

Game Theory



Giacomo Bonanno

GAME THEORY. 2nd Edition

Copyright © 2015, 2018 Giacomo Bonanno, all rights reserved

You are free to redistribute this book in pdf format. If you make use of any part of this book you must give appropriate credit to the author. You may not remix, transform, or build upon the material without permission from the author. You may not use the material for commercial purposes.

This book is also available in print in two volumes:

Volume 1: Basic Concepts ISBN-13: 978-1983604638 ISBN-10: 1983604631

Volume 2: Advanced Topics ISBN-13: 978-1985862517 ISBN-10: 1985862514

Kindle Direct Publishing (formerly CreateSpace Independent Publishing)

For details see: http://faculty.econ.ucdavis.edu/faculty/bonanno/GT_Book.html

Giacomo Bonanno is Professor of Economics at the
University of California, Davis

<http://faculty.econ.ucdavis.edu/faculty/bonanno/>

Preface to the first edition, 2015

After teaching game theory (at both the undergraduate and graduate level) at the University of California, Davis for 25 years, I decided to organize all my teaching material in a textbook. There are many excellent textbooks in game theory and there is hardly any need for a new one. However, there are two distinguishing features of this textbook: (1) it is open access and thus free,¹ and (2) it contains an unusually large number of exercises with complete and detailed answers.

I tried to write the book in such a way that it would be accessible to anybody with minimum knowledge of mathematics (high-school level algebra and some elementary notions of probability) and no prior knowledge of game theory. However, the book is intended to be rigorous and it includes several proofs. I believe it is appropriate for an advanced undergraduate class in game theory and also for a first-year graduate-level class.

I expect that there will be some typos and (hopefully minor) mistakes. If you come across any typos or mistakes, I would be grateful if you could inform me: I can be reached at gfbonanno@ucdavis.edu. I will maintain an updated version of the book on my web page at

<http://www.econ.ucdavis.edu/faculty/bonanno/>

I also intend to add, some time in the future, a further collection of exercises and exam questions with detailed answers. Details will appear on my web page.

I am very grateful to Elise Tidrick for meticulously going through each chapter of the book and for suggesting numerous improvements. Her insightful and constructive comments have considerably enhanced this book.

I would also like to thank Nicholas Bowden, Lester Lusher, Burkhard Schipper, Matthieu Stigler, Sukjoon Lee, Minfei Xu and Pedro Paulo Funari for pointing out typos.

¹There may be several other free textbooks on game theory available. The only one I am aware of is the excellent book by Ariel Rubinstein and Martin Osborne, MIT Press, 1994, which can be downloaded for free from Ariel Rubinstein's web page: <http://arielrubinstein.tau.ac.il>. In my experience this book is too advanced for an undergraduate class in game theory.

Preface to the second edition, 2018

This second edition introduces

- ◇ substantial reformatting (having translated into \LaTeX the previous file written in Microsoft Word),
- ◇ the addition of 15 new exercises (bringing the total to 180),
- ◇ improved exposition, with additional text and examples, and enhanced topics,
- ◇ the availability of a printed version of the book, split into two volumes (the first covering the basic concepts and the second dealing with advanced topics).

I am indebted to Elise Tidrick for the enormous amount of work she put into helping me with this new edition. Not only did she design the covers of the two volumes with original artwork, but she spent hours helping me translate the original file into \LaTeX code, taking care of reformatting the figures and nudging and instructing me on how to modify the \LaTeX output so as to ensure a smooth flow in the reading of the book.

I would like to thank Dr. Chula Kooanantkul for pointing out several typos and Mathias Legrand for making the latex template used for this book available for free (the template was downloaded from <http://www.latextemplates.com/template/the-legrand-orange-book>).

Contents

1	Introduction	11
----------	---------------------------	-----------

I **Games with Ordinal Payoffs**

2	Ordinal Games in Strategic Form	17
2.1	Game frames and games	17
2.2	Strict and weak dominance	24
2.3	Second-price auction	29
2.4	The pivotal mechanism	32
2.5	Iterated deletion procedures	35
2.5.1	IDSDS	35
2.5.2	IDWDS	37
2.6	Nash equilibrium	39
2.7	Games with infinite strategy sets	44
2.8	Proofs of theorems	46
2.9	Exercises	50
2.9.1	Exercises for Section 2.1: Game frames and games	50
2.9.2	Exercises for Section 2.2: Strict/weak dominance	51
2.9.3	Exercises for Section 2.3: Second price auction	52
2.9.4	Exercises for Section 2.4: The pivotal mechanism	54
2.9.5	Exercises for Section 2.5: Iterated deletion procedures	55
2.9.6	Exercises for Section 2.6: Nash equilibrium	57

2.9.7	Exercises for Section 2.7: Games with infinite strategy sets	57
2.10	Solutions to exercises	59
3	Perfect-information Games	75
3.1	Trees, frames and games	75
3.2	Backward induction	80
3.3	Strategies in perfect-information games	83
3.4	Relationship between backward induction and other solutions	86
3.5	Perfect-information games with two players	90
3.6	Exercises	94
3.6.1	Exercises for Section 3.1: Trees, frames and games	94
3.6.2	Exercises for Section 3.2: Backward induction	95
3.6.3	Exercises for Section 3.3: Strategies in perfect-information games	96
3.6.4	Exercises for Section 3.5: Two-player games	98
3.7	Solutions to exercises	101
4	General Dynamic Games	117
4.1	Imperfect Information	117
4.2	Strategies	124
4.3	Subgames	126
4.4	Subgame-perfect equilibrium	128
4.5	Games with chance moves	134
4.6	Exercises	140
4.6.1	Exercises for Section 4.1: Imperfect information	140
4.6.2	Exercises for Section 4.2: Strategies	141
4.6.3	Exercises for Section 4.3: Subgames	141
4.6.4	Exercises for Section 4.4: Subgame-perfect equilibrium	142
4.6.5	Exercises for Section 4.5: Games with chance moves	144
4.7	Solutions to exercises	148

II

Games with Cardinal Payoffs

5	Expected Utility Theory	169
5.1	Money lotteries and attitudes to risk	169
5.2	Expected utility: theorems	171
5.3	Expected utility: the axioms	177
5.4	Exercises	184
5.4.1	Exercises for Section 5.1: Money lotteries and attitudes to risk	184
5.4.2	Exercises for Section 5.2: Expected utility theory	185
5.4.3	Exercises for Section 5.3: Expected utility axioms	186

5.5	Solutions to Exercises	187
6	Strategic-form Games	193
6.1	Strategic-form games with cardinal payoffs	193
6.2	Mixed strategies	196
6.3	Computing the mixed-strategy Nash equilibria	201
6.4	Strict dominance and rationalizability	206
6.5	Exercises	210
6.5.1	Exercises for Section 6.1: Strategic-form games with cardinal payoffs	210
6.5.2	Exercises for Section 6.2: Mixed strategies	211
6.5.3	Exercises for Section 6.3: Computing the mixed-strategy Nash equilibria	212
6.5.4	Exercises for Section 6.4: Strict dominance and rationalizability	214
6.6	Solutions to exercises	216
7	Extensive-form Games	227
7.1	Behavioral strategies in dynamic games	227
7.2	Subgame-perfect equilibrium revisited	237
7.3	Problems with the notion of subgame-perfect equilibrium	240
7.4	Exercises	242
7.4.1	Exercises for section 7.1: Behavioral strategies in dynamic games	242
7.4.2	Exercises for section 7.2: Subgame-perfect equilibrium revisited	244
7.5	Solutions to exercises	249

III Knowledge, Common Knowledge, Beliefs

8	Common Knowledge	261
8.1	Individual knowledge	261
8.2	Interactive knowledge	264
8.3	Common knowledge	271
8.4	Exercises	276
8.4.1	Exercises for Section 8.1: Individual knowledge	276
8.4.2	Exercises for Section 8.2: Interactive knowledge	279
8.4.3	Exercises for Section 8.3: Common knowledge	281
8.5	Solutions to Exercises	284
9	Adding Beliefs to Knowledge	295
9.1	Sets and probability: brief review	295
9.1.1	Sets	295
9.1.2	Probability	297

9.2	Probabilistic beliefs	298
9.3	Conditional probability and Bayes' rule	300
9.3.1	Conditional probability	300
9.4	Changing beliefs in response to new information	303
9.4.1	Belief updating	303
9.4.2	Belief revision	305
9.5	Harsanyi consistency of beliefs or like-mindedness	308
9.6	Agreeing to disagree	314
9.7	Proof of the Agreement Theorem	322
9.8	Exercises	323
9.8.1	Exercises for Section 9.1: Sets and probability	323
9.8.2	Exercises for Section 9.2: Probabilistic beliefs	324
9.8.3	Exercises for Section 9.3: Conditional probability and Bayes' rule	324
9.8.4	Exercises for Section 9.4: Changing beliefs in response to new information 327	
9.8.5	Exercises for Section 9.5: Harsanyi consistency of beliefs or like-mindedness 330	
9.8.6	Exercises for Section 9.6: Agreeing to disagree	332
9.9	Solutions to Exercises	333
10	Rationality	347
10.1	Models of strategic-form games	347
10.2	Common knowledge of rationality in strategic-form games	352
10.3	Common knowledge of rationality in extensive-form games	354
10.4	Proofs of Theorems	356
10.5	Exercises	358
10.5.1	Exercises for Section 10.1: Model of strategic-form games	358
10.5.2	Exercises for Section 10.2: Common knowledge of rationality in strategic-form games	360
10.5.3	Exercises for Section 10.3: Common knowledge of rationality in extensive-form games	361
10.6	Solutions to Exercises	362

IV

Refinements of Subgame Perfection

11	Weak Sequential Equilibrium	369
11.1	Assessments and sequential rationality	369
11.2	Bayesian updating at reached information sets	376
11.3	A first attempt: Weak sequential equilibrium	378
11.4	Exercises	387
11.4.1	Exercises for Section 11.1: Assessments and sequential rationality	387

11.4.2	Exercises for Section 11.2: Bayesian updating at reached information sets	389
11.4.3	Exercises for Section 11.3: Weak sequential equilibrium	390
11.5	Solutions to Exercises	393
12	Sequential Equilibrium	403
12.1	Consistent assessments	403
12.2	Sequential equilibrium	408
12.3	Is ‘consistency’ a good notion?	411
12.4	Exercises	413
12.4.1	Exercises for Section 12.1: Consistent assessments	413
12.4.2	Exercises for Section 12.2: Sequential equilibrium	414
12.5	Solutions to Exercises	418
13	Perfect Bayesian Equilibrium	429
13.1	Belief revision and AGM consistency	429
13.2	Bayesian consistency	433
13.3	Perfect Bayesian equilibrium	436
13.4	Adding independence	440
13.4.1	Weak independence	440
13.4.2	Strong independence	444
13.5	Characterization of SE in terms of PBE	446
13.6	History-based definition of extensive-form game	449
13.7	Proofs	452
13.8	Exercises	454
13.8.1	Exercises for Section 13.1: Belief revision and AGM consistency	454
13.8.2	Exercises for Section 13.2: Bayesian consistency	455
13.8.3	Exercises for Section 13.3: Perfect Bayesian equilibrium	457
13.8.4	Exercises for Section 13.4: Adding independence	458
13.8.5	Exercises for Section 13.5: Characterization of sequential equilibrium in terms of PBE	459
13.9	Solutions to Exercises	463

V

Incomplete Information

14	Static Games	483
14.1	Interactive situations with incomplete information	483
14.2	One-sided complete information	485
14.3	Two-sided incomplete information	492
14.4	Multi-sided incomplete information	496

14.5 Exercises	499
14.5.1 Exercises for Section 14.2: One-sided incomplete information.	499
14.5.2 Exercises for Section 14.3: Two-sided incomplete information.	501
14.5.3 Exercises for Section 14.4: Multi-sided incomplete information.	504
14.6 Solutions to Exercises	508
15 Dynamic Games	521
15.1 One-sided incomplete information	521
15.2 Multi-sided incomplete information	538
15.3 Exercises	544
15.3.1 Exercises for Section 15.1: One-sided incomplete information	544
15.3.2 Exercises for Section 15.2: Multi-sided incomplete information	548
15.4 Solutions to Exercises	551
16 The Type-Space Approach	567
16.1 Types of players	567
16.2 Types that know their payoffs	568
16.3 The general case	571
16.4 Exercises	574
16.4.1 Exercises for Section 16.2: Types that know their own payoffs	574
16.4.2 Exercises for Section 16.3: The general case	576
16.5 Solutions to Exercises	578
Index	589

1. Introduction

The discipline of game theory was pioneered in the early 20th century by mathematicians Ernst Zermelo (1913) and John von Neumann (1928). The breakthrough came with John von Neumann and Oskar Morgenstern's book, *Theory of Games and Economic Behavior*, published in 1944. This was followed by important work by John Nash (1950-51) and Lloyd Shapley (1953). Game theory had a major influence on the development of several branches of economics (industrial organization, international trade, labor economics, macroeconomics, etc.). Over time the impact of game theory extended to other branches of the social sciences (political science, international relations, philosophy, sociology, anthropology, etc.) as well as to fields outside the social sciences, such as biology, computer science, logic, etc. In 1994 the Nobel Memorial prize in economics was given to three game theorists, John Nash, John Harsanyi and Reinhard Selten, for their theoretical work in game theory which was very influential in economics. At the same time, the US Federal Communications Commission was using game theory to help it design a \$7-billion auction of the radio spectrum for personal communication services (naturally, the bidders used game theory too!). The Nobel Memorial prize in economics was awarded to game theorists three more times: in 2005 to Robert Aumann and Thomas Schelling, in 2007 to Leonid Hurwicz, Eric Maskin and Roger Myerson and in 2012 to Lloyd Shapley and Alvin Roth.

Game theory provides a formal language for the representation and analysis of *interactive situations*, that is, situations where several “entities”, called *players*, take actions that affect each other. The nature of the players varies depending on the context in which the game theoretic language is invoked: in evolutionary biology (see, for example, John Maynard Smith, 1982) players are non-thinking living organisms;¹ in computer science

¹Evolutionary game theory has been applied not only to the analysis of animal and insect behavior but also to studying the “most successful strategies” for tumor and cancer cells (see, for example, Gerstung *et al.*, 2011).

(see, for example, Shoham-Leyton-Brown, 2008) players are artificial agents; in behavioral game theory (see, for example, Camerer, 2003) players are “ordinary” human beings, etc. Traditionally, however, game theory has focused on interaction among intelligent, sophisticated and rational individuals. For example, Robert Aumann describes game theory as follows:

“Briefly put, game and economic theory are concerned with the interactive behavior of *Homo rationalis* – rational man. *Homo rationalis* is the species that always acts both purposefully and logically, has well-defined goals, is motivated solely by the desire to approach these goals as closely as possible, and has the calculating ability required to do so.” (Aumann, 1985, p. 35.)

This book is concerned with the traditional interpretation of game theory.

Game theory is divided into two main branches. The first is *cooperative game theory*, which assumes that the players can communicate, form coalitions and sign binding agreements. Cooperative game theory has been used, for example, to analyze voting behavior and other issues in political science and related fields.

We will deal exclusively with the other main branch, namely *non-cooperative game theory*. Non-cooperative game theory models situations where the players are either unable to communicate or are able to communicate but cannot sign binding contracts. An example of the latter situation is the interaction among firms in an industry in an environment where antitrust laws make it illegal for firms to reach agreements concerning prices or production quotas or other forms of collusive behavior.

The book is divided into five parts. The printed version of the book is split into two volumes. **Volume 1** covers the basic concepts and encompasses Chapters 1-7 (Parts I and II), while **Volume 2** is devoted to advanced topics, encompassing Chapters 8-16 (Parts III to V).

Part I deals with games with *ordinal* payoffs, that is, with games where the players’ preferences over the possible outcomes are only specified in terms of an ordinal ranking (outcome o is better than outcome o' or o is just as good as o'). Chapter 2 covers strategic-form games, Chapter 3 deals with dynamic games with perfect information and Chapter 4 with general dynamic games with (possibly) imperfect information.

Part II is devoted to games with *cardinal* payoffs, that is, with games where the players’ preferences extend to uncertain prospects or lotteries: players are assumed to have a consistent ranking of the set of lotteries over basic outcomes. Chapter 5 reviews the theory of expected utility, Chapter 6 discusses the notion of mixed strategy in strategic-form games and of mixed-strategy Nash equilibrium, while Chapter 7 deals with mixed strategies in dynamic games.

Parts III, IV and V cover a number of advanced topics.

Part III deals with the notions of knowledge, common knowledge and belief. Chapter 8 explains how to model what an individual knows and what she is uncertain about and how to extend the analysis to the interactive knowledge of several individuals (e.g. what Individual 1 knows about what Individual 2 knows about some facts or about the state of knowledge of Individual 1). The chapter ends with the notion of common knowledge. Chapter 9 adds probabilistic beliefs to the knowledge structures of the previous chapter

and discusses the notions of Bayesian updating, belief revision, like-mindedness and the possibility of “agreeing to disagree”. Chapter 10 uses the interactive knowledge-belief structures of the previous two chapters to model the players’ state of mind in a possible play of a given game and studies the implications of common knowledge of rationality in strategic-form games.

Part IV focuses on dynamic (or extensive-form) games and on the issue of how to refine the notion of subgame-perfect equilibrium (which was introduced in Chapters 4 and 7). Chapter 11 introduces a simple notion, called weak sequential equilibrium, which achieves some desirable goals (such as the elimination of strictly dominated choices) but fails to provide a refinement of subgame-perfect equilibrium. Chapter 12 explains the more complex notion of sequential equilibrium, which is extensively used in applications of game theory. That notion, however, leaves much to be desired from a practical point of view (it is typically hard to show that an equilibrium is indeed a sequential equilibrium) and also from a conceptual point of view (it appeals to a topological condition, whose interpretation is not clear). Chapter 13 introduces an intermediate notion, called perfect Bayesian equilibrium, whose conceptual justification is anchored in the so called AGM theory of belief revision, extensively studied in philosophy and computer science, which was pioneered by Carlos Alchourrón (a legal scholar), Peter Gärdenfors (a philosopher) and David Makinson (a logician) in 1985. In Chapter 13 we also provide an alternative characterization of sequential equilibrium based on the notion of perfect Bayesian equilibrium, which is free of topological conditions.

Part V deals with the so-called “theory of games of incomplete information”, which was pioneered by John Harsanyi (1967-68). This theory is usually explained using the so-called “type-space” approach suggested by Harsanyi. However, we follow a different approach: the so-called “state-space” approach, which makes use of the interactive knowledge-belief structures developed in Part III. We find this approach both simpler and more elegant. For completeness, in Chapter 16 we explain the commonly used type-based structures and show how to convert a state-space structure into a type-space structure and *vice versa* (the two approaches are equivalent). Chapter 14 deals with situations of incomplete information that involve static (or strategic-form) games, while Chapter 15 deals with situations of incomplete information that involve dynamic (or extensive-form) games.

At the end of each section of each chapter the reader is invited to try the exercises for that section. All the exercises are collected in the penultimate section of the chapter, followed by a section containing complete and detailed answers for each exercise. For each chapter, the set of exercises culminates in a “challenging question”, which is more difficult and more time consuming than the other exercises. In game theory, as in mathematics in general, it is essential to test one’s understanding of the material by attempting to solve exercises and problems. The reader is encouraged to attempt solving exercises after the introduction of every new concept.

The spacing in this book does not necessarily follow conventional formatting standards. Rather, it is the editor’s intention that each step is made plain in order for the student to easily follow along and quickly discover where he/she may grapple with a complete understanding of the material.



