt the decision-theoretic, i.e., ex-ante, point of view. However, if we adopt the ex-ante point of view for one player, we must not adopt, at the same time, the ex-post point of view for any other player.

Working with a common state space, Ω , is just like placing characters on a board that is for all to see. If I set up my pieces, everybody can observe what happens. Of course, the same holds true for me regarding the other players. Now, if I change my mind and move my pieces to another place, the others can change their minds, too. So why should I expect that my own strategy, i.e., formation of pieces, has no influence on the formation of the pieces of any other player, although everyone can see everything? Put another way, why should it be possible to consider the strategy of one player fixed, while the strategy of another player is variable? I think that this is possible only if we leave the objectivistic framework and make use of the subjectivistic approach.

¹⁵² The given viewpoints have nothing to do with our usual notions of “a priori” and “a posteriori.”

What is the main difference between Aumann's objectivistic approach and the subjectivistic approach advocated throughout this book? In Aumann's model of strategic conflict, a player can decide how to assign his actions to each state of the world. Nonetheless, whatever assignment he makes, i.e., strategy he performs, his opponents are always informed about that strategy—although, in general, they do not know his particular action. Indeed, according to Aumann (1987, p. 6),

“each ω includes a specification of which action is chosen by each player of G at that state ω . Conditional on a given ω , everybody knows everything; but in general, nobody knows which is really the true ω .”

Hence, the players might be uncertain about the true state of the world, but they cannot be uncertain about what happens *if* any state of the world obtains. Since the players have a common state space, they must know the strategy of each other!

By contrast, in the subjectivistic framework each player has his *own* state space, Ω , and the possible strategies of his opponents represent events. His strategy is not part of his own state space—it is part of the state spaces of the other players. Further, a player cannot produce any wrong assignment on Ω . The descriptions that are associated with each state of the world, ω , are not optional to his opponents. The player *himself* creates them by counterfactual reasoning and so each label that is assigned to some $\omega \in \Omega$ is correct by definition. All associations just take place in the mind of the corresponding player and reflect his personal uncertainty—given the private information that he has obtained a posteriori.

7.5 Minimax Solution

Obviously, von Neumann's (1928) seminal work on parlor games is a cornerstone in game theory. Together with Zermelo (1913), von Neumann's contribution is one of the first formal attempts to solve strategic conflicts. It contains his celebrated minimax theorem, which can be used to analyze 2-person zero-sum games.¹⁵³ Nonetheless, von Neumann discusses also zero-sum games with $n > 2$ players and uses the notion of coalition. This means that he lays the foundation for cooperative game theory, but here I assume that the number of players is $n = 2$. Hence, the players have strictly opposing interests and so the games that we take into consideration are noncooperative.

7.5.1 Von Neumann's Minimax Theorem

The proof of the following theorem can be found in von Neumann (1928).

Theorem 12 (Minimax theorem). *Let $A \subset \mathbb{R}^r$ and $B \subset \mathbb{R}^s$ with $A, B \neq \emptyset$ be compact and convex. Further, let $f: A \times B \rightarrow \mathbb{R}$ be some continuous function such that $f(x, \cdot)$ is convex*

¹⁵³ Recall that Zermelo (1913) focuses on Chess (see Section 5.4), which is a zero-sum game, too.

for all $x \in A$ and $f(\cdot, y)$ is concave for all $y \in B$. Then we have that

$$\max_{x \in A} \min_{y \in B} f(x, y) = \min_{y \in B} \max_{x \in A} f(x, y). \quad (7.1)$$

The minimax theorem reminds us of the so-called max-min inequality

$$\sup_{x \in A} \inf_{y \in B} f(x, y) \leq \inf_{y \in B} \sup_{x \in A} f(x, y),$$

which holds true for any real-valued function f on a (nonempty) domain $A \times B$ with $A \subset \mathbb{R}^r$ and $B \subset \mathbb{R}^s$.¹⁵⁴ Hence, the max-min inequality turns into an equality if the (sufficient) conditions that are stated by the minimax theorem are met.

7.5.2 2-Person Zero-Sum Games

A 2-person zero-sum game can be characterized like this:

- There are two players with (nonempty) action sets $A \subseteq \mathbb{R}^r$ and $B \subseteq \mathbb{R}^s$.
- There is a function $f : A \times B \rightarrow \mathbb{R}$ that assigns Player 1 the payoff $f(x, y)$ and Player 2 the payoff $-f(x, y)$.

Hence, Player 1 can control only the first argument of f , whereas Player 2 is able to control only its second argument. The reason why we speak about a “zero-sum” game is because the payoffs of the players sum up to zero for all $(x, y) \in A \times B$.

Player 1 aims at maximizing his payoff function $f(\cdot, y)$ by choosing some action $x \in A$. However, he assumes that Player 2 will adapt his position, accordingly. More precisely, he thinks that Player 2 will *not* hold still if he alternates his action but rather search for a best response to each (alternative) action of Player 1. This explains why Player 1 can achieve *at most* the payoff $\max_{x \in A} \min_{y \in B} f(x, y)$ if Player 2 always searches for a best response to Player 1. The same arguments apply, *mutatis mutandis*, to Player 2, who aims at maximizing his own objective function $-f(x, \cdot)$. We conclude that Player 1 can realize *at least* $\min_{y \in B} \max_{x \in A} f(x, y)$ if Player 2 acts in the same way, i.e., maximizes his own objective function, given that Player 1 always searches for a best response to Player 2.

In my opinion, von Neumann not only provides one of the first formal treatments of game theory. He also appears to be the first one who, essentially, discusses *coherent* games—even if only implicitly. Why is the considered game coherent? This is because we assume that the players can act and react to one another, at the same time, where to “act and react” means to choose an element from an Euclidean space. Hence, the players are able to interact. This is opposite to Nash’s approach to game theory, which is based on the strategic-independence assumption.

¹⁵⁴ In fact, from $\inf_y f(x, y) \leq \inf_y \sup_x f(x, y)$ it follows that $\sup_x \inf_y f(x, y) \leq \inf_y \sup_x f(x, y)$.

However, I fear that von Neumann would vehemently reject my interpretation of his minimax approach. He categorically denies interaction but explicitly presumes that the players have imperfect information (von Neumann, 1928, p. 302):

“Die Spieler S_1, S_2 wählen irgendwelche der Zahlen [...] ohne die Wahl des anderen zu kennen.”

This means that “The players S_1, S_2 choose any numbers [...] without knowing the choice of the other,” where each number represents a specific action, i.e., an element of A or B , respectively. Similar statements can be found in von Neumann (1928, pp. 299–300). Obviously, he presumes that the zero-sum game is *static*.

His motivation follows Murphy’s law (see Section 2.4.2): “Anything that can go wrong will go wrong!” According to von Neumann (1928, p. 302), Player 1 does not know what Player 2 is going to do and so he is prepared for the worst. More precisely, the worst thing that can happen if he chooses Action $x \in A$ is that Player 2 chooses some action $y \in B$ such that Player 1 receives only the payoff $\min_{y \in B} f(x, y)$. Now, Player 1 thinks about his opportunities and comes to the conclusion that, by choosing some action $x \in A$ that maximizes $\min_{y \in B} f(x, y)$, he can achieve *at least* the payoff $\max_{x \in A} \min_{y \in B} f(x, y)$ —irrespective of whatever Player 2 does! Hence, the static interpretation of a 2-person zero-sum game renders $\max_{x \in A} \min_{y \in B} f(x, y)$ a *lower* bound, whereas the coherent interpretation makes it an *upper* bound for the payoff of Player 1.

Analogously, Player 2 comes to the same conclusion and so he chooses some action $y \in B$ that minimizes $\max_{x \in A} f(x, y)$. In this way, he gains at least the payoff $-\min_{y \in B} \max_{x \in A} f(x, y)$ —irrespective of whatever Player 1 does. This means that Player 1 achieves *at most* $\min_{y \in B} \max_{x \in A} f(x, y)$, given that Player 2 acts in the described way. Now, we observe that $\min_{y \in B} \max_{x \in A} f(x, y)$ represents an *upper* bound in the static framework, whereas it is a *lower* bound in the coherent model.

It is more or less a matter of taste whether we analyze a 2-person zero-sum game from the static or the coherent perspective. I prefer the coherent perspective because the static one presumes that the players are extremely pessimistic, which is not very convincing.¹⁵⁵ Hence, as is already done in Section 6.6.2, I call $f(x, y)$ and $-f(x, y)$ “pay-offs” and the arguments of f , i.e., x and y , “actions.” However, this does not mean that the following analysis is based on the assumption that the game is coherent. It can very well be static or even dynamic. This will be resumed in the next section.

Definition 16 (Proper game). A 2-person zero-sum game is said to be proper if and only if it satisfies the conditions of the minimax theorem.

Assume that each player performs his best-response strategy. The minimax theorem can be interpreted thus: In a proper (coherent) 2-person zero-sum game it does not harm to abandon our assumption that the players act simultaneously, i.e., at Time 0.

¹⁵⁵ Apparently, I share this opinion with Ellsberg (1956).

In fact, if Player 1 acts first, he achieves the payoff on the left-hand side of the minimax equality (7.1) and if Player 2 acts first, Player 1 realizes the payoff on the right-hand side of the minimax equality.¹⁵⁶ However, due to the minimax theorem, the two results coincide and so

$$v := \max_{x \in A} \min_{y \in B} f(x, y) = \min_{y \in B} \max_{x \in A} f(x, y)$$

can be considered the value, with respect to Player 1, of the corresponding game.

So far we have discussed only the *value* of a (proper) 2-person zero-sum game. However, it is still unclear whether or not there *exists* a so-called minimax solution. Note that the following definition does not require the game to be proper.

Definition 17 (Minimax solution). A solution $(x^*, y^*) \in A \times B$ of a 2-person zero-sum game is said to be minimax if and only if it is a minimax point, i.e.,

$$(x^*, y^*) \in \mathcal{M} := \arg \max_{x \in A} \min_{y \in B} f(x, y) \cap \arg \min_{y \in B} \max_{x \in A} f(x, y).$$

I say that Player 1 acts in the “minimax sense” if and only if he maximizes his objective function, $f(\cdot, y)$, given that Player 2 gives a best response to every action $x \in A$ of Player 1. Analogously, Player 2 acts in the “minimax sense” if and only if he maximizes his own objective function, $-f(x, \cdot)$, given that Player 1 gives a best response to every action $y \in B$ of Player 2. Hence, a minimax solution is a solution in which both players act in the “minimax sense.”

If a 2-person zero-sum game has a minimax solution, the minimax equality holds true. Indeed, if (x^*, y^*) is a minimax point, we must have that

$$\max_{x \in A} \min_{y \in B} f(x, y) = f(x^*, y^*) = \min_{y \in B} \max_{x \in A} f(x, y).$$

Thus, if there *exists* a solution in which both players act in the “minimax sense,” the minimax equality must be satisfied *by definition*. Further, it is clear that we must obtain the same value, v , for each minimax point of the game.

The minimax equality can be derived also like this: The strict inequality

$$\max_{x \in A} \min_{y \in B} f(x, y) < \min_{y \in B} \max_{x \in A} f(x, y)$$

tells us that the greatest payoff of Player 1 is lower than his lowest payoff, given that *both* players act in the minimax sense. This result is clearly impossible and so the game cannot have any minimax solution. Thus, if the game *has* some minimax solution, von Neumann’s minimax equality must hold true.¹⁵⁷

To sum up, we can derive the minimax equality, without making use of the minimax theorem at all, just by assuming that the game *has* a minimax solution.

156 In von Neumann and Morgenstern (1953, Section 14.2), the former situation is called the minorant game, whereas the latter is said to be the majorant game.

157 Note that the max-min inequality precludes that “<” turns into “>.”

Theorem 13 (Minimax solution). *A solution of a proper 2-person zero-sum game is minimax if and only if it is an equilibrium.*

Proof. Let (x^*, y^*) be a minimax solution. Player 1 maximizes his payoff function given the action of Player 2 (“ $\max_{x \in A} f(x, y)$ ”) and vice versa (“ $\min_{y \in B} f(x, y)$ ”). Hence (x^*, y^*) is an equilibrium. Conversely, let (x^*, y^*) be an equilibrium. Thus, for all $x \in A$, we have that

$$f(x^*, y^*) \geq f(x, y^*) \geq \min_{y \in B} f(x, y).$$

Thus, Player 1 cannot do better by moving from x^* to any other action $x \in A$, provided that Player 2 gives a best response to every action of Player 1. The same argument holds true, mutatis mutandis, for Player 2 and so (x^*, y^*) represents a minimax solution. \square

Hence, in a proper 2-person zero-sum game, it makes no difference at all whether we speak about a minimax solution or an equilibrium. Now, our main observation is that the existence of an equilibrium and thus of a minimax solution is always guaranteed by Theorem 7. This means that the minimax equality must be satisfied. Note that Nash’s original existence theorem (Nash, 1951) cannot be applied to our general definition of a 2-person zero-sum game, since Nash presumes that each payoff function f_i is multilinear, not only quasiconcave in x_i .

The next results make use only of Theorem 13. Corollary 2 asserts that the existence of a minimax solution is not only sufficient but also necessary for the minimax equality, whereas Corollary 3 implies that every proper 2-person zero-sum game has either exactly one or infinitely many minimax solutions.

Corollary 2. *Every proper 2-person zero-sum game has a minimax solution if and only if the minimax equality is satisfied.*

Proof. The “only-if part” has already been discussed and is trivial. Thus, let us turn to the “if part”: Consider some points

$$(x_1^*, y_1^*) \in \arg \max_{x \in A} \min_{y \in B} f(x, y) \quad \text{and} \quad (x_2^*, y_2^*) \in \arg \min_{y \in B} \max_{x \in A} f(x, y).$$

Obviously, we have that

$$f(x_1^*, y) \geq f(x_1^*, y_1^*) \quad \text{and} \quad f(x, y_2^*) \leq f(x_2^*, y_2^*)$$

for all $x \in A$ and $y \in B$, which implies that

$$f(x_1^*, y_1^*) \leq f(x_1^*, y_2^*) \leq f(x_2^*, y_2^*).$$

If the minimax equality is satisfied, i.e., $f(x_1^*, y_1^*) = f(x_2^*, y_2^*)$, we have that

$$f(x_1^*, y_1^*) = f(x_1^*, y_2^*) = f(x_2^*, y_2^*).$$

Thus, it holds that

$$f(x_1^*, y) \geq f(x_1^*, y_2^*) \quad \text{and} \quad f(x, y_2^*) \leq f(x_1^*, y_2^*)$$

for all $x \in A$ and $y \in B$. Hence, (x_1^*, y_2^*) is an equilibrium, i.e., a minimax solution. \square

Corollary 3. *If a proper 2-person zero-sum game has two different minimax solutions, it has infinitely many minimax solutions.*

Proof. Let $(x_1^*, y_1^*) \neq (x_2^*, y_2^*)$ be two minimax solutions, i.e., equilibria, of the game. Then we have that

$$f(x_1^*, y_1^*) \geq f(x, y_1^*) \quad \text{and} \quad f(x_2^*, y) \geq f(x_2^*, y_2^*)$$

as well as

$$f(x_1^*, y_1^*) \leq f(x_1^*, y) \quad \text{and} \quad f(x, y_2^*) \leq f(x_2^*, y_2^*)$$

for all actions x and y . Moreover, it is clear that $f(x_1^*, y_1^*) = f(x_2^*, y_2^*) = v$, which implies that

$$f(x_1^*, y_1^*) = f(x_1^*, y_2^*) = f(x_2^*, y_1^*) = f(x_2^*, y_2^*) = v.$$

Now, let x^* and y^* be any convex combinations of x_1^* and x_2^* as well as of y_1^* and y_2^* . This means that $x^* = \alpha x_1^* + (1 - \alpha)x_2^*$ and $y^* = \beta y_1^* + (1 - \beta)y_2^*$ with $0 \leq \alpha, \beta \leq 1$. Since (x_1^*, y_1^*) is an equilibrium and $f(x, \cdot)$ is convex, we have that

$$v = f(x_1^*, y_1^*) \leq f(x_1^*, y^*) \leq \underbrace{\beta f(x_1^*, y_1^*)}_{=v} + (1 - \beta) \underbrace{f(x_1^*, y_2^*)}_{=v} = v$$

and thus $f(x_1^*, y^*) = v$. Similar arguments imply that

$$f(x_2^*, y^*) = f(x^*, y_1^*) = f(x^*, y_2^*) = v.$$

Finally, since $f(\cdot, y)$ is concave, we have that

$$f(x^*, y^*) \geq \underbrace{\alpha f(x_1^*, y^*)}_{=v} + (1 - \alpha) \underbrace{f(x_2^*, y^*)}_{=v} = v$$

and, since $f(x, \cdot)$ is convex, it follows that

$$f(x^*, y^*) \leq \underbrace{\beta f(x^*, y_1^*)}_{=v} + (1 - \beta) \underbrace{f(x^*, y_2^*)}_{=v} = v,$$

which means that $f(x^*, y^*) = v$. Suppose that $y_1^* \neq y_2^*$ and assume that $0 < \beta < 1$, which means that $y^* \neq y_1^*, y_2^*$. For each action $x \in A$ it holds that

$$f(x, y^*) \leq \underbrace{\beta f(x, y_1^*)}_{\leq v} + (1 - \beta) \underbrace{f(x, y_2^*)}_{\leq v} \leq v = f(x^*, y^*)$$

and for each action $y \in B$ we have that

$$f(x^*, y) \geq \underbrace{\alpha f(x_1^*, y)}_{\geq v} + (1 - \alpha) \underbrace{f(x_2^*, y)}_{\geq v} \geq v = f(x^*, y^*).$$

We conclude that (x^*, y^*) is an equilibrium and thus a minimax solution that differs both from (x_1^*, y_1^*) and from (x_2^*, y_2^*) . We can apply the same arguments in the case of $x_1^* \neq x_2^*$. Thus, we can find an infinite number of minimax solutions in the convex hull of $\{(x_1^*, y_1^*), (x_2^*, y_1^*), (x_1^*, y_2^*), (x_2^*, y_2^*)\}$.¹⁵⁸ \square

Our previous arguments are *not* based on the minimax theorem at all. Theorem 7 reveals that every proper 2-person zero-sum game has an equilibrium and thus a minimax solution, (x^*, y^*) , which implies that von Neumann's minimax equality holds true just by definition. Further, Player 1 achieves the payoff $v = f(x^*, y^*)$, which means that Player 2 obtains the payoff $-v$. Hence, we can always assign a proper 2-person zero-sum game an unambiguous value, v , just because it *has* an equilibrium. Similar arguments can be found in Kakutani (1941) and von Neumann (1937) without referring to (Nash) equilibrium.

This completes our general discussion of 2-person zero-sum games. For more details on that topic see, e.g., Appendix 2 in Luce and Raiffa (1957).

7.5.3 Some Examples

In most practical applications of von Neumann's minimax theorem, one considers a finite (static or dynamic) 2-person zero-sum game in normal form, which allows the players to perform mixed strategies in the sense of Nash (1951).¹⁵⁹

Let $U \in \mathbb{R}^{r \times s}$ be the decision matrix of Player 1, which means that $-U'$ is the decision matrix of Player 2. Hence,

- $x'Uy = f(x, y)$ is the expected utility of Player 1 and
- $y'(-U')x = -x'Uy = -f(x, y)$ is the expected utility of Player 2;

where x and y represent the (pure or mixed) strategies of the players.¹⁶⁰

Note that the functions f and $-f$, which quantify the expected utilities of the players, are bilinear forms. Moreover, the strategies x and y belong to the standard simplex in \mathbb{R}^r and \mathbb{R}^s , respectively. This means that every finite 2-person zero-sum game in normal form is proper if we allow the players to perform mixed strategies.

¹⁵⁸ The convex combination $\alpha\beta(x_1^*, y_1^*) + (1 - \alpha)\beta(x_2^*, y_1^*) + \alpha(1 - \beta)(x_1^*, y_2^*) + (1 - \alpha)(1 - \beta)(x_2^*, y_2^*)$, in fact, produces the desired minimax solution (x^*, y^*) .

¹⁵⁹ Actually, this was already the main motivation of von Neumann (1928).

¹⁶⁰ Hence, each player has an infinite number of strategies unless $r = 1$ or $s = 1$. In Section 7.5.2, these strategies were called “actions” and the expected utilities were said to be “payoffs.”