Games with Ordinal Payoffs

2 Ordinal Games in Strategic Form 17

- 2.1 Game frames and games
- 2.2 Strict and weak dominance
- 2.3 Second-price auction
- 2.4 The pivotal mechanism
- 2.5 Iterated deletion procedures
- 2.6 Nash equilibrium
- 2.7 Games with infinite strategy sets
- 2.8 Proofs of theorems
- 2.9 Exercises
- 2.10 Solutions to exercises

3 Perfect-information Games 75

- 3.1 Trees, frames and games
- 3.2 Backward induction
- 3.3 Strategies in perfect-information games
- 3.4 Relationship between backward induction and other solutions
- 3.5 Perfect-information games with two players
- 3.6 Exercises
- 3.7 Solutions to exercises

4 General Dynamic Games 117

- 4.1 Imperfect Information
- 4.2 Strategies
- 4.3 Subgames
- 4.4 Subgame-perfect equilibrium
- 4.5 Games with chance moves
- 4.6 Exercises
- 4.7 Solutions to exercises

2. Ordinal Games in Strategic Form

2.1 Game frames and games

Game theory deals with interactive situations where two or more individuals, called players, make decisions that jointly determine the final outcome. To see an example point your browser to the following video:¹

<https://www.youtube.com/watch?v=tBtr8-VMj0E>.

In this video each of two players, Sarah and Steve, has to pick one of two balls: inside one ball appears the word ‘split’ and inside the other the word ‘steal’ (each player is first asked to secretly check which of the two balls in front of him/her is the split ball and which is the steal ball). They make their decisions simultaneously. The possible outcomes are shown in Figure 2.1, where each row is labeled with a possible choice for Sarah and each column with a possible choice for Steven. Each cell in the table thus corresponds to a possible pair of choices and the resulting outcome is written inside the cell.

		Steven	
		Split	Steal
Sarah	Split	Sarah gets \$50,000 Steven gets \$50,000	Sarah gets nothing Steven gets \$100,000
	Steal	Sarah gets \$100,000 Steven gets nothing	Sarah gets nothing Steven gets nothing

Figure 2.1: The *Golden Balls* “game”

What should a rational player do in such a situation? It is tempting to reason as follows.

¹The video shows an excerpt from *Golden Balls*, a British daytime TV game show. If you search for ‘Split or Steal’ on youtube.com you will find several instances of this game.

Let us focus on Sarah’s decision problem. She realizes that her decision alone is not sufficient to determine the outcome; she has no control over what Steven will choose to do. However, she can envision two scenarios: one where Steven chooses *Steal* and the other where he chooses *Split*.

- If Steven decides to *Steal*, then it does not matter what Sarah does, because she ends up with nothing, no matter what she chooses.
- If Steven picks *Split*, then Sarah will get either \$50,000 (if she also picks *Split*) or \$100,000 (if she picks *Steal*).

Thus Sarah should choose *Steal*.

The above argument, however, is not valid because it is based on an implicit and unwarranted assumption about how Sarah ranks the outcomes; namely, it assumes that Sarah is selfish and greedy, which may or may not be true. Let us denote the outcomes as follows:

- o_1 : Sarah gets \$50,000 and Steven gets \$50,000.
- o_2 : Sarah gets nothing and Steven gets \$100,000.
- o_3 : Sarah gets \$100,000 and Steven gets nothing.
- o_4 : Sarah gets nothing and Steven gets nothing.

Table 2.1: Names for the outcomes shown in Figure 2.1

If, indeed, Sarah is *selfish and greedy* – in the sense that, in evaluating the outcomes, she focuses exclusively on what she herself gets and prefers more money to less – then her ranking of the outcomes is as follows: $o_3 \succ o_1 \succ o_2 \sim o_4$ (which reads ‘ o_3 is better than o_1 , o_1 is better than o_2 and o_2 is just as good as o_4 ’). But there are other possibilities. For example, Sarah might be *fair-minded* and view the outcome where both get \$50,000 as better than all the other outcomes. For instance, her ranking could be $o_1 \succ o_3 \succ o_2 \succ o_4$; according to this ranking, besides valuing fairness, she also displays *benevolence* towards Steven, in the sense that – when comparing the two outcomes where she gets nothing, namely o_2 and o_4 – she prefers the one where at least Steven goes home with some money. If, in fact, Sarah is fair-minded and benevolent, then the logic underlying the above argument would yield the opposite conclusion, namely that she should choose *Split*.

Thus we cannot presume to know the answer to the question “What is the rational choice for Sarah?” if we don’t know what her preferences are. It is a common mistake (unfortunately one that even game theorists sometimes make) to reason under the assumption that players are selfish and greedy. This is, typically, an unwarranted assumption. Research in experimental psychology, philosophy and economics has amply demonstrated that many people are strongly motivated by considerations of fairness. Indeed, fairness seems to motivate not only humans but also primates, as shown in the following video:² http://www.ted.com/talks/frans_de_waal_do_animals_have_morals.

The situation illustrated in Figure 2.1 is not a game as we have no information about the preferences of the players; we use the expression *game-frame* to refer to it. In the case

²Also available at <https://www.youtube.com/watch?v=GcJxRqTs5nk>

where there are only two players and each player has a small number of possible choices (also called *strategies*), a game-frame can be represented – as we did in Figure 2.1 – by means of a table, with as many rows as the number of possible strategies of Player 1 and as many columns as the number of strategies of Player 2; each row is labeled with one strategy of Player 1 and each column with one strategy of Player 2; inside each cell of the table (which corresponds to a pair of strategies, one for Player 1 and one for Player 2) we write the corresponding outcome.

Before presenting the definition of game-frame, we remind the reader of what the Cartesian product of two or more sets is. Let S_1 and S_2 be two sets. Then the Cartesian product of S_1 and S_2 , denoted by $S_1 \times S_2$, is the set of *ordered pairs* (x_1, x_2) where x_1 is an element of S_1 ($x_1 \in S_1$) and x_2 is an element of S_2 ($x_2 \in S_2$). For example, if $S_1 = \{a, b, c\}$ and $S_2 = \{D, E\}$ then

$$S_1 \times S_2 = \{(a, D), (a, E), (b, D), (b, E), (c, D), (c, E)\}.$$

The definition extends to the general case of n sets ($n \geq 2$): an element of $S_1 \times S_2 \times \dots \times S_n$ is an ordered n -tuple (x_1, x_2, \dots, x_n) where, for each $i = 1, \dots, n$, $x_i \in S_i$.

The definition of game-frame is as follows

Definition 2.1.1 A *game-frame in strategic form* is a list of four items (a quadruple) $\langle I, (S_1, S_2, \dots, S_n), O, f \rangle$ where:

- $I = \{1, 2, \dots, n\}$ is a set of *players* ($n \geq 2$).
- (S_1, S_2, \dots, S_n) is a list of sets, one for each player. For every Player $i \in I$, S_i is the set of *strategies* (or possible choices) of Player i . We denote by S the Cartesian product of these sets: $S = S_1 \times S_2 \times \dots \times S_n$; thus an element of S is a list $s = (s_1, s_2, \dots, s_n)$ consisting of one strategy for each player. We call S the set of *strategy profiles*.
- O is a set of *outcomes*.
- $f : S \rightarrow O$ is a function that associates with every strategy profile s an outcome $f(s) \in O$.

Using the notation of Definition 2.1.1, the situation illustrated in Figure 2.1 is the following game-frame in strategic form:

- $I = \{1, 2\}$ (letting Sarah be Player 1 and Steven Player 2),
- $(S_1, S_2) = (\{Split, Steal\}, \{Split, Steal\})$; thus $S_1 = S_2 = \{Split, Steal\}$, so that the set of strategy profiles is $S = \{(Split, Split), (Split, Steal), (Steal, Split), (Steal, Steal)\}$,
- O is the set of outcomes listed in Table 2.1,
- f is the following function:

$$\begin{array}{cccccc} s : & (Split, Split) & (Split, Steal) & (Steal, Split) & (Steal, Steal) & \\ f(s) : & o_1 & o_2 & o_3 & o_4 & \end{array}$$

(that is, $f(Split, Split) = o_1$, $f(Split, Steal) = o_2$, etc.).

From a game-frame one obtains a *game* by adding, for each player, her preferences over (or ranking of) the possible outcomes. We use the notation shown in Table 2.2. For example, if M denotes ‘Mexican food’ and J denotes ‘Japanese food’, then $M \succ_{Alice} J$ means that Alice prefers Mexican food to Japanese food and $M \sim_{Bob} J$ means that Bob is indifferent between the two.

Notation	Interpretation
$o \succsim_i o'$	Player i considers outcome o to be <i>at least as good as</i> o' (that is, either better than or just as good as)
$o \succ_i o'$	Player i considers outcome o to be <i>better than</i> o' (that is, she prefers o to o')
$o \sim_i o'$	Player i considers outcome o to be <i>just as good as</i> o' (that is, she is indifferent between o and o')

Table 2.2: Notation for preference relations

- R** The “at least as good” relation \succsim is sufficient to capture also strict preference \succ and indifference \sim . In fact, starting from \succsim , one can define strict preference as follows: $o \succ o'$ if and only if $o \succsim o'$ and $o' \not\sucsim o$ and one can define indifference as follows: $o \sim o'$ if and only if $o \succsim o'$ and $o' \succsim o$.

We will assume throughout this book that the “at least as good” relation \succsim_i of Player i – which embodies her preferences over (or ranking of) the outcomes – is complete (for every two outcomes o_1 and o_2 , either $o_1 \succsim_i o_2$ or $o_2 \succsim_i o_1$, or both) and transitive (if $o_1 \succsim_i o_2$ and $o_2 \succsim_i o_3$ then $o_1 \succsim_i o_3$).³

There are (at least) four ways of representing, or expressing, a complete and transitive preference relation over (or ranking of) a set of outcomes. For example, suppose that $O = \{o_1, o_2, o_3, o_4, o_5\}$ and that we want to represent the following ranking (expressing the preferences of a given individual): o_3 is better than o_5 , which is just as good as o_1 , o_1 is better than o_4 , which, in turn, is better than o_2 (thus, o_3 is the best outcome and o_2 is the worst outcome). We can represent this ranking in one of the following ways:

- As a subset of $O \times O$ (the interpretation of $(o, o') \in O \times O$ is that o is at least as good as o'):

$$\begin{aligned} & \{(o_1, o_1), (o_1, o_2), (o_1, o_4), (o_1, o_5) \\ & (o_2, o_2), \\ & (o_3, o_1), (o_3, o_2), (o_3, o_3), (o_3, o_4), (o_3, o_5), \\ & (o_4, o_2), (o_4, o_4), \\ & (o_5, o_1), (o_5, o_2), (o_5, o_4), (o_5, o_5)\} \end{aligned}$$

- By using the notation of Table 2.2: $o_3 \succ o_5 \sim o_1 \succ o_4 \succ o_2$.

³Transitivity of the “at least as good” relation implies transitivity of the indifference relation (if $o_1 \sim o_2$ and $o_2 \sim o_3$ then $o_1 \sim o_3$) as well as transitivity of the strict preference relation (not only in the sense that (1) if $o_1 \succ o_2$ and $o_2 \succ o_3$ then $o_1 \succ o_3$, but also that (2) if $o_1 \succ o_2$ and $o_2 \sim o_3$ then $o_1 \succ o_3$ and (3) if $o_1 \sim o_2$ and $o_2 \succ o_3$ then $o_1 \succ o_3$).

- By listing the outcomes in a column, starting with the best at the top and proceeding down to the worst, thus using the convention that if outcome o is listed above outcome o' then o is preferred to o' , while if o and o' are written next to each other (on the same row), then they are considered to be just as good:

best o_3
 o_1, o_5
 o_4
 worst o_2

- By assigning a number to each outcome, with the convention that *if the number assigned to o is greater than the number assigned to o'* then o is preferred to o' , and if two outcomes are assigned the same number then they are considered to be just as

good. For example, we could choose the following numbers: $\begin{matrix} o_1 & o_2 & o_3 & o_4 & o_5 \\ 6 & 1 & 8 & 2 & 6 \end{matrix}$.

Such an assignment of numbers is called a *utility function*. A useful way of thinking of utility is as an “index of satisfaction”: the higher the index the better the outcome; however, this suggestion is just to aid memory and should be taken with a grain of salt because a utility function *does not measure anything* and, furthermore, as explained below, the actual numbers used as utility indices are completely arbitrary.⁴

Definition 2.1.2 Given a complete and transitive ranking \succsim of a finite set of outcomes O , a function $U : O \rightarrow \mathbb{R}$ (where \mathbb{R} denotes the set of real numbers)^a is said to be an *ordinal utility function that represents the ranking \succsim* if, for every two outcomes o and o' , $U(o) > U(o')$ if and only if $o \succ o'$ and $U(o) = U(o')$ if and only if $o \sim o'$. The number $U(o)$ is called the *utility of outcome o* .^b

^aThe notation $f : X \rightarrow Y$ is used to denote a function which associates with every $x \in X$ an element $y = f(x)$ with $y \in Y$.

^bThus, $o \succsim o'$ if and only if $U(o) \geq U(o')$.

R Note that the statement “for Alice the utility of Mexican food is 10” is in itself a meaningless statement; on the other hand, what would be a meaningful statement is “for Alice the utility of Mexican food is 10 and the utility of Japanese food is 5”, because such a statement conveys the information that she prefers Mexican food to Japanese food. However, the two numbers 10 and 5 have no other meaning besides the fact that 10 is greater than 5: for example, we **cannot** infer from these numbers that she considers Mexican food twice as good as Japanese food. The reason for this is that we could have expressed the same fact, namely that she prefers Mexican food to Japanese food, by assigning utility 100 to Mexican food and -25 to Japanese food, or with any other two numbers (as long as the number assigned to Mexican food is larger than the number assigned to Japanese food).

⁴Note that assigning a utility of 1 to an outcome o does not mean that o is the “first choice”. Indeed, in this example a utility of 1 is assigned to the worst outcome: o_2 is the worst outcome because it has the lowest utility (which happens to be 1, in this example).

It follows from the above remark that there is an infinite number of utility functions that represent the same ranking. For instance, the following are equivalent ways of representing the ranking $o_3 \succ o_1 \succ o_2 \sim o_4$ (f , g and h are three out of the many possible utility functions):

<i>outcome</i> \rightarrow	o_1	o_2	o_3	o_4
<i>utility function</i> \downarrow				
f :	5	2	10	2
g :	0.8	0.7	1	0.7
h :	27	1	100	1

Utility functions are a particularly convenient way of representing preferences. In fact, by using utility functions one can give a more condensed representation of games, as explained in the last paragraph of the following definition.

Definition 2.1.3 An ordinal game in strategic form is a quintuple

$\langle I, (S_1, \dots, S_n), O, f, (\succsim_1, \dots, \succsim_n) \rangle$ where:

- $\langle I, (S_1, \dots, S_n), O, f \rangle$ is a game-frame in strategic form (Definition 2.1.1) and
- for every Player $i \in I$, \succsim_i is a complete and transitive ranking of the set of outcomes O .

If we replace each ranking \succsim_i with a utility function U_i that represents it, and we assign, to each strategy profile s , Player i 's utility of $f(s)$ (recall that $f(s)$ is the outcome associated with s) then we obtain a function $\pi_i : S \rightarrow \mathbb{R}$ called Player i 's *payoff function*. Thus $\pi_i(s) = U_i(f(s))$.^a Having done so, we obtain a triple $\langle I, (S_1, \dots, S_n), (\pi_1, \dots, \pi_n) \rangle$ called a *reduced-form ordinal game in strategic form* ('reduced-form' because some information is lost, namely the specification of the possible outcomes).

^aNote that, in this book, the symbol π is not used to denote the irrational number used to compute the circumference and area of a circle, but rather as the Greek letter for 'p' which stands for 'payoff'.

For example, take the game-frame illustrated in Figure 2.1, let Sarah be Player 1 and Steven Player 2 and name the possible outcomes as shown in Table 2.1. Let us add the information that both players are selfish and greedy (that is, **Player 1**'s ranking is $o_3 \succ_1 o_1 \succ_1 o_2 \sim_1 o_4$ and **Player 2**'s ranking is $o_2 \succ_2 o_1 \succ_2 o_3 \sim_2 o_4$) and let us represent their rankings with the following utility functions (note, again, that the choice of numbers 2, 3 and 4 for utilities is arbitrary: any other three numbers would do):

<i>outcome</i> \rightarrow	o_1	o_2	o_3	o_4
<i>utility function</i> \downarrow				
U_1 (Player 1):	3	2	4	2
U_2 (Player 2):	3	4	2	2

Then we obtain the reduced-form game shown in Figure 2.2, where in each cell the first number is the payoff of Player 1 and the second number is the payoff of Player 2.

		Player 2 (Steven)	
		Split	Steal
Player 1 (Sarah)	Split	3 3	2 4
	Steal	4 2	2 2

Figure 2.2: One possible game based on the game-frame of Figure 2.1

On the other hand, if we add to the game-frame of Figure 2.1 the information that **Player 1** is fair-minded and benevolent (that is, her ranking is $o_1 \succ_1 o_3 \succ_1 o_2 \succ_1 o_4$), while Player 2 is selfish and greedy and represent these rankings with the following utility functions:

<i>outcome</i> →	o_1	o_2	o_3	o_4
<i>utility function</i> ↓				
U_1 (Player 1):	4	2	3	1
U_2 (Player 2):	3	4	2	2

then we obtain the reduced-form game shown in Figure 2.3.

		Player 2 (Steven)	
		Split	Steal
Player 1 (Sarah)	Split	4 3	2 4
	Steal	3 2	1 2

Figure 2.3: Another possible game based on the game-frame of Figure 2.1

In general, a player will act differently in different games, even if they are based on the same game-frame, because her incentives and objectives (as captured by her ranking of the outcomes) will be different. For example, one can argue that in the game of Figure 2.2 a rational Player 1 would choose *Steal*, while in the game of Figure 2.3 the rational choice for Player 1 is *Split*.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 2.9.1 at the end of this chapter.

2.2 Strict and weak dominance

In this section we define two relations on the set of strategies of a player. Before introducing the formal definition, we shall illustrate these notions with an example. The first relation is called “strict dominance”. Let us focus our attention on one player, say Player 1, and select two of her strategies, say a and b . We say that a **strictly** dominates b (for Player 1) if, for every possible strategy profile of the other players, strategy a of Player 1, in conjunction with the strategies selected by the other players, yields a payoff for Player 1 which is **greater than** the payoff associated with strategy b (in conjunction with the strategies selected by the other players). For example, consider the following two-player game, where only the payoffs of Player 1 are shown:

		Player 2		
		<i>E</i>	<i>F</i>	<i>G</i>
Player 1	<i>A</i>	3 ...	2 ...	1 ...
	<i>B</i>	2 ...	1 ...	0 ...
	<i>C</i>	3 ...	2 ...	1 ...
	<i>D</i>	2 ...	0 ...	0 ...

Figure 2.4: A game showing only the payoffs of Player 1

In this game for Player 1 strategy A strictly dominates strategy B :

- if Player 2 selects E then A in conjunction with E gives Player 1 a payoff of 3, while B in conjunction with E gives her only a payoff of 2,
- if Player 2 selects F then A in conjunction with F gives Player 1 a payoff of 2, while B in conjunction with F gives her only a payoff of 1,
- if Player 2 selects G then A in conjunction with G gives Player 1 a payoff of 1, while B in conjunction with G gives her only a payoff of 0.

In the game of Figure 2.4 we also have that A strictly dominates D and C strictly dominates D ; however, it is not the case that B strictly dominates D because, in conjunction with strategy E of Player 2, B and D yield the same payoff for Player 1.

The second relation is called “weak dominance”. The definition is similar to that of strict dominance, but we replace ‘greater than’ with ‘greater than or equal to’ while insisting on at least one strict inequality: a **weakly** dominates b (for Player 1) if, for every possible strategy profile of the other players, strategy a of Player 1, in conjunction with the strategies selected by the other players, yields a payoff for Player 1 which is **greater than or equal to** the payoff associated with strategy b (in conjunction with the strategies selected by the other players) and, furthermore, there is at least one strategy profile of the other players against which strategy a gives a larger payoff to Player 1 than strategy b . In the example of Figure 2.4, we have that, while it is not true that B strictly dominates D , it is true that B weakly dominates D :

- if Player 2 selects E , then B in conjunction with E gives Player 1 the same payoff as D in conjunction with E (namely 2),
- if Player 2 selects F , then B in conjunction with F gives Player 1 a payoff of 1, while D in conjunction with F gives her only a payoff of 0,
- if Player 2 selects G then B in conjunction with G gives Player 1 the same payoff as D in conjunction with G (namely 0).

In order to give the definitions in full generality we need to introduce some notation. Recall that S denotes the set of strategy profiles, that is, an element s of S is an ordered list of strategies $s = (s_1, \dots, s_n)$, one for each player. We will often want to focus on one player, say Player i , and view s as a pair consisting of the strategy of Player i and the remaining strategies of all the other players. For example, suppose that there are three players and the strategy sets are as follows: $S_1 = \{a, b, c\}$, $S_2 = \{d, e\}$ and $S_3 = \{f, g\}$. Then one possible strategy profile is $s = (b, d, g)$ (thus $s_1 = b$, $s_2 = d$ and $s_3 = g$). If we focus on, say, Player 2 then we will denote by s_{-2} the sub-profile consisting of the strategies of the players *other than 2*: in this case $s_{-2} = (b, g)$. This gives us an alternative way of denoting s , namely as (s_2, s_{-2}) . Continuing our example where $s = (b, d, g)$, letting $s_{-2} = (b, g)$, we can denote s also by (d, s_{-2}) and we can write the result of replacing Player 2's strategy d with her strategy e in s by (e, s_{-2}) ; thus $(d, s_{-2}) = (b, d, g)$ while $(e, s_{-2}) = (b, e, g)$. In general, given a Player i , we denote by S_{-i} the set of strategy profiles of the players other than i (that is, S_{-i} is the Cartesian product of the strategy sets of the other players; in the above example we have that $S_{-2} = S_1 \times S_3 = \{a, b, c\} \times \{f, g\} = \{(a, f), (a, g), (b, f), (b, g), (c, f), (c, g)\}$). We denote an element of S_{-i} by s_{-i} .

Definition 2.2.1 Given an ordinal game in strategic form, let i be a Player and a and b two of her strategies ($a, b \in S_i$). We say that, for Player i ,

- a *strictly dominates* b (or b is *strictly dominated* by a) if, in every situation (that is, no matter what the other players do), a gives Player i a payoff which is *greater* than the payoff that b gives. Formally: for every $s_{-i} \in S_{-i}$, $\pi_i(a, s_{-i}) > \pi_i(b, s_{-i})$.^a
- a *weakly dominates* b (or b is *weakly dominated* by a) if, in every situation, a gives Player i a payoff which is *greater than or equal to* the payoff that b gives and, furthermore, there is at least one situation where a gives a greater payoff than b . Formally: for every $s_{-i} \in S_{-i}$, $\pi_i(a, s_{-i}) \geq \pi_i(b, s_{-i})$ and there exists an $\bar{s}_{-i} \in S_{-i}$ such that $\pi_i(a, \bar{s}_{-i}) > \pi_i(b, \bar{s}_{-i})$.^b
- a is *equivalent to* b if, in every situation, a and b give Player i the *same* payoff. Formally: for every $s_{-i} \in S_{-i}$, $\pi_i(a, s_{-i}) = \pi_i(b, s_{-i})$.^c

^aOr, stated in terms of rankings instead of payoffs, $f(a, s_{-i}) \succ_i f(b, s_{-i})$ for every $s_{-i} \in S_{-i}$.

^bOr, stated in terms of rankings, $f(a, s_{-i}) \succeq_i f(b, s_{-i})$, for every $s_{-i} \in S_{-i}$, and there exists an $\bar{s}_{-i} \in S_{-i}$ such that $f(a, \bar{s}_{-i}) \succ_i f(b, \bar{s}_{-i})$.

^cOr, stated in terms of rankings, $f(a, s_{-i}) \sim_i f(b, s_{-i})$, for every $s_{-i} \in S_{-i}$.

For example, in the game of Figure 2.5 (which reproduces Figure 2.4), we have that

- A strictly dominates B.
- A and C are equivalent.
- A strictly dominates D.
- B is strictly dominated by C.
- B weakly (but not strictly) dominates D.
- C strictly dominates D.

		Player 2		
		<i>E</i>	<i>F</i>	<i>G</i>
Player 1	<i>A</i>	3 ...	2 ...	1 ...
	<i>B</i>	2 ...	1 ...	0 ...
	<i>C</i>	3 ...	2 ...	1 ...
	<i>D</i>	2 ...	0 ...	0 ...

Figure 2.5: Copy of Figure 2.4

- R Note that if strategy a strictly dominates strategy b then it also satisfies the conditions for weak dominance, that is, ‘ a strictly dominates b ’ implies ‘ a weakly dominates b ’. Throughout the book the expression ‘ a weakly dominates b ’ will be interpreted as ‘ a dominates b weakly but not strictly’.

The expression ‘ a dominates b ’ can be understood as ‘ a is better than b ’. The next term we define is ‘dominant’ which can be understood as ‘best’. Thus one cannot meaningfully say “ a dominates” because one needs to name another strategy that is dominated by a ; for example, one would have to say “ a dominates b ”. On the other hand, one *can* meaningfully say “ a is dominant” because it is like saying “ a is best”, which means “ a is better than every other strategy”.

Definition 2.2.2 Given an ordinal game in strategic form, let i be a Player and a one of her strategies ($a \in S_i$). We say that, for Player i ,

- a is a *strictly dominant* strategy if a strictly dominates every other strategy of Player i .
- a is a *weakly dominant* strategy if, for every other strategy x of Player i , one of the following is true: either (1) a weakly dominates x or (2) a is equivalent to x .

For example, in the game shown in Figure 2.5, *A* and *C* are both weakly dominant strategies for Player 1. Note that if a player has two or more strategies that are weakly dominant, then any two of those strategies must be equivalent. On the other hand, there can be at most one strictly dominant strategy.

R The reader should convince herself/himself that the definition of weakly dominant strategy given in Definition 2.2.2 is equivalent to the following: $a \in S_i$ is a weakly dominant strategy for Player i if and only if, for every $s_{-i} \in S_{-i}$, $\pi_i(a, s_{-i}) \geq \pi_i(s_i, s_{-i})$ for every $s_i \in S_i$.⁵

In accordance with the convention established earlier, the expression ‘ a is a weakly dominant strategy’ will have the default interpretation ‘ a is a *weakly but not strictly dominant strategy*’.

Note: if you claim that, for some player, “strategy x is (weakly or strictly) dominated” then you ought to name another strategy of that player that dominates x . Saying “ x is dominated” is akin to saying “ x is worse”: worse than what? On the other hand, claiming that strategy x is weakly dominant is akin to claiming that it is best, that is, better than, or just as good as, any other strategy.

Definition 2.2.3 Given an ordinal game in strategic form, let $s = (s_1, \dots, s_n)$ be a strategy profile. We say that^a

- s is a *strict dominant-strategy profile* if, for every Player i , s_i is a strictly-dominant strategy.
- s is a *weak dominant-strategy profile* if, for every Player i , s_i is a weakly-dominant strategy and, furthermore, for at least one Player j , s_j is not a strictly-dominant strategy.

^aIn the literature, the expression ‘dominant-strategy equilibrium’ is often used instead of ‘dominant-strategy profile’, but the latter seems to be more appropriate.

If we refer to a strategy profile as a dominant-strategy profile, without qualifying it as weak or strict, then the default interpretation will be ‘weak’.

In the game of Figure 2.6 (which reproduces Figure 2.2), *Steal* is a weakly dominant strategy for each player and thus $(Steal, Steal)$ is a weak dominant-strategy profile.

		Player 2 (Steven)	
		Split	Steal
Player 1 (Sarah)	Split	3 3	2 4
	Steal	4 2	2 2

Figure 2.6: Copy of Figure 2.2

		Player 2 (Steven)	
		Split	Steal
Player 1 (Sarah)	Split	4 3	2 4
	Steal	3 2	1 2

Figure 2.7: Copy of Figure 2.3

In the game of Figure 2.7 (which reproduces Figure 2.3), *Split* is a strictly dominant strategy for Player 1, while *Steal* is a weakly (but not strictly) dominant strategy for Player 2 and thus $(Split, Steal)$ is a weak dominant-strategy profile.

⁵Or, stated in terms of rankings, for every $s_{-i} \in S_{-i}$, $f(a, s_{-i}) \succsim_i f(s_i, s_{-i})$ for every $s_i \in S_i$.

The *Prisoner's Dilemma* is an example of a game with a strict dominant-strategy profile. For a detailed account of the history of this game and an in-depth analysis of it see <http://plato.stanford.edu/entries/prisoner-dilemma> or http://en.wikipedia.org/wiki/Prisoner's_dilemma.

An instance of the Prisoner's Dilemma is the following situation. Doug and Ed work for the same company and the annual party is approaching. They know that they are the only two candidates for the best-worker-of-the-year prize and at the moment they are tied; however, only one person can be awarded the prize and thus, unless one of them manages to outperform the other, nobody will receive the prize. Each chooses between exerting *Normal* effort or *Extra* effort (that is, work overtime) before the party. The corresponding game-frame is shown in Figure 2.8.

		Player 2 (Ed)	
		Normal effort	Extra effort
Player 1 (Doug)	Normal effort	O_1	O_2
	Extra effort	O_3	O_4

- O_1 : nobody gets the prize and nobody sacrifices family time
- O_2 : Ed gets the prize and sacrifices family time, Doug does not
- O_3 : Doug gets the prize and sacrifices family time, Ed does not
- O_4 : nobody gets the prize and both sacrifice family time

Figure 2.8: The Prisoner's Dilemma game-frame

Suppose that both **Doug** and **Ed** are willing to sacrifice family time to get the prize, but otherwise value family time; furthermore, they are envious of each other, in the sense that they prefer nobody getting the prize to the other person's getting the prize (even at the personal cost of sacrificing family time). That is, their rankings are as follows: $O_3 \succ_{\text{Doug}} O_1 \succ_{\text{Doug}} O_4 \succ_{\text{Doug}} O_2$ and $O_2 \succ_{\text{Ed}} O_1 \succ_{\text{Ed}} O_4 \succ_{\text{Ed}} O_3$. Using utility functions with values from the set $\{0, 1, 2, 3\}$ we can represent the game in reduced form as shown in Figure 2.9. In this game exerting extra effort is a strictly dominant strategy for every player; thus *(Extra effort, Extra effort)* is a strict dominant-strategy profile.

Definition 2.2.4 Given an ordinal game in strategic form, let o and o' be two outcomes. We say that o is *strictly Pareto superior* to o' if every player prefers o to o' (that is, if $o \succ_i o'$, for every Player i). We say that o is *weakly Pareto superior* to o' if every player considers o to be at least as good as o' and at least one player prefers o to o' (that is, if $o \succeq_i o'$, for every Player i and there is a Player j such that $o \succ_j o'$).

In reduced-form games, this definition can be extended to strategy profiles as follows. If s and s' are two strategy profiles, then s is *strictly Pareto superior* to s' if $\pi_i(s) > \pi_i(s')$ for every Player i and s is *weakly Pareto superior* to s' if $\pi_i(s) \geq \pi_i(s')$ for every Player i and, furthermore, there is a Player j such that $\pi_j(s) > \pi_j(s')$.

		Player 2 (Ed)	
		Normal effort	Extra effort
Player 1 (Doug)	Normal effort	2 2	0 3
	Extra effort	3 0	1 1

Figure 2.9: The Prisoner's Dilemma game

For example, in the Prisoner's Dilemma game of Figure 2.9, outcome o_1 is strictly Pareto superior to o_4 or, in terms of strategy profiles, $(Normal\ effort, Normal\ effort)$ is strictly Pareto superior to $(Extra\ effort, Extra\ effort)$.

When a player has a strictly dominant strategy, it would be irrational for that player to choose any other strategy, since she would be guaranteed a lower payoff in every possible situation (that is, no matter what the other players do). Thus in the Prisoner's Dilemma individual rationality leads to $(Extra\ effort, Extra\ effort)$ despite the fact that both players would be better off if they both chose *Normal effort*. It is obvious that if the players could reach a *binding* agreement to exert normal effort then they would do so; however, the underlying assumption in non-cooperative game theory is that such agreements are not possible (e.g. because of lack of communication or because such agreements are illegal or cannot be enforced in a court of law, etc.). Any non-binding agreement to choose *Normal effort* would not be viable: if one player expects the other player to stick to the agreement, then he will gain by cheating and choosing *Extra effort* (on the other hand, if a player does not believe that the other player will honor the agreement then he will gain by deviating from the agreement herself). The Prisoner's Dilemma game is often used to illustrate a conflict between individual rationality and collective rationality: $(Extra\ effort, Extra\ effort)$ is the individually rational profile while $(Normal\ effort, Normal\ effort)$ would be the collectively rational one.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 2.9.2 at the end of this chapter.

2.3 Second-price auction

The *second-price auction*, or *Vickrey auction*, is an example of a game that has a weak dominant-strategy profile. It is a "sealed-bid" auction where bidders submit bids without knowing the bids of the other participants in the auction. The object which is auctioned is then assigned to the bidder who submits the highest bid (the winner), but the winner pays not her own bid but rather the second-highest bid, that is the highest bid among the bids that remain after removing the winner's own bid. Tie-breaking rules must be specified for selecting the winner when the highest bid is submitted by two or more bidders (in which case the winner ends up paying her own bid, because the second-highest bid is equal to the winner's bid). We first illustrate this auction with an example:

Two oil companies bid for the right to drill a field. The possible bids are \$10 million, \$20 million, . . . , \$50 million. In case of ties the winner is Player 2 (this was decided earlier by tossing a coin). Let us take the point of view of Player 1. Suppose that Player 1 ordered a geological survey and, based on the report, concludes that the oil field would generate a profit of \$30 million. Suppose also that Player 1 is indifferent between any two outcomes where the oil field is given to Player 2 and prefers to get the oil field herself if and only if it has to pay not more than \$30 million for it; furthermore, getting the oil field for \$30 million is just as good as not getting it. Then we can take as utility function for Player 1 the net gain to Player 1 from the oil field (defined as profits from oil extraction minus the price paid for access to the oil field) if Player 1 wins, and zero otherwise.

		Player 2				
		\$10M	\$20M	\$30M	\$40M	\$50M
<b style="color: red;">Player 1 (value \$30M)	\$10M	0	0	0	0	0
	\$20M	20	0	0	0	0
	\$30M	20	10	0	0	0
	\$40M	20	10	0	0	0
	\$50M	20	10	0	-10	0

Figure 2.10: A second-price auction where, in case of ties, the winner is Player 2

In Figure 2.10 we have written inside each cell only the payoff of Player 1. For example, why is Player 1's payoff 20 when it bids \$30M and Player 2 bids \$10M? Since Player 1's bid is higher than Player 2's bid, Player 1 is the winner and thus the drilling rights are assigned to Player 1; hence Player 1 obtains something worth \$30M and pays, not its own bid of \$30M, but the bid of Player 2, namely \$10M; it follows that Player 1's net gain is $$(30 - 10)M = $20M$.

It can be verified that for Player 1 submitting a bid equal to the value it assigns to the object (namely, a bid of \$30M) is a weakly dominant strategy: it always gives Player 1 the largest of the payoffs that are possible, given the bid of the other player. This does not imply that it is the only weakly dominant strategy; indeed, in this example bidding \$40M is also a weakly dominant strategy for Player 1 (in fact, it is equivalent to bidding \$30M).

Now we can describe the second-price auction in more general terms. Let $n \geq 2$ be the number of bidders. We assume that all non-negative numbers are allowed as bids and that the tie-breaking rule favors the player with the lowest index among those who submit the highest bid: for example, if the highest bid is \$250 and it is submitted by Players 5, 8 and 10, then the winner is Player 5. We shall denote the possible outcomes as pairs (i, p) , where i is the winner and p is the price that the winner has to pay. Finally we denote by b_i the bid of Player i . We start by describing the case where there are only two bidders and then generalize to the case of an arbitrary number of bidders. We denote the set of non-negative numbers by $[0, \infty)$.

The case where $n = 2$: in this case we have that $I = \{1, 2\}$, $S_1 = S_2 = [0, \infty)$, $O = \{(i, p) : i \in \{1, 2\}, p \in [0, \infty)\}$ and $f : S \rightarrow O$ is given by

$$f((b_1, b_2)) = \begin{cases} (1, b_2) & \text{if } b_1 \geq b_2 \\ (2, b_1) & \text{if } b_1 < b_2 \end{cases}.$$

The case where $n \geq 2$: in the general case the second-price auction is the following game-frame:

- $I = \{1, \dots, n\}$.
- $S_i = [0, \infty)$ for every $i = 1, \dots, n$. We denote an element of S_i by b_i .
- $O = \{(i, p) : i \in I, p \in [0, \infty)\}$.
- $f : S \rightarrow O$ is defined as follows. Let $H(b_1, \dots, b_n) \subseteq I$ be the set of bidders who submit the highest bid: $H(b_1, \dots, b_n) = \{i \in I : b_i \geq b_j \text{ for all } j \in I\}$ and let $\hat{i}(b_1, \dots, b_n)$ be the smallest number in the set $H(b_1, \dots, b_n)$, that is, the winner of the auction. Finally, let b^{\max} denote the maximum bid and $b^{\text{second}}(b_1, \dots, b_n)$ the second-highest bid,⁶ that is,

$$b^{\max}(b_1, \dots, b_n) = \text{Max}\{b_1, \dots, b_n\}$$

$$b^{\text{second}}(b_1, \dots, b_n) = \text{Max}(\{b_1, \dots, b_n\} \setminus \{b^{\max}(b_1, \dots, b_n)\}).$$

Then $f(b_1, \dots, b_n) = (\hat{i}(b_1, \dots, b_n), b^{\text{second}}(b_1, \dots, b_n))$.

How much should a player bid in a second-price auction? Since what we have described is a game-frame and not a game, we cannot answer the question unless we specify the player's preferences over the set of outcomes O . Let us say that Player i in a second-price auction is *selfish and greedy* if she only cares about whether or not she wins and – conditional on winning – prefers to pay less; furthermore, she prefers winning to not winning if and only if she has to pay less than the true value of the object for her, which we denote by v_i , and is indifferent between not winning and winning if she has to pay v_i . Thus the ranking of a selfish and greedy player is as follows (together with everything that follows from transitivity):

$$\begin{aligned} (i, p) \succ_i (i, p') & \quad \text{if and only if } p < p' \\ (i, p) \succ_i (j, p') & \quad \text{for all } j \neq i \text{ and for all } p', \text{ if and only if } p < v_i \\ (i, v_i) \sim_i (j, p') & \quad \text{for all } j \neq i \text{ and for all } p' \\ (j, p) \sim_i (k, p') & \quad \text{for all } j \neq i, k \neq i \text{ and for all } p \text{ and } p'. \end{aligned}$$

An ordinal utility function that represents these preferences is:⁷

$$U_i(j, p) = \begin{cases} v_i - p & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$$

⁶For example, if $n = 5, b_1 = \$10, b_2 = \$14, b_3 = \$8, b_4 = \14 and $b_5 = \$14$ then $H(\$10, \$14, \$8, \$14, \$14) = \{2, 4, 5\}$, $\hat{i}(\$10, \$14, \$8, \$14, \$14) = 2$, $b^{\max}(\$10, \$14, \$8, \$14, \$14) = \14 and $b^{\text{second}}(\$10, \$14, \$8, \$14, \$14) = \14 .

⁷Of course there are many more. For example, also the following utility function represents those preferences: $U_i(j, p) = \begin{cases} 2^{(v_i - p)} & \text{if } i = j \\ 1 & \text{if } i \neq j \end{cases}$

Using this utility function we get the following payoff function for Player i :

$$\pi_i(b_1, \dots, b_n) = \begin{cases} v_i - b^{\text{second}}(b_1, \dots, b_n) & \text{if } i = \hat{i}(b_1, \dots, b_n) \\ 0 & \text{if } i \neq \hat{i}(b_1, \dots, b_n) \end{cases}$$

We can now state the following theorem. The proof is given in Section 2.8.

Theorem 2.3.1 — Vickrey, 1961. In a second-price auction, if Player i is selfish and greedy then it is a weakly dominant strategy for Player i to bid her true value, that is, to choose $b_i = v_i$.

Note that, for a player who is not selfish and greedy, Theorem 2.3.1 is not true. For example, if a player has the same preferences as above for the case where she wins, but, conditional on not winning, prefers the other player to pay as much as possible (she is spiteful) or as little as possible (she is generous), then bidding her true value is no longer a dominant strategy.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 2.9.3 at the end of this chapter.