7.5.3.1 Rock-Paper-Scissors

A typical example of a finite (static, one-shot) 2-person zero-sum game is Rock-Paper-Scissors (see Section 4.4.1). Its payoff matrix is given in Table 7.10.

Table 7.10: Payoff matrix of Rock-Paper-Scissors.

Ann	Bob		
	Rock	Paper	Scissors
Rock	(0, 0)	(-1, 1)	(1, -1)
Paper	(1, -1)	(0, 0)	(-1, 1)
Scissors	(-1, 1)	(1, -1)	(0, 0)

The solution (x^*, y^*) with $x^* = y^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ represents a Nash equilibrium and thus it is minimax. It is the *unique* minimax solution of this game. Nonetheless, this solution is quite unsatisfactory because we are not able to predict the particular actions of the players. Nash (1951) and von Neumann (1928) tell us that the players will randomize their actions, which has been extensively discussed in Section 7.2.4. In Section 7.4 we have discovered another, more subtle, understanding of “mixed strategy”: Correlated equilibrium suggests that both Ann and Bob make a definite choice between Rock, Paper, and Scissors on the basis of their private information. That is, each player considers his own action deterministic, but since his information is imperfect, the action of the opponent appears to be stochastic.

In contrast to the traditional, i.e., equilibrium-based, approach, we have seen in Section 4.4.1 that the subjectivistic approach is able to predict the actions of the players unless their subjective probability distribution is uniform. The optimal action of a player depends on the modes of his subjective probability distribution. The decision rule goes like this:

- If the single mode is Rock, choose Paper.
- If the single mode is Paper, choose Scissors.
- If the single mode is Scissors, choose Rock.
- If the two modes are Rock and Paper, choose Paper.
- If the two modes are Rock and Scissors, choose Rock.
- If the two modes are Paper and Scissors, choose Scissors.
- If the three modes are Rock, Paper, and Scissors, choose anything.

Hence, in most cases, the optimal choice of each player is uniquely determined by his subjective probability measure. Moreover, since each player makes a definite choice, the solution of the game cannot be a mixed-strategy Nash equilibrium and thus a minimax solution.

7.5.3.2 A 3×2 Game

In order to demonstrate Corollary 3, I would like to analyze a 3×2 zero-sum game that is characterized by the payoff matrix

$$\begin{bmatrix} (1, -1) & (1, -1) \\ (1, -1) & (1, -1) \\ (-1, 1) & (-1, 1) \end{bmatrix}.$$

We can see that

- (x_1^*, y_1^*) with $x_1^* = (1, 0, 0)$ and $y_1^* = (1, 0)$ as well as
- (x_2^*, y_2^*) with $x_2^* = (0, 1, 0)$ and $y_2^* = (0, 1)$

are two possible minimax solutions.

Now, in order to find another minimax solution, we might choose any convex combination of the first and second action of Player 1, i.e., combine the first and second row of the payoff matrix. Further, regarding Player 2 we might combine the first and second column in an arbitrary way. Each solution that is constructed in such a way is minimax and so this game contains an infinite number of minimax solutions.

7.6 Refinement

Despite its huge significance in game theory, Nash equilibrium is often considered insufficient or implausible. Colman (2004) points out the possibility of multiple and payoff-inefficient Nash equilibria (Harsanyi and Selten, 1988). Moreover, Nash equilibria can be trembling-hand imperfect (Selten, 1975) or even subgame imperfect (Selten, 1965, 1973). Hence, Nash equilibrium is not a sufficient condition for rationality.

The problem of multiple solutions is ubiquitous in traditional game theory. Nash equilibrium is generally assumed to be a valid but insufficient solution concept. In the literature, one can find a huge number of procedures that aim at a refinement of Nash equilibrium. Refinement means to eliminate some Nash equilibria, or any other solutions, in order to obtain a smaller set of solutions that are considered reasonable.¹⁶¹

At best, the remaining set of reasonable solutions is a singleton.

Refinement criteria such as payoff or risk dominance, perfectness, properness, stability, etc., are developed in order to eliminate all Nash equilibria that are considered implausible (Govindan and Wilson, 2008; Harsanyi and Selten, 1988; Kreps and Wilson, 1982; Myerson, 1978; Selten, 1975). Here, I will concentrate on payoff efficiency, risk dominance, and (trembling-hand) perfectness. Subgame perfectness has already been illustrated in Section 7.2.2.3.

I will explain the refinement procedures by means of the game show, which has already been discussed, e.g., in Section 4.2.3. Its payoff matrix is given in Table 7.11. This

¹⁶¹ In Chapter 6, we already applied a refinement procedure for coherent games.

game has two pure-strategy and infinitely many mixed-strategy Nash equilibria. They are represented by the intersection points of the best-response curves in Figure 7.12:

- Ann chooses Up and Bob chooses Left.
- Ann chooses Down and Bob chooses Right.
- Ann chooses Up with probability $0 < p \leq \frac{9}{10}$, whereas Bob chooses Right.

Table 7.11: Payoff matrix of the game show.

Ann	Bob	
	Left	Right
Up	(2, 1)	(9, 0)
Down	(1, -1)	(9, 8)




Figure 7.12: Best-response curves and Nash equilibria of the game show (blue point on the upper left and blue line on the right).

This result is somewhat unsatisfying. I think that in game and decision theory, at least, we should try to make concrete predictions about the players' actions, or to provide clear recommendations—depending on whether our approach is descriptive or prescriptive.¹⁶² Thus, we should aim at achieving a *unique* answer. Are we able to reduce the set of rational solutions?

7.6.1 Payoff Efficiency

Payoff efficiency plays a major role when solving coherent games (see Chapter 6). The following definition refers to Nash equilibria, not to rational solutions of a coherent game, and it goes back to Harsanyi and Selten (1988).

Definition 18 (Payoff efficiency, II). A Nash equilibrium is payoff efficient if and only if there is no other Nash equilibrium in which all players have a higher expected payoff.

¹⁶² Aumann (1985) discusses the goals of traditional game theory. Further, the distinction between descriptive and prescriptive decision theories was explained in Section 2.2.

Harsanyi and Selten (1988) suggest to eliminate all Nash equilibria that are payoff inefficient. For example, in the game show the Nash equilibrium Up/Left is clearly payoff inefficient and thus it can be eliminated. The basic idea is that Ann and Bob will agree only on payoff-efficient Nash equilibria. This need not be accomplished by communication. If the players have common knowledge, they could *tacitly* agree that Up/Left is an unacceptable solution for both of them.

However, even under the very strong assumption of common knowledge, there is no guarantee that players (in a static game) will, in fact, avoid the actions of a payoff-inefficient Nash equilibrium. In fact, Up and Left are rationalizable actions (Bernheim, 1984; Pearce, 1984): Ann may choose Up if she believes that Bob chooses Left, which is reasonable if he believes that Ann chooses Up, etc.¹⁶³ The problem is that common knowledge does not imply that the players *know* the action of the other. Moreover, we have already seen in Figure 4.1 that Up/Left can very well be a rational solution from a subjectivistic point of view if we drop the common-knowledge assumption.

Another drawback of this refinement procedure is that games typically possess more than one payoff-efficient Nash equilibrium. That is, in most cases we still do not know which solution is “right” and have to apply additional refinement procedures in order to come to some conclusion. However, applying different refinement procedures one after another makes not much sense. Each method has its own philosophy and different philosophies can very well contradict each other, as we will see below.

7.6.2 Risk Dominance

Now, I turn to another refinement procedure called risk dominance, which goes back to Harsanyi and Selten. As Selten (1995) points out, they developed three different approaches to risk dominance, but Harsanyi and Selten (1988) describe only the third one. Here, I concentrate on their first approach, which is presented by Selten (1995).

It is clear that the solution Down/Right is preferable both for Ann and for Bob. However, even if they agree on this solution backstage, nobody knows whether the other player deviates when it comes down to it. Thus, Ann might ask herself the following question: “What happens if Bob deviates to Left with probability π but sticks to Right with probability $1 - \pi$? Is it still optimal for me to adhere to Down?” If Bob deviates with probability π , Ann’s expected utilities for Up and for Down are

$$EU_{\text{Ann}}(\text{Up}) = \pi 2 + (1 - \pi) 9$$

and

$$EU_{\text{Ann}}(\text{Down}) = \pi 1 + (1 - \pi) 9.$$

163 The concept of rationalizability will be explained in more detail in Section 8.1.

Hence, it is optimal for Ann to choose Down if and only if

$$\pi 2 + (1 - \pi) 9 \leq \pi 1 + (1 - \pi) 9,$$

which leads us to the conclusion that Ann's critical threshold for π is $\pi_{\text{Ann}}^* = 0$. Thus, if Bob deviates with any probability $\pi > 0$, Ann prefers to choose Up. Well, this is not very surprising, since Up dominates Down in the weak sense.

Now, we change our perspective to Bob. If Ann deviates with probability π , his expected utilities are

$$\text{EU}_{\text{Bob}}(\text{Left}) = \pi 1 + (1 - \pi)(-1)$$

and

$$\text{EU}_{\text{Bob}}(\text{Right}) = \pi 0 + (1 - \pi) 8.$$

This means that Bob considers Right optimal if and only if

$$\pi 1 + (1 - \pi)(-1) \leq \pi 0 + (1 - \pi) 8.$$

It follows that Bob's critical threshold for π is $\pi_{\text{Bob}}^* = 0.9$, i.e., he will prefer Left only if Ann deviates with probability greater than 90%.

Selten (1995) calls the critical thresholds of Ann, i.e., $\pi_{\text{Ann}}^* = 0$, and of Bob, i.e., $\pi_{\text{Bob}}^* = 0.9$, diagonal probabilities. He writes that "The diagonal probability $\pi_i(\varphi, \psi)$ is a natural index of player i 's individual stability at φ against deviations of other players to ψ ," where φ and ψ represent the two Nash equilibria that are compared with one another. In our case, φ corresponds to Down/Right and ψ is Up/Left. Selten's stability index of Down/Right is the product of π_{Ann}^* and π_{Bob}^* :

$$\Pi_{\text{Down/Right}} := \pi_{\text{Ann}}^* \pi_{\text{Bob}}^* \cdot^{164}$$

The diagonal probabilities of the game show and thus also its stability indices can easily be deduced from the change points of the best-response diagram in Figure 7.12. If we choose Down/Right as a starting point, π_{Bob}^* just corresponds to the change point of the black, i.e., Bob's, best-response curve, whereas π_{Ann}^* is the change point of the red, i.e., Ann's, best-response curve. Hence, $\Pi_{\text{Down/Right}}$ is the rectangular area on the lower right, which is zero in our case. Further, it is quite easy to see that the stability index of each mixed-strategy Nash equilibrium of the game show (see the blue line in Figure 7.12) with respect to Up/Left is zero, too. By contrast, the stability index $\Pi_{\text{Up/Left}}$ equals the rectangular area on the upper left of the best-response diagram (see Figure 7.13), i.e.,

$$\Pi_{\text{Up/Left}} = (1 - \pi_{\text{Ann}}^*)(1 - \pi_{\text{Bob}}^*) = (1 - 0)(1 - 0.9) = 0.1.$$




Figure 7.13: The stability index of Up/Left, $\Pi_{\text{Up/Left}}$, corresponds to the blue area.

Risk dominance suggests that we should discard all Nash equilibria that have a lower stability index compared with another Nash equilibrium. To be more precise, the comparison is two-fold: We compare φ with respect to ψ as well as ψ with respect to φ and favor the equilibrium point with the higher stability index. The Nash equilibrium with the higher stability index is said to be risk dominant.

The motivation of this terminology is that the risk-dominant equilibrium is less vulnerable to deviations. For example, consider a strict (pure-strategy) Nash equilibrium, like Up/Left in the game show, a starting point. The bigger the rectangular area that can be found at the corresponding corner of the best-response diagram, the more space is left for the deviation probability π . In the game show, Up/Left risk-dominates Down/Right and also all mixed-strategy Nash equilibria. Thus, we can analyze the overall situation, just as usual, by taking a look at the best-response diagram, which immediately reveals that a corner solution is more “stable” the further the change points are away from that corner.

For example, the upper left corner of the best-response diagram in Figure 7.12 marks the solution Up/Left. The first change point is 10 percentage points away and the next one is 100 percentage points away. Hence, the stability index of this corner solution amounts to $0.1 \cdot 1 = 0.1$. The lower right of the best-response diagram marks the solution Down/Right. The first change point is 0 percentage points away and the next one is 90 percentage points away. We conclude that the corresponding stability index is $0 \cdot 0.9 = 0$, i.e., Down/Right is not stable at all.

From the viewpoint of subjectivistic game theory, risk dominance turns out to be a familiar concept. I would say that the stability index of a Nash equilibrium is just a measure of how stable the solution is, concerning the *subjective* probabilities of the players. More precisely, the solution Up/Left is quite stable, compared with Down/Right, because Ann has an incentive to choose Up whenever she thinks that Bob chooses Left with probability greater than $\pi_{\text{Ann}}^* = 0$ and Bob has an incentive to choose Left whenever he thinks that Ann chooses Up with probability greater than $\pi_{\text{Bob}} = 0.9$. Hence, the “space” of (subjective) probabilities that lead to the solution

164 Actually, Selten (1995) uses influence weights w_{Ann} and w_{Bob} to define $(\pi_{\text{Ann}}^*)^{w_{\text{Ann}}} (\pi_{\text{Bob}}^*)^{w_{\text{Bob}}}$ as a stability index. Here, I set $w_{\text{Ann}} = w_{\text{Bob}} = 1$ for the sake of simplicity.

Up/Left, i.e., $\Pi_{\text{Up/Left}} = 0.1$, is greater than the space of probabilities that lead to the solution Down/Right, i.e., $\Pi_{\text{Down/Right}} = 0$.

This line of argument is purely subjectivistic and it precisely reflects the basic idea of Harsanyi and Selten (1988, p. 83) as well as Selten (1995). However, in my opinion, risk dominance need not be restricted to Nash equilibria. We can apply the same principle to any other (corner) solution of a 2×2 game.¹⁶⁵ Recall that p is Bob's subjective probability that Ann chooses Up, whereas q is Ann's subjective probability that Bob chooses Left. Now, let us make the following definition of stability¹⁶⁶:

Definition 19 (Stability). Consider any 2×2 game and let φ be some corner solution in the best-response diagram. The stability of φ corresponds to the (Lebesgue) area of the set

$$\{(p, q) \in [0, 1]^2 : \text{The solution } \varphi \text{ is rational with } p \text{ and } q\}.$$

For example, reconsider the chicken game, whose payoff matrix can be found in Table 7.12. In Section 4.3.1.1, I have argued that both players will swerve, although this is not a Nash equilibrium. As we can see in Figure 7.14, the stability of this solution equals $0.9^2 = 0.81$, which is essentially greater than the stability $0.1 \cdot 0.9 = 0.09$ of the pure-strategy Nash equilibria Straight/Swerve and Swerve/Straight.

Table 7.12: Payoff matrix of the chicken game.

Andy	Bob	
	Swerve	Straight
Swerve	(1, 1)	(0, 2)
Straight	(2, 0)	(-9, -9)




Figure 7.14: The stability of the solution Swerve/Swerve (green point) corresponds to the green area, whereas the stability of Straight/Swerve and Swerve/Straight (blue points) equals the blue areas.

¹⁶⁵ A solution in the subjectivistic framework is always a corner solution because the players do not randomize their actions (see Section 7.2.4).

¹⁶⁶ There exist many notions of “stability,” which typically refer to Nash equilibria. Some authors even consider Nash equilibrium per se a “stable” solution, i.e., they ignore any further requirement.

Thus, Swerve/Swerve risk-dominates the two pure-strategy Nash equilibria and the pure-strategy Nash equilibria risk-dominate the mixed-strategy Nash equilibrium of the chicken game. However, it is well-known that risk dominance is not transitive: If a Nash equilibrium φ risk-dominates some Nash equilibrium ψ and ψ risk-dominates a Nash equilibrium ϕ , it does not follow that φ risk-dominates ϕ . This poses no problem in our subjectivistic framework because a “mixed-strategy solution” does not exist at all. Hence, we just need to take corner solutions into account. Further, the stability of a corner solution, as defined above, is not a binary measure. That is, in contrast to Selten’s stability index, we need no reference point, ψ , to quantify the stability of φ .

7.6.3 Perfectness

Selten (1975) starts to distinguish between perfect and subgame-perfect Nash equilibria. Every perfect Nash equilibrium is subgame perfect, but the converse is not true. Perfect Nash equilibria are often called trembling-hand perfect.

Consider any finite game Γ in normal form. A perturbed game, $\Gamma_n = (\Gamma, \eta_n)$, is a game Γ in which the players must perform completely mixed strategies, i.e., assign each action a positive probability. The lower probability bounds of their strategies are contained in the (positive) vector η_n . Now, fix any Nash equilibrium of Γ_n for $n = 1, 2, \dots$ and suppose that the sequence $\{\eta_n\}$ converges to zero. Every limit of Nash equilibria of the perturbed games, $\Gamma_1, \Gamma_2, \dots$, is a Nash equilibrium of the original game Γ (Selten, 1975, Lemma 3). However, a Nash equilibrium of Γ need not be the limit of some Nash equilibria in $\{\Gamma_n\}$ with $\eta_n \rightarrow \mathbf{0}$.¹⁶⁷

Definition 20 (Perfectness). A Nash equilibrium of a finite game Γ in normal form is perfect if and only if it is the limit of some Nash equilibria in $\{\Gamma_n\}$ with $\eta_n \rightarrow \mathbf{0}$.

It is clear that every Nash equilibrium in which all equilibrium strategies are completely mixed is perfect. Further, a Nash equilibrium is perfect only if the equilibrium strategy of each player is not weakly dominated by another strategy.¹⁶⁸

For example, the payoff-efficient Nash equilibrium Down/Right in the game show is imperfect because Up dominates Down in the weak sense. If Bob chooses Left with positive probability q , it is always better for Ann to go up. Indeed, it does not matter at all how small $q > 0$ is. Bob could simply be nervous or distracted by the studio lighting, etc. Due to the same reasons, also the mixed-strategy Nash equilibria are trembling-hand imperfect.¹⁶⁹ By contrast, the pure-strategy Nash equilibrium Up/Left

¹⁶⁷ Here, $\mathbf{0}$ denotes a vector of zeros.

¹⁶⁸ However, the converse is not true, i.e., the absence of weakly dominated strategies is not sufficient for a perfect Nash equilibrium.

¹⁶⁹ Note that the equilibrium strategies are not *completely* mixed.

is not affected by a “trembling hand”: We can always find some $p, q < 1$ such that Up remains to be optimal for Ann and Left remains to be optimal for Bob. Hence, all Nash equilibria except for Up/Left are imperfect and thus, according to Selten (1975), they should be discarded.

Selten’s “model of slight mistakes” describes incomplete rationality.¹⁷⁰ As Selten (1975, p. 35) points out, “There cannot be any mistakes if the players are absolutely rational.” He considers rationality as a limiting case of incomplete rationality. A Nash equilibrium that breaks down if some player deviates by an *arbitrarily* small amount from the equilibrium strategy is considered implausible because the other player must fear to perform a suboptimal strategy if his opponent makes the slightest mistake.

Well, at least for me, it is hard to understand why players should not be rational in the context of traditional game theory, which is based substantially on the assumption that players act like robots and apply a random generator in order to make their choices. However, if we leave the objectivistic framework and do no longer assume that the players can perform a mixed strategy, Selten’s principal idea appears to be entirely perspicuous: If Ann is not *convinced* that Bob goes right, she will go up. Whether Bob is rational or not makes no difference at all. Of course, as already mentioned above, he could be rational but still choose the wrong action just by mistake. We do not have to explain *why* Ann doubts that Bob goes right. Whenever she fears that *anything* can go wrong with Bob, she will go up.

7.7 Conclusion

Nash equilibrium is still the basic paradigm of noncooperative game theory. It is based on the assumption that all players use a random generator in order to make their choices. The random generator allows them to apply a mixed strategy, but it can be used also to perform a pure strategy, in which case some action is generated with probability 1. The players make their decisions, i.e., choose their strategies, independently and the actions that are generated by their strategies are stochastically independent. The strategic-independence assumption, i.e., the assumption that the players choose their strategies independently, implies that no player knows the strategy of any other.