2.4 The pivotal mechanism

An article in the *Davis Enterprise* (the local newspaper in Davis, California) on January 12, 2001 started with the following paragraph:

“By consensus, the Davis City Council agreed Wednesday to order a communitywide public opinion poll to gauge how much Davis residents would be willing to pay for a park tax and a public safety tax.”

Opinion polls of this type are worthwhile only if there are reasons to believe that the people who are interviewed will respond honestly. But will they? If I would like more parks and believe that the final tax I will have to pay is independent of the amount I state in the interview, I would have an incentive to overstate my willingness to pay, hoping to swing the decision towards building a new park. On the other hand, if I fear that the final tax will be affected by the amount I report, then I might have an incentive to understate my willingness to pay.

The *pivotal mechanism*, or *Clarke mechanism*, is a game designed to give the participants an incentive to report their true willingness to pay.

A public project, say to build a park, is under consideration. The cost of the project is $\$C$. There are n individuals in the community. If the project is carried out, individual i ($i = 1, \dots, n$) will have to pay $\$c_i$ (with $c_1 + c_2 + \dots + c_n = C$); these amounts are specified as part of the project. Note that we allow for the possibility that some individuals might have to contribute a larger share of the total cost C than others (e.g. because they live closer to the projected park and would therefore benefit more from it). Individual i has an initial wealth of $\$\bar{m}_i > 0$. If the project is carried out, individual i receives benefits from

it that she considers equivalent to receiving $\$v_i$. Note that for some individual i , v_i could be negative, that is, the individual could be harmed by the project (e.g. because she likes peace and quiet and a park close to her home would bring extra traffic and noise). We assume that individual i has the following utility-of-wealth function:

$$U_i(\$m) = \begin{cases} m & \text{if the project is not carried out} \\ m + v_i & \text{if the project is carried out} \end{cases}$$

The socially efficient decision is to carry out the project if and only if $\sum_{i=1}^n v_i > C$ (recall that \sum is the summation sign: $\sum_{i=1}^n v_i$ is a short-hand for $v_1 + v_2 + \dots + v_n$). For example, suppose that $n = 2$, $\bar{m}_1 = 50$, $\bar{m}_2 = 60$, $v_1 = 19$, $v_2 = -15$, $C = 6$, $c_1 = 6$, $c_2 = 0$. In this case $\sum_{i=1}^n v_i = 19 - 15 = 4 < C = 6$ hence the project should not be carried out. To see this consider the following table:

	If the project is not carried out	If the project is carried out
Utility of Individual 1	50	$50 + 19 - 6 = 63$
Utility of Individual 2	60	$60 - 15 = 45$

If the project is carried out, Individual 1 has a utility gain of 13, while Individual 2 has a utility loss of 15. Since the loss is greater than the gain, we have a Pareto inefficient situation. Individual 2 could propose the following alternative to Individual 1: let us not carry out the project and I will pay you \$14. Then Individual 1's wealth and utility would be $50 + 14 = 64$ and Individual 2's wealth and utility would be $60 - 14 = 46$ and thus they would both be better off.

Thus Pareto efficiency requires that the project be carried out if and only if $\sum_{i=1}^n v_i > C$. This would be a simple decision for the government if it knew the v_i 's. But, typically, these values are private information to the individuals. Can the government find a way to induce the individuals to reveal their true valuations? It seems that in general the answer is No: those who gain from the project would have an incentive to overstate their potential gains, while those who suffer would have an incentive to overstate their potential losses.

Influenced by Vickrey's work on second-price auctions, Clarke suggested the following procedure or game. Each individual i is asked to submit a number w_i which will be interpreted as the gross benefit (if positive) or harm (if negative) that individual i associates with the project. Note that, in principle, individual i can lie and report a value w_i which is different from the true value v_i . Then the decision will be:

$$\text{Carry out the project?} \begin{cases} \text{Yes} & \text{if } \sum_{j=1}^n w_j > C \\ \text{No} & \text{if } \sum_{j=1}^n w_j \leq C \end{cases}$$

However, this is not the end of the story. Each individual will be classified as either not pivotal or pivotal.

$$\text{Individual } i \text{ is not pivotal if } \begin{cases} \text{either } \left(\sum_{j=1}^n w_j > C \text{ and } \sum_{j \neq i} w_j > \sum_{j \neq i} c_j \right) \\ \text{or } \left(\sum_{j=1}^n w_j \leq C \text{ and } \sum_{j \neq i} w_j \leq \sum_{j \neq i} c_j \right) \end{cases}$$

and she is pivotal otherwise. In other words, individual i is pivotal if the decision about the project that would be made in the restricted society resulting from removing individual i is different from the decision that is made when she is included. If an individual is not pivotal then she has to pay no taxes. If individual i is pivotal then she has to pay a tax in the amount of

$$\left| \sum_{j \neq i} w_j - \sum_{j \neq i} c_j \right|, \text{ that is, the absolute value of } \sum_{j \neq i} w_j - \sum_{j \neq i} c_j$$

(recall that the absolute value of a is equal to a , if a is positive, and to $-a$, if a is a negative; for instance, $|4| = 4$ and $|-4| = -(-4) = 4$).

For example, let $n = 3$, $C = 10$, $c_1 = 3$, $c_2 = 2$, $c_3 = 5$.

Suppose that they state the following benefits/losses (which may or may not be the true ones): $w_1 = -1, w_2 = 8, w_3 = 3$.

Then $\sum_{i=1}^3 w_i = 10 = C$.

Thus the project will not be carried out. Who is pivotal? The answer is provided in Figure 2.11.

Individual $i =$	$\sum w_j$ (including i)	$\sum c_j$ (including i)	Decision	$\sum w_j$ $j \neq i$ (without i)	$\sum c_j$ $j \neq i$ (without i)	Decision	Pivotal?	Tax
1	10	10	No	$8 + 3 = 11$	$2 + 5 = 7$	Yes	Yes	$11 - 7 = 4$
2	10	10	No	$-1 + 3 = 2$	$3 + 5 = 8$	No	No	0
3	10	10	No	$-1 + 8 = 7$	$3 + 2 = 5$	Yes	Yes	$7 - 5 = 2$

Figure 2.11: Example of pivotal mechanism

It may seem that, since it involves paying a tax, being pivotal is a bad thing and one should try to avoid it. It is certainly possible for individual i to make sure that she is not pivotal: all she has to do is to report $w_i = c_i$; in fact, if $\sum_{j \neq i} w_j > \sum_{j \neq i} c_j$ then adding c_i

to both sides yields $\sum_{j=1}^n w_j > C$ and if $\sum_{j \neq i} w_j \leq \sum_{j \neq i} c_j$ then adding c_i to both sides yields

$\sum_{j=1}^n w_j \leq C$. It is not true, however, that it is best to avoid being pivotal. The following

example shows that one can gain by being truthful even if it involves being pivotal and thus having to pay a tax. Let $n = 4, C = 15, c_1 = 5, c_2 = 0, c_3 = 5$ and $c_4 = 5$.

Suppose that $\bar{m}_1 = 40$ and $v_1 = 25$.

Imagine that you are Individual 1 and, for whatever reason, you expect the following reports by the other individuals: $w_2 = -40, w_3 = 15$ and $w_4 = 20$.

If you report $w_1 = c_1 = 5$ then you ensure that you are not pivotal.

In this case $\sum_{j=1}^4 w_j = 5 - 40 + 15 + 20 = 0 < C = 15$ and thus the project is not carried out and your utility is equal to $\bar{m}_1 = 40$. If you report truthfully, that is, you report $w_1 = v_1 = 25$ then $\sum_{j=1}^4 w_j = 25 - 40 + 15 + 20 = 20 > C = 15$ and the project is carried out; furthermore, you are pivotal and have to pay a tax t_1 equal to

$$\left| \sum_{j=2}^4 w_j - \sum_{j=2}^4 c_j \right| = |(-40 + 15 + 20) - (0 + 5 + 5)| = |-15| = 15$$

and your utility will be $\bar{m}_1 + v_1 - c_1 - t_1 = 40 + 25 - 5 - 15 = 45$; hence you are better off. Indeed, the following theorem states that no individual can ever gain by lying.

The proof of Theorem 2.4.1 is given in Section 2.8.

Theorem 2.4.1 — Clarke, 1971. In the pivotal mechanism (under the assumed preferences) truthful revelation (that is, stating $w_i = v_i$) is a weakly dominant strategy for every Player i .

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 2.9.4 at the end of this chapter.

2.5 Iterated deletion procedures

If in a game a player has a (weakly or strictly) dominant strategy then the player ought to choose that strategy: in the case of strict dominance, choosing any other strategy guarantees that the player will do worse and in the case of weak dominance, no other strategy can give a better outcome, no matter what the other players do. Unfortunately, games that have a dominant-strategy profile are not very common. What should a player do when she does not have a dominant strategy? We shall consider two iterative deletion procedures that can help solve some games.

2.5.1 IDSDS

The Iterated Deletion of Strictly Dominated Strategies (IDSDS) is the following procedure or algorithm. Given a finite ordinal strategic-form game G , let G^1 be the game obtained by removing from G , for every Player i , those strategies of Player i (if any) that are strictly dominated in G by some other strategy; let G^2 be the game obtained by removing from G^1 , for every Player i , those strategies of Player i (if any) that are strictly dominated in G^1 by some other strategy, and so on. Let G^∞ be the output of this procedure. Since the initial game G is finite, G^∞ will be obtained in a finite number of steps.

Figure 2.12 illustrates this procedure. If G^∞ contains a single strategy profile (this is **not** the case in the example of Figure 2.12), then we call that strategy profile the *iterated strict dominant-strategy profile*. If G^∞ contains two or more strategy profiles then we refer to those strategy profiles merely as the *output of the IDSDS procedure*. For example, in the game of Figure 2.12 the output of the IDSDS procedure is the set of strategy profiles $\{(A, e), (A, f), (B, e), (B, f)\}$.

What is the significance of the output of the IDSDS procedure? Consider game G of Figure 2.12. Since, for Player 2, h is strictly dominated by g , if Player 2 is rational she will not play h . Thus, if Player 1 believes that Player 2 is rational then he believes that Player 2 will not play h , that is, he restricts attention to game G^1 ; since, in G^1 , D is strictly dominated by C for Player 1, if Player 1 is rational he will not play D . It follows that if Player 2 believes that Player 1 is rational and that Player 1 believes that Player 2 is rational, then Player 2 restricts attention to game G^2 where rationality requires that Player 2 not play g , etc. It will be shown in a later chapter that if there is common knowledge of rationality,⁸ then only strategy profiles that survive the IDSDS procedure can be played; the converse is also true: any strategy profile that survives the IDSDS procedure is compatible with common knowledge of rationality.

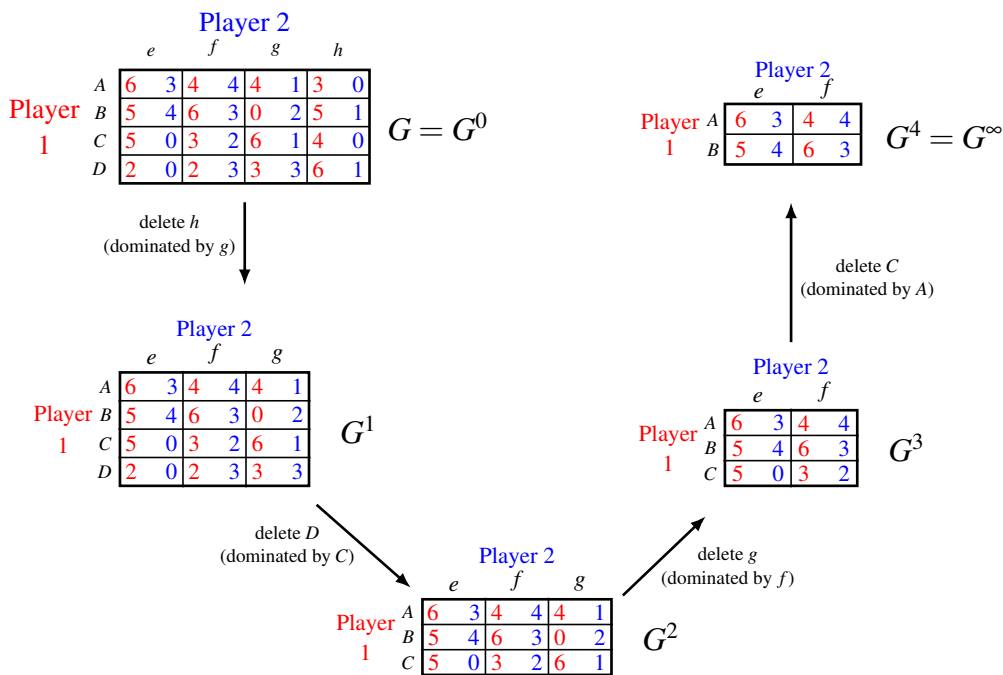


Figure 2.12: An example of the IDSDS procedure

R In finite games, the order in which strictly dominated strategies are deleted is irrelevant, in the sense that any sequence of deletions of strictly dominated strategies leads to the same output.

⁸An event E is commonly known if everybody knows E and everybody knows that everybody knows E and everybody knows that everybody knows that everybody knows E , and so on.

2.5.2 IDWDS

The *Iterated Deletion of Weakly Dominated Strategies* (IDWDS) is a weakening of IDSDS in that it allows the deletion also of *weakly* dominated strategies. However, this procedure has to be defined carefully, since in this case the order of deletion *can* matter. To see this, consider the game shown in Figure 2.13.

		Player 2	
		L	R
Player 1	A	4 0	0 0
	T	3 2	2 2
	M	1 1	0 0
	B	0 0	1 1

Figure 2.13: A strategic-form game with ordinal payoffs

Since M is strictly dominated by T for Player 1, we can delete it and obtain the reduced game shown in Figure 2.14

		Player 2	
		L	R
Player 1	A	4 0	0 0
	T	3 2	2 2
	B	0 0	1 1

Figure 2.14: The game of Figure 2.13 after deletion of strategy M

Now L is weakly dominated by R for Player 2. Deleting L we are left with the reduced game shown in Figure 2.15.

		Player 2	
		R	
Player 1	A	0 0	
	T	2 2	
	B	1 1	

Figure 2.15: The game of Figure 2.14 after deletion of strategy L

Now A and B are strictly dominated by T . Deleting them we are left with $\boxed{(T, R)}$, with corresponding payoffs $(2, 2)$.

Alternatively, going back to the game of Figure 2.13, we could note that B is strictly dominated by T ; deleting B we are left with

		Player 2			
		L	R		
Player 1	A	4	0	0	0
	T	3	2	2	2
	M	1	1	0	0

Figure 2.16: The game of Figure 2.13 after deletion of strategy B

Now R is weakly dominated by L for Player 2. Deleting R we are left with the reduced game shown in Figure 2.17.

		Player 2	
		L	
Player 1	A	4	0
	T	3	2
	M	1	1

Figure 2.17: The game of Figure 2.16 after deletion of strategy R

Now T and M are strictly dominated by A and deleting them leads to (A, L) with corresponding payoffs $(4, 0)$. Since one order of deletion leads to (T, R) with payoffs $(2, 2)$ and the other to (A, L) with payoffs $(4, 0)$, the procedure is not well defined.

Definition 2.5.1 — IDWDS. In order to avoid the problem illustrated above, the IDWDS procedure is defined as follows: *at every step identify, for every player, all the strategies that are weakly (or strictly) dominated and then delete all such strategies in that step.* If the output of the IDWDS procedure is a single strategy profile then we call that strategy profile the *iterated weak dominant-strategy profile* (otherwise we just use the expression ‘output of the IDWDS procedure’).

For example, the IDWDS procedure when applied to the game of Figure 2.13 leads to the set of strategy profiles shown in Figure 2.18.⁹

		Player 2			
		<i>L</i>	<i>R</i>		
Player 1	<i>A</i>	4	0	0	0
	<i>T</i>	3	2	2	2

Figure 2.18: The output of the IDWDS procedure applied to the game of Figure 2.13

Hence the game of Figure 2.13 does not have an iterated weak dominant-strategy profile.

The interpretation of the output of the IDWDS procedure is not as simple as that of the IDSDS procedure: certainly common knowledge of rationality is not sufficient. In order to delete weakly dominated strategies one needs to appeal not only to rationality but also to some notion of caution: a player should not completely rule out any of her opponents' strategies. However, this notion of caution is in direct conflict with the process of deletion of strategies. In this book we shall not address the issue of how to justify the IDWDS procedure.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 2.9.5 at the end of this chapter.

2.6 Nash equilibrium

Games where either the IDSDS procedure or the IDWDS procedure leads to a unique strategy profile are not very common. How can one then “solve” games that are not solved by either procedure? The notion of Nash equilibrium offers a more general alternative. We first define Nash equilibrium for a two-player game.

Definition 2.6.1 Given an ordinal game in strategic form with two players, a strategy profile $s^* = (s_1^*, s_2^*) \in S_1 \times S_2$ is a Nash equilibrium if the following two conditions are satisfied:

1. for every $s_1 \in S_1$, $\pi_1(s_1^*, s_2^*) \geq \pi_1(s_1, s_2^*)$ (or stated in terms of outcomes and preferences, $f(s_1^*, s_2^*) \succeq_1 f(s_1, s_2^*)$), and
2. for every $s_2 \in S_2$, $\pi_2(s_1^*, s_2^*) \geq \pi_2(s_1^*, s_2)$ (or, $f(s_1^*, s_2^*) \succeq_2 f(s_1^*, s_2)$).

⁹Note that, in general, the output of the IDWDS procedure is a subset of the output of the IDSDS procedure (not necessarily a proper subset). The game of Figure 2.13 happens to be a game where the two procedures yield the same outcome.

For example, in the game of Figure 2.19 there are two Nash equilibria: (T,L) and (B,C) .

		Player 2					
		L		C		R	
Player 1	T	3	2	0	0	1	1
	M	3	0	1	5	4	4
	B	1	0	2	3	3	0

Figure 2.19: A strategic-form game with ordinal payoffs

There are several possible interpretations of this definition:

“No regret” interpretation: s^* is a Nash equilibrium if there is no player who, after observing the opponent’s choice, regrets his own choice (in the sense that he could have done better with a different strategy of his, given the observed strategy of the opponent).

“Self-enforcing agreement” interpretation: imagine that the players are able to communicate before playing the game and reach a non-binding agreement expressed as a strategy profile s^* ; then no player will have an incentive to deviate from the agreement (if she believes that the other player will follow the agreement) if and only if s^* is a Nash equilibrium.

“Viable recommendation” interpretation: imagine that a third party makes a public recommendation to each player on what strategy to play; then no player will have an incentive to deviate from the recommendation (if she believes that the other players will follow the recommendation) if and only if the recommended strategy profile is a Nash equilibrium.

“Transparency of reason” interpretation: if players are all “equally rational” and Player 2 reaches the conclusion that she should play y , then Player 1 must be able to duplicate Player 2’s reasoning process and come to the same conclusion; it follows that Player 1’s choice of strategy is not rational unless it is a strategy x that is optimal against y . A similar argument applies to Player 2’s choice of strategy (y must be optimal against x) and thus (x,y) is a Nash equilibrium.

It should be clear that all of the above interpretations are mere rewording of the formal definition of Nash equilibrium in terms of the inequalities given in Definition 2.6.1.

The generalization of Definition 2.6.1 to games with more than two players is straightforward.

Definition 2.6.2 Given an ordinal game in strategic form with n players, a strategy profile $s^* \in S$ is a Nash equilibrium if the following n inequalities are satisfied: for every Player $i = 1, \dots, n$,

$$\pi_i(s^*) \geq \pi_i(s_1^*, \dots, s_{i-1}^*, s_i, s_{i+1}^*, \dots, s_n^*) \text{ for all } s_i \in S_i.$$

The reader should convince himself/herself that a (weak or strict) dominant strategy equilibrium is a Nash equilibrium and the same is true of a (weak or strict) iterated dominant-strategy profile.