Nonetheless, Nash equilibrium tells us that each player is able to choose a strategy that is optimal with respect to the given strategies of his opponents, although those strategies are not evident to him. The principal idea of Nash equilibrium is typically justified by the transparency of reason, which states that the players can step into the shoes of each other in order to deduce his strategy. Indeed, if the conjectures of the players about the strategy of each other are correct and all players choose

¹⁷⁰ In fact, this model goes back to Harsanyi (Selten, 1975, Footnote 2).

an optimal strategy, the solution of the game must be a Nash equilibrium. Nonetheless, the precise reasons why each player should be able to deduce not only the optimal but even the *actual* strategies of his opponents typically remain vague and unclear.

Bacharach's transparency of reason is often explained by complete information and common knowledge. The common-knowledge assumption makes great demands on the private information and mental capabilities of all players, which go far beyond the notion of rationality in the subjectivistic framework. In most real-life situations, players suffer from incomplete information, and even if they *have* complete information, it turns out that the common-knowledge assumption is violated. In the next chapter it will be shown that the players need not be able to deduce the strategy of each other even if everybody has complete information and the common-knowledge assumption is satisfied. Hence, Bacharach's transparency of reason, which serves as an anchor point for the equilibrium doctrine, does *not* follow from complete information and common knowledge. This means that complete information and common knowledge do not imply that the rational solution of a game is a Nash equilibrium.

Moreover, there exist games in which the transparency of reason cannot take place at all from a subjectivistic point of view. For example, in a typical discoordination game like Matching Pennies it is impossible for both players to possess only one optimal action, to deduce the optimal action of the adversary, and to be convinced that he performs this action. Simply put, it is impossible for the players to deduce the action of one another unless some player *doubts* that the other acts in an optimal way. This holds true for all finite 2-person games without a pure-strategy Nash equilibrium.

Players suffer from incomplete information in most real-life situations. Harsanyi's Bayesian model transforms incomplete information into imperfect information by assuming that Nature chooses the type of each player at random from some set of types before the game starts. Thus, everybody is, seemingly, able to apply the transparency of reason in order to deduce the strategy of each type of player. This is essentially based on the common-prior assumption, which claims that all subjects have the same unconditional probabilities. However, we already know that the transparency of reason does not follow from complete information and common knowledge. This means that the rational solution of an ordinary game with complete information need not be a Nash equilibrium and thus also the rational solution of a Bayesian game need not be a (Bayesian) Nash equilibrium—even if the players agree on the Bayesian model and have common knowledge.

We cannot really observe that rational subjects choose their actions at random. In real life, players make a definite choice—they do not use a random generator. Aumann avoids the randomization paradigm by introducing asymmetric information. According to Aumann's model of strategic conflict, the actions of the players appear to be stochastic—but not because the players actually make their choices at random.

The reason is that the players have imperfect information and so they consider the strategies of their opponents random variables. This allows him to drop the stochastic-independence assumption of Nash equilibrium.

Aumann's model of strategic conflict marks a cornerstone in traditional game theory and it serves as a fundament for many other models that have been developed so far. However, it has two conceptual drawbacks:

1. In order to justify correlated equilibrium, Aumann makes use of the Harsanyi doctrine, which claims that all players have the same prior. This assumption thwarts Savage's principal idea, namely that a subjective probability measure reflects the individual preferences of a rational subject. It implies that the probability measure is objective.
2. Aumann presumes also that the players have a common state space. Hence, they can very well be uncertain about the action of each other, but they cannot be uncertain about their *strategies*. This is because every random variable that takes place in a state space must be evident to all subjects that share the same state space. Thus, each player knows the strategies of the others and so we should expect that the strategic-independence assumption is violated in Aumann's model of strategic conflict. However, in order to justify correlated equilibrium, we have to assume that the players act independently.

Why should a rational solution in Aumann's model of strategic conflict be an equilibrium at all if everybody knows the strategy of each other, which means that the players are able to act in a *coherent* way? Hence, they can interact, but interaction goes beyond correlation and any other form of stochastic dependence. Indeed, we have seen in Chapter 6 that the rational solution of a coherent game need not be an equilibrium.

The solution of a finite 2-person zero-sum game is a Nash equilibrium if and only if it is minimax. If the game is coherent, we can justify the assumption that the players perform their best-response strategies. This means that both players know that the other gives a best response—irrespective of whatever strategy they decide to perform. Hence, the resulting solution is minimax, and although the strategic-independence assumption is clearly violated, it represents a Nash equilibrium. However, this is true just because the players have strictly opposing interests, in which case performing the best-response strategy seems to make sense, but we have seen in Chapter 6 that the best-response pattern turns out to be erroneous in most other coherent games. We conclude that finite 2-person zero-sum games represent an exceptional case.

A rational solution of a static, dynamic, or coherent game need not be a (Nash) equilibrium, but also the converse is true. This means that a Nash equilibrium need not be a rational solution. In order to eliminate all unreasonable Nash equilibria of a game, we can apply some of the many refinement procedures that can be found in the literature. However, the problem is that refinement comes to opposing results. For

example, payoff efficiency *discards* the Nash equilibrium Up/Left, whereas risk dominance and trembling-hand perfectness *propagate* Up/Left in the game show. It is quite unsatisfactory that we come to mutually exclusive conclusions after applying different refinement procedures. Moreover, although refinement is interesting in its own right, I think that the overall approach does not address the root of the problem: A rational solution need not be a Nash equilibrium. This will be discussed in the next chapter.

8 The Epistemic Approach

The solution of a strategic conflict depends essentially on what a player thinks about the others and what he thinks that the others think about the others, etc. This is the domain of epistemology, a branch of philosophy that is concerned with knowledge and belief. The epistemic approach is relatively new in game theory. In contrast to the traditional approach, it scrutinizes the strategic reasoning of the players by referring to subjective rather than objective probabilities.

A discussion of epistemic game theory can be found in Battigalli and Bonanno (1999), Board (2002), Brandenburger (1992, 2007), Dekel and Gul (1997), Dekel and Siniscalchi (2015) as well as Geanakoplos (1992). Some monographs on the subject matter are written by Brandenburger (2014), de Bruin (2010), and Perea (2012). The latter work marks the first textbook on that topic and is highly recommended for those who want to start studying epistemic game theory.

In contrast to traditional game theory, the epistemic program typically deals with belief rather than knowledge. Perea (2012, p. 66) points out that

“the term ‘knowledge’ is too strong to describe a player’s state of mind in a game. [...] Player i can at best have a belief about the rationality of player j , and this belief may very well be wrong!”

In this chapter, we still focus on static games. A player’s belief not only contains his opinion about the rationality of his opponents but also a subjective probability distribution regarding each other’s action. Epistemic game theory makes use of so-called belief hierarchies, which specify the players’ beliefs about each other, their beliefs about each other’s belief about each other, etc. Hence, the epistemic approach describes the strategic reasoning of the players and demonstrates that a rational solution of a game need not be a Nash equilibrium.

This work goes very much in the same direction: I assume that the players are rational. Hence, they possess some utility functions and subjective probabilities regarding the actions (or strategies) of each other. However, subjectivistic game theory does not require us to assume that any player thinks that the others are rational, too, and that the others have any specific beliefs about each other, etc. The subjectivistic approach differs from the epistemic approach in that the former guarantees that every rational player *has* a unique conjecture, i.e., a subjective probability distribution concerning the potential actions of his opponents, whereas the latter aims at explaining *how* the players come to their conclusions.

More precisely, epistemic game theory clarifies under which circumstances the subjective probability distributions of the players are consistent with their belief hierarchies, whereas subjectivistic game theory does not scrutinize the players’ reasoning about any other. Of course, also in the subjectivistic context we can think about each player’s individual reasons for having a specific prior—which, in fact, has often been done throughout this book—but our subjectivistic arguments are by far more un-

<https://doi.org/10.1515/9783110596106-008>

sophisticated compared with epistemic game theory. It appears that epistemic game theory is a spin-off of subjectivistic game theory that is concerned with the question of whether a specific action is not only optimal but also *reasonable*.

8.1 Rationalizability

We have already seen in Section 7.2.2.2 that Nash equilibrium is not a necessary condition for rationality, i.e., $R \neq N$. This is observed also by Bernheim (1984) and Pearce (1984), who show that finite games can have rational solutions that go far beyond Nash equilibrium—although the traditional assumption of common knowledge is satisfied! Bernheim (1984) concentrates on static games, whereas Pearce (1984) treats also dynamic games. In any case, the corresponding strategies are called rationalizable.¹⁷¹

Here, I focus on static n -person games in normal form. I provide only a brief and informal discussion of rationalizability. Nonetheless, its basic idea—as well as its strong connection to the subjectivistic approach—shall become clear to the reader.

8.1.1 Basic Idea

The theory of rationalizability is based on conjectures, which have already been used in Section 7.2 in the context of Nash equilibrium. Hence, let us assume that Player i has a conjecture, c_{ij} , about each other player $j \neq i$, which shall express his personal uncertainty about his actions. However, this makes sense only if we suppose that Player i thinks that the actions of his opponents are stochastically independent. Otherwise, he would have to take the joint probability distribution of their actions into account—not only their marginal distributions. Bernheim (1984, p. 1014) writes:

“A question arises here as to whether an agent’s probabilistic conjectures can allow for correlation between the choices of other players. In a purely non-cooperative framework, such a correlation would be nonsensical: the choices of any two agents are by definition independent events; they cannot affect each other.”

Thus, he presumes that (the players think that) the actions are stochastically independent. However, Aumann (1974, 1987) argues that the action of a player can very well depend stochastically on the action of another, which leads him to the notion of correlated equilibrium (see Section 7.4). Thus, Brandenburger and Dekel (2014) drop the stochastic-independence assumption of rationalizability. Further, Brandenburger and Friedenberg (2014) distinguish between “intrinsic” correlation, where the belief hierarchies of the players are assumed to depend on each other, and Aumann’s “extrinsic”

¹⁷¹ The term “rationalizability” goes back to Bernheim (1984).

correlation, typically attributed to a physical source of information or signal.¹⁷² I will come back to this point in Section 8.2.2, where the stochastic-independence assumption is dropped, but here I follow the original approach to rationalizability, which is based on the standard assumption of strategic *and* stochastic independence. Moreover, rationalizability presumes that the players have complete information and common knowledge (Pearce, 1984).

A strategy of a player is said to be rationalizable if and only if we can find some conjectures for this player about the other players that (i) are consistent with the common-knowledge assumption and (ii) render the given strategy a best response to his opponents. All strategies of a player that are not rationalizable can be eliminated. A solution of the game is considered rationalizable if and only if the associated strategies are rationalizable.

8.1.2 Solution Concept

In order to find the rationalizable strategies of a game, we can start with Player 1 and remove all of his strategies that can never be a best response. To be more precise, we remove a strategy if we cannot find any conjectures of Player 1, about Player 2, 3, \dots , n , such that the given strategy is optimal. Now, we go further by eliminating all strategies of Player 2 that can never be a best response to Player 1, 3, \dots , n , etc. This is done also for all other players, after which we hopefully end up with a smaller set of strategies. Then we start again with Player 1 and remove all strategies that can never be a best response by taking only those strategies of the other players into account that have not been eliminated in the first round, etc. This procedure is repeated until we are no longer able to eliminate any strategy. The remaining strategies are rationalizable.

A pure strategy that is strictly dominated by some mixed strategy can never be a best response. Thus, we can remove all pure strategies of a player that are strictly dominated by a mixed strategy. However, in general, the converse is not true. This means that a pure strategy that can never be a best response need not be strictly dominated by a mixed strategy—unless we consider a 2-person game (Bernheim, 1984, p. 1016) or allow for correlation (Brandenburger and Dekel, 2014).

For example, consider a 3-person game in which Ann, Andy, and Bob play against each other. Ann and Andy can choose only between Action a and Action b , whereas the action set of Bob is $\{1, 2, 3, 4\}$. It is assumed that all players obtain the same payoffs, which are presented in Table 8.1. The numbers on the upper left of the payoff matrices indicate the available actions of Bob.

Action 1 is not strictly dominated by any convex combination of Action 2, 3, and 4.¹⁷³ This can be seen like this: Let p be the probability that Bob chooses Action

¹⁷² It is worth emphasizing that both forms of correlation represent *stochastic* dependence, whereas interaction, which was discussed throughout Chapter 6, refers to *strategic* dependence.

¹⁷³ Thus, Action 1 cannot be strictly dominated by any mixed strategy that involves Action 1.

Table 8.1: Payoff matrices of a 3-person game between Ann, Andy, and Bob.

1			2			3			4		
Ann			Ann			Ann			Ann		
Andy			Andy			Andy			Andy		
<i>a</i>	<i>b</i>		<i>a</i>	<i>b</i>		<i>a</i>	<i>b</i>		<i>a</i>	<i>b</i>	
<i>a</i>	(1.5)	(0)	<i>a</i>	(3)	(0)	<i>a</i>	(0)	(0)	<i>a</i>	(1)	(1)
<i>b</i>	(0)	(1.5)	<i>b</i>	(0)	(0)	<i>b</i>	(0)	(3)	<i>b</i>	(1)	(1)

2 and q be the probability that he chooses Action 3, which means that he chooses Action 4 with probability $1 - p - q$. Suppose that Action 1 is strictly dominated by some convex combination of the other actions. Then we must have that

$$p3 + (1 - p - q) > 1.5$$

$$1 - p - q > 0$$

$$q3 + (1 - p - q) > 1.5.$$

The second inequality immediately implies that $p + q < 1$, whereas the first and third inequality yield

$$p2 - q > 0.5$$

$$q2 - p > 0.5.$$

Hence, we must have that

$$p2 + q(-1) > 0.5 \quad \text{and} \quad p(-1) + q2 > 0.5.$$

The reader can easily verify that the first inequality can be satisfied only for $p > 0.5$, but then the second inequality must be violated. Similarly, if the second inequality is satisfied, we must have that $p < 0.5$, in which case the first one is violated. Hence, Action 1 cannot be strictly dominated by any mixed strategy.

It remains to show that Action 1 cannot be a best response. Suppose that Action 1 is a best response. Now, let p be the probability that Ann chooses Action a and q be the probability that Andy chooses Action a . Then we must have that

$$1.5pq + 1.5(1 - p)(1 - q) \geq 3pq$$

$$1.5pq + 1.5(1 - p)(1 - q) \geq 3(1 - p)(1 - q)$$

$$1.5pq + 1.5(1 - p)(1 - q) \geq 1.$$

The first inequality implies that $p + q \leq 1$. From the second one we conclude that $p + q \geq 1$. Thus, we must have that $p + q = 1$ and so the third inequality reads

$$p(1 - p) \geq \frac{1}{3},$$

but it holds that $p(1 - p) \leq \frac{1}{4}$ for all $0 \leq p \leq 1$.

Hence, Action 1 cannot be a best response of Bob to Ann and Andy although this action, i.e., pure strategy, is not strictly dominated by any mixed strategy! The problem is that rationalizability does not allow for stochastic dependence. If Bob were allowed to believe that the actions of Ann and Andy are correlated, then we would come to the conclusion that Action 1 *can* be a best response and thus it would not be possible to eliminate this strategy. This point will be explained in more detail in Section 8.2.3.

8.1.3 Some Examples

8.1.3.1 An Anti-Coordination Game

Consider the chicken game, which is *the* anti-coordination game per se. Its payoff matrix is given in Table 8.2. According to the algorithm described in the previous section, we cannot eliminate any strategy in this game and thus all (pure and mixed) strategies are rationalizable. How can we explain that this startling result is consistent with common knowledge?

Table 8.2: Payoff matrix of the chicken game.

Andy	Bob	
	Swerve	Straight
Swerve	(1, 1)	(0, 2)
Straight	(2, 0)	(-9, -9)

Andy considers Straight optimal if he believes that Bob swerves, while Bob will prefer to swerve if he thinks that Andy goes straight, etc. Thus, we have already created a periodic sequence, i.e., a cycle, of optimal choices, which can be repeated infinitely many times. Hence, the pure strategies that are involved in that cycle, i.e., Straight for Andy and Swerve for Bob, are rationalizable. The same principle leads us to the conclusion that it is also rationalizable for Andy to swerve and for Bob to go straight. It may even happen that both players go straight, although they are rational and know that the other is rational, too, and know that the other knows that the other is rational, etc. Hence, common knowledge does not imply the transparency of reason and so the solution of the game need not be a Nash equilibrium!

Now, suppose that Andy decides to swerve with probability 90 % and to go straight with probability 10 %. This strategy belongs to the mixed-strategy Nash equilibrium (see Figure 7.2). An equilibrium strategy, let it be pure or mixed, is always rationalizable: Andy performs the equilibrium strategy because he thinks that Bob performs the equilibrium strategy, and Bob performs the equilibrium strategy because he thinks that Andy performs the equilibrium, etc. Thus, every Nash equilibrium is rationalizable.

Further, since the mixed-strategy Nash equilibrium of the chicken game consists only of (completely) mixed strategies, *all* pure and mixed strategies of Andy and Bob

are rationalizable. It makes no difference at all whether Andy performs his equilibrium strategy or any other strategy because he always achieves the same expected payoff, and the same holds true, *mutatis mutandis* for Bob. We conclude that all solutions of the chicken game are consistent with the common-knowledge assumption.

8.1.3.2 A Discoordination Game

The payoff matrix of Matching Pennies, which is a typical discoordination game, can be found in Table 8.3. Since we cannot eliminate any strategy by applying the aforementioned algorithm, it turns out that all strategies are rationalizable. For example, Ann may choose Heads because she thinks that Bob chooses Heads, which is optimal for Bob if he thinks that she chooses Tails, which is optimal for Ann if she thinks that Bob chooses Tails, and this is optimal for Bob if he thinks that she chooses Heads, etc. This completes our cycle of optimal choices and demonstrates that all pure strategies are rationalizable both for Ann and for Bob.

Table 8.3: Payoff matrix of Matching Pennies.

Ann	Bob	
	Heads	Tails
Heads	(1, -1)	(-1, 1)
Tails	(-1, 1)	(1, -1)

In the same way, we can explain why every mixed strategy is rationalizable, too. It can be seen from Figure 4.6 that Matching Pennies contains precisely one Nash equilibrium, in which both players decide to choose Heads and Tails each with probability 50 %. This means that every mixed strategy is consistent with the common-knowledge assumption. We conclude that anything can happen in this game, depending on the individual conjectures of the players.

8.1.3.3 A Coordination Game

Let us reconsider the reunion game, which is a typical coordination game. Its payoffs are given in Table 8.4. Even here we cannot eliminate any pure or mixed strategy by applying the aforementioned algorithm, and thus all strategies in this game turn out to be rationalizable.

Table 8.4: Payoff matrix of the reunion game.

Ann	Bob	
	Station	Museum
Station	(1, 1)	(0, 0)
Museum	(0, 0)	(1, 1)

Ann may consider Station optimal if she thinks that Bob goes to the station and this is optimal for Bob, too, if he thinks that Ann prefers to go to the station, etc. We can apply the same principle to Museum and come to the conclusion that also Museum is rationalizable both for Ann and for Bob. Hence, although it is common knowledge that both players are rational, it can very well happen that Ann goes to the station, whereas Bob goes to the museum.

Instead, Ann might decide to throw a coin, in which case she would go to the station and to the museum each with probability 50 %. This (mixed) strategy is rationalizable, too, because it is an equilibrium strategy. Thus, Ann's other mixed strategies are rationalizable as well, and the same conclusion can be drawn, *mutatis mutandis*, for Bob. Once again, we conclude that *all* pure- and mixed-strategy solutions are rationalizable in this game.

8.1.3.4 The Prisoners' Dilemma

Now, consider the prisoners' dilemma in its standard form as a counterexample. Its penalty matrix can be found in Table 8.5. Since Deny is strictly dominated by Confess, no rational prisoner who takes only the penalties into account will ever prefer to deny. In fact, the entries in the second row of the penalty matrix are always better for Mary than the entries in the first row. Thus, we cannot find any conjecture of Mary about Joe such that Deny appears to be at least as good as Confess. These arguments hold true, *mutatis mutandis*, for Joe.

Table 8.5: Penalty matrix of the prisoners' dilemma.

Mary	Joe	
	Deny	Confess
Deny	(1, 1)	(5, 0)
Confess	(0, 5)	(4, 4)

We conclude that Deny is not rationalizable, which enables us to eliminate the first row and the first column of the penalty matrix. Hence, rationalizability leads us to the typical, i.e., noncooperative, solution of the prisoners' dilemma.