Definition 2.6.3 Consider an ordinal game in strategic form, a Player i and a strategy profile $\bar{s}_{-i} \in S_{-i}$ of the players other than i . A strategy $s_i \in S_i$ of Player i is a *best reply* (or *best response*) to \bar{s}_{-i} if $\pi_i(s_i, \bar{s}_{-i}) \geq \pi_i(s'_i, \bar{s}_{-i})$, for every $s'_i \in S_i$.

For example, in the game of Figure 2.20, for Player 1 there are two best replies to L , namely M and T , while the unique best reply to C is B and the unique best reply to R is M ; for Player 2 the best reply to T is L , the best reply to M is C and the best reply to B is C .

		Player 2					
		L		C		R	
Player 1	T	3	2	0	0	1	1
	M	3	0	1	5	4	4
	B	1	0	2	3	3	0

Figure 2.20: A strategic-form game with ordinal payoffs

- R** Using the notion of best reply, an alternative definition of Nash equilibrium is as follows: $\bar{s} \in S$ is a Nash equilibrium if and only if, for every Player i , $\bar{s}_i \in S_i$ is a best reply to $\bar{s}_{-i} \in S_{-i}$.

A quick way to find the Nash equilibria of a two-player game is as follows: in each column of the table underline the largest payoff of Player 1 in that column (if there are several instances, underline them all) and in each row underline the largest payoff of Player 2 in that row; if a cell has both payoffs underlined then the corresponding strategy profile is a Nash equilibrium. Underlining of the maximum payoff of Player 1 in a given column identifies the best reply of Player 1 to the strategy of Player 2 that labels that column and similarly for Player 2. This procedure is illustrated in Figure 2.21, where there is a unique Nash equilibrium, namely (B, E) .

		Player 2							
		<i>E</i>		<i>F</i>		<i>G</i>		<i>H</i>	
Player 1	<i>A</i>	<u>4</u>	0	3	2	2	<u>3</u>	4	1
	<i>B</i>	<u>4</u>	<u>2</u>	2	1	1	<u>2</u>	0	<u>2</u>
	<i>C</i>	3	<u>6</u>	<u>5</u>	5	<u>3</u>	1	<u>5</u>	0
	<i>D</i>	2	<u>3</u>	3	2	1	2	3	<u>3</u>

Figure 2.21: A strategic-form game with ordinal payoffs

Exercise 2.3 in Section 2.9.1 explains how to represent a three-player game by means of a set of tables. In a three-player game the procedure for finding the Nash equilibria is the same, with the necessary adaptation for Player 3: in each cell underline the payoff of Player 3 if and only if her payoff is the largest of all her payoffs in the same cell across different tables. This is illustrated in Figure 2.22, where there is a unique Nash equilibrium, namely (B, R, W) .

		Player 2					
		<i>L</i>			<i>R</i>		
Player 1	<i>T</i>	0	0	<u>0</u>	2	<u>8</u>	<u>6</u>
	<i>B</i>	<u>5</u>	3	<u>2</u>	<u>3</u>	<u>4</u>	<u>2</u>

Player 3 chooses *W*

		Player 2					
		<i>L</i>			<i>R</i>		
Player 1	<i>T</i>	0	0	<u>0</u>	<u>1</u>	<u>2</u>	5
	<i>B</i>	<u>1</u>	<u>6</u>	1	0	0	1

Player 3 chooses *E*

Figure 2.22: A three-player game with ordinal payoffs

Unfortunately, when the game has too many players or too many strategies – and it is thus impossible or impractical to represent it as a set of tables – there is no quick procedure for finding the Nash equilibria: one must simply apply the definition of Nash equilibrium. This is illustrated in the following example.

■ **Example 2.1** There are 50 players. A benefactor asks them to simultaneously and secretly write on a piece of paper a request, which must be a multiple of \$10 up to a maximum of \$100 (thus the possible strategies of each player are \$10, \$20, ..., \$90, \$100). He will then proceed as follows: if not more than 10% of the players ask for \$100 then he will grant every player's request, otherwise every player will get nothing. Assume that every player is selfish and greedy (only cares about how much money she gets and prefers more money to less). What are the Nash equilibria of this game? There are several:

- every strategy profile where 7 or more players request \$100 is a Nash equilibrium (everybody gets nothing and no player can get a positive amount by unilaterally changing her request, since there will still be more than 10% requesting \$100; on the other hand, convince yourself that a strategy profile where exactly 6 players request \$100 is not a Nash equilibrium),
- every strategy profile where exactly 5 players request \$100 and the remaining players request \$90 is a Nash equilibrium.

Any other strategy profile is not a Nash equilibrium: (1) if fewer than 5 players request \$100, then a player who requested less than \$100 can increase her payoff by switching to a request of \$100, (2) if exactly 5 players request \$100 and among the remaining players there is one who is not requesting \$90, then that player can increase her payoff by increasing her request to \$90. ■

We conclude this section by noting that, since so far we have restricted attention to ordinal games, there is no guarantee that an arbitrary game will have at least one Nash equilibrium. An example of a game that has no Nash equilibria is the *Matching Pennies* game. This is a simultaneous two-player game where each player has a coin and decides whether to show the Heads face or the Tails face. If both choose *H* or both choose *T* then Player 1 wins, otherwise Player 2 wins. Each player strictly prefers the outcome where she herself wins to the alternative outcome. The game is illustrated in Figure 2.23.

		Player 2			
		<i>H</i>	<i>T</i>		
Player 1	<i>H</i>	1	0	0	1
	<i>T</i>	0	1	1	0

Figure 2.23: The matching pennies game

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 2.9.6 at the end of this chapter.

2.7 Games with infinite strategy sets

Games where the strategy set of one or more players is infinite cannot be represented using a table or set of tables. However, all the concepts introduced in this chapter can still be applied. In this section we will focus on the notion of Nash equilibrium. We start with an example.

■ **Example 2.2** There are two players. Each player has to write down a real number (not necessarily an integer) greater than or equal to 1; thus the strategy sets are $S_1 = S_2 = [1, \infty)$. Payoffs are as follows (π_1 is the payoff of Player 1, π_2 the payoff of Player 2, x is the number written by Player 1 and y the number written by Player 2):

$$\pi_1(x, y) = \begin{cases} x - 1 & \text{if } x < y \\ 0 & \text{if } x \geq y \end{cases} \quad \text{and} \quad \pi_2(x, y) = \begin{cases} y - 1 & \text{if } x > y \\ 0 & \text{if } x \leq y \end{cases}$$

What are the Nash equilibria of this game? ■

There is only one Nash equilibrium, namely (1,1) with payoffs (0,0). First of all, we must show that (1,1) is indeed a Nash equilibrium.

If Player 1 switched to some $x > 1$ then her payoff would remain 0: $\pi_1(x, 1) = 0$, for all $x \in [1, \infty)$ and the same is true for Player 2 if he unilaterally switched to some $y > 1$: $\pi_2(1, y) = 0$, for all $y \in [1, \infty)$.

Now we show that no other pair (x, y) is a Nash equilibrium.

Consider first an arbitrary pair (x, y) with $x = y > 1$. Then $\pi_1(x, y) = 0$, but if Player 1 switched to an \hat{x} strictly between 1 and x ($1 < \hat{x} < x$) her payoff would be $\pi_1(\hat{x}, y) = \hat{x} - 1 > 0$ (recall that, by hypothesis, $x = y$).

Now consider an arbitrary (x, y) with $x < y$. Then $\pi_1(x, y) = x - 1$, but if Player 1 switched to an \hat{x} strictly between x and y ($x < \hat{x} < y$) her payoff would be $\pi_1(\hat{x}, y) = \hat{x} - 1 > x - 1$.

The argument for ruling out pairs (x, y) with $y < x$ is similar.

Note the interesting fact that, for Player 1, $x = 1$ is a weakly dominated strategy: indeed it is weakly dominated by any other strategy: $x = 1$ guarantees a payoff of 0 for Player 1, while any $\hat{x} > 1$ would yield a positive payoff to Player 1 in some cases (against any $y > \hat{x}$) and 0 in the remaining cases. The same is true for Player 2. Thus in this game there is *a unique Nash equilibrium where the strategy of each player is weakly dominated!*

[Note: the rest of this section makes use of calculus. The reader who is not familiar with calculus should skip this part.]

We conclude this section with an example based on the analysis of competition among firms proposed by Augustine Cournot in a book published in 1838. In fact, Cournot is the one who invented what we now call ‘Nash equilibrium’, although his analysis was restricted to a small class of games. Consider $n \geq 2$ firms which produce an identical product. Let q_i be the quantity produced by Firm i ($i = 1, \dots, n$). For Firm i the cost of producing q_i units of output is $c_i q_i$, where c_i is a positive constant. For simplicity we will restrict attention to the case of two firms ($n = 2$) and identical cost functions: $c_1 = c_2 = c$. Let Q be total industry output, that is, $Q = q_1 + q_2$. The price at which each firm can sell each unit of output is given by the inverse demand function $P = a - bQ$ where a and b are positive constants. Cournot assumed that each firm was only interested in its own profit and preferred higher profit to lower profit (that is, each firm is “selfish and greedy”).

The profit function of Firm 1 is given by

$$\pi_1(q_1, q_2) = Pq_1 - cq_1 = [a - b(q_1 + q_2)]q_1 - cq_1 = (a - c)q_1 - b(q_1)^2 - bq_1q_2.$$

Similarly, the profit function of Firm 2 is given by

$$\pi_2(q_1, q_2) = (a - c)q_2 - b(q_2)^2 - bq_1q_2$$

Cournot defined an equilibrium as a pair (\bar{q}_1, \bar{q}_2) that satisfies the following two inequalities:

$$\pi_1(\bar{q}_1, \bar{q}_2) \geq \pi_1(q_1, \bar{q}_2), \text{ for every } q_1 \geq 0 \quad (\clubsuit)$$

$$\pi_2(\bar{q}_1, \bar{q}_2) \geq \pi_2(\bar{q}_1, q_2) \text{ for every } q_2 \geq 0. \quad (\diamond)$$

Of course, this is the same as saying that (\bar{q}_1, \bar{q}_2) is a Nash equilibrium of the game where the players are the two firms, the strategy sets are $S_1 = S_2 = [0, \infty)$ and the payoff functions are the profit functions. How do we find a Nash equilibrium? First of all, note that the profit functions are differentiable. Secondly note that (\clubsuit) says that, having fixed the value of q_2 at \bar{q}_2 , the function $\pi_1(q_1, \bar{q}_2)$ – viewed as a function of q_1 alone – is maximized at the point $q_1 = \bar{q}_1$. A necessary condition for this (if $\bar{q}_1 > 0$) is that the derivative of this function be zero at the point $q_1 = \bar{q}_1$, that is, it must be that $\frac{\partial \pi_1}{\partial q_1}(\bar{q}_1, \bar{q}_2) = 0$. This condition is also sufficient since the second derivative of this function is always negative ($\frac{\partial^2 \pi_1}{\partial q_1^2}(q_1, q_2) = -2b$ for every (q_1, q_2)). Similarly, by (\diamond) , it must be that $\frac{\partial \pi_2}{\partial q_2}(\bar{q}_1, \bar{q}_2) = 0$. Thus the Nash equilibrium is found by solving the system of two equations

$$\begin{cases} \frac{\partial \pi_1}{\partial q_1}(q_1, q_2) = a - c - 2bq_1 - bq_2 = 0 \\ \frac{\partial \pi_2}{\partial q_2}(q_1, q_2) = a - c - 2bq_2 - bq_1 = 0 \end{cases}$$

The solution is $\bar{q}_1 = \bar{q}_2 = \frac{a-c}{3b}$. The corresponding price is $\bar{P} = a - b\left(2\frac{a-c}{3b}\right) = \frac{a+2c}{3}$ and the corresponding profits are $\pi_1\left(\frac{a-c}{3b}, \frac{a-c}{3b}\right) = \pi_2\left(\frac{a-c}{3b}, \frac{a-c}{3b}\right) = \frac{(a-c)^2}{9b}$.

For example, if $a = 25$, $b = 2$, $c = 1$ then the Nash equilibrium is given by (4,4) with corresponding profits of 32 for each firm. The analysis can easily be extended to the case of more than two firms.

The reader who is interested in further exploring the topic of competition among firms can consult any textbook on Industrial Organization.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 2.9.7 at the end of this chapter.

2.8 Proofs of theorems

Theorem [Vickrey, 1961] In a second-price auction, if Player i is selfish and greedy then it is a weakly dominant strategy for Player i to bid her true value, that is, to choose $b_i = v_i$.

Proof. In order to make the notation simpler and the argument more transparent, we give the proof for the case where $n = 2$. We shall prove that bidding v_1 is a weakly dominant strategy for Player 1 (the proof for Player 2 is similar). Assume that Player 1 is selfish and greedy. Then we can take her payoff function to be as follows:

$$\pi_1(b_1, b_2) = \begin{cases} v_1 - b_2 & \text{if } b_1 \geq b_2 \\ 0 & \text{if } b_1 < b_2 \end{cases}$$

We need to show that, whatever bid Player 2 submits, Player 1 cannot get a higher payoff by submitting a bid different from v_1 . Two cases are possible (recall that b_2 denotes the actual bid of Player 2, which is unknown to Player 1).

Case 1: $b_2 \leq v_1$. In this case, bidding v_1 makes Player 1 the winner and his payoff is $v_1 - b_2 \geq 0$. Consider a different bid b_1 . If $b_1 \geq b_2$ then Player 1 is still the winner and his payoff is still $v_1 - b_2 \geq 0$; thus such a bid is as good as (hence not better than) v_1 . If $b_1 < b_2$ then the winner is Player 2 and Player 1 gets a payoff of 0. Thus such a bid is also not better than v_1 .

Case 2: $b_2 > v_1$. In this case, bidding v_1 makes Player 2 the winner and thus Player 1 gets a payoff of 0. Any other bid $b_1 < b_2$ gives the same outcome and payoff. On the other hand, any bid $b_1 \geq b_2$ makes Player 1 the winner, giving him a payoff of $v_1 - b_2 < 0$, thus making Player 1 worse off than with a bid of v_1 .



Theorem [Clarke, 1971] In the pivotal mechanism (under the assumed preferences) truthful revelation (that is, stating $w_i = v_i$) is a weakly dominant strategy for every Player i .

Proof. Consider an individual i and possible statements w_j for $j \neq i$. Several cases are possible.

Case 1: $\sum_{j \neq i} w_j > \sum_{j \neq i} c_j$ and $v_i + \sum_{j \neq i} w_j > c_i + \sum_{j \neq i} c_j = C$. Then

	decision	i 's tax	i 's utility
if i states v_i	Yes	0	$\bar{m}_i + v_i - c_i$
if i states w_i such that $w_i + \sum_{j \neq i} w_j > C$	Yes	0	$\bar{m}_i + v_i - c_i$
if i states w_i such that $w_i + \sum_{j \neq i} w_j \leq C$	No	$\sum_{j \neq i} w_j - \sum_{j \neq i} c_j$	$\bar{m}_i - \left(\sum_{j \neq i} w_j - \sum_{j \neq i} c_j \right)$

Individual i cannot gain by lying if and only if

$$\bar{m}_i + v_i - c_i \geq \bar{m}_i - \left(\sum_{j \neq i} w_j - \sum_{j \neq i} c_j \right), \quad \text{i.e. if and only if} \quad v_i + \sum_{j \neq i} w_j \geq C,$$

which is true by our hypothesis.

Case 2: $\sum_{j \neq i} w_j > \sum_{j \neq i} c_j$ and $v_i + \sum_{j \neq i} w_j \leq c_i + \sum_{j \neq i} c_j = C$. Then

	decision	i 's tax	i 's utility
if i states v_i	No	$\sum_{j \neq i} w_j - \sum_{j \neq i} c_j$	$\bar{m}_i - \left(\sum_{j \neq i} w_j - \sum_{j \neq i} c_j \right)$
if i states w_i such that $w_i + \sum_{j \neq i} w_j \leq C$	No	$\sum_{j \neq i} w_j - \sum_{j \neq i} c_j$	$\bar{m}_i - \left(\sum_{j \neq i} w_j - \sum_{j \neq i} c_j \right)$
if i states w_i such that $w_i + \sum_{j \neq i} w_j > C$	Yes	0	$\bar{m}_i + v_i - c_i$

Individual i cannot gain by lying if and only if $\bar{m}_i - \left(\sum_{j \neq i} w_j - \sum_{j \neq i} c_j \right) \geq \bar{m}_i + v_i - c_i$,
i.e. if and only if $v_i + \sum_{j \neq i} w_j \leq C$, which is true by our hypothesis.

Case 3: $\sum_{j \neq i} w_j \leq \sum_{j \neq i} c_j$ and $v_i + \sum_{j \neq i} w_j \leq c_i + \sum_{j \neq i} c_j = C$. Then

	decision	i 's tax	i 's utility
if i states v_i	No	0	\bar{m}_i
if i states w_i such that $w_i + \sum_{j \neq i} w_j \leq C$	No	0	\bar{m}_i
if i states w_i such that $w_i + \sum_{j \neq i} w_j > C$	Yes	$\left(\sum_{j \neq i} c_j - \sum_{j \neq i} w_j \right)$ (recall that $\sum_{j \neq i} w_j \leq \sum_{j \neq i} c_j$)	$\bar{m}_i + v_i - c_i - \left(\sum_{j \neq i} c_j - \sum_{j \neq i} w_j \right)$

Individual i cannot gain by lying if and only if $\bar{m}_i \geq \bar{m}_i + v_i - c_i - \left(\sum_{j \neq i} c_j - \sum_{j \neq i} w_j \right)$,
i.e. if and only if $v_i + \sum_{j \neq i} w_j \leq C$, which is true by our hypothesis.

Case 4: $\sum_{j \neq i} w_j \leq \sum_{j \neq i} c_j$ and $v_i + \sum_{j \neq i} w_j > c_i + \sum_{j \neq i} c_j = C$. Then

	decision	i 's tax	i 's utility
if i states v_i	Yes	$\left(\sum_{j \neq i} c_j - \sum_{j \neq i} w_j \right)$ (recall that $\sum_{j \neq i} w_j \leq \sum_{j \neq i} c_j$)	$\bar{m}_i + v_i - c_i - \left(\sum_{j \neq i} w_j - \sum_{j \neq i} c_j \right)$
if i states w_i such that $w_i + \sum_{j \neq i} w_j > C$	Yes	$\left(\sum_{j \neq i} c_j - \sum_{j \neq i} w_j \right)$ (recall that $\sum_{j \neq i} w_j \leq \sum_{j \neq i} c_j$)	$\bar{m}_i + v_i - c_i - \left(\sum_{j \neq i} w_j - \sum_{j \neq i} c_j \right)$
if i states w_i such that $w_i + \sum_{j \neq i} w_j \leq C$	No	0	\bar{m}_i

Individual i cannot gain by lying if and only if $\bar{m}_i + v_i - c_i - \left(\sum_{j \neq i} w_j - \sum_{j \neq i} c_j \right) \geq \bar{m}_i$,
i.e. if and only if $v_i + \sum_{j \neq i} w_j \geq C$, which is true by our hypothesis.

Since we have covered all the possible cases, the proof is complete. ■

2.9 Exercises

2.9.1 Exercises for Section 2.1: Game frames and games

The answers to the following exercises are in Section 2.10 at the end of this chapter.

Exercise 2.1 Antonia and Bob cannot decide where to go to dinner. Antonia proposes the following procedure: she will write on a piece of paper either the number 2 or the number 4 or the number 6, while Bob will write on his piece of paper either the number 1 or 3 or 5. They will write their numbers secretly and independently. They then will show each other what they wrote and choose a restaurant according to the following rule: if the sum of the two numbers is 5 or less, they will go to a Mexican restaurant, if the sum is 7 they will go to an Italian restaurant and if the number is 9 or more they will go to a Japanese restaurant.

(a) Let Antonia be Player 1 and Bob Player 2. Represent this situation as a game frame, first by writing out each element of the quadruple of Definition 2.1.1 and then by using a table (label the rows with Antonia's strategies and the columns with Bob's strategies, so that we can think of Antonia as choosing the row and Bob as choosing the column).

(b) Suppose that Antonia and Bob have the following preferences (where M stands for 'Mexican', I for 'Italian' and J for 'Japanese'):

for Antonia: $M \succ_{Antonia} I \succ_{Antonia} J$; for Bob: $I \succ_{Bob} M \succ_{Bob} J$.

Using utility function with values 1, 2 and 3 represent the corresponding reduced-form game as a table.

Exercise 2.2 Consider the following two-player game-frame where each player is given a set of cards and each card has a number on it. The players are Antonia (Player 1) and Bob (Player 2). Antonia's cards have the following numbers (one number on each card): 2, 4 and 6, whereas Bob's cards are marked 0, 1 and 2 (thus different numbers from the previous exercise). Antonia chooses one of her own cards and Bob chooses one of his own cards: this is done without knowing the other player's choice. The outcome depends on the sum of the points of the chosen cards. If the sum of the points on the two chosen cards is greater than or equal to 5, Antonia gets \$10 minus that sum; otherwise (that is, if the sum is less than 5) she gets nothing; furthermore, if the sum of points is an odd number, Bob gets as many dollars as that sum; if the sum of points turns out to be an even number and is less than or equal to 6, Bob gets \$2; otherwise he gets nothing. (The money comes from a third party.)

(a) Represent the game-frame described above by means of a table. As in the previous exercise, assign the rows to Antonia and the columns to Bob.

(b) Using the game-frame of part (a) obtain a reduced-form game by adding the information that each player is selfish and greedy. This means that each player only cares about how much money he/she gets and prefers more money to less.

Exercise 2.3 Alice (Player 1), Bob (Player 2), and Charlie (Player 3) play the following simultaneous game. They are sitting in different rooms facing a keyboard with only one key and each has to decide whether or not to press the key. Alice wins if the number of people who press the key is odd (that is, all three of them or only Alice or only Bob or only Charlie), Bob wins if exactly two people (he may be one of them) press the key and Charlie wins if nobody presses the key.

- (a) Represent this situation as a game-frame. Note that we can represent a three-player game with a *set of tables*: Player 1 chooses the row, Player 2 chooses the column and Player 3 chooses the table (that is, we label the rows with Player 1's strategies, the columns with Player 2's strategies and the tables with Player 3's strategies).
- (b) Using the game-frame of part (a) obtain a reduced-form game by adding the information that each player prefers winning to not winning and is indifferent between any two outcomes where he/she does not win. For each player use a utility function with values from the set $\{0, 1\}$.
- (c) Using the game-frame of part (a) obtain a reduced-form game by adding the information that (1) each player prefers winning to not winning, (2) Alice is indifferent between any two outcomes where she does not win, (3) conditional on not winning, Bob prefers if Charlie wins rather than Alice, (4) conditional on not winning, Charlie prefers if Bob wins rather than Alice. For each player use a utility function with values from the set $\{0, 1, 2\}$.

2.9.2 Exercises for Section 2.2: Strict/weak dominance

The answers to the following exercises are in Section 2.10 at the end of this chapter.

Exercise 2.4 There are two players. Each player is given an unmarked envelope and asked to put in it either nothing or \$300 of his own money or \$600 of his own money. A referee collects the envelopes, opens them, gathers all the money, then adds 50% of that amount (using his own money) and divides the total into two equal parts which he then distributes to the players.

- (a) Represent this game frame with two alternative tables: the first table showing in each cell the amount of money distributed to Player 1 and the amount of money distributed to Player 2, the second table showing the change in wealth of each player (money received minus contribution).
- (b) Suppose that Player 1 has some animosity towards the referee and ranks the outcomes in terms of how much money the referee loses (the more, the better), while Player 2 is selfish and greedy and ranks the outcomes in terms of her own net gain. Represent the corresponding game using a table.
- (c) Is there a strict dominant-strategy profile?

Exercise 2.5 Consider again the game of Part (b) of Exercise 2.1.

- (a) Determine, for each player, whether the player has *strictly* dominated strategies.
- (b) Determine, for each player, whether the player has *weakly* dominated strategies.

Exercise 2.6 There are three players. Each player is given an unmarked envelope and asked to put in it either nothing or \$3 of his own money or \$6 of his own money. A referee collects the envelopes, opens them, gathers all the money and then doubles the amount (using his own money) and divides the total into three equal parts which he then distributes to the players.

For example, if Players 1 and 2 put nothing and Player 3 puts \$6, then the referee adds another \$6 so that the total becomes \$12, divides this sum into three equal parts and gives \$4 to each player.

Each player is selfish and greedy, in the sense that he ranks the outcomes exclusively in terms of his net change in wealth (what he gets from the referee minus what he contributed).

- (a) Represent this game by means of a set of tables. (Do not treat the referee as a player.)
- (b) For each player and each pair of strategies determine if one of the two dominates the other and specify if it is weak or strict dominance.
- (c) Is there a strict dominant-strategy profile?

2.9.3 Exercises for Section 2.3: Second price auction

The answers to the following exercises are in Section 2.10 at the end of this chapter.

Exercise 2.7 For the second-price auction partially illustrated in Figure 2.10 – reproduced below (recall that the numbers are the payoffs of Player 1 only) – complete the representation by adding the payoffs of Player 2, assuming that Player 2 assigns a value of \$50M to the field and, like Player 1, ranks the outcomes in terms of the net gain from the oil field (defined as profits minus the price paid, if Player 2 wins, and zero otherwise).

		Player 2				
		\$10M	\$20M	\$30M	\$40M	\$50M
Player 1	\$10M	0	0	0	0	0
	\$20M	20	0	0	0	0
	\$30M	20	10	0	0	0
	\$40M	20	10	0	0	0
	\$50M	20	10	0	-10	0

Exercise 2.8 Consider the following “third-price” auction. There are $n \geq 3$ bidders. A single object is auctioned and Player i values the object $\$v_i$, with $v_i > 0$. The bids are simultaneous and secret.

The utility of Player i is: 0 if she does not win and $(v_i - p)$ if she wins and pays $\$p$.

Every non-negative number is an admissible bid. Let b_i denote the bid of Player i .

The winner is the highest bidder. In case of ties the bidder with the lowest index among those who submitted the highest bid wins (e.g. if the highest bid is $\$120$ and it is submitted by players 6, 12 and 15, then the winner is Player 6). The losers don't get anything and don't pay anything. The winner gets the object and pays the **third** highest bid, which is defined as follows.

Let i be the winner and fix a Player j such that

$$b_j = \max(\{b_1, \dots, b_n\} \setminus \{b_i\})$$

[note: if

$$\max(\{b_1, \dots, b_n\} \setminus \{b_i\})$$

contains more than one element, then we pick any one of them]. Then the third price is defined as

$$\max(\{b_1, \dots, b_n\} \setminus \{b_i, b_j\}).$$

For example, if $n = 3$ and the bids are $b_1 = 30$, $b_2 = 40$ and $b_3 = 40$ then the winner is Player 2 and she pays $\$30$. If $b_1 = b_2 = b_3 = 50$ then the winner is Player 1 and she pays $\$50$. For simplicity, let us restrict attention to the case where $n = 3$ and $v_1 > v_2 > v_3 > 0$. Does Player 1 have a weakly dominant strategy in this auction? ■

2.9.4 Exercises for Section 2.4: The pivotal mechanism

The answers to the following exercises are in Section 2.10 at the end of this chapter.

Exercise 2.9 The pivotal mechanism is used to decide whether a new park should be built. There are 5 individuals. According to the proposed project, the cost of the park would be allocated as follows:

Individual	1	2	3	4	5
Share of cost	$c_1 = \$30$	$c_2 = \$25$	$c_3 = \$25$	$c_4 = \$15$	$c_5 = \$5$

For every individual $i = 1, \dots, 5$, let v_i be the perceived gross benefit (if positive; perceived gross loss, if negative) from having the park built. The v_i 's are as follows:

Individual	1	2	3	4	5
Gross benefit	$v_1 = \$60$	$v_2 = \$15$	$v_3 = \$55$	$v_4 = -\$25$	$v_5 = -\$20$

(Thus the net benefit (loss) to individual i is $v_i - c_i$). Individual i has the following utility of wealth function (where m_i denotes the wealth of individual i):

$$U_i(\$m_i) = \begin{cases} m_i & \text{if the project is not carried out} \\ m_i + v_i & \text{if the project is carried out} \end{cases}$$

Let \bar{m}_i be the initial endowment of money of individual i and assume that \bar{m}_i is large enough that it exceeds c_i plus any tax that the individual might have to pay.

(a) What is the Pareto-efficient decision: to build the park or not?

Assume that the pivotal mechanism is used, so that each individual i is asked to state a number w_i which is going to be interpreted as the gross benefit to individual i from carrying out the project. There are no restrictions on the number w_i : it can be positive, negative or zero. Suppose that the individuals make the following announcements:

Individual	1	2	3	4	5
Stated benefit	$w_1 = \$70$	$w_2 = \$10$	$w_3 = \$65$	$w_4 = -\$30$	$w_5 = -\$5$

(b) Would the park be built based on the above announcements?

(c) Using the above announcements and the rules of the pivotal mechanism, fill in the following table:

Individual	1	2	3	4	5
Pivotal?					
Tax					

- (d) As you know, in the pivotal mechanism each individual has a dominant strategy. If all the individuals played their dominant strategies, would the park be built?
- (e) Assuming that all the individuals play their dominant strategies, find out who is pivotal and what tax (if any) each individual has to pay?
- (f) Show that if every other individual reports his/her true benefit, then it is best for Individual 1 to also report his/her true benefit.

2.9.5 Exercises for Section 2.5: Iterated deletion procedures

The answers to the following exercises are in Section 2.10 at the end of this chapter.

Exercise 2.10 Consider again the game of Part (b) of Exercise 2.1.

- (a) Apply the IDSDS procedure (Iterated Deletion of Strictly Dominated Strategies).
- (b) Apply the IDWDS procedure (Iterated Deletion of Weakly Dominated Strategies).

Exercise 2.11 Apply the IDSDS procedure to the game shown in Figure 2.24. Is there a strict iterated dominant-strategy profile?

		Player 2					
		<i>d</i>		<i>e</i>		<i>f</i>	
Player 1	<i>a</i>	8	6	0	9	3	8
	<i>b</i>	3	2	2	1	4	3
	<i>c</i>	2	8	1	5	3	1

Figure 2.24: A strategic-form game with ordinal payoffs

Exercise 2.12 Consider the following game. There is a club with three members: Ann, Bob and Carla. They have to choose which of the three is going to be president next year. Currently, Ann is the president. Each member is both a candidate and a voter. Voting is as follows: each member votes for one candidate (voting for oneself is allowed); if two or more people vote for the same candidate then that person is chosen as the next president; if there is complete disagreement, in the sense that there is exactly one vote for each candidate, then the person for whom Ann voted is selected as the next president.

- Represent this voting procedure as a game frame, indicating inside each cell of each table which candidate is elected.
- Assume that the players' preferences are as follows: $Ann \succ_{Ann} Carla \succ_{Ann} Bob$, $Carla \succ_{Bob} Bob \succ_{Bob} Ann$, $Bob \succ_{Carla} Ann \succ_{Carla} Carla$. Using utility values 0, 1 and 2, convert the game frame into a game.
- Apply the IDWDS to the game of Part (b). Is there an iterated weak dominant-strategy profile?
- Does the extra power given to Ann (in the form of tie-breaking in case of complete disagreement) benefit Ann?

Exercise 2.13 Consider the game shown in Figure 2.25.

- Apply the IDSDS procedure. Is there an iterated strict dominant-strategy profile?
- Apply the IDWDS procedure. Is there an iterated weak dominant-strategy profile?

		Player 2					
		<i>D</i>		<i>E</i>		<i>F</i>	
Player 1	<i>a</i>	2	3	2	2	3	1
	<i>b</i>	2	0	3	1	1	0
	<i>c</i>	1	4	2	0	0	4

Figure 2.25: A strategic-form game with ordinal payoffs

2.9.6 Exercises for Section 2.6: Nash equilibrium

The answers to the following exercises are in Section 2.10 at the end of this chapter.

Exercise 2.14 Find the Nash equilibria of the game of Exercise 2.2. ■

Exercise 2.15 Find the Nash equilibria of the games of Exercise 2.3 (b) and (c). ■

Exercise 2.16 Find the Nash equilibria of the game of Exercise 2.4 (b). ■

Exercise 2.17 Find the Nash equilibria of the game of Exercise 2.6. ■

Exercise 2.18 Find the Nash equilibria of the game of Exercise 2.7. ■

Exercise 2.19 Find a Nash equilibrium of the game of Exercise 2.8 for the case where

$$n = 3 \quad \text{and} \quad v_1 > v_2 > v_3 > 0$$

(there are several Nash equilibria: you don't need to find them all). ■

Exercise 2.20 Find the Nash equilibria of the game of Exercise 2.12 (b). ■

Exercise 2.21 Find the Nash equilibria of the game of Exercise 2.13. ■

2.9.7 Exercises for Section 2.7: Games with infinite strategy sets

The answers to the following exercises are in Section 2.10 at the end of this chapter.

Exercise 2.22 Consider a simultaneous n -player game where each Player i chooses an effort level $a_i \in [0, 1]$. The payoff to Player i is given by

$$\pi_i(a_1, \dots, a_n) = 4 \min\{a_1, \dots, a_n\} - 2a_i$$

(interpretation: efforts are complementary and each player's cost per unit of effort is 2).

- (a) Find all the Nash equilibria and prove that they are indeed Nash equilibria.
- (b) Are any of the Nash equilibria Pareto efficient?
- (c) Find a Nash equilibrium where each player gets a payoff of 1. ■

Exercise 2.23 — ★★★ Challenging Question ★★★.

The Mondevil Corporation operates a chemical plant, which is located on the banks of the Sacramento river. Downstream from the chemical plant is a group of fisheries. The Mondevil plant emits by-products that pollute the river, causing harm to the fisheries. The profit Mondevil obtains from operating the chemical plant is $\$ \Pi > 0$.

The harm inflicted on the fisheries due to water pollution is equal to $\$ L > 0$ of lost profit [without pollution the fisheries' profit is $\$ A$, while with pollution it is $\$(A - L)$]. Suppose that the fisheries collectively sue the Mondevil Corporation. It is easily verified in court that Mondevil's plant pollutes the river. However, the values of Π and L cannot be verified by the court, although they are commonly known to the litigants.

Suppose that the court requires the Mondevil attorney (Player 1) and the fisheries' attorney (Player 2) to play the following litigation game. Player 1 is asked to announce a number $x \geq 0$, which the court interprets as a claim about the plant's profits. Player 2 is asked to announce a number $y \geq 0$, which the court interprets as the fisheries' claim about their profit loss. The announcements are made simultaneously and independently.

Then the court uses Posner's nuisance rule to make its decision (R. Posner, *Economic Analysis of Law*, 9th edition, 1997). According to the rule, if $y > x$, then Mondevil must shut down its chemical plant. If $x \geq y$, then the court allows Mondevil to operate the plant, but the court also requires Mondevil to pay the fisheries the amount y . Note that the court cannot force the attorneys to tell the truth: in fact, it would not be able to tell whether or not the lawyers were reporting truthfully. Assume that the attorneys want to maximize the payoff (profits) of their clients.

- (a) Represent this situation as a strategic-form game by describing the strategy set of each player and the payoff functions.
- (b) Is it a dominant strategy for the Mondevil attorney to make a truthful announcement (i.e. to choose $x = \Pi$)? [Prove your claim.]
- (c) Is it a dominant strategy for the fisheries' attorney to make a truthful announcement (i.e. to choose $y = L$)? [Prove your claim.]
- (d) For the case where $\Pi > L$ (recall that Π and L denote the true amounts), find **all** the Nash equilibria of the litigation game. [Prove that what you claim to be Nash equilibria are indeed Nash equilibria and that there are no other Nash equilibria.]
- (e) For the case where $\Pi < L$ (recall that Π and L denote the true amounts), find **all** the Nash equilibria of the litigation game. [Prove that what you claim to be Nash equilibria are indeed Nash equilibria and that there are no other Nash equilibria.]
- (f) Does the court rule give rise to a Pareto efficient outcome? [Assume that the players end up playing a Nash equilibrium.]

2.10 Solutions to exercises

Solution to Exercise 2.1.

- (a) $I = \{1,2\}$, $S_1 = \{2,4,6\}$, $S_2 = \{1,3,5\}$, $O = \{M,I,J\}$
 (where M stands for ‘Mexican’, I for ‘Italian’ and J for ‘Japanese’).

The set of strategy profiles is

$$S = \{(2,1), (2,3), (2,5), (4,1), (4,3), (4,5), (6,1), (6,3), (6,5)\};$$

the outcome function is:

$$f(2,1) = f(2,3) = f(4,1) = M,$$

$$f(2,5) = f(4,3) = f(6,1) = I \text{ and}$$

$$f(4,5) = f(6,3) = f(6,5) = J.$$

The representation as a table is shown in Figure 2.26.

		Player 2 (Bob)		
		1	3	5
Player 1 (Antonia)	2	M	M	I
	4	M	I	J
	6	I	J	J

Figure 2.26: The game-frame for Exercise 2.1 (a)

- (b) Using values 1, 2 and 3, the utility functions are as follows, where U_1 is the utility function of Player 1 (Antonia) and U_2 is the utility function of Player 2 (Bob):

$$\begin{pmatrix} & M & I & J \\ U_1: & 3 & 2 & 1 \\ U_2: & 2 & 3 & 1 \end{pmatrix}$$

The reduced-form game is shown in Figure 2.27. □

		Player 2 (Bob)					
		1	3	5			
Player 1 (Antonia)	2	3	2	3	2	2	3
	4	3	2	2	3	1	1
	6	2	3	1	1	1	1

Figure 2.27: The game for Exercise 2.1 (b)

Solution to Exercise 2.2.

- (a) The game-frame is shown in Figure 2.28.
- (b) When the outcomes are sums of money and Player i is selfish and greedy, we can take the following as Player i 's utility function: $U_i(\$x) = x$ (other utility functions would do too: the only requirement is that the utility of a larger sum of money is larger than the utility of a smaller sum of money). Thus the reduced-form game is shown in Figure 2.29. \square

		Player 2 (Bob)		
		0	1	2
Player 1 (Antonia)	2	Antonia gets nothing Bob gets \$2	Antonia gets nothing Bob gets \$3	Antonia gets nothing Bob gets \$2
	4	Antonia gets nothing Bob gets \$2	Antonia gets \$5 Bob gets \$5	Antonia gets \$4 Bob gets \$2
	6	Antonia gets \$4 Bob gets \$2	Antonia gets \$3 Bob gets \$7	Antonia gets \$2 Bob gets nothing

Figure 2.28: The game-frame for Exercise 2.2 (a)

		Player 2 (Bob)		
		0	1	2
Player 1 (Antonia)	2	0 2	0 3	0 2
	4	0 2	5 5	4 2
	6	4 2	3 7	2 0

Figure 2.29: The game for Exercise 2.2 (b)

Solution to Exercise 2.3.

- (a) The game-frame is shown in Figure 2.30.
- (b) The reduced-form game is shown in Figure 2.31.
- (c) The reduced-form game is shown in Figure 2.32. For Alice we chose 1 and 0 as utilities, but one could also use 2 and 1 or 2 and 0. \square

		Bob	
		<i>P</i>	<i>not P</i>
Alice	<i>P</i>	Alice wins	Bob wins
	<i>not P</i>	Bob wins	Alice wins

Charlie: *P*

		Bob	
		<i>P</i>	<i>not P</i>
Alice	<i>P</i>	Bob wins	Alice wins
	<i>not P</i>	Alice wins	Charlie wins

Charlie: *not P*

Figure 2.30: The game-frame for Exercise 2.3 (a)

		Bob	
		<i>P</i>	<i>not P</i>
Alice	<i>P</i>	1 0 0	0 1 0
	<i>not P</i>	0 1 0	1 0 0

Charlie: *P*

		Bob	
		<i>P</i>	<i>not P</i>
Alice	<i>P</i>	0 1 0	1 0 0
	<i>not P</i>	1 0 0	0 0 1

Charlie: *not P*

Figure 2.31: The reduced-form game for Exercise 2.3 (b)

		Bob	
		<i>P</i>	<i>not P</i>
Alice	<i>P</i>	1 0 0	0 2 1
	<i>not P</i>	0 2 1	1 0 0

Charlie: *P*

		Bob	
		<i>P</i>	<i>not P</i>
Alice	<i>P</i>	0 2 1	1 0 0
	<i>not P</i>	1 0 0	0 1 2

Charlie: *not P*

Figure 2.32: The reduced-form game for Exercise 2.3 (c)

Solution to Exercise 2.4.