8.1.4 General Remarks

8.1.4.1 Nash Equilibria

The proof of the following theorem is trivial and thus it can be skipped. Recall that we are still speaking about a finite game in normal form.

Theorem 14 (Rationalizability, I). *Every strategy that is part of a Nash equilibrium is rationalizable.*

Further, it is not difficult to see that every *action* that is part of a Nash equilibrium is rationalizable. In fact, all actions that have been deemed rationalizable in the previous examples were part of some Nash equilibrium. Nonetheless, a game can contain rationalizable actions that are *not* part of a Nash equilibrium.

For example, let us consider a game with the payoff matrix in Table 8.6. Its best-response diagram is given in Figure 8.1. As we can see, there is precisely one pure-strategy Nash equilibrium, which can be found on the upper left of the best-response diagram. Moreover, there are infinitely many mixed-strategy Nash equilibria, i.e., (p, q) with $\frac{1}{2} \leq p < 1$ and $q = 1$.

Table 8.6: Payoff matrix of a 2×2 game.

Ann	Bob	
	Left	Right
Up	(3, 1)	(1, 0)
Down	(3, 1)	(0, 2)




Figure 8.1: Best-response curves of the 2×2 game.

Right is not part of any Nash equilibrium, but it is rationalizable: Bob may choose Right if he assumes that Ann chooses Up and Down each with probability $\frac{1}{2}$. Ann may choose Up and Down each with probability $\frac{1}{2}$ if she assumes that Bob chooses Left. Bob may choose Left if he assumes that Ann chooses Up and Down each with probability $\frac{1}{2}$, etc. Note that Down is rationalizable, too, because this action is part of a mixed-strategy Nash equilibrium. Hence, the fact that Down is *weakly* dominated by Up does not imply that Down must be discarded under the common-knowledge assumption!

Together with Nash's existence theorem (see Section 7.2), Theorem 14 implies that every finite game has a rationalizable solution. The following theorem asserts that a rationalizable solution must be a Nash equilibrium if we cannot find any other

rationalizable solution. This is an immediate consequence of Theorem 14, too, together with Nash's existence theorem. Thus, I skip also the proof of the following theorem.

Theorem 15 (Rationalizability, II). *If a rationalizable solution is unique, it must be a Nash equilibrium.*

Hence, if the common-knowledge assumption leads us to one and *only one* (rationalizable) solution, this must be a Nash equilibrium. However, the problem is that most finite games possess many rationalizable solutions and some of them are not Nash equilibria. Put another way, common knowledge does not imply Bacharach's transparency of reason (see Section 7.2.3).

8.1.4.2 Mixed Strategies

In order to find the rationalizable strategies of a player, we must eliminate all pure strategies that are strictly dominated by any other strategy.¹⁷⁴ This sounds trivial, but it is not as easy as it seems: When searching for strict dominance, we must take also mixed strategies into account. However, does this imply that a player really compares a pure strategy with a mixed strategy? The answer is “No”!

For example, consider a game with the payoff matrix in Table 8.7. At first glance, we cannot eliminate anything, but this is a fallacy. Down is not strictly dominated by Up or by Middle. However, it is strictly dominated by a convex combination of Up and Middle: Ann could choose Up and Middle each with probability $\frac{1}{2}$ in order to obtain the expected utilities 4 and 6, which is strictly better than 2 and 5.

Table 8.7: Payoff matrix of a 3×2 game.

Ann	Bob	
	Left	Right
Up	(1, 3)	(8, 2)
Middle	(7, 2)	(4, 1)
Down	(2, -1)	(5, 0)

The fact that Down is strictly dominated by Ann's mixed strategy just implies that it is always better to choose either the *pure* strategy Up or the *pure* strategy Middle. Thus, we cannot find any conjecture of Ann about Bob such that Down turns into an optimal action. Figure 8.2 reveals that for each conjecture of Ann, which is denoted by

¹⁷⁴ Remember that this procedure is necessary but, in general, not sufficient to find all rationalizable solutions of a game (see Section 8.1.2).

the probability $0 \leq q \leq 1$ that Bob chooses Left, it is either better to choose Up or Middle *alone*. Hence, rationalizability does not require us to assume that the players perform any mixed strategy, i.e., that they randomize their actions. In particular, it is clear that all actions that are part of a rationalizable strategy are rationalizable, too (Bernheim, 1984, p. 1016). This means that the players may always choose a definite action.




Figure 8.2: Down is not rationalizable.

The fact that either Up or Middle is better than Down is not caused by accident: If some action is strictly dominated by a mixed strategy, then there must exist a better action. This can be seen as follows: If Action $a \in A$ is strictly dominated by a mixed strategy s , then the expected utility of the mixed strategy, $EU(s)$, must be greater than $EU(a)$. Note that this holds for *all* conjectures c_{ij} of Player i about each other player $j \neq i$. Hence, the mixed strategy must contain some action $b \in A$ such that $EU(b) \geq EU(s) > EU(a)$. Thus, Action a cannot be a best response, since there is always a better action.

To sum up, Down is not rationalizable and thus it can be eliminated. After we have eliminated Down, we can see that Right is strictly dominated by Left. Hence, we can eliminate Right, too, but then we can discard also Up. We conclude that the only rationalizable solution of this game is Middle/Left, which is also a Nash equilibrium. This example demonstrates that in some cases it is possible to derive a unique rational solution by applying the assumption of common knowledge although no action is strictly dominated by another.

8.1.5 The Subjectivistic Interpretation

The concept of rationalizability reminds us of the subjectivistic approach to game theory, which is propagated in this book. In fact, Bernheim (1984) refers to Savage (1954) when motivating his work. Similarly, Pearce (1984) points out that players lack an objective probability distribution. Instead, their decisions are based on their (subjective) priors. Indeed, the mere fact that the players have conjectures, i.e., do not *know* the

(pure or mixed) strategy of any other player, suggests that Bernheim and Pearce leave the objectivistic framework. Another crucial observation is that rationalizability derives all rationalizable solutions only by requiring that the choice of each player is optimal from his *own*, i.e., individual, perspective but not from the perspective of an outside observer. This distinguishes rationalizability from traditional game theory, which suggests that all players come to the *same* conclusion regarding the optimal strategy of each other by common knowledge.

A rational player need not and will not randomize his actions (see Section 7.2.4). In the case in which there is one and only one optimal action, he will choose precisely this action. Otherwise, he may freely decide between his optimal actions, but then it makes no sense at all to flip a coin or to apply any other random generator, since each optimal action yields the same expected utility. Hence, we may consider the conjecture of a player his subjective probability distribution on the action set of another player, whereas the conjectures mentioned in Section 7.2.3 suggest that the players randomize their actions. In the context of Nash equilibrium, Player i thinks that Player j uses a random generator with discrete probability distribution c_{ij} and he believes that the random generators produce stochastically independent outcomes. By contrast, rationalizability does not require the players to randomize their actions at all. This means that c_{ij} reflects the uncertainty of Player i regarding the *definite* action of Player j .

Despite the many similarities, rationalizability differs in several aspects from the subjectivistic approach presented in this book:

1. We are not restricted to games in which the players think that the actions of their opponents are stochastically independent.
2. We do not have to assume that the players have complete information or, even more, common knowledge.
3. We are not restricted to static or dynamic games. This means that we may consider also coherent games.

8.2 Reasonability

8.2.1 Common Belief

Epistemic game theory has extended the general idea of rationalizability, which is based, either implicitly (Bernheim, 1984) or explicitly (Pearce, 1984), on the assumption of common knowledge. However, this assumption is quite strong and we can obtain similar results by making the weaker assumption of common belief.

Consider any static n -person game and let G_i be a list of propositions of Player i regarding the number of players, their action sets, and their payoffs, which means that all players are rational. In addition, G_i shall imply that each player chooses some op-

timal action.¹⁷⁵ Common belief means that Player $i = 1, 2, \dots, n$ believes that G_i is true, he believes that all players believe that G_i is true, he believes that all players believe that all players believe that G_i is true and so on, ad infinitum. This is called “common” belief although the players can have their own lists of propositions, G_1, G_2, \dots, G_n , and those lists can very well differ from each other.

More precisely, let G be some list of propositions, i.e., an entire description of the game. It is clear that every list of propositions is a proposition, too. Hence, let us define $\beta^0(G) := G$. Further, let $\beta(G)$ be the proposition that all players believe that the proposition G is true. Then $\beta^2(G)$ means that all players believe that all players believe that G is true, etc. Common belief states that Player $i = 1, 2, \dots, n$ believes that the proposition $\beta^k(G_i)$ is true for $k = 0, 1, \dots$, which implies that the player believes that his own proposition, G_i , is true.¹⁷⁶ Thus, common belief states that each player believes in common belief in his *own* proposition.

8.2.1.1 The Mallorca Game

The concept of common belief shall be demonstrated by the Mallorca game, which has already been discussed in Section 4.4.2. The following exposition is based on Perea (2012, Section 3.1). Table 8.8 contains the market shares of Ann—depending on her and Bob’s decision. We simply obtain Bob’s decision matrix by exchanging Ann with Bob and thus we may focus on Ann without loss of generality.

Table 8.8: Ann’s decision matrix of the Mallorca game.

Ann	Bob						
	a	b	c	d	e	f	g
a	$\frac{6}{12}$	$\frac{1}{12}$	$\frac{2}{12}$	$\frac{3}{12}$	$\frac{4}{12}$	$\frac{5}{12}$	$\frac{6}{12}$
b	$\frac{11}{12}$	$\frac{6}{12}$	$\frac{3}{12}$	$\frac{4}{12}$	$\frac{5}{12}$	$\frac{6}{12}$	$\frac{7}{12}$
c	$\frac{10}{12}$	$\frac{9}{12}$	$\frac{6}{12}$	$\frac{5}{12}$	$\frac{6}{12}$	$\frac{7}{12}$	$\frac{8}{12}$
d	$\frac{9}{12}$	$\frac{8}{12}$	$\frac{7}{12}$	$\frac{6}{12}$	$\frac{7}{12}$	$\frac{8}{12}$	$\frac{9}{12}$
e	$\frac{8}{12}$	$\frac{7}{12}$	$\frac{6}{12}$	$\frac{5}{12}$	$\frac{6}{12}$	$\frac{9}{12}$	$\frac{10}{12}$
f	$\frac{7}{12}$	$\frac{6}{12}$	$\frac{5}{12}$	$\frac{4}{12}$	$\frac{3}{12}$	$\frac{6}{12}$	$\frac{11}{12}$
g	$\frac{6}{12}$	$\frac{5}{12}$	$\frac{4}{12}$	$\frac{3}{12}$	$\frac{2}{12}$	$\frac{1}{12}$	$\frac{6}{12}$

¹⁷⁵ Recall that the mere fact that a player is rational does not imply that his *choice* is optimal.

¹⁷⁶ Equivalently, we could have assumed that the player believes that $\beta^k(G_i)$ is true for $k = 1, 2, \dots$ because “If I believe that I believe in God, then I believe in God.” This follows from Savage’s postulates of rational choice, which have been discussed in Section 1.3.

If Ann believes that Bob chooses Location d , her best choice is d , too. By contrast, if she thinks that Bob chooses Location g , she prefers f , etc. However, a closer look into Ann's decision matrix reveals that g is strictly dominated by f . Hence, Ann would never prefer to choose g . Since Ann thinks that Bob is rational, and that he has the same decision matrix, she comes to the conclusion that he will never choose g , too. The same argument applies to a , which is dominated by b . Thus, Ann is convinced that Bob abandons both Location a and Location g .

After a and b have been eliminated from Bob's action set, it turns out that f is strictly dominated by e and b is strictly dominated by c . Hence, Ann will eliminate also b and f . Further, she thinks that Bob thinks that she is rational, and that she has the same decision matrix, which means that she thinks that he thinks that she eliminates a and g , but then b and f are strictly dominated by c and e also from Bob's perspective. Thus, Ann believes that he will eliminate b and f , too. If we go further with the same arguments, we finally come to the conclusion that Ann is convinced that Bob chooses Location d . Due to the same logic, also Bob should be convinced that Ann chooses Location d . Hence, according to epistemic game theory, i.e., given that the assumption of common belief is satisfied, the only rational solution of the Mallorca game is (d, d) .

8.2.1.2 Common Belief vs. Knowledge

What is the difference between common belief and common knowledge? Common belief differs from common knowledge in that it allows the players to have “wrong beliefs,” i.e., beliefs that do not match reality (Perea, 2012, p. 66). By contrast, common knowledge requires evidence and, by its very definition, says that the players cannot fail at all. This assumption is much stronger than common belief.

For example, in the Mallorca game common knowledge implies that Ann and Bob know that they *prefer* Location d , whereas common belief allows them to be wrong in their opinions. Things become even more intricate if we would presume that the players know that they *choose* an optimal strategy, that they know that they know that they choose an optimal strategy, etc. In this case, common knowledge would imply that Ann and Bob know the other's action. However, the Mallorca game is static and so the private information sets of the players are trivial, which means that nobody can know the choice of the other! This issue has already been elaborated in Section 7.1.2.

Thus, in my opinion, the overall construction of common knowledge is quite problematic. Admittedly, it is still impossible for a player to *believe* that he performs any specific action: Belief always refers to an event, but the action of the player cannot be an event from his own perspective.¹⁷⁷ However, common belief does not suffer from

¹⁷⁷ Remember that a subject believes that $F \in \mathcal{F}$ happens if and only if he assigns F probability 1.

the same problems as common knowledge. This can be seen like this: If I know that you know that it is sunny, then I must know that it is sunny. By contrast, if I just believe that you believe that it is sunny, then I need not believe that it is sunny. Thus, Ann might very well believe that Bob believes that she chooses Location d without believing herself that she chooses Location d .

Hence, in a static game, we can unscrupulously say that “I believe that you *choose* Action a ,” but the assertion “I know that you choose Action a ” cannot be true in that context. On the contrary, saying that “I know that you *prefer* Action a ” poses no problem at all and, of course, “I believe that you prefer Action a ” is even more harmless. However, the latter two statements do not imply that the speaker assigns probability 1 to Action a , since it can very well happen that somebody prefers an action but still performs another one. By contrast, the very first statement, i.e., “I believe that you choose Action a ,” conveys that the speaker considers each action other than a null.

8.2.1.3 Subjectivistic vs. Epistemic Game Theory

Despite the many obvious similarities between the subjectivistic and the epistemic approach, in general, subjectivistic game theory does not come to the same conclusions as epistemic game theory.¹⁷⁸ The reason is that subjectivistic game theory does not require common belief regarding any aspect of the game. Strategic conflicts that occur in real life are often characterized by the fact that the players’ personal opinions about the game diverge. Typically, the players have incomplete information because they do not know the utility functions of the others. Indeed, epistemic game theory is aware of the problem of incomplete information. Just like in traditional game theory, it usually transforms incomplete information into imperfect information and considers the game Bayesian (see Section 7.3). However, my arguments go a little bit further.

In general, the players do not have the same beliefs about all relevant aspects of the game, i.e., their own (Bayesian) models, G_1, G_2, \dots, G_n , differ from each other. We could even assume that some player doubts that another player is rational at all.¹⁷⁹ Thus, it is hard to imagine why a (rational) player should believe that his opponents believe in his *own* model and that they believe that the others believe in his own model, etc. For example, Ann or Bob might not agree on the rules of the game or they could have different opinions regarding its payoffs, etc. We could imagine also that Ann and Bob doubt that the other is rational. Nonetheless, if we make the additional assumption of common belief in Section 4.4.2, the epistemic solution of the Mallorca game, i.e., (d, d) , coincides with the subjectivistic solution.

To sum up, the subjectivistic approach can be considered more general than the epistemic approach and the former comes to the same conclusions as the latter under the common-belief assumption. Finally, the reader must decide for himself whether

¹⁷⁸ The subjectivistic solution of the Mallorca game has been elaborated in Section 4.4.2.

¹⁷⁹ Of course, in this case his list of propositions does not imply the rationality of the other player.

or not this assumption can really be justified for the specific situation, i.e., strategic conflict, at hand.

8.2.2 Stochastic Dependence

Players are usually uncertain about the actions of their opponents. Hence, they assign a positive probability to more than one action of some opponent and so their conjectures are mixed. In a multiperson game, i.e., $n > 2$, a player need not think that the actions of his opponents are *stochastically* independent, although he may assume that their actions are *strategically* independent. Thus, even in a static game, it can very well happen that their choices are conditionally independent but unconditionally dependent. This can be clarified by simple considerations of subjectivistic decision theory.

For example, suppose that Ann is playing a game with Andy and Bob in which the payoffs depend on whether they arrive at work by car or by bicycle. Their choices shall be strategically independent, which means that no player knows which vehicle the others are going to choose. Ann thinks that Andy and Bob each chooses Car with probability 90 % if it rains and Bicycle with probability 60 % if it does not rain. Their choices are conditionally independent both on the event “Rain” and on the event “No rain.” Ann assigns Rain probability 30 %. Her situation is illustrated in Figure 8.3, which clearly reveals that the actions of Andy and Bob are conditionally independent, i.e., they do not depend on one another *given* any specific weather condition




Figure 8.3: Ann's event tree in a 3-person game.

The same situation can be expressed, equivalently, by the event tree in Figure 8.4. The probability that Andy decides to go to work by car is

$$P(\text{Andy chooses Car}) = 0.3 \cdot 0.9 + 0.7 \cdot 0.4 = 0.55$$

and it can easily be seen that $P(\text{Bob chooses Car}) = 0.55$, too. Further, the probability of Rain conditional on the event that Andy chooses Car can be calculated by the Bayes rule:

$$\begin{aligned} P(\text{Rain} \mid \text{Andy chooses Car}) &= \frac{P(\text{Andy chooses Car} \mid \text{Rain}) \cdot P(\text{Rain})}{P(\text{Andy chooses Car})} \\ &= \frac{0.9 \cdot 0.3}{0.55} = 0.4909. \end{aligned}$$

In the same way, we obtain the probability that it rains given that Andy decides to go to work by bicycle:

$$\begin{aligned} P(\text{Rain} \mid \text{Andy chooses Bicycle}) &= \frac{P(\text{Andy chooses Bicycle} \mid \text{Rain}) \cdot P(\text{Rain})}{P(\text{Andy chooses Bicycle})} \\ &= \frac{0.1 \cdot 0.3}{0.45} = 0.0667. \end{aligned}$$




Figure 8.4: Ann's equivalent event tree in the 3-person game.

We conclude that the probability that Bob chooses Car given that Andy chooses Car is

$$\begin{aligned} P(\text{Bob chooses Car} \mid \text{Andy chooses Car}) &= 0.4909 \cdot 0.9 + 0.5091 \cdot 0.4 \\ &= 0.6455, \end{aligned}$$

whereas the probability that Bob chooses Car given that Andy chooses Bicycle is

$$\begin{aligned} P(\text{Bob chooses Car} \mid \text{Andy chooses Bicycle}) &= 0.0667 \cdot 0.9 + 0.9333 \cdot 0.4 \\ &= 0.4334. \end{aligned}$$

Hence, the choices of Andy and Bob are (unconditionally) *dependent* from Ann's perspective. We can ignore such kind of complications by focusing on 2-person games.

8.2.3 Optimality vs. Rationality

Perea (2012, Definition 2.4.5) distinguishes between optimal and rational actions:

- An action is said to be optimal if and only if there is no better action, given the player's conjecture about his opponents.¹⁸⁰
- An action is said to be rational if and only if it is optimal for *some* conjecture about his opponents.

Remember that Location *d* is the only *rationalizable* action for a player in the Mallorca game. This is because we can eliminate *a, b, c, e, f, g* by an iterated application of the dominance principle. Epistemic game theory explains why we should eliminate all and *only* all strictly dominated actions in order to deduce the set of reasonable solutions. Recall that this is not true, in general, for the concept of rationalizability, which allows us to eliminate *more* actions than those that are strictly dominated by another strategy (see Section 8.1.2). However, it is obvious that this makes no difference in the Mallorca game.

The following theorem presumes that the game is static and that the players believe in strategic independence, which is a standard assumption of epistemic game theory.

Theorem 16 (Dominance). *Consider any finite game in normal form. An action is rational if and only if it is not strictly dominated by any pure or mixed strategy.*

Proof. Let *A* be any real-valued matrix. Then exactly one of the following two statements is satisfied¹⁸¹:

- There exists some nonzero vector $x \geq \mathbf{0}$ such that $Ax = \mathbf{0}$.¹⁸²
- There exists some vector *y* such that $y'A > \mathbf{0}$.