- (a) The tables are shown in Figure 2.33.
- (b) For Player 1 we can take as his payoff the total money lost by the referee and for Player 2 her own net gain as shown in Figure 2.34.
- (c) For Player 1 contributing \$600 is a strictly dominant strategy and for Player 2 contributing \$0 is a strictly dominant strategy. Thus $(\$600, \$0)$ is the strict dominant-strategy profile. \square

Distributed money:

		Player 2					
		0		300		600	
Player 1	0	0	0	225	225	450	450
	300	225	225	450	450	675	675
	600	450	450	675	675	900	900

.....

Net amounts:

		Player 2					
		0		300		600	
Player 1	0	0	0	225	-75	450	-150
	300	-75	225	150	150	375	75
	600	-150	450	75	375	300	300

Figure 2.33: The tables for Exercise 2.4 (a)

		Player 2					
		0		300		600	
Player 1	0	0	0	150	-75	300	-150
	300	150	225	300	150	450	75
	600	300	450	450	375	600	300

Figure 2.34: The game for Exercise 2.4 (b)

Solution to Exercise 2.5. The game under consideration is shown in Figure 2.35.

- (a) For Player 1, 6 is strictly dominated by 2. There is no other strategy which is strictly dominated. Player 2 does not have any strictly dominated strategies.
- (b) For Player 1, 6 is weakly dominated by 4 (and also by 2, since strict dominance implies weak dominance); 4 is weakly dominated by 2. Player 2 does not have any weakly dominated strategies. \square

		Player 2 (Bob)					
		1		3		5	
Player 1 (Antonia)	2	3	2	3	2	2	3
	4	3	2	2	3	1	1
	6	2	3	1	1	1	1

Figure 2.35: The game for Exercise 2.5

Solution to Exercise 2.6.

- (a) The game under consideration is shown in Figure 2.36.
- (b) For Player 1, 0 strictly dominates 3 and 6, 3 strictly dominates 6 (the same is true for every player). Thus 0 is a strictly dominant strategy.
- (c) The strict dominant-strategy profile is (0,0,0) (everybody contributes 0). □

		Player 2								
		0		3		6				
Player 1	0	0	0	2	-1	2	4	-2	4	
	3	-1	2	2	1	1	4	3	0	6
	6	-2	4	4	0	3	6	2	2	8
		Player 3: 0								

		Player 2								
		0		3		6				
Player 1	0	2	2	-1	4	1	1	6	0	3
	3	1	4	1	3	3	3	5	2	5
	6	0	6	3	2	5	5	4	4	7
		Player 3: 3								

		Player 2								
		0		3		6				
Player 1	0	4	4	-2	6	3	0	8	2	2
	3	3	6	0	5	5	2	7	4	4
	6	2	8	2	4	7	4	6	6	6
		Player 3: 6								

Figure 2.36: The game for Exercise 2.6

Solution to Exercise 2.7. The game under consideration is shown in Figure 2.37. \square

		Player 2 (value \$50M)									
		\$10M		\$20M		\$30M		\$40M		\$50M	
Player 1 (value \$30M)	\$10M	0	40	0	40	0	40	0	40	0	40
	\$20M	20	0	0	30	0	30	0	30	0	30
	\$30M	20	0	10	0	0	20	0	20	0	20
	\$40M	20	0	10	0	0	0	0	10	0	10
	\$50M	20	0	10	0	0	0	-10	0	0	0

Figure 2.37: The game for Exercise 2.7

Solution to Exercise 2.8. No. Suppose, by contradiction, that \hat{b}_1 is a weakly dominant strategy for Player 1. It cannot be that $\hat{b}_1 > v_1$, because when $b_2 = b_3 = \hat{b}_1$ Player 1 wins and pays \hat{b}_1 , thereby obtaining a payoff of $v_1 - \hat{b}_1 < 0$, whereas bidding 0 would give him a payoff of 0.

It cannot be that $\hat{b}_1 = v_1$ because when $b_2 > \hat{b}_1$ and $b_3 < v_1$ the auction is won by Player 2 and Player 1 gets a payoff of 0, while a bid of Player 1 greater than b_2 would make him the winner with a payoff of $v_1 - b_3 > 0$.

Similarly, it cannot be that $\hat{b}_1 < v_1$ because when $b_2 > v_1$ and $b_3 < v_1$ then with \hat{b}_1 the auction is won by Player 2 and Player 1 gets a payoff of 0, while a bid greater than b_2 would make him the winner with a payoff of $v_1 - b_3 > 0$. \square

Solution to Exercise 2.9.

(a) Since $\sum_{i=1}^5 v_i = 85 < \sum_{i=1}^5 c_i = 100$ the Pareto efficient decision is not to build the park.

(b) Since $\sum_{i=1}^5 w_i = 110 > \sum_{i=1}^5 c_i = 100$ the park would be built.

(c) Individuals 1 and 3 are pivotal and each of them has to pay a tax of \$30. The other individuals are not pivotal and thus are not taxed.

(d) For each individual i it is a dominant strategy to report v_i and thus, by Part (a), the decision will be the Pareto efficient one, namely not to build the park.

(e) When every individual reports truthfully, Individuals 4 and 5 are pivotal and Individual 4 has to pay a tax of \$25, while individual 5 has to pay a tax of \$10. The others are not pivotal and do not have to pay a tax.

(f) Assume that all the other individuals report truthfully; then if Individual 1 reports truthfully, he is not pivotal, the project is not carried out and his utility is \bar{m}_1 . Any other w_1 that leads to the same decision (not to build the park) gives him the same utility.

If, on the other hand, he chooses a w_1 that leads to a decision to build the park, then Individual 1 will become pivotal and will have to pay a tax $t_1 = 45$ with a utility

of $\bar{m}_1 + v_1 - c_1 - t_1 = \bar{m}_1 + 60 - 30 - 45 = \bar{m}_1 - 15$, so that he would be worse off relative to reporting truthfully. \square

Solution to Exercise 2.10. The game under consideration is shown in Figure 2.38.

		Player 2 (Bob)					
		1		3		5	
Player 1 (Antonia)	2	3	2	3	2	2	3
	4	3	2	2	3	1	1
	6	2	3	1	1	1	1

Figure 2.38: The game for Exercise 2.10

- (a) The first step of the procedure eliminates 6 for Player 1. After this step the procedure stops and thus the output is as shown in Figure 2.39.

		Player 2 (Bob)					
		1		3		5	
Player 1 (Antonia)	2	3	2	3	2	2	3
	4	3	2	2	3	1	1

Figure 2.39: The output of the IDWDS procedure applied to the game of Figure 2.38

- (b) The first step of the procedure eliminates 4 and 6 for Player 1 and nothing for Player 2. The second step of the procedure eliminates 1 and 3 for Player 2. Thus the output is the strategy profile (2,5), which constitutes the iterated weak dominant-strategy profile of this game. \square

Solution to Exercise 2.11. The game under consideration is shown in Figure 2.40.

In this game c is strictly dominated by b ;

- after deleting c , d becomes strictly dominated by f ;
- after deleting d , a becomes strictly dominated by b ;
- after deleting a , e becomes strictly dominated by f ;
- deletion of e leads to only one strategy profile, namely (b, f) .

Thus (b, f) is the iterated strict dominant-strategy profile. \square

		Player 2					
		<i>d</i>		<i>e</i>		<i>f</i>	
Player 1	<i>a</i>	8	6	0	9	3	8
	<i>b</i>	3	2	2	1	4	3
	<i>c</i>	2	8	1	5	3	1

Figure 2.40: The game for Exercise 2.11

Solution to Exercise 2.12.

(a) The game-frame under consideration is shown in Figure 2.41.

		BOB					BOB					BOB		
		<i>A</i>	<i>B</i>	<i>C</i>			<i>A</i>	<i>B</i>	<i>C</i>			<i>A</i>	<i>B</i>	<i>C</i>
A N N	<i>A</i>	<i>A</i>	<i>A</i>	<i>A</i>	A N N	<i>A</i>	<i>A</i>	<i>B</i>	<i>A</i>	A N N	<i>A</i>	<i>A</i>	<i>C</i>	
	<i>B</i>	<i>A</i>	<i>B</i>	<i>B</i>		<i>B</i>	<i>B</i>	<i>B</i>	<i>B</i>		<i>B</i>	<i>B</i>	<i>B</i>	<i>C</i>
	<i>C</i>	<i>A</i>	<i>C</i>	<i>C</i>		<i>C</i>	<i>C</i>	<i>B</i>	<i>C</i>		<i>C</i>	<i>C</i>	<i>C</i>	<i>C</i>
		CARLA: <i>A</i>					CARLA: <i>B</i>					CARLA: <i>C</i>		

Figure 2.41: The game-frame for Exercise 2.12 (a)

(b) The game under consideration is shown in Figure 2.42.

(c) For Ann, both *B* and *C* are weakly dominated by *A*; for Bob, *A* is weakly dominated by *C*; for Carla, *C* is weakly dominated by *B*.

Thus in the first step of the IDWDS we delete *B* and *C* for Ann, *A* for Bob and *C* for Carla.

Hence the game reduces to the Figure 2.43. In this game, for Bob, *C* is weakly dominated by *B* and for Carla, *A* is weakly dominated by *B*.

Thus in the second and final step of the IDWDS we delete *C* for Bob and *A* for Carla and we are left with a unique strategy profile, namely (A, B, B) , that is, Ann votes for herself and Bob and Carla vote for Bob. This is the iterated weak dominant-strategy profile.

(d) The elected candidate is Bob, who is Ann's least favorite; thus the extra power given to Ann (tie breaking in case of total disagreement) turns out to be detrimental for Ann! \square

		BOB		
		A	B	C
CARLA: A	A	2,0,1	2,0,1	2,0,1
	N	2,0,1	0,1,2	0,1,2
	N	2,0,1	1,2,0	1,2,0

		BOB		
		A	B	C
CARLA: B	A	2,0,1	0,1,2	2,0,1
	N	0,1,2	0,1,2	0,1,2
	N	1,2,0	0,1,2	1,2,0

		BOB		
		A	B	C
CARLA: C	A	2,0,1	2,0,1	1,2,0
	N	0,1,2	0,1,2	1,2,0
	N	1,2,0	1,2,0	1,2,0

Figure 2.42: The game for Exercise 2.12 (b)

		CARLA	
		A	B
BOB	B	0 1	1 2
	C	0 1	0 1

Figure 2.43: The reduced game for Exercise 2.12 (c)

Solution to Exercise 2.13. The game under consideration is shown in Figure 2.44.

		Player 2					
		D		E		F	
Player 1	a	2	3	2	2	3	1
	b	2	0	3	1	1	0
	c	1	4	2	0	0	4

Figure 2.44: The game for Exercise 2.13

- (a) The output of the IDSDS is shown in Figure 2.45 (first delete c and then F). Thus there is no iterated strict dominant-strategy profile.
- (b) The output of the IDWDS is (b, E) (in the first step delete c and F , the latter because it is weakly dominated by D ; in the second step delete a and in the third step delete D). Thus (b, E) is the iterated weak dominant-strategy profile. □

		Player 2			
		D		E	
Player 1	a	2	3	2	2
	b	2	0	3	1

Figure 2.45: The output of the IDSDS procedure applied to the game of Figure 2.44

Solution to Exercise 2.14. The game under consideration is shown in Figure 2.46. There is only one Nash equilibrium, namely (4,1) with payoffs (5,5). □

		Player 2 (Bob)					
		0		1		2	
Player 1 (Antonia)	2	0	2	0	3	0	2
	4	0	2	5	5	4	2
	6	4	2	3	7	2	0

Figure 2.46: The game for Exercise 2.14

Solution to Exercise 2.15. The game of Exercise 2.3 (b) is shown in Figure 2.47. This game has only one Nash equilibrium, namely (*not P*, *P*, *not P*). The game of Exercise 2.3

		Bob					
		P			not P		
Alice	P	1	0	0	0	1	0
	not P	0	1	0	1	0	0

Charlie: P

.....

		Bob					
		P			not P		
Alice	P	0	1	0	1	0	0
	not P	1	0	0	0	0	1

Charlie: not P

Figure 2.47: The first game for Exercise 2.15

(c) is shown in Figure 2.48. This game does not have any Nash equilibria. □

		Bob					
		<i>P</i>			<i>not P</i>		
Alice	<i>P</i>	1	0	0	0	2	1
	<i>not P</i>	0	2	1	1	0	0
Charlie: <i>P</i>							
.....							
		Bob					
		<i>P</i>			<i>not P</i>		
Alice	<i>P</i>	0	2	1	1	0	0
	<i>not P</i>	1	0	0	0	1	2
Charlie: <i>not P</i>							

Figure 2.48: The second game for Exercise 2.15

Solution to Exercise 2.16. The game under consideration is shown in Figure 2.49. This game has only one Nash equilibrium, namely (600,0). □

		Player 2					
		0	300		600		
Player 1	0	0	0	150	-75	300	-150
	300	150	225	300	150	450	75
	600	300	450	450	375	600	300

Figure 2.49: The game for Exercise 2.16

Solution to Exercise 2.17. The game under consideration is shown in Figure 2.50. This game has only one Nash equilibrium, namely (0,0,0). □

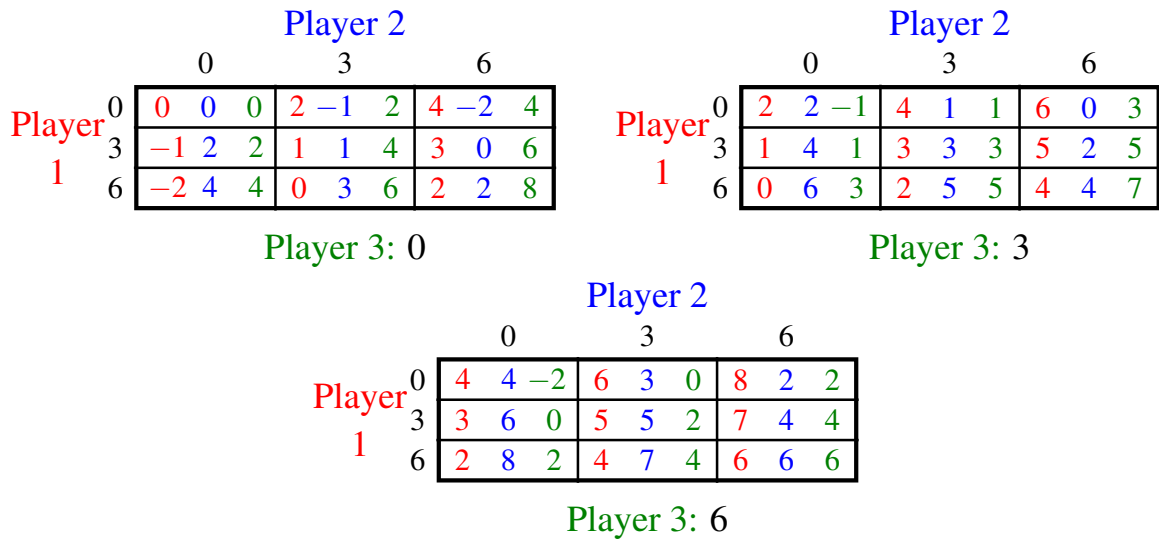


Figure 2.50: The game for Exercise 2.17

Solution to Exercise 2.18. The game under consideration is shown in Figure 2.51. This game has 15 Nash equilibria:

- (10,30), (10,40), (10,50), (20,30), (20,40), (20,50), (30,30), (30,40), (30,50), (40,40), (40,50), (50,10), (50,20), (50,30), (50,50). □

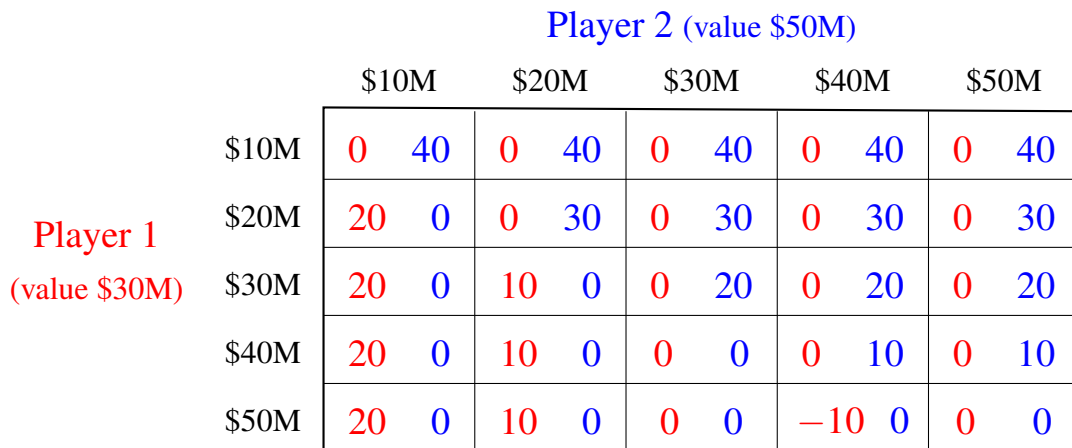


Figure 2.51: The game for Exercise 2.18

Solution to Exercise 2.19. A Nash equilibrium is $b_1 = b_2 = b_3 = v_1$ (with payoffs $(0,0,0)$). Convince yourself that this is indeed a Nash equilibrium.

There are many more Nash equilibria: for example, any triple (b_1, b_2, b_3) with $b_2 = b_3 = v_1$ and $b_1 > v_1$ is a Nash equilibrium (with payoffs $(0,0,0)$) and so is any triple (b_1, b_2, b_3) with $b_2 = b_3 = v_2$ and $b_1 \geq v_2$ (with payoffs $(v_1 - v_2, 0, 0)$). \square

Solution to Exercise 2.20. The game under consideration is shown in Figure 2.52. There are 5 Nash equilibria: (A,A,A) , (B,B,B) , (C,C,C) , (A,C,A) and (A,B,B) . \square

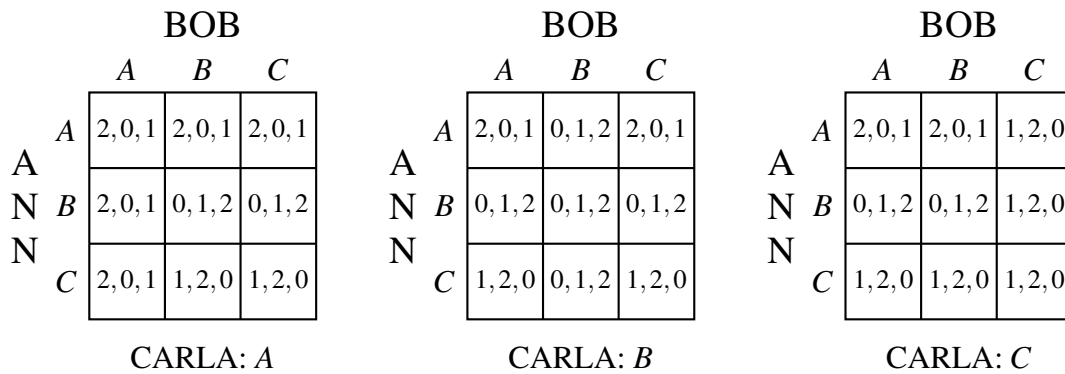


Figure 2.52: The game for Exercise 2.20

Solution to Exercise 2.21. The game under consideration is shown in Figure 2.53. There are 2 Nash equilibria: (a,D) and (b,E) . \square




Figure 2.53: The game for Exercise 2.21

Solution to Exercise 2.22.

- (a) For every
- $e \in [0, 1]$
- ,
- (e, e, \dots, e)
- is a Nash equilibrium.