180 Now, I speak about the “conjecture” and not about the “conjectures” of the player. This is because in the context of reasonability, each player possesses a joint probability distribution about the others' actions but not $n - 1$ marginal probability distributions about each other's action.

181 This result is referred to as Gordan's theorem (Bernstein, 2009, Fact 4.11.16).

182 The symbol $\mathbf{0}$ denotes a vector of zeros.

Fix some player and let Action a be the last element of his action set. We prove that Action a is strictly dominated by a pure or mixed strategy if and only if it is irrational. We may treat the game like a 2-person game by aggregating all combinations of actions of the opponents in one action set of a fictional adversary. Let U be the decision matrix of the player except for the last row, which shall be denoted by u . Define $\Delta := U - \mathbf{1}u$ and $A := [\Delta \mathbf{I}]$.¹⁸³ Hence, Action a is rational if and only if there exists some nonzero vector $x \geq \mathbf{0}$ such that $Ax = \mathbf{0}$. If Action a is irrational, we must have that $y'A > \mathbf{0}$. This means that Action a is strictly dominated by a mixed strategy that assigns each row of U a positive probability. Conversely, if Action a is strictly dominated by a mixed strategy that assigns each row of U a positive probability, we have a vector y such that $y'A > \mathbf{0}$, which means that Action a is irrational. If y assigns some rows of U probability 0, we can eliminate the corresponding rows of A and columns of the identity matrix \mathbf{I} . Let \tilde{A} be the reduced matrix such that $\tilde{y}'\tilde{A} > \mathbf{0}$. Hence, $\tilde{A}\tilde{x} = \mathbf{0}$ cannot have any solution with $\tilde{x} \geq \mathbf{0}$ and $\tilde{x} \neq \mathbf{0}$, which means that $Ax = \mathbf{0}$ cannot have any solution with $x \geq \mathbf{0}$ and $x \neq \mathbf{0}$ either. Thus, Action a must be irrational. \square

8.2.4 Reasonable Solutions

We have clarified both the notion of optimal and the notion of rational action, but what constitutes a *reasonable* action from the viewpoint of epistemic game theory? First of all, common belief presumes that the action of the player is rational, but this is by far not enough. It implies also that the player believes that each other is rational, too. Thus, he should consider each irrational action of his opponents null. Moreover, he believes that all other players believe that their opponents are rational, which means that the player believes that they consider all irrational actions of the other players null, too, and so on, ad infinitum.

Suppose that $G_1 = G_2 = \dots = G_n =: G$. Theorem 16 asserts that an action is rational if and only if it is not strictly dominated by any pure or mixed strategy. Hence, in order to find all reasonable actions of the game, first of all we have to eliminate all strictly dominated actions. This step is motivated by Proposition G , which implies that all players are rational. However, common belief contains also the proposition $\beta(G)$, i.e., all players believe that all players are rational. Thus, we must eliminate also all actions that are strictly dominated by taking only those actions into account that survive after the first step, etc. In fact, common belief states that (the players think that) $\beta^k(G)$ is true for $k = 0, 1, \dots$

Hence, the corresponding solution procedure is easy (Perea, 2012, Chapter 3.7):

1. Eliminate all actions of all players that are strictly dominated by taking all possible actions of the other players into account.

¹⁸³ Here, $\mathbf{1}$ denotes a vector of ones and \mathbf{I} is an identity matrix.

2. Eliminate all actions of all players that are strictly dominated by taking only the residual actions of the other players into account.
3. Repeat this procedure until there is no more action that can be eliminated.

All actions that survive are reasonable and, conversely, all reasonable actions survive. This holds true for any number $n > 1$ of players. Note, however, that the resulting actions need not be *rationalizable* in the case of $n > 2$.

It is worth emphasizing that an action is considered reasonable only from the perspective of some specific player although the algorithm is supposed to eliminate the irrational actions of *all* players. However, the iteration takes place only in the mind of the corresponding player and is independent of the belief of any other player. However, here we make the simplifying assumption that $G_1 = G_2 = \dots = G_n$ and thus our result is equally valid for Player $i = 1, 2, \dots, n$.

Now, a solution is said to be reasonable if and only if the actions that are part of that solution are reasonable for each single player.

For example, in the Mallorca game, the only reasonable choice is Location d because all other actions can be eliminated by an iterated elimination of strictly dominated actions. This holds true both for Ann and for Bob, since we assume that their individual propositions about the game, i.e., G_{Ann} and G_{Bob} , are identical.

What is the main difference between reasonability and rationalizability?

1. Reasonability is based on common belief, whereas rationalizability requires common knowledge.
2. The former allows the actions of the players to be stochastically dependent, whereas the latter presumes that they are stochastically independent.

Since reasonability allows for stochastic dependence, it is more conservative than rationalizability. Hence, if the number of players is greater than 2, it can happen that a solution is reasonable but not rationalizable although the assumption of common knowledge is satisfied. However, irrespective of the number of players, a rationalizable solution is always reasonable under the assumption of common knowledge, in which case we have that $G_1 = G_2 = \dots = G_n = G \equiv \Gamma$. Hence, reasonability can be considered a generalization of rationalizability.

8.2.5 Belief Hierarchies

8.2.5.1 The Epistemic Model

The preceding explanations about common belief are quite general. Epistemic game theory concretises the situation of the players by using a semantic model. The basic epistemic model for an n -person game in normal form goes like this¹⁸⁴: Every player

184 It is still presumed that the players choose their actions independently, although each player may consider the actions of his opponents stochastically dependent.

$i \in \{1, 2, \dots, n\}$ has a finite action set $A_i \neq \emptyset$ and a utility function $u_i : \times_{j=1}^n A_j \rightarrow \mathbb{R}$. Each action $a_i \in A_i$ of every player i is associated with some probability distribution on $\times_{j \neq i} A_j$, where each action $a_j \in A_j$ with $j \neq i$ is associated with some probability distribution on $\times_{k \neq j} A_k$ and so on, ad infinitum.

I should explain the peculiarity of this model in more detail: Consider Player i and some action $a_i \in A_i$. The probability distribution that is associated with Action a_i represents the conjecture of Player i regarding the actions of his opponents, given that Player i performs Action a_i . Now, i.e., at the next level, each action of Player $j \neq i$ is associated with some probability distribution that reflects Player i 's belief about the conjecture of Player j regarding the actions of his opponents, given that Player j performs the corresponding action, etc. It is worth emphasizing that the probability distributions that are obtained at the second level need *not* coincide with the first-level conjectures of the opponents of Player i . Hence, the conjectures at all subsequent levels take place only in the mind of Player i and, in general, they do not reflect the actual conjectures of his opponents.

The epistemic model is sometimes called (interactive) belief system because it represents the belief hierarchies of the players. A belief hierarchy specifies the conjecture of a player, his belief regarding the conjectures of his opponents, his belief about their beliefs regarding the conjectures of their opponents, etc. Hence, a belief hierarchy is an infinite tree of beliefs. The epistemic model enables us to deduce the belief hierarchies of the players in a quite elegant way. We can find many variants and extensions of this model in the literature. For example, we could consider the player's action, conjecture, and payoffs part of some *type* in order to transform incomplete information into imperfect information à la Harsanyi (1967–1968),¹⁸⁵ etc., but here I assume for the sake of simplicity that the players have identical opinions regarding their payoffs. The interested reader may consult the contributions mentioned at the beginning of this chapter for a more sophisticated treatment of the subject matter.

Now, we can apply the terminology that we have used so far to the epistemic model. Consider some player $i \in \{1, 2, \dots, n\}$ and any action $a_i \in A_i$:

- The player believes that all players are rational if and only if he is convinced that his opponents choose a rational action.
- He believes that all players believe that all players are rational if and only if he is convinced that his opponents believe that all players are rational.
- He believes in common belief in rationality if and only if he believes that all players are rational, he believes that all players believe that all players are rational and so on, ad infinitum.

Action a_i is reasonable if and only if we can find a belief hierarchy for which (i) this action is optimal for Player i and (ii) he believes in common belief in rationality.

185 The overall approach has been explained in Section 7.3.

Finally, as we already know, a solution of the game is reasonable if and only if the actions that are part of that solution are reasonable for each single player.

8.2.5.2 Beliefs Diagrams

We can illustrate belief hierarchies in a convenient way by using beliefs diagrams (Perea, 2012). In such a diagram, the conjectures of the player who is taken into consideration are represented by arrows that start on the left-hand side and point to some actions in the middle of the diagram. More precisely, each arrow points to some action of another player that is not considered null. Conversely, each action with positive probability has an ingoing arrow. This shall be demonstrated by the Mallorca game. A beliefs diagram of this game is depicted in Figure 8.5.¹⁸⁶




Figure 8.5: A beliefs diagram of the Mallorca game.

Ann never considers Location *a* optimal because *a* is strictly dominated by *b*. The same holds true for Location *g*, which is strictly dominated by *f*. Hence, *a* and *g* have no outgoing arrows in the beliefs diagram, which means that *a* and *g* cannot be optimal for *any* conjecture of Ann about Bob. By contrast, *b* and *f* each have one outgoing arrow: Location *b* is optimal if Ann believes that Bob chooses *a*, whereas Location *f* is optimal if she thinks that he chooses *g*, etc. Beliefs diagrams are usually not unique. For example, Ann considers Location *d* optimal if she believes that Bob chooses *d*, too, but Location *d* would be optimal also if she would expect that Bob chooses *c* or *e*.

The conjecture of a player about the actions of his opponents is called his first-order belief, his belief regarding the conjectures of his opponents about the actions of the other players represents his second-order belief, etc.

¹⁸⁶ It can be found also in Perea (2012, p. 69).

The arrows on the left-hand side of Figure 8.5 illustrate Ann’s first-order belief, whereas the arrows on the right-hand side represent her second-order belief. Arrow *A* meets Arrow *B*. This means that Location *c* makes sense for Ann if she believes that Bob prefers Location *b*, which is true if he believes that Ann prefers Location *a*, but this makes no sense if she believes that Bob believes that she is rational. The problem is that Bob’s outgoing arrow for Location *b*, i.e., Arrow *B*, leads to Location *a*, which can never be optimal for Ann. In fact, the only combination of arrows in Figure 8.5 that does not suffer from the aforementioned problem is $d \rightarrow d \rightarrow d$, in which case we have created a cycle of optimal choices. Thus, Location *d* turns out to be a reasonable action for Ann. Are we able to create any other beliefs diagram in order to justify another action? The answer is “No”! Since *d* is the only action that survives an iterated elimination of strictly dominated actions, it is the only reasonable action of Ann. We can apply the same arguments to Bob and come to the conclusion that (d, d) is the unique reasonable solution of the Mallorca game.

Now, let us reconsider the 3×2 game from Section 8.1.4.2. Its payoff matrix can be found in Table 8.9. The action set of Ann is $A = \{\text{Up, Middle, Down}\}$, whereas $B = \{\text{Left, Right}\}$ represents the action set of Bob. Figure 8.6 contains a beliefs diagram of this game. Down is strictly dominated by a convex combination of Up and Middle, which means that Down is an irrational choice. Hence, we cannot find any outgoing arrow for Down.

Table 8.9: Payoff matrix of a 3×2 game.

Ann	Bob	
	Left	Right
Up	(1, 3)	(8, 2)
Middle	(7, 2)	(4, 1)
Down	(2, -1)	(5, 0)




Figure 8.6: A beliefs diagram of the 3×2 game.

By contrast, we can easily find two outgoing arrows each for Up and for Middle. For example, Ann considers Up optimal if she thinks that Bob chooses Left with probability $\frac{1}{4}$ and Right with probability $\frac{3}{4}$ (see Figure 8.2). This is illustrated by the two outgoing arrows from Up to Left and to Right. Note that the subjective probabilities are ignored in the beliefs diagram. In fact, the actual probabilities of Ann's conjecture are not relevant at all in order to judge whether an action is rational or not. In a similar way, we can find two outgoing arrows from Middle to Left and to Right, since Ann considers Middle optimal, e.g., if she thinks that Bob chooses Left with probability $\frac{1}{2}$ and Right with probability $\frac{1}{2}$. Hence, both Up and Middle are rational for Ann.

Further, Left is optimal for Bob, e.g., if he thinks that Ann chooses Up with probability $\frac{1}{2}$ and Middle with probability $\frac{1}{2}$, whereas Right turns out to be optimal, e.g., if Bob thinks that Ann chooses Middle with probability $\frac{1}{4}$ and Down with probability $\frac{3}{4}$. Hence, we can assign Left an outgoing arrow to Up and another one to Middle as well as Right an outgoing arrow to Middle and another one to Down. Note that these considerations still take place in Ann's mind—not in Bob's!

We conclude that both Left and Right are rational for Bob. However, Ann's first- and second-order beliefs, which are depicted in Figure 8.6, are inconsistent with the assumption of common belief. The reason is that she believes that Bob assigns Down a positive probability, but she should expect that he is aware of the fact that Down is irrational for her and thus consider this action null in the second order.

Does Figure 8.6 contain any belief hierarchy at all? The answer is "No" because it does not specify all third-order beliefs of Ann. We can specify the missing beliefs by restarting from the left. More precisely, in this case we would have to start with Down because Ann assigns Down a positive probability in the second order, but Down has no outgoing arrow on the left-hand side of Figure 8.6. For this purpose, let us simply draw a dashed arrow from Down to Right (Perea, 2012, p. 82), which can be seen in Figure 8.7. This arrow makes our sequence of actions periodic and now the beliefs diagram contains a belief hierarchy.




Figure 8.7: Inconsistent belief hierarchy of Ann in the 3×2 game.

However, if Ann is convinced that Bob chooses Right, then choosing Down is sub-optimal. Hence, the given belief hierarchy is inconsistent. In fact, Down can *never* be optimal, since it is always better for Ann to choose either Up or Middle (see Figure 8.2). Thus, Down marks an irrational action and so the given belief hierarchy is not only inconsistent with the assumption of common belief—we even cannot find *any* consistent belief hierarchy in which Ann chooses Down! Remember that this issue arose only because Right points to Down, but the problem is that we cannot find any conjecture for Bob that makes Right an optimal choice if the outgoing arrows point only to Up or to Middle. This is because Left strictly dominates Right if we eliminate Down. Hence, we need an outgoing arrow to Down in order to make Right an optimal choice. However, then we cannot close our cycle of optimal actions, since Down is irrational for Ann.

We conclude that Up and Middle cannot be a reasonable choice for Ann unless we are able to drop some outgoing arrow to Right, which means that Ann must be convinced that Bob chooses Left. Well, in this case Up cannot be optimal for Ann and so her only reasonable action is Middle, which can have only one outgoing arrow to Left. Due to the same reasons, Ann cannot believe that Bob assigns Up a positive probability and so we obtain the belief hierarchy in Figure 8.8. This belief hierarchy constitutes a periodic sequence of optimal choices, i.e., each arrow meets another, and thus it is consistent with the assumption of common belief.




Figure 8.8: Consistent belief hierarchy of Ann in the 3×2 game.

In the same way, we are able to conclude that Bob's only reasonable action is Left and his belief hierarchy can be immediately deduced from Figure 8.8. In fact, the same solution has already been derived in Section 8.1.4.2 by an iterated elimination of strictly dominated actions.

Finally, let us reconsider the 2×2 game from Section 8.1.4.1, whose payoff matrix can be found in Table 8.10. The beliefs diagram in Figure 8.9 demonstrates that all actions in this game are reasonable. Up is always optimal for Ann because Up weakly dominates Down. Hence, we can assign Up two outgoing arrows, i.e., one arrow to Left and one to Right. By contrast, Down, is optimal for Ann only if she is convinced that

Bob chooses Left. For this reason, there can be only one outgoing arrow, which must lead to Left.

Table 8.10: Payoff matrix of a 2×2 game.

Ann	Bob	
	Left	Right
Up	(3, 1)	(1, 0)
Down	(3, 1)	(0, 2)




Figure 8.9: Consistent belief hierarchy of Ann in the 2×2 game.

Further, Bob considers Left optimal, e.g., if he believes that Ann chooses Up with probability $\frac{1}{2}$ and Down with probability $\frac{1}{2}$. The same holds true for Right and so we have two outgoing arrows from Left as well as from Right, which lead us to Up and to Down. Now, each action has an arrow that leads us to an arrow that leads us to another arrow, etc. This means that all actions are reasonable and thus each solution of this game is reasonable, too. The same result can be obtained by verifying that this game contains no *strictly* dominated action.

8.3 Conclusion

Rationalizability shows that a rational solution need not be a Nash equilibrium even if the players have common knowledge. Hence, the transparency of reason is not a logical consequence of common knowledge. To be more precise, the transparency of reason claims that the players are able to deduce the strategy of each other, but common knowledge enables them only to deduce the strategies that cannot be optimal for any other. However, in general, they are not able to deduce the optimal strategies of their adversaries and all the less each other's *actual* strategy. This means that common knowledge is not enough in order to justify the transparency of reason. However, if the basic assumptions of Bernheim and Pearce are satisfied, each solution that is considered rational in our subjectivistic framework must also be rationalizable.

Rationalizability creates a wide range of possible solutions, which seems to be a highly desirable escape from the equilibrium doctrine. Hence, the concept of rationalizability is contrary to refinement. However, the problem is that now we are left with even *more* ambiguity about the actual behavior of rational subjects in a strategic conflict. Rationalizability is not specific enough in order to generate useful predictions about the actual behavior of rational players.

For example, in the anti-coordination, discoordination, and coordination games that we discussed in Section 8.1.3, each single action is rationalizable. Only in the prisoners' dilemma we are able to suggest an unambiguous solution, namely that the players fail to cooperate. However, this is a standard result of traditional game theory and it is just the reason why we call it "dilemma."

In most strategic conflicts we cannot find many actions that are not rationalizable and although there exist some exceptions, e.g., the prisoners' dilemma, I doubt that the common-knowledge assumption is satisfied at all. The assumptions of complete information and common knowledge seem to be very restrictive. I can hardly imagine any real-life situation in which these standard assumptions of traditional game theory are satisfied, and experimental studies show that people do not make use of the transparency of reason in order to deduce the actions of each other.

If we drop the basic assumptions of rationalizability, we can justify actions that are no longer rationalizable in the sense of Bernheim and Pearce but still reasonable in the sense of epistemic game theory. The main difference between reasonability and rationalizability is that the former is based on common belief and allows the actions of the players to be stochastically dependent, whereas the latter requires common knowledge and presumes that the actions are stochastically independent. Reasonability simplifies the overall solution procedure because we can find all reasonable actions by eliminating all strictly dominated actions, iteratively. However, by eliminating only strictly dominated actions we allow more actions to survive. Hence, reasonability is even *more* conservative than rationalizability and thus it is a weaker solution concept.

The set of solutions of a strategic conflict that can be justified without common belief and without specifying the subjective probabilities of the players is usually greater than the set of reasonable solutions. However, the set of *rational* solutions, which consists of all solutions in which the players choose an optimal action given their specific priors, is much smaller in most case studies. The subjectivistic approach typically leads us to a *unique* rational solution. For this purpose, the game theorist has to justify the subjective probabilities (and payoffs) of the players. This can often be done in a relatively simple way by putting oneself in the position of each player. However, in order to accomplish this goal, we must go beyond the formal description of the game and take the specific situations of the players, i.e., the basic storyline, into account.

Epistemic game theory is a fantastic playground. It shows how far we can go with mathematics in order to describe human reasoning. However, I fear that this is much too far for most practical applications. On the one hand, common belief is a very strong

assumption and, on the other hand, it is not strong enough in order to make clear predictions. I think that we should aim for a methodology that (i) enables us to describe the real-life behavior of rational subjects and (ii) comes to unambiguous results, i.e., unique rational solutions. In my opinion, this goal can be accomplished only by subjectivistic game theory, which specifies the subjective probabilities of each player. Nonetheless, epistemic game theory can be considered a special branch of subjectivistic game theory, which is based on the (additional) assumption of common knowledge or, at least, of common belief. Thus, it is characterized by belief hierarchies, which do not appear in the pure subjectivistic framework.

9 Theory of Moves

9.1 Motivation

Nash equilibrium plays a predominant role in game theory. Most contributions that are made in the literature during the 20th century propagate only equilibrium solutions. We have already seen that common knowledge is not convincing enough to support Nash equilibrium. However, we can try to apply alternative arguments, like von Neumann's minimax theorem, the self-enforcement argument, Aumann's (or Harsanyi's) Bayesian arguments, etc., in order to justify the equilibrium doctrine.

Risse (2000) vividly discusses the typical arguments in favor of Nash equilibrium. He comes to the conclusion that "All of these arguments either fail entirely or have a very limited scope." Interestingly, he also alludes to interaction:

"One agent judges deviating reasonable if the other one deviates as well. The other one figures this out. Suppose mutual deviation is profitable for him as well. Since he knows that the first agent would deviate if he did, he deviates. In this way, two agents transparent to each other could abandon an equilibrium. The point is that a Nash Equilibrium, by definition, only discourages uni-lateral deviation."

This is precisely my motivation behind analyzing coherent games. Similarly, Bjerring (1978) observes that

"'knowing' an opponent means that every particular interaction with him should be viewed as one in a series of interactions, and that series itself has an important unity overlooked by traditional game theory."

He further indicates that "knowing an opponent" might mean

"having had past interactions with him, having learned some of his likely ways of behaving in certain circumstances, having built up mutually comfortable topics of conversation, or 'patterns' of interacting, and so on."

In fact, we have already observed the pattern of interaction described by Bjerring in the iterated prisoners' dilemma (see Section 5.5) and in the quasicohherent games that have been discussed in Chapter 6.

Bjerring's arguments are retrospective, i.e., they refer to the history of actions that have been performed by the players in the past. By contrast, Brams (1994, p. 4) mentions that players should "think ahead in deciding whether or not to move." This means that they need to consider "the consequences of a series of moves and countermoves from the resulting outcome." Hence, in contrast to Bjerring (1978), Brams' point of view is prospective. Moreover, Brams (1994, pp. 22–23) elaborates:

<https://doi.org/10.1515/9783110596106-009>

“The standard theory [...] does not raise questions about the rationality of moving or departing from outcomes — at least beyond an immediate departure, à la Nash. [...] The question then becomes whether a player, by departing from an outcome, can do better not just in an immediate or myopic sense but, instead, in an extended or nonmyopic sense.”

Risse’s argument is myopic but allows for *multilateral* deviations, whereas Brams considers alternating moves and countermoves around the payoff matrix of a normal-form game. His theory of moves (TOM) is based on an extensive-form analysis with backward induction. Its goal is to determine where play will end up after starting from any initial state.¹⁸⁷ In the long run, even *unilateral* deviations can lead to nonmyopic equilibria that are not predicted by standard game theory.

To sum up, the aforementioned authors note that the decisions of the players can depend on one another, in which case the rational solution of the game may not be a Nash equilibrium.¹⁸⁸ This is one of the key findings of this work, too.

In this chapter, I discuss the TOM. This was one of my main motivations to analyze strategic conflicts in which the players are able to interact. In fact, we will see that the TOM and the solution concept for coherent games (see Chapter 6) are similar in character. However, I provide only a brief introduction to Brams’ theory of moves. The interested reader should study Brams and Mattli (1993), Brams (1994) as well as Brams (1997) for more details on that topic. Some, more contemporary, contributions that make use of the TOM are, e.g., Brams (2003, 2007, 2011) and Brams (2018).

9.2 Basic Rules of TOM

TOM focuses on 2×2 games and thus we may call the players, as usual, “Ann” and “Bob.” It is assumed that Ann and Bob have an *ordinal* utility function, which means that they are able to assign each outcome of the game a rank number.¹⁸⁹ To be more precise, Brams (1994, p. 1) presumes that the players can *strictly* rank the possible outcomes from worst (“1”) to best (“4”) and calls such a game strict ordinal. Hence, the payoffs shall be considered rank numbers rather than utilities in the sense of expected-utility theory. This approach does not require Savage’s axioms of rational choice and so it may happen that a player has no (unique) subjective probability measure. However, probabilities and mixed strategies play no role at all in the context of TOM.

According to Brams (1994, p. 23), there exist 78 nonisomorphic 2×2 strict ordinal games. There remain 57 “conflict games” after excluding all “cooperative games,” which have a mutually best outcome, i.e., “(4, 4)” (Brams, 1994, p. 1). The conflict games are provided in the appendix of Brams (1994).

¹⁸⁷ I thank Steve Brams very much for clarifying this point in a personal communication.

¹⁸⁸ More precisely, Bjerring (1978) refers to Harsanyi’s (1967–1968) Bayesian Nash equilibrium, whereas Brams (1994) and Risse (2000) focus on the ordinary Nash equilibrium (Nash, 1951).

¹⁸⁹ Here, the term “outcome” is synonymous to “solution” in the subjectivistic framework.

Rules of play (Brams, 1994, p. 24): The game starts in some *initial* state, i.e., an intersection of a row and a column of the 2×2 payoff matrix. Either player can switch, unilaterally, his action and thereby change the state of the game. If Ann starts making a move, the new state will be in the same column as the initial state. By contrast, if Bob starts making a move, the new state will be in the same row as the initial state. After the first move has been made, the other player can proceed further by making the next (unilateral) move, in which case the game, once again, comes into a new state. The alternating moves and countermoves continue until the player whose turn it is to move next decides not to move further. Then the game terminates and the final state represents its outcome.

The rules of play state how the players can make their moves and countermoves, and they precisely define how the game comes to an end, provided that anyone has made some move right from the start. However, they do not say anything about which player moves first and whether or not somebody starts moving at all. I will come back to this question soon. Before that, I should highlight some important aspects of TOM:

- The players move alternately and backtracking is impossible.
- In the initial state, either player can start to move, but in each other state only the player who did not move before is able to move further.
- The players obtain their payoffs only at the end of the game.

We start in some initial state because every game has some history, which has brought the players into the current position (Brams, 1994, p. 26). We will see that the outcome of the game *essentially* depends on its initial state.

The next rules of TOM are concerned with the question of whether or not some player moves from the initial state and which player it is.

Rule of rational termination (Brams, 1994, p. 27): A player moves from the initial state only if the final state is better for him.

Hence, a player does not move from the initial state if his move would lead him to a worse (final) state or if it would return him to the initial state. Note that if the player stops moving further in some state *after* the initial state, stopping means that the game comes to an end. By contrast, stopping in the initial state does not inevitably end the game, since the other player can decide to move instead.

The rule of rational termination states that a player, actually, *does not* move from the initial state if the final state is not better for him. However, it is important to note also that a player would stop moving *if* he had moved from the initial state and were back to square one after all moves and countermoves.¹⁹⁰

190 I thank Steve Brams very much for giving me the opportunity to discuss this question with him.

Brams (1994, p. 28) presumes that the players have complete information. This means that everybody knows the payoff matrix and the rules of the game. In my opinion, the overall logic requires us to assume also that the players have, at least, bounded knowledge, which will become clear after the next rule. Hence, let us assume that the players have *common* knowledge.

Two-sidedness rule (Brams, 1994, p. 28): A player decides whether to move or not from every state by backward induction, taking his own forthcoming moves and the countermoves of his adversary into account. If the final state is better than the initial state for one player only, that player starts moving from the initial state.¹⁹¹

Hence, if Ann and Bob do not consider the final state better than the initial state, the rule of rational termination tells us that they will not move at all and so the initial state of the game immediately turns into the final state. If only one player has an incentive to move, i.e., to generate an outcome that is better for him than the initial state, this one starts moving. However, the two-sidedness rule fails to specify who moves first if both players consider the final state better than the initial state.

Here, I refer only to the *basic* rules of TOM. Brams (1994) extends the rules of TOM by taking order, moving, and threat power into consideration. A player is said to have order power if and only if he can dictate which player moves first from the initial state, provided that it is preferable for both players to move first, which solves the indeterminacy problem discussed in the previous paragraph. The precise meanings of moving and threat power are explained by Brams (1994) in Chapter 4 and Chapter 5. In the subsequent analysis, I ignore any form of power.

9.3 Standard Example

9.3.1 The Solution According to TOM

The overall solution concept can be best illustrated by means of a simple example. Here, I choose the standard example of Brams (1994, p. 19), which is illustrated in Table 9.1. The appendix in Brams (1994) reveals that this is the penultimate game out of all 57 possible conflict games. In a strict ordinal game, we can identify each state with the corresponding tuple of payoffs. For example, the state on the upper right of the payoff matrix is denoted by “(4, 2).”

Suppose that the game has started in the initial state (2, 4) and assume that the alternating moves and countermoves have been made counterclockwise until Ann and

¹⁹¹ Brams (1994, p. 28) writes: “If it is rational for one player to move and the other player not to move from the initial state, then the player who moves takes precedence [...]” Due to the rule of rational termination, it is rational for some player to move only if the final state is better for him.


Table 9.1: Standard example of TOM (Game 56).

Ann	Bob	
	Left	Right
Up	(2, 4)	(4, 2)
Down	(1, 1)	(3, 3)


Bob are in State (4, 2). It is Bob's turn and he is thinking about moving further or not. Bob knows that Ann will stop the game if he moves from right to left, since then the players are back to square one. Thus, he will move left because 4 is greater than 2. This being said, Ann knows that Bob moves left if she moves up in State (3, 3) and so she will *not* move up because otherwise she would end up with 2 instead of 3. Further, in State (1, 1), Bob knows that Ann knows that he moves left in State (4, 2) and so he knows that she will stop in State (3, 3). Hence, he moves right in State (1, 1). Finally, in the initial state, Ann knows that Bob knows that she knows that he moves left in State (4, 2), which means that she knows that he will move right in State (1, 1). For this reason, she will move down in the initial state (2, 4) and, after Bob's countermove from left to right, the game stops in the (final) state (3, 3).

Now, it is time to step into the shoes of Bob. If he starts moving right in the initial state, (2, 4), the possible moves and countermoves are made clockwise. Assume that the players end up in State (1, 1) after their moves and countermoves. Ann knows that the game ends in State (2, 4) if she moves up and so she will move up because 2 is greater than 1. In State (3, 3), Bob knows that Ann moves up in State (1, 1) and so he will move left. Further, in State (4, 2), Ann knows that Bob moves left in State (3, 3), in which case she moves up in State (1, 1). For this reason, she will *not* move down in State (4, 2) because then she would sacrifice her best payoff, 4, in order to obtain only 2. Finally, in the initial state (2, 4), Bob knows that Ann does not move down if he moves right. This means that he holds out in the initial state, which contains the best payoff from his own perspective.¹⁹²

Ann's backward induction can be illustrated like this (Brams, 1994, p. 29):



Here, "A" stands for Ann and "B" stands for Bob. Further, the underline indicates the final state of the game if Ann moves first. Similarly, Bob's backward induction reads (Brams, 1994, p. 30):



¹⁹² Actually, we could have seen immediately, i.e., without backward induction, that Bob will stay in the initial state simply because he cannot do better elsewhere.

The two-sidedness rule implies that Ann will start moving down from the initial state, after which Bob moves right in order to reach the final state (3, 3) (Brams, 1994, p. 31).

Note that (3, 3) is not a (pure-strategy) Nash equilibrium. Brams calls this state a nonmyopic equilibrium (NME). It is *not* considered the initial state of a new instance of the game, in which case both players would have the possibility to move again and thus to reach another state.

Now, we can proceed further in order to analyze what happens if the game starts in the initial state (4, 2). In this case, the backward induction of Ann is:

$$\begin{array}{ccccccc} & A & & B & & A & & B \\ (4, 2) & \nrightarrow & (3, 3) & \rightarrow & (1, 1) & \rightarrow & (2, 4) & \nrightarrow & (4, 2) \end{array}$$

This result is not surprising at all because Ann has already achieved her best payoff in the initial state and so she clearly has no incentive to move anywhere. Moreover, Bob's reasoning goes like this:

$$\begin{array}{ccccccc} & B & & A & & B & & A \\ (4, 2) & \nrightarrow & (2, 4) & \rightarrow & (1, 1) & \rightarrow & (3, 3) & \rightarrow & (4, 2) \end{array}$$

Hence, Bob will not move either and so we have already reached the final state of the game, i.e., (4, 2) (Brams, 1994, p. 31).

If the initial state of the game is (3, 3), Ann's strategic reasoning is:

$$\begin{array}{ccccccc} & A & & B & & A & & B \\ (3, 3) & \nrightarrow & (4, 2) & \rightarrow & (2, 4) & \rightarrow & (1, 1) & \rightarrow & (3, 3) \end{array}$$

Hence, Ann will hold out in the initial state. Further, Bob's backward induction is:

$$\begin{array}{ccccccc} & B & & A & & B & & A \\ (3, 3) & \rightarrow & (1, 1) & \rightarrow & (2, 4) & \nrightarrow & (4, 2) & \nrightarrow & (3, 3) \end{array}$$

Thus, he will move left, after which Ann moves up and the game stops in the final state (2, 4) (Brams, 1994, p. 32). Interestingly, TOM predicts that Ann will not move from the initial state (3, 3) because she cannot improve her situation by moving up, whereas Bob achieves his best payoff, 4, by moving left. However, if Bob moves left, the game ends up with a payoff of 2 for Ann, which is worse than 3. Thus, it would be better for Ann to move up, although then she will come back to square one! Nonetheless, the rule of rational termination forbids Ann to move up in this situation.

Finally, if the game starts in the initial state (1, 1), the backward induction for Ann reads:

$$\begin{array}{ccccccc} & A & & B & & A & & B \\ (1, 1) & \rightarrow & (2, 4) & \nrightarrow & (4, 2) & \nrightarrow & (3, 3) & \nrightarrow & (1, 1) \end{array}$$

Hence, Ann will move up, provided that Bob stays in the initial state. Moreover, Bob's backward induction is:

$$\begin{array}{ccccccc} & B & & A & & B & & A \\ (1, 1) & \rightarrow & (3, 3) & \nrightarrow & (4, 2) & \rightarrow & (2, 4) & \nrightarrow & (1, 1) \end{array}$$

Thus, Bob will move right if Ann holds out in the initial state. The problem is that now *both* players have an incentive to move, but the basic rules of TOM do not say anything about the outcome of the game in such a situation:

- The rule of rational termination just states that a player will move *only* if the final state is better for him but *not if* it is better.
- The two-sidedness rule requires a player to move away from the initial state if he ends up in a better position and the other one does *not* improve by moving away. However, this does not reflect the current situation, since both Ann and Bob can do better by moving away from the initial state.

It would even be better for Ann not to move up if she knows that Bob moves right because then she obtains 3, which is greater than 2. Conversely, it would be better for Bob not to move right, given that he knows that Ann moves up, since then he gets 4 rather than 3. However, without order power, nobody can force the other to move first.

If the initial state is (1, 1), Brams (1994, p. 32) considers the outcome of the game indeterminate, since it can happen that either Ann moves up, in which case the players end up in the state (2, 4), or Bob moves right so that the outcome is (3, 3). Moreover, according to the remarks made by Brams (1994, p. 32), some player *must* move away from the initial state, but we do not know who it is. Thus, both (2, 4) and (3, 3) are considered nonmyopic equilibria (Brams, 1994, p. 33).

We conclude that the nonmyopic equilibria of the game are (2, 4), (4, 2), and (3, 3). However, we can find only one Nash equilibrium: (2, 4). Note that Down is strictly dominated by Up. For this reason, the nonmyopic equilibrium (3, 3) is ruled out, categorically, by traditional game theory.

9.3.2 The Coherent Solution

To the best of my knowledge, TOM is the only systematic attempt to solve strategic conflicts with perfect information—besides the solution concept presented in Chapter 6. Here, “perfect information” is not understood in the sense of von Neumann and Morgenstern (1953), who require only that the players in a dynamic game know the complete history of the game (including the moves of Nature) before making their next moves. Here, perfect information means that the players know (also) the *forthcoming* moves of their opponents. Thus, we should ask ourselves how the standard example of TOM, i.e., Game 56, can be explained if it is treated as a coherent game. Do we come to similar or even to the same conclusions?

The rational plays of Game 56 are given in Figure 9.1. Each rational solution is marked blue. On the lower right we can find also the three nonmyopic equilibria. As we can see, each rational solution of this game is a NME and vice versa. Only the solution (1, 1) cannot be rational. In that case, each player receives his worst payoff and, since the game is strict ordinal, it is always better to deviate.




Figure 9.1: Rational plays of Game 56 and its nonmyopic equilibria (lower right).

Ann receives her best payoff in $(4, 2)$ and so she would like to perform Strategy 2, i.e., to choose Down if and only if Bob chooses Left, in order to force Bob to choose Right. By contrast, Bob receives his best payoff in $(2, 4)$. Thus, he would like to perform Strategy 1, i.e., to choose categorically Left, in order to force Ann to choose Up. It is worth emphasizing that the favorable strategies of Ann and Bob are not their best-response strategies: The best-response strategy of Ann is Strategy 1, whereas the best-response strategy of Bob is Strategy 3.

If Ann and Bob decide to perform their favorable strategies, then the resulting solution, i.e., $(1, 1)$, is irrational (see the top left of Figure 9.2). Hence, this cannot be the play if we presume that both players are rational. The same holds true if both players perform their best-response strategies (see the top right of Figure 9.2).

We conclude that, by choosing his own favorable strategy, each player can force the other *not* to choose his favorable strategy. There is only one play in which no player outdoes the other. It can be found on the bottom of Figure 9.2: Ann performs Strategy 4, i.e., she chooses categorically Down, whereas Bob performs Strategy 3, which means that he goes left if and only if Ann goes up. In this case, the players end up in $(3, 3)$.




Figure 9.2: Irrational plays (top) and reasonable play (bottom) of Game 56.

This is the only reasonable solution of this game, provided that we assume that both players act under equal conditions.¹⁹³

It should be emphasized that finding the reasonable solution of a coherent game is typically a question of refinement and thus requires us to provide a meaningful argument that justifies our conclusion. Here, our conclusion is justified by the idea that $(3, 3)$ is a compromise between $(2, 4)$ and $(4, 2)$, which can be achieved because both players can force the other to act in the desired way. Of course, we could ignore this argument and apply any other refinement procedure as well. However, the *rational* solution of Game 56 must always be $(2, 4)$, $(4, 2)$, or $(3, 3)$, i.e., it must be a NME.

9.4 TOM vs. Coherent Games

The games that are taken into consideration by TOM, which we may call “TOM games” in the subsequent analysis, differ in many aspects from coherent games:

1. The moves and countermoves in a TOM game take place at different time points. A coherent game is instantaneous, i.e., it starts and stops at Time 0.
2. In a TOM game, the players make their choices alternately, i.e., the decision of one player *precludes* the decision of the other. By contrast, in a coherent game, the players make their choices simultaneously and so the decision of one player *includes* the decision of the other.

¹⁹³ Interestingly, this is also the outcome predicted Brams (1994, p. 55) after analyzing the so-called anticipation game that is associated with Game 56.

3. Deviations in a TOM game, i.e., moves and countermoves, can be made only unilaterally. The players in a coherent game can deviate multilaterally. Moreover, the moves and countermoves in a TOM game can actually happen, whereas deviations in a coherent game are only hypothetical.
4. The basic rules of TOM determine the players' behavior, i.e., when they move and when not, right from the start.¹⁹⁴ On the contrary, in a coherent game we do not specify the players' behavior in advance but assume that the players are rational and have perfect information.
5. TOM games require us to fix some initial state, whereas coherent games do not have any starting point.
6. TOM is based on the assumption of complete information and common or, at least, bounded knowledge. By contrast, coherent games do not require us to assume anything but perfect information.
7. In a TOM game, the action set of each player is {Stay, Move}, whereas the typical action sets of a 2×2 coherent game are {Up, Down} and {Left, Right}.
8. Coherent games are not required to be 2×2 .¹⁹⁵ In principle, we can imagine any finite number of players with finite action sets, and in many cases coherent games have a rational solution even if their action sets are infinite.
9. TOM presumes that the players have an ordinal utility function, whereas coherent games have been discussed in the context of subjectivistic game theory, which is based on the assumption that the players are rational in the sense of Savage (1972) and thus possess a cardinal utility function. However, in fact, we need only rank numbers in order to solve a coherent game.¹⁹⁶ Nonetheless, TOM refers to strict ordinal games, but coherent games do not require that the preference orders of the players are strict ordinal, i.e., we may consider any arbitrary number of ties.¹⁹⁷
10. Players in a TOM game are able to stop the game, deliberately, by deciding not to move. In a coherent game, this is impossible—the game just stops automatically after the players have made their choices.

Despite the large number of differences between TOM and coherent games, there is one important aspect in which both games coincide: Each player knows the strategy of the other. Put another way, both players have perfect information.

In the standard example of TOM, i.e., Game 56, the set of nonmyopic equilibria coincides with the set of rational solutions. This is not always the case for the 57 conflict games, which are listed in the appendix of Brams (1994). However, the following theorem implies that an irrational solution cannot be a NME.

194 The fact that there exist some indeterminacies shall be ignored here.

195 Actually, TOM is not restricted to 2×2 games either, which is clarified by Brams (1994) in the introduction, but the basic rules of TOM refer to 2×2 games.

196 Since the players in a coherent game have perfect information, probabilities play no role at all.

197 This enables us to analyze 2×2 games like Matching Pennies and Split or Steal.

Theorem 17 (Nonmyopic equilibrium). *Every TOM game has a NME and each NME is a rational solution.*

Proof. The fact that every TOM game has a NME is trivial (Brams, 1994, p. 33). Let

$$\begin{bmatrix} (u_{11}, v_{11}) & (u_{12}, v_{12}) \\ (u_{21}, v_{21}) & (u_{22}, v_{22}) \end{bmatrix}$$

be the payoff matrix of the TOM game and assume, without loss of generality, that Up/Left is a NME. Suppose that Up/Left is also an irrational solution and thus

- $u_{21} > u_{11}$ and $u_{22} > u_{11}$ or
- $v_{12} > v_{11}$ and $v_{22} > v_{11}$.

The fact that Up/Left is a NME implies that Player 1 or Player 2 will not move away from the upper left when arriving at this state after any number of moves and countermoves. Assume that the initial state that leads us to the NME Up/Left coincides with Up/Left. Then both players do not move away from the initial state although the inequalities mentioned above reveal that it is certainly better for some player to move away. Hence, the initial state cannot coincide with Up/Left, which means that Player 1 decides to move from the lower left to the upper left or Player 2 decides to move from the upper right to the upper left. If Player 1 decides to move from the lower left to the upper left, we have that $u_{11} > u_{21}$, i.e., $v_{12} > v_{11}$ and $v_{22} > v_{11}$, but then it makes no sense for Player 2 to stop moving. By contrast, if Player 2 decides to move from the upper right to the upper left, we have that $v_{11} > v_{12}$, i.e., $u_{21} > u_{11}$ and $u_{22} > u_{11}$, and so it makes no sense for Player 1 to stop moving. Thus, Up/Left cannot be an irrational solution. \square

Table 9.2 reveals that 31 out of all 57 conflict games have one NME, 24 games have two nonmyopic equilibria, and two games have three nonmyopic equilibria. Theorem 17 guarantees that all nonmyopic equilibria are rational solutions. Hence, searching for the nonmyopic equilibria of a conflict game represents a refinement procedure, provided that we consider the game coherent. However, it turns out that most rational solutions are nonmyopic equilibria: Only 15 out of 100 rational solutions fail to be a nonmyopic equilibrium. Table 9.2 does not take any further refinement into account. It contains the numbers of rational (and irrational), not the numbers of reasonable (and unreasonable) solutions of the conflict games.

Table 9.2: Numbers of nonmyopic equilibria (“NME”) and rational solutions (“R”) of all 57 conflict games, where \neg means “not.”

	R	\neg R	Σ		R	\neg R	Σ		R	\neg R	Σ
NME	31	0	31	NME	48	0	48	NME	6	0	6
\neg NME	10	83	93	\neg NME	5	43	48	\neg NME	0	2	2
Σ	41	83	124	Σ	53	43	96	Σ	6	2	8
31 games with 1 NME				24 games with 2 NMEs				2 games with 3 NMEs			

9.5 Conclusion

To the best of my knowledge, besides the solution concept for coherent games that is presented in Chapter 6, only TOM describes a method to solve strategic conflicts in which the players are able to foresee the action of their adversary. This explains why TOM can be considered a refinement procedure for games in which the players have perfect information, although TOM games differ in many aspects from coherent games. In a TOM game, a NME need not be a Nash equilibrium and it can even happen that a Nash equilibrium is not a NME, which is shown in the appendix of Brams (1994). Similar conclusions have been made for coherent games: Every Nash equilibrium represents a rational solution, but a rational solution need not be a Nash equilibrium. If we consider a TOM game coherent, it turns out that each NME is a rational solution, but once again the converse is not true. No matter how we look at it: Focusing on Nash equilibria turns out to be implausible if the players have perfect information.

The basic rules of TOM refer to 2×2 strict ordinal games, and the 57 conflict games, which are presented by Brams (1994), can be understood as political conflicts that have brought the players into some initial state. Now, the players have to decide whether to stay or to move further until the game reaches some final state. TOM explains whether or not some player prefers to leave the initial state and it also predicts the final state, i.e., the resulting outcome, of the game, provided that the players act according to the basic rules of TOM. The solution concept for coherent games is by far less specific than TOM and thus it comes to more general conclusions unless we apply a refinement procedure. This clarifies why the set of rational solutions of a TOM game can be greater than the set of nonmyopic equilibria, provided that we consider the game coherent. However, it turns out that the number of rational solutions of all 57 conflict games that are not a NME is relatively small.

If we believe that the basic rules of TOM are satisfied, it is more appropriate to solve a strict ordinal 2×2 game by calculating its nonmyopic equilibria. Otherwise, we can use the solution concept for coherent games, which has been elaborated in Chapter 6. The quintessence of this chapter is that the presented solution concepts are similar in character. In particular, they do not contradict each other. TOM is just stronger than the solution concept for coherent games without refinement. However, refinement typically enables us to find a unique rational solution, which usually holds true also for TOM if we specify some initial state.

Bibliography

- Adams, D. (1979): *The Hitchhiker's Guide to the Galaxy*, Pan Books.
- Anscombe, F., Aumann, R. (1963): "A definition of subjective probability," *Annals of Mathematical Statistics* **34**, pp. 199–205.
- Arrow, K. (1971): "The theory of risk aversion," in Y. J. Saatio (editor), "Aspects of the Theory of Risk Bearing," Markham, pp. 90–109.
- Aumann, R. (1974): "Subjectivity and correlation in randomized strategies," *Journal of Mathematical Economics* **1**, pp. 67–96.
- Aumann, R. (1976): "Agreeing to disagree," *The Annals of Statistics* **4**, pp. 1236–1239.
- Aumann, R. (1985): "What is game theory trying to accomplish?" in K. Arrow, S. Honkapohja (editors), "Frontiers of Economics," Basil Blackwell, pp. 5–46.
- Aumann, R. (1987): "Correlated equilibrium as an expression of Bayesian rationality," *Econometrica* **55**, pp. 1–18.
- Aumann, R. (1995): "Backward induction and common knowledge of rationality," *Games and Economic Behavior* **8**, pp. 6–19.
- Aumann, R. (1998): "Common priors: a reply to Gul," *Econometrica* **66**, pp. 929–938.
- Aumann, R. (1999): "Interactive epistemology I: knowledge," *International Journal of Game Theory* **28**, pp. 263–300.
- Aumann, R., Dreze, J. (2009): "Assessing strategic risk," *American Economic Journal: Microeconomics* **1**, pp. 1–16.
- Axelrod, R. (1984): *The Evolution of Cooperation*, Basic Books.
- Axelrod, R., Hamilton, W. (1981): "The evolution of cooperation," *Science* **211**, pp. 1390–1396.
- Bacharach, M. (1987): "A theory of rational decision in games," *Erkenntnis* **27**, pp. 17–55.
- Battigalli, P., Bonanno, G. (1999): "Recent results on belief, knowledge and the epistemic foundations of game theory," *Research in Economics* **53**, pp. 149–225.
- Bellman, R. (1957): *Dynamic Programming*, Princeton University Press.
- Bernheim, D. (1984): "Rationalizable strategic behavior," *Econometrica* **52**, pp. 1007–1028.
- Bernstein, D. (2009): *Matrix Mathematics: Theory, Facts, and Formulas*, Princeton University Press, 2nd edition.
- Bicchieri, C., Dalla Chiara, M. (1992): *Knowledge, Belief, and Strategic Interaction*, Cambridge University Press.
- Bjerring, A. (1978): "The 'tracing procedure' and a theory of rational interaction," in C. Hooker, J. Leach, E. McClennen (editors), "Foundations and Applications of Decision Theory, Volume I: Theoretical Foundations," Reidel, pp. 1–54.
- Board, O. (2002): "Knowledge, beliefs, and game-theoretic solution concepts," *Oxford Review of Economic Policy* **18**, pp. 418–432.
- Board, O. (2004): "Dynamic interactive epistemology," *Games and Economic Behavior* **49**, pp. 49–80.
- Bolker, E. (1967): "A simultaneous axiomatization of utility and subjective probability," *Philosophy of Science* **34**, pp. 333–340.
- Bonanno, G. (2015): "Counterfactuals and the prisoner's dilemma," in M. Peterson (editor), "The Prisoner's Dilemma," Cambridge University Press, pp. 133–155.
- Borel, E. (1924): "A propos d'un traité de probabilités," *Revue Philosophique de la France et de l'Étranger* **98**, pp. 321–336.
- Brams, S. (1994): *Theory of Moves*, Cambridge University Press, 3rd edition.
- Brams, S. (1997): "The rationality of surprise: unstable Nash equilibria and the theory of moves," in N. Geva, A. Mintz (editors), "Decisionmaking on War and Peace: The Cognitive Rational

<https://doi.org/10.1515/9783110596106-010>

- Debate," Lynne Rienner Publishers, pp. 103–131.
- Brams, S. (2003): *Biblical Games: Game Theory and the Hebrew Bible*, MIT Press, 2nd edition.
- Brams, S. (2007): *Superior Beings: If They Exist, How Would We Know?*, Springer, 2nd edition.
- Brams, S. (2011): *Game Theory and the Humanities: Bridging Two Worlds*, MIT Press.
- Brams, S. (2018): *Divine Games: Game Theory and the Undecidability of a Superior Being*, MIT Press.
- Brams, S., Maitli, W. (1993): "Theory of moves: overview and examples," *Conflict Management and Peace Science* **12**, pp. 1–39.
- Brandenburger, A. (1992): "Knowledge and equilibrium in games," *Journal of Economic Perspectives* **6**, pp. 83–101.
- Brandenburger, A. (2007): "The power of paradox: some recent developments in interactive epistemology," *International Journal of Game Theory* **35**, pp. 465–492.
- Brandenburger, A. (2014): *The Language of Game Theory: Putting Epistemics into the Mathematics of Games*, World Scientific.
- Brandenburger, A., Dekel, E. (2014): "Rationalizability and correlated equilibria," in A. Brandenburger (editor), "The Language of Game Theory: Putting Epistemics into the Mathematics of Games," Chapter 3, World Scientific, pp. 43–57.
- Brandenburger, A., Friedenberg, A. (2014): "Intrinsic correlation in games," in A. Brandenburger (editor), "The Language of Game Theory: Putting Epistemics into the Mathematics of Games," Chapter 4, World Scientific, pp. 59–111.
- Briggs, R. (2017): "Normative theories of rational choice: expected utility," Technical report, Stanford University.
- Burdzy, K. (2016): *Resonance*, Imperial College Press.
- Camerer, C., Weber, M. (1992): "Recent developments in modeling preferences: uncertainty and ambiguity," *Journal of Risk and Uncertainty* **5**, pp. 325–370.
- Chai, J., Li, C., Wakker, P., Wang, T., Yang, J. (2016): "Reconciling Savage's and Luce's modeling of uncertainty: the best of both worlds," *Journal of Mathematical Psychology* **75**, pp. 10–18.
- Colman, A. (2004): "Reasoning about strategic interaction: solution concepts in game theory," in K. Manktelow, M. Chung (editors), "Psychology of Reasoning: Theoretical and Historical Perspectives," Psychology Press, pp. 287–308.
- Cournot, A. (1897): *Researches into the Mathematical Principles of the Theory of Wealth*, Macmillan.
- Dalla Chiara, M. (1992): "Possible worlds, counterfactuals, and epistemic operators," in C. Bicchieri, M. Dalla Chiara (editors), "Knowledge, Belief, and Strategic Interaction," Cambridge University Press, pp. 155–166.
- de Bruin, B. (2010): *Explaining Games: The Epistemic Programme in Game Theory*, Springer.
- de Finetti, B. (1937): "La prévision: ses lois logiques, ses sources subjectives," *Annales de l'Institut Henri Poincaré* **7**, pp. 1–68.
- Debreu, G. (1952): "A social equilibrium existence theorem," *Proceedings of the National Academy of Sciences of the United States of America* **38**, pp. 886–893.
- Dekel, E., Gul, F. (1997): "Rationality and knowledge in game theory," in D. Kreps, K. Wallis (editors), "Advances in Economics and Econometrics: Theory and Applications (Seventh World Congress)," Cambridge University Press, pp. 87–172.
- Dekel, E., Siniscalchi, M. (2015): "Epistemic game theory," in H. Young, S. Zamir (editors), "Handbook of Game Theory with Economic Applications," Elsevier, pp. 619–702.
- Domotor, Z. (1978): "Axiomatization of Jeffrey utilities," *Synthese* **39**, pp. 165–210.
- Eisenführ, F., Weber, M. (1994): *Rationales Entscheiden*, Springer, 2nd edition.
- Ellsberg, D. (1956): "Theory of the reluctant duelist," *The American Economic Review* **46**, pp. 909–923.
- Ellsberg, D. (1961): "Risk, ambiguity, and the Savage axioms," *Quarterly Journal of Economics* **75**, pp. 643–669.

- Fan, K. (1952): "Fixed-point and minimax theorems in locally convex topological linear spaces," *Proceedings of the National Academy of Sciences of the United States of America* **38**, pp. 121–126.
- Fishburn, P. (1967): "Preference-based definitions of subjective probability," *Annals of Mathematical Statistics* **38**, pp. 1605–1617.
- Fishburn, P. (1969): "A general theory of subjective probabilities and expected utilities," *Annals of Mathematical Statistics* **40**, pp. 1419–1429.
- Fishburn, P. (1970): *Utility Theory for Decision Making*, RAC.
- Fishburn, P. (1981): "Subjective expected utility theory: an overview of normative theories," *Theory and Decision* **13**, pp. 139–199.
- Fishburn, P. (1991): "Nontransitive preferences in decision theory," *Journal of Risk and Uncertainty* **4**, pp. 113–134.
- Geanakoplos, J. (1992): "Common knowledge," *Management Science* **6**, pp. 53–82.
- Gibbard, A., Harper, W. (1981): "Counterfactuals and two kinds of expected utility," in W. Harper, R. Stalnaker, G. Pearce (editors), *Ifs: Conditionals, Belief, Decision, Chance, and Time*, Reidel, pp. 153–190.
- Gibbons, R. (1992): *A Primer in Game Theory*, Prentice Hall.
- Gilboa, I. (1987): "Expected utility with pure subjective non-additive probabilities," *Journal of Mathematical Economics* **16**, pp. 65–88.
- Glicksberg, I. (1952): "A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points," *Proceedings of the American Mathematical Society* **3**, pp. 170–174.
- Govindan, S., Wilson, R. (2008): "Refinements of Nash equilibrium," in S. Durlauf, L. Blume (editors), *The New Palgrave Dictionary of Economics*, Palgrave Macmillan, 2nd edition, pp. 1–14.
- Gul, F. (1998): "A comment on Aumann's Bayesian view," *Econometrica* **66**, pp. 923–927.
- Güth, W., Schmittberger, R., Schwarze, B. (1982): "An experimental analysis of ultimatum bargaining," *Journal of Economic Behavior & Organization* **3**, pp. 367–388.
- Harper, W., Stalnaker, R., Pearce, G. (1981): *Ifs: Conditionals, Belief, Decision, Chance, and Time*, Reidel.
- Harsanyi, J. (1967–1968): "Games with incomplete information played by 'Bayesian' players, I–III," *Management Science* **14**, pp. 159–182, 320–334, 486–502.
- Harsanyi, J., Selten, R. (1988): *A General Theory of Equilibrium Selection in Games*, MIT Press.
- Hotelling, H. (1929): "Stability in competition," *The Economic Journal* **39**, pp. 41–57.
- Jeffrey, R. (1965a): "Ethics and the logic of decision," *The Journal of Philosophy* **62**, pp. 528–539.
- Jeffrey, R. (1965b): *The Logic of Decision*, McGraw-Hill.
- Jeffrey, R. (1981): "The logic of decision defended," *Synthese* **48**, pp. 473–492.
- Jeffrey, R. (1983): *The Logic of Decision*, The University of Chicago Press, 2nd edition.
- Kadane, J., Larkey, P. (1982): "Subjective probability and the theory of games," *Management Science* **28**, pp. 113–120.
- Kahneman, D. (2011): *Thinking, Fast and Slow*, Penguin.
- Kahneman, D., Tversky, A. (1979): "Prospect theory: an analysis of decision under risk," *Econometrica* **47**, pp. 263–291.
- Kakutani, S. (1941): "A generalization of Brouwer's fixed point theorem," *Duke Mathematical Journal* **8**, pp. 457–459.
- Kant, I. (1787): *Critique of Pure Reason*, Cambridge University Press, translated and edited by P. Guyer and A. W. Wood, 1998.
- Karni, E. (2017): "States of Nature and the nature of states," *Economics and Philosophy* **33**, pp. 73–90.
- Knight, F. (1921): *Risk, Uncertainty and Profit*, Hart, Schaffner and Marx.

- Kolmogoroff, A. (1933): *Grundbegriffe der Wahrscheinlichkeitsrechnung*, Springer.
- Kopylov, I. (2007): "Subjective probabilities on 'small' domains," *Journal of Economic Theory* **133**, pp. 236–265.
- Kreps, D., Wilson, R. (1982): "Sequential equilibria," *Econometrica* **50**, pp. 863–894.
- Kydland, F., Prescott, E. (1977): "Rules rather than discretion: the inconsistency of optimal plans," *Journal of Political Economy* **85**, pp. 473–492.
- Lave, L. (1962): "An empirical approach to the prisoners' dilemma game," *Quarterly Journal of Economics* **76**, pp. 424–436.
- Levi, I. (1997): *The Covenant of Reason: Rationality and the Commitments of Thought*, Cambridge University Press.
- Levi, I. (2008): "Feasibility," in C. Bicchieri, M. Dalla Chiara (editors), "Knowledge, Belief, and Strategic Interaction," Cambridge University Press, pp. 1–21.
- Lewis, D. (1969): *Convention: A Philosophical Study*, Harvard University Press.
- Lewis, D. (1973): *Counterfactuals*, Blackwell.
- Lewis, D. (1981): "Probabilities of conditionals and conditional probabilities," in W. Harper, R. Stalnaker, G. Pearce (editors), "Ifs: Conditionals, Belief, Decision, Chance, and Time," Reidel, pp. 129–147.
- Liu, Y., Price, H. (2018): "Ramsey and Joyce on deliberation and prediction," Technical report, PhilSci Archive, <http://philsci-archive.pitt.edu/14972/>.
- Loewenstein, G., Prelec, D. (1992): "Anomalies in intertemporal choice: evidence and an interpretation," *Quarterly Journal of Economics* **107**, pp. 573–597.
- Luce, R. (2000): *Utility of Gains and Losses: Measurement-Theoretical and Experimental Approaches*, Lawrence Erlbaum Publishers.
- Luce, R., Krantz, D. (1971): "Conditional expected utility," *Econometrica* **39**, pp. 253–271.
- Luce, R., Raiffa, H. (1957): *Games and Decisions: Introduction and Critical Survey*, Dover, 3rd edition.
- Machina, M. (2003): "States of the world and the state of decision theory," in D. Meyer (editor), "The Economics of Risk," W. E. Upjohn Institute for Employment Research, pp. 17–49.
- Mengel, F. (2018): "Risk and temptation: a meta-study on prisoner's dilemma games," *Economic Journal* **128**, pp. 3182–3209.
- Morris, S. (1995): "The common prior assumption in economic theory," *Economics and Philosophy* **11**, pp. 227–253.
- Moulin, H. (1988): *Axioms of Cooperative Decision Making*, Cambridge University Press.
- Myerson, R. (1978): "Refinements of the Nash equilibrium concept," *International Journal of Game Theory* **7**, pp. 73–80.
- Myerson, R. (1991): *Game Theory: Analysis of Conflict*, Harvard University Press.
- Nagel, R. (1995): "Unraveling in guessing games: an experimental study," *American Economic Review* **85**, pp. 1313–1326.
- Nash, J. (1950a): "The bargaining problem," *Econometrica* **18**, pp. 155–162.
- Nash, J. (1950b): "Equilibrium points in n -person games," *Proceedings of the National Academy of Sciences of the United States of America* **36**, pp. 48–49.
- Nash, J. (1951): "Non-cooperative games," *Annals of Mathematics* **54**, pp. 286–295.
- Nash, J. (1953): "Two-person cooperative games," *Econometrica* **21**, pp. 128–140.
- Nozick, R. (1969): "Newcomb's problem and two principles of choice," in N. Rescher (editor), "Essays in Honor of Carl G. Hempel: A Tribute on the Occasion of his Sixty-Fifth Birthday," Springer, pp. 114–146.
- Oosterbeek, H., Sloof, R., van de Kuilen, G. (2004): "Cultural differences in ultimatum game experiments: evidence from a meta-analysis," *Experimental Economics* **7**, pp. 171–188.
- Pearce, D. (1984): "Rationalizable strategic behavior and the problem of perfection," *Econometrica* **52**, pp. 1029–1050.

- Perea, A. (2012): *Epistemic Game Theory: Reasoning and Choice*, Cambridge University Press.
- Peters, H. (2015): *Game Theory – A Multi-Leveled Approach*, Springer, 2nd edition.
- Pfanzagl, J. (1967): "Subjective probability derived from the Morgenstern-von Neumann utility concept," in M. Shubik (editor), "Essays in Mathematical Economics," Princeton University Press, pp. 237–251.
- Pfanzagl, J. (1968): *Theory of Measurement*, Wiley.
- Pratt, J. (1964): "Risk aversion in the small and in the large," *Econometrica* **32**, pp. 122–136.
- Pratt, J., Raiffa, H., Schlaifer, R. (1964): "The foundations of decision under uncertainty: an elementary exposition," *Journal of the American Statistical Association* **59**, pp. 353–375.
- Raiffa, H. (1968): *Decision Analysis. Introductory Lectures on Choices under Uncertainty*, Addison-Wesley.
- Ramsey, F. (1926): "Truth and probability," in R. Braithwaite (editor), "The Foundations of Mathematics and other Logical Essays," Harcourt, Brace and Company, pp. 156–198, 1931.
- Rieck, C. (2016): *Spieltheorie: Eine Einführung*, Rieck, C., 15th edition.
- Risse, M. (2000): "What is rational about Nash equilibria?" *Synthese* **124**, pp. 361–384.
- Rodriguez Marine, G. (1995): *The Role of Counterfactuals in the Foundations of Equilibrium Concepts in Game Theory*, Ph.D. thesis, University of California.
- Rubinstein, A. (1991): "Comments on the interpretation of game theory," *Econometrica* **59**, pp. 909–924.
- Samet, D. (2005): "Counterfactuals in Wonderland," *Games and Economic Behavior* **51**, pp. 537–541.
- Savage, L. (1951): "The theory of statistical decision," *Journal of the American Statistical Association* **46**, pp. 55–67.
- Savage, L. (1954): *The Foundations of Statistics*, Wiley.
- Savage, L. (1972): *The Foundations of Statistics*, Dover, 2nd edition.
- Schelling, T. (1980): *The Strategy of Conflict*, Harvard University Press.
- Schmeidler, D. (1989): "Subjective probability and expected utility without additivity," *Econometrica* **57**, pp. 571–587.
- Schwalbe, U., Walker, P. (2001): "Zermelo and the early history of game theory," *Games and Economic Behavior* **34**, pp. 123–137.
- Selten, R. (1965): "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit," *Zeitschrift für die gesamte Staatswissenschaft* **121**, pp. 301–324, 667–689.
- Selten, R. (1973): "A simple model of imperfect competition, where 4 are few and 6 are many," *International Journal of Game Theory* **2**, pp. 141–201.
- Selten, R. (1975): "Re-examination of the perfectness concept for equilibrium points in extensive games," *International Journal of Game Theory* **4**, pp. 25–55.
- Selten, R. (1995): "An axiomatic theory of a risk dominance measure for bipolar games with linear incentives," *Games and Economic Behavior* **8**, pp. 213–263.
- Selten, R. (2001): "Die konzeptionellen Grundlagen der Spieltheorie einst und jetzt," Technical report 2/2001, University of Bonn.
- Selten, R., Nagel, R. (1998): "Das Zahlenwahlspiel – Ergebnisse und Hintergrund," *Spektrum der Wissenschaft* **2**, pp. 16–22.
- Shin, H. (1992): "Counterfactuals and a theory of equilibrium in games," in C. Bicchieri, M. Dalla Chiara (editors), "Knowledge, Belief, and Strategic Interaction," Cambridge University Press, pp. 397–413.
- Slovic, P., Tversky, A. (1974): "Who accepts Savage's axiom?" *Behavioral Science* **19**, pp. 368–373.
- Smith, J. (1982): *Evolution and the Theory of Games*, Cambridge University Press.
- Spence, M. (1973): "Job market signaling," *Quarterly Journal of Economics* **87**, pp. 355–374.
- Spohn, W. (1977): "Where Luce and Krantz do really generalize Savage's decision model," *Erkenntnis* **11**, pp. 113–134.

- Stalnaker, R. (1981a): "Indicative conditionals," in W. Harper, R. Stalnaker, G. Pearce (editors), "Ifs: Conditionals, Belief, Decision, Chance, and Time," Reidel, pp. 193–210.
- Stalnaker, R. (1981b): "A theory of conditionals," in W. Harper, R. Stalnaker, G. Pearce (editors), "Ifs: Conditionals, Belief, Decision, Chance, and Time," Reidel, pp. 41–55.
- Stalnaker, R. (1996): "Knowledge, belief and counterfactual reasoning in games," *Economics and Philosophy* **12**, pp. 133–163.
- Strotz, R. (1955): "Myopia and inconsistency in dynamic utility maximization," *Review of Economic Studies* **23**, pp. 165–180.
- Suppes, P. (1969): "The role of subjective probability and utility in decision-making," in "Studies in the Methodology and Foundations of Science: Selected Papers from 1951 to 1969," Springer, pp. 87–104.
- van den Assem, M., van Dolder, D., Thaler, R. (2012): "Split or steal? Cooperative behavior when the stakes are large," *Management Science* **58**, pp. 2–20.
- von Neumann, J. (1928): "Zur Theorie der Gesellschaftsspiele," *Mathematische Annalen* **100**, pp. 295–320.
- von Neumann, J. (1937): "Über ein ökonomisches Gleichungssystem und eine Verallgemeinerung des Brouwerschen Fixpunktsatzes," in K. Menger (editor), "Ergebnisse eines mathematischen Kolloquiums," Vienna, p. 83.
- von Neumann, J., Morgenstern, O. (1944): *Theory of Games and Economic Behavior*, Princeton University Press.
- von Neumann, J., Morgenstern, O. (1953): *Theory of Games and Economic Behavior*, Princeton University Press, 3rd edition.
- Wakker, P. (1989): *Additive Representations of Preferences: A New Foundation of Decision Analysis*, Kluwer Academic Publishers.
- Wakker, P. (1990): "Under stochastic dominance Choquet-expected utility and anticipated utility are identical," *Theory and Decision* **29**, pp. 119–132.
- Wakker, P. (1993a): "Clarification of some mathematical misunderstandings about Savage's foundations of statistics, 1954," *Mathematical Social Sciences* **25**, pp. 199–202.
- Wakker, P. (1993b): "Unbounded utility for Savage's 'Foundations of Statistics,' and other models," *Mathematics of Operations Research* **18**, pp. 446–485.
- Wakker, P. (2001): "Testing and characterizing properties of nonadditive measures through violations of the sure-thing principle," *Econometrica* **69**, pp. 1039–1059.
- Wakker, P. (2010): *Prospect Theory: For Risk and Ambiguity*, Cambridge University Press.
- Wald, A. (1950): *Statistical Decision Functions*, Wiley.
- Zambrano, E. (2004): "Counterfactual reasoning and common knowledge of rationality in normal form games," *Topics in Theoretical Economics* **4**, Article 8.
- Zermelo, E. (1913): "Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels," in "Proceedings of the Fifth International Congress of Mathematicians," Cambridge University Press, pp. 501–504.

Games

Games against Nature

- Dictator Game 155
- Ellsberg's Paradox 67
- Homeowner's Problem 42
- Horse Race 47
- Meeting 48
- Newcomb's Paradox 58, 65, 99, 190, 233
- Odysseus and the Sirens 81
- Oil-Wildcatting Problem 83
- Omelette Example 4, 29, 42
- Roulette 19, 70
- Two-Urns Problem 63, 66
- Weather Bet 51
- Weekend Problem 74, 81

Strategic games

- Assurance Game *see* Stag Hunt
- Aumann's Example 262
- Battle of the Sexes 137, 209
- Bench Game 225
- Cat-and-Mouse Game 131, 205
- Chess 97, 105, 160
- Chicken Game 124, 202, 242, 280, 291
- Cournot Duopoly 234
- De-Montmort Game 130
- Entrepreneurship 114
- Friend or Foe 255, 259

- Game of Numbers 135
- Game Show 117, 194, 240, 275
- Golden Balls *see* Split or Steal
- Guessing Game 235
- Hawk-Dove Game *see* Chicken Game
- Invade or Retreat 127
- Iterated Prisoners' Dilemma 170, 186
- Mallorca Game 142, 298, 303, 307
- Matching Pennies 129, 203, 292
- Monitoring Game 248
- Nuclear Threat 150, 156
- Odds and Evens 95, 129
- Peace-War Game 180
- Penalty Shoot-Out 133, 158
- Prisoners' Dilemma 144, 150, 170, 211, 293
- Rebel Without a Cause 116
- Reunion Game 134, 292
- Rock-Paper-Scissors 140, 274
- Sheriff's Dilemma 254, 257
- Split or Steal 148, 172, 214
- Stag Hunt 139
- TOM Game ("Game 56") 318
- Trapped in a Cave 199
- Trust Dilemma *see* Stag Hunt
- Ultimatum Game 153, 243
- Win-Win Situation 136, 208

Index

A 220

C 4

S 239

W 31

E 19

EU 22

Γ 232

Ω 3

P 18

P_I 23

\mathcal{F} 4

\mathcal{N} 13

S 4

ω 4

ω_0 4

$<$ 9

$<_I$ 23

\leq 9

\leq_I 23

\sim 9

\sim_I 23

$>$ 9

$>_I$ 23

\geq 9

\geq_I 23

n 91

p 119

p^* 124

q 119

q^* 124

u_I 24

w 31

w_0 31

42 135

a posteriori 23, 41, 80

a priori 23, 41, 80

act *see* Savage act

action 37, 92, 101, 105, 106, 183, 191

action set 42

action space 220

actual world 31

additivity 18

alternative world 33

altruism 94

ambiguity aversion 15, 67, 71, 121

anti-coordination game 124

antisymmetric binary relation 7

asymmetric binary relation 6

asymmetric information 262

awareness 27

Axelrod-Hamilton experiment 171,
179

backward induction 23, 55, 77, 245

bargaining solution 198, 239

Bayes postulate *see* **P9**

Bayes rule 25, 302

Bayes theorem 24

Bayesian game 254

Bayesian Nash equilibrium 254

Bayesian rationality 25, 76, 78, 260

belief 56, 92, 287

belief hierarchy 287, 306, 309

belief system *see* epistemic model

beliefs diagram 307

Bellman postulate *see* **P8**

Bellman rationality 24, 76, 78

best-response curve 120, 192, 278

best-response diagram 105, 120, 124

best-response strategy 199, 201, 208

bimatrix *see* payoff matrix

binary relation 6

binding agreement 94, 212

Boolean algebra 4, 22, 260

bounded knowledge 232

càdlàg 184, 186, 207

càglàd 186

Cerberus 34

chance node 37

Circe 82

coherent game 102, 183, 220, 268

common belief 135, 297, 298

common knowledge 202, 232, 233,
241, 252, 277, 288

common-prior assumption 254, 258,
260, 265

common-state-space assumption 51,
112, 118, 265

communication 60, 203

complete binary relation 7

complete information 155, 190, 202,
232, 252

completely mixed strategy 281

composite act 11

conflict game 316

conjecture 246

connex binary relation 7

consequence 4, 40

consistency 27

constant act 13, 47

constant-sum game *see* zero-sum
game

continuous-time game 116, 150, 186

convention 135

conviction 27, 53, 56, 57, 60, 78, 118,
135, 149, 164, 241

cooperation 109, 147, 170, 174, 200,
212

cooperative game theory 93, 109, 239

coordination game 134, 146

correlated equilibrium 259, 261, 274,
288

counterfactual 32, 35

counterfactual reasoning 29, 54, 59,
67, 70, 74, 93, 94, 183, 244, 266,
267

credible threat 157, 179, 212, 215

critical threshold 124, 243

D1 12

D2 13

D3 13

D4 16

D5 17

decision matrix 5, 40, 44

decision node 36

decision tree 36, 42, 45, 98, 109

defection 140, 170, 174

deliberation crowds out prediction
see Levi's postulate

descriptive decision theory 19, 28

descriptive game theory 92

deterrence 157

diagonal probability 278

diligence 27

discoordination game 128

dominance 62, 66

dominance principle 15, 61, 66, 107,
115, 121, 214

dynamic game 97, 105, 153

effectiveness 27

empty proposition 31

end node 37

epistemic model 305

equilibrium 220, 223, 271

equivalence 9

essential event 13

event 4, 37

event node 35

event tree 35, 42

evidence *see* information

evolutionary game theory 251

- ex ante 41, 266
- ex post 41, 266
- expected utility 20, 22
- extensive form 105, 109
- extensive-form game *see* dynamic game
- extreme pessimism 121

- favorable strategy 199, 201, 203, 213
- finite decision problem 38, 77
- finite game 106, 160, 165, 221, 238, 240, 281, 288, 293, 303
- focal point 135, 136, 211
- forward deduction 77, 245
- free will *see* freedom of choice
- freedom of choice 219, 233
- full scenario matrix *see* scenario matrix

- game 91
- game against Nature 91, 109
- game tree 97, 100, 103, 105, 109, 243
- Gordan's theorem 303

- Hades 34
- Harsanyi doctrine *see* common-prior assumption
- Heracles 34, 36
- heterogeneous scenario 43
- hierarchy of rationality 25
- homo economicus 156, 171
- homogeneous scenario 43

- implication *see* material conditional
- indicative conditional 32, 36
- information 23, 24, 60
- information function 217
- information graph 219, 226
- interaction 92, 102, 104, 109, 183, 196, 212
- irreflexive binary relation 6
- irregular decision problem 78

- isomorphic games 123, 129, 198, 225

- knowledge 56, 93, 186, 245

- Laplace experiment 50, 86
- law of large numbers 87, 130, 251
- Levi's postulate 49–52, 54, 55, 59, 113, 118, 233, 247
- losing position 162

- majorant game 270
- material conditional 30
- max-min inequality 268
- maximin rule 52, 54
- minimax equality 268, 273
- minimax point *see* minimax solution
- minimax solution 270, 275
- minimax theorem 142, 268
- minimax-regret rule 53
- minorant game 270
- mixed strategy 54, 142, 238, 247, 259, 262
- monotonic conditional 35
- moving power 318
- multiple-solutions problem 80, 103, 121, 141, 241, 245, 247, 275
- Murphy's law 52, 269
- mutual knowledge 232

- Nash equilibrium 124, 130, 164, 198, 201, 215, 238, 288, 291, 321
- negatively transitive binary relation 7
- negligible event *see* null event
- NME *see* nonmyopic equilibrium
- noncooperative game theory 93, 109, 238
- nonmyopic equilibrium 320, 321
- nonstrict order 8
- normal form 105, 109
- normal-form game *see* static game
- normative decision theory 10, 15, 29
- null event 13

objective probability 19, 27, 28, 67,
130, 248, 251, 260, 265
objectivity 27, 254
Odysseus 82
omniscient being 58
one-shot game 95, 100
optimal action 303
order 8
order power 318, 321

P1 10

P2 11

P3 13

P4 15

P5 16

P6 16

P7 17

P8 23

P9 24

parallel game *see* static game

partial order 8

partition 12

payoff efficiency 198, 208, 213, 276,
281

payoff matrix 95, 105, 109, 116

perfect information 112, 153, 190, 202,
219, 252, 262

perfect Nash equilibrium 281

perturbed game 281

pessimism 53, 74, 121, 258

planning horizon 177

play 192, 199

population game 257

positive decision theory 29

posterior 23, 56, 79, 253, 254

power 125, 138, 205, 228, 318

prediction crowds out deliberation
see Levi's postulate

preference relation 9

prescriptive decision theory 28

prescriptive game theory 92

prior 22, 25, 56, 79, 254, 260, 265, 284,
287

private information set 22

probability measure 18

proper game 269

proposition 31

pure strategy 238

qualitative probability 16

quantitative probability *see*
probability measure

quasicoherent game 184

quasiconcave 223

randomization 53, 130, 142, 243, 249,
259, 262, 274, 280, 296

rational action 303

rational play 194

rational solution 92, 118, 119, 124,
129, 133, 194, 220, 246

rationalizability 288

rationalizable solution 289

rationalizable strategy 289

reaction 92, 183

reasonable action 288, 304, 306

reasonable play 199

reasonable solution 198, 305, 307

reduced decision matrix 52, 63, 66

reduced scenario matrix 43, 63

refinement 197, 275

reflexive binary relation 6

reflexive conditional 35, 233

regret 53

representation theorem 3, 19, 22, 28,
75, 225

response curve 104, 188

response diagram 103, 191, 207

response function 217

response graph 217, 223, 226

restricted Savage act 11

risk 67

- risk dominance 277, 279
- root 103
- Savage act 4
- Savage rationality 17, 76
- scenario 38, 42, 46, 75
- scenario matrix 42
- scenario set 42, 47
- scepticism 140
- self-commitment 83
- self-enforcement 199, 215, 247
- self-fulfilling prophecy 116, 241
- sequential decision problem 38, 75
- sequential game *see* dynamic game
- serial game *see* dynamic game
- set of possible worlds 31
- σ -additivity 18
- σ -algebra 4, 6
- signaling 60
- simple decision problem 38
- simple function 20
- simple Savage act 20
- simultaneous game *see* static game
- Sirens 81
- solution 91
- stability 280
- stability index 278
- stamina 173
- state of the world 3, 6, 49
- state space 3, 51, 55, 267
- static game 94, 105, 111, 269
- stochastic independence 238, 301
- strategic form *see* normal form
- strategic game 91, 109
- strategic independence 111, 116, 231, 239, 263, 266, 301, 303
- strategy 38, 42, 75, 101, 105, 106, 191
- strict dominance 108, 110, 234, 236
- strict Nash equilibrium 239, 279
- strict order 8
- strict ordinal game 316
- strict partial order 8
- strict total order 8
- strict weak order 9
- subgame-perfect Nash equilibrium 244, 281
- subjective probability 18–20, 22, 67
- substantive conditional 31, 34, 233
- superdominance 66
- sure-thing principle 12–15, 61, 69
- surprise 23, 79, 175, 245
- symmetric binary relation 6
- tat for tit 174
- The Answer to the Ultimate Question of Life, the Universe, and Everything *see* “42”
- theory of moves 316
- thinking fast 134
- thinking slowly 134
- threat power 318
- time-inconsistency paradox 80
- tit for tat 145, 157, 172, 174, 201, 203, 212, 215
- TOM *see* theory of moves
- total order 8
- transitive binary relation 7
- transitive conditional 35
- transparency of reason 231, 249, 252, 295
- trembling-hand perfect Nash equilibrium *see* perfect Nash equilibrium
- trivial information 22
- trivial information set 23
- true state of the world 4
- trust 140, 150, 263
- type 253, 306
- unbeatable position 162
- unbounded knowledge *see* common knowledge
- uncertainty 67
- util 132, 149

utility function 19, 22, 25

vacuous action 41

value 270, 273

weak dominance 107, 110, 278, 281,
294

weak order 9, 10, 12

well-defined decision problem 5

win-win situation 136, 147

winning position 161

Wojtyla 33, 35

world 3

Zermelo's Chess theorem 161

zero-sum game 140, 142, 225, 268