The payoff of Player i is $\pi_i(e, e, \dots, e) = 2e$; if player i increases her effort to $a > e$ (of course, this can only happen if $e < 1$), then her payoff decreases to $4e - 2a$ and if she decreases her effort to $a < e$ (of course, this can only happen if $e > 0$), then her payoff decreases to $2a$.

There is no Nash equilibrium where two players choose different levels of effort.

Proof: suppose there is an equilibrium (a_1, a_2, \dots, a_n) where $a_i \neq a_j$ for two players i and j .

Let $a_{\min} = \min\{a_1, \dots, a_n\}$ and let k be a player such that $a_k > a_{\min}$ (such a player exists by our supposition).

Then the payoff to player k is $\pi_k = 4a_{\min} - 2a_k$ and if she reduced her effort to a_{\min} her payoff would increase to $2a_{\min}$.

- (b) Any symmetric equilibrium with $e < 1$ is Pareto inefficient, because all the players would be better off if they collectively switched to $(1, 1, \dots, 1)$. On the other hand, the symmetric equilibrium $(1, 1, \dots, 1)$ is Pareto efficient.
- (c) The symmetric equilibrium $(\frac{1}{2}, \frac{1}{2}, \dots, \frac{1}{2})$. □

Solution to Exercise 2.23.

- (a) The strategy sets are
- $S_1 = S_2 = [0, \infty)$
- . The payoff functions are as follows:
- ¹⁰

$$\pi_1(x, y) = \begin{cases} \Pi - y & \text{if } x \geq y \\ 0 & \text{if } y > x \end{cases} \quad \text{and} \quad \pi_2(x, y) = \begin{cases} A - L + y & \text{if } x \geq y \\ A & \text{if } y > x \end{cases}$$

- (b) Yes, for player 1 choosing
- $x = \Pi$
- is a weakly dominant strategy.

Proof. Consider an arbitrary y . We must show that $x = \Pi$ gives at least a high a payoff against y as any other x . Three cases are possible.

Case 1: $y < \Pi$. In this case $x = \Pi$ or any other x such that $x \geq y$ yields $\pi_1 = \Pi - y > 0$, while any $x < y$ yields $\pi_1 = 0$.

Case 2: $y = \Pi$. In this case 1's payoff is zero no matter what x he chooses.

Case 3: $y > \Pi$. In this case $x = \Pi$ or any other x such that $x < y$ yields $\pi_1 = 0$, while any $x \geq y$ yields $\pi_1 = \Pi - y < 0$.

- (c) No, choosing
- $y = L$
- is not a dominant strategy for Player 2. For example, if
- $x > L$
- then choosing
- $y = L$
- yields
- $\pi_2 = A$
- while choosing a
- y
- such that
- $L < y \leq x$
- yields
- $\pi_2 = A - L + y > A$
- .

¹⁰We have chosen to use accounting profits as payoffs. Alternatively, one could take as payoffs the changes in profits relative to the initial situation, namely

$$\pi_1(x, y) = \begin{cases} -y & \text{if } x \geq y \\ -\Pi & \text{if } y > x \end{cases} \quad \text{and} \quad \pi_2(x, y) = \begin{cases} y & \text{if } x \geq y \\ L & \text{if } y > x. \end{cases}$$

The answers are the same, whichever choice of payoffs one makes.

- (d) **Suppose that $\Pi > L$.** If (x, y) is a Nash equilibrium **with** $x \geq y$ then it must be that $y \leq \Pi$ (otherwise Player 1 could increase its payoff by reducing x below y) and $y \geq L$ (otherwise Player 2 would be better off by increasing y above x).

Thus it must be $L \leq y \leq \Pi$, which is possible, given our assumption.

However, it cannot be that $x > y$, because Player 2 would be getting a higher payoff by increasing y to x .

Thus it must be $x \leq y$, which (together with our hypothesis that $x \geq y$) implies that $x = y$. Thus the following are Nash equilibria:

all the pairs (x, y) with $L \leq y \leq \Pi$ and $x = y$.

Now consider pairs (x, y) **with** $x < y$. Then it cannot be that $y < \Pi$, because Player 1 could increase its payoff by increasing x to y . Thus it must be $y \geq \Pi$ (hence, by our supposition that $\Pi > L$, $y > L$). Furthermore, it must be that $x \leq L$ (otherwise Player 2 could increase its profits by reducing y to (or below) x). Thus

(x, y) with $x < y$ is a Nash equilibrium if and only if $x \leq L$ and $y \geq \Pi$.

- (e) **Suppose that $\Pi < L$.** For the same reasons given above, an equilibrium with $x \geq y$ requires $L \leq y \leq \Pi$. However, this is not possible given that $\Pi < L$. Hence,

there is no Nash equilibrium (x, y) with $x \geq y$.

Thus we must restrict attention to pairs (x, y) **with** $x < y$. As explained above, it must be that $y \geq \Pi$ and $x \leq L$. Thus,

(x, y) with $x < y$ is a Nash equilibrium if and only if $\Pi \leq y$ and $x \leq L$.

- (f) Pareto efficiency requires that the chemical plant be shut down if $\Pi < L$ and that it remain operational if $\Pi > L$.

Now, when $\Pi < L$ all the equilibria have $x < y$ which leads to shut-down, hence a Pareto efficient outcome.

When $\Pi > L$, there are two types of equilibria: one where $x = y$ and the plant remains operational (a Pareto efficient outcome) and the other where $x < y$ in which case the plant shuts down, yielding a Pareto inefficient outcome. \square

3. Perfect-information Games

3.1 Trees, frames and games

Often interactions are not simultaneous but sequential. For example, in the game of Chess the two players, White and Black, take turns moving pieces on the board, having full knowledge of the opponent's (and their own) past moves. Games with sequential interaction are called *dynamic games* or *games in extensive form*. This chapter is devoted to the subclass of dynamic games characterized by *perfect information*, namely the property that, whenever it is her turn to move, a player knows all the preceding moves.

Perfect-information games are represented by means of rooted directed trees.

Definition 3.1.1 A *rooted directed tree* consists of a set of nodes and directed edges joining them.

- The root of the tree has no directed edges leading to it (has indegree 0), while every other node has exactly one directed edge leading to it (has indegree 1).
- There is a unique path (that is, a unique sequence of directed edges) leading from the root to any other node. A node that has no directed edges out of it (has outdegree 0) is called a *terminal node*, while every other node is called a *decision node*.
- We shall denote the set of nodes by X , the set of decision nodes by D and the set of terminal nodes by Z . Thus $X = D \cup Z$.

Definition 3.1.2 A *finite extensive form (or frame) with perfect information* consists of the following items.

- A finite rooted directed tree.
- A set of players $I = \{1, \dots, n\}$ and a function that assigns one player to every decision node.
- A set of actions A and a function that assigns one action to every directed edge, satisfying the restriction that no two edges out of the same node are assigned the same action.
- A set of outcomes O and a function that assigns an outcome to every terminal node.

■ **Example 3.1** Amy (Player 1) and Beth (Player 2) have decided to dissolve a business partnership whose assets have been valued at \$100,000. The charter of the partnership prescribes that the senior partner, Amy, make an offer concerning the division of the assets to the junior partner, Beth. The junior partner can *Accept*, in which case the proposed division is implemented, or *Reject*, in which case the case goes to litigation.

- Litigating involves a cost of \$20,000 in legal fees for each partner and the typical verdict assigns 60% of the assets to the senior partner and the remaining 40% to the junior partner.
- Suppose, for simplicity, that there is no uncertainty about the verdict (how to model uncertainty will be discussed in a later chapter). Suppose also that there are only two possible offers that Amy can make: a 50-50 split or a 70-30 split.

This situation can be represented as a finite extensive form with perfect information as shown in Figure 3.1. Each outcome is represented as two sums of money: the top one is what Player 1 gets and the bottom one what Player 2 gets. ■




Figure 3.1: A perfect-information extensive form representing the situation described in Example 3.1

What should we expect the players to do in the game of Figure 3.1? Consider the following reasoning, which is called *backward induction* reasoning, because it starts from the end of the game and proceeds backwards towards the root:

- If Player 2 (the junior partner) is offered a 50-50 split then, if she accepts, she will get \$50,000, while, if she rejects, she will get \$20,000 (the court-assigned 40% minus legal fees in the amount of \$20,000); thus, if rational, she will accept.
- Similarly, if Player 2 is offered a 70-30 split then, if she accepts, she will get \$30,000, while, if she rejects, she will get \$20,000 (the court-assigned 40% minus legal fees in the amount of \$20,000); thus, if rational, she will accept.
- Anticipating all of this, Player 1 realizes that, if she offers a 50-50 split then she will end up with \$50,000, while if she offers a 70-30 split then she will end up with \$70,000; thus, if Player 1 is rational and believes that Player 2 is rational, she will offer a 70-30 split and Player 2, being rational, will accept.

The above reasoning suffers from the same flaw as the reasoning described in Chapter 2: it is not a valid argument because it is based on an implicit assumption about how Player 2 ranks the outcomes, which may or may not be correct. For example, Player 2 may feel that she worked as hard as her senior partner and the only fair division is a 50-50 split; indeed she may feel so strongly about this that – if offered an unfair 70-30 split – she would be willing to sacrifice \$10,000 in order to “teach a lesson to Player 1”; in other words, she ranks outcome o_4 above outcome o_3 .

Using the terminology introduced in Chapter 2, we say that the situation represented in Figure 3.1 is not a game but a *game-frame*. In order to convert that frame into a game we need to add a ranking of the outcomes for each player.

Definition 3.1.3 A *finite extensive game with perfect information* is a finite extensive form with perfect information together with a ranking \succsim_i of the set of outcomes O , for every player $i \in I$.

As usual, it is convenient to represent the ranking of Player i by means of an ordinal utility function $U_i : O \rightarrow \mathbb{R}$. For example, take the extensive form of Figure 3.1 and assume that **Player 1** is selfish and greedy, that is, her ranking is:

best o_3
 o_1 (or, in the alternative notation, $o_3 \succ_1 o_1 \succ_1 o_2 \sim_1 o_4$).
 worst o_2, o_4

while **Player 2** is concerned with fairness and her ranking is:

best o_1
 o_2, o_4 (or, in the alternative notation, $o_1 \succ_2 o_2 \sim_2 o_4 \succ_2 o_3$)
 worst o_3

Then we can represent the players' preferences using the following utility functions:

outcome \rightarrow utility function \downarrow	o_1	o_2	o_3	o_4
U_1 (Player 1)	2	1	3	1
U_2 (Player 2)	3	2	1	2

and replace each outcome in Figure 3.1 with a pair of utilities or payoffs, as shown in

Figure 3.2, thereby obtaining one of the many possible games based on the frame of Figure 3.1.

Now that we have a game (rather than just a game-frame), we can apply the backward-induction reasoning and conclude that Player 1 will offer a 50-50 split, anticipating that Player 2 would reject the offer of a 70-30 split, and Player 2 will accept Player 1's 50-50 offer. The choices selected by the backward-induction reasoning have been highlighted in Figure 3.2 by doubling the corresponding edges.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 3.6.1 at the end of this chapter.




Figure 3.2: A perfect-information game based on the frame of Figure 3.1

3.2 Backward induction

The backward-induction reasoning mentioned above can be formalized as an algorithm for solving any finite perfect-information game. We say that a node is *marked* if a utility vector is associated with it. Initially all and only the terminal nodes are marked; the following procedure provides a way of marking all the nodes.

Definition 3.2.1 The *backward-induction algorithm* is the following procedure for solving a finite perfect-information game:

1. Select a decision node x whose immediate successors are all marked. Let i be the player who moves at x . Select a choice that leads to an immediate successor of x with the highest payoff (or utility) for Player i (highest among the utilities associated with the immediate successors of x). Mark x with the payoff vector associated with the node that follows the selected choice.
2. Repeat the above step until all the nodes have been marked.

Note that, since the game is finite, the above procedure is well defined. In the initial steps one starts at those decision nodes that are followed only by terminal nodes, call them penultimate nodes. After all the penultimate nodes have been marked, there will be unmarked nodes whose immediate successors are all marked and thus the step can be repeated. Note also that, in general, at a decision node there may be several choices that maximize the payoff of the player who moves at that node. If that is the case, then the procedure requires that *one* such choice be selected. This arbitrary selection may lead to the existence of several backward-induction solutions.

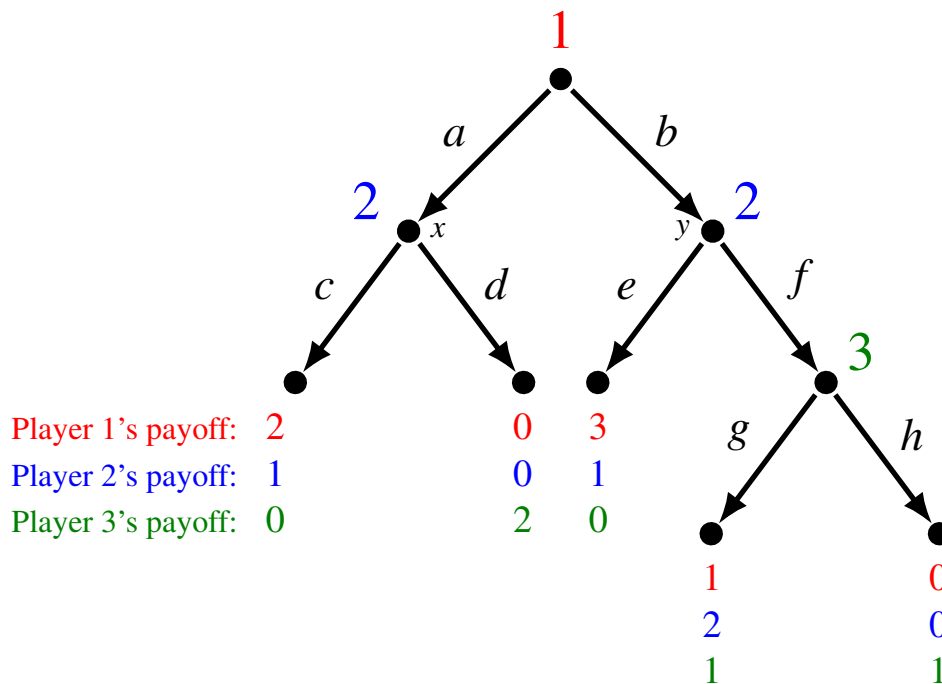


Figure 3.3: A perfect-information game with multiple backward-induction solutions.

For example, in the game of 3.3 starting at node x of Player 2 we select choice c (since it gives Player 2 a higher payoff than d). Then we move on to Player 3's node and we find that both choices there are payoff maximizing for Player 3; thus there are two ways to proceed, as shown in the next two figures.

In Figure 3.4 we show the steps of the backward-induction algorithm with the selection of choice g , while Figure 3.5 shows the steps of the algorithm with the selection of choice h . As before, the selected choices are shown by double edges. In Figures 3.4 and 3.5 the marking of nodes is shown explicitly, but later on we will represent the backward-induction solution more succinctly by merely highlighting the selected choices.




Figure 3.4: One possible output of the backward-induction algorithm applied to the game of Figure 3.3




Figure 3.5: Another possible output of the backward-induction algorithm applied to the game of Figure 3.3

How should one define the output of the backward-induction algorithm and the notion of backward-induction solution? What kind of objects are they? Before we answer this question we need to introduce the notion of strategy in a perfect-information game.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 3.6.2 at the end of this chapter.

3.3 Strategies in perfect-information games

A strategy for a player in a perfect-information game is a complete, contingent plan on how to play the game. Consider, for example, the game shown in Figure 3.6 (which reproduces 3.3) and let us focus on Player 2.




Figure 3.6: Copy of the game of Figure 3.3

Before the game is played, Player 2 does not know what Player 1 will do and thus a complete plan needs to specify what she will do if Player 1 decides to play a and what she will do if Player 1 decides to play b . A possible plan, or strategy, is “if Player 1 chooses a then I will choose c and if Player 1 chooses b then I will choose e ”, which we can denote more succinctly as (c, e) . The other possible plans, or strategies, for Player 2 are (c, f) , (d, e) and (d, f) . The formal definition of strategy is as follows.

Definition 3.3.1 A *strategy* for a player in a perfect-information game is a list of choices, one for each decision node of that player.

For example, suppose that Player 1 has three decision nodes in a given game: at one node she has three possible choices, a_1, a_2 and a_3 , at another node she has two possible choices, b_1 and b_2 , and at the third node she has four possible choices, c_1, c_2, c_3 and c_4 . Then a strategy for Player 1 in that game can be thought of as a way of filling in three blanks:

$$\left(\underbrace{\hspace{2cm}}_{\text{one of } a_1, a_2, a_3}, \underbrace{\hspace{2cm}}_{\text{one of } b_1, b_2}, \underbrace{\hspace{2cm}}_{\text{one of } c_1, c_2, c_3, c_4} \right).$$

Since there are 3 choices for the first blank, 2 for the second and 4 for the third, the total number of possible strategies for Player 1 in this case would be $3 \times 2 \times 4 = 24$. One strategy is (a_2, b_1, c_1) , another strategy is (a_1, b_2, c_4) , etc.

It should be noted that the notion of strategy involves redundancies. To see this, consider the game of Figure 3.7. In this game a possible strategy for Player 1 is (a, g) , which means that Player 1 is planning to choose a at the root of the tree and would choose g at her other node. But if Player 1 indeed chooses a , then her other node will *not* be reached and thus why should Player 1 make a plan on what to do there? One could justify this redundancy in the notion of strategy in a number of ways:

1. Player 1 is so cautious that she wants her plan to cover also the possibility that she might make mistakes in the implementation of parts of her plan (in this case, she allows for the possibility that – despite her intention to play a – she might end up playing b), or
2. we can think of a strategy as a set of instructions given to a third party on how to play the game on Player 1's behalf, in which case Player 1 might indeed worry about the possibility of mistakes in the implementation and thus want to cover all contingencies.

An alternative justification relies on a different interpretation of the notion of strategy: not as a plan of Player 1 but as a belief in the mind of Player 2 concerning what Player 1 would do. For the moment we will set this issue aside and simply use the notion of strategy as given in Definition 4.2.1.




Figure 3.7: A perfect-information game

Using Definition 4.2.1, one can associate with every perfect-information game a strategic-form (or normal-form) game: a strategy profile determines a unique terminal node that is reached if the players act according to that strategy profile and thus a unique vector of payoffs. Figure 3.8 shows the strategic-form associated with the perfect-information game of Figure 3.7, with the Nash equilibria highlighted.

		Player 2							
		<i>ce</i>	<i>cf</i>	<i>de</i>	<i>df</i>				
Player 1	<i>ag</i>	2	1	2	1	0	0	0	0
	<i>ah</i>	2	1	2	1	0	0	0	0
	<i>bg</i>	3	1	1	2	3	1	1	2
	<i>bh</i>	3	1	1	0	3	1	1	0

Figure 3.8: The strategic form of the perfect-information game of Figure 3.7 with the Nash equilibria highlighted

Because of the redundancy discussed above, the strategic form also displays redundancies: in this case the top two rows are identical.

Armed with the notion of strategy, we can now revisit the notion of backward-induction solution. Figure 3.9 shows the two backward-induction solutions of the game of Figure 3.7.




Figure 3.9: The backward-induction solutions of the game of Figure 3.7

It is clear from the definition of backward-induction algorithm (Definition 3.2.1) that the procedure selects a choice at every decision node and thus yields a strategy profile for the entire game: the backward-induction solution shown in Panel (a) of Figure 3.9 is the strategy profile $((a, g), (c, f))$, while the backward-induction solution shown in Panel (b) is the strategy profile $((b, h), (c, e))$. Both of them are Nash equilibria of the corresponding strategic-form game, but not all the Nash equilibria correspond to backward-induction solutions. The relationship between the two concepts is explained in the next section.

- R** A backward-induction solution is a strategy profile. Since strategies contain a description of what a player actually does and also of what the player would do in circumstances that do not arise, one often draws a distinction between the backward-induction *solution* and the backward-induction *outcome* which is defined as the sequence of actual moves. For example, the backward-induction outcome associated with the solution $((a, g), (c, f))$ is the play ac with corresponding payoff $(2, 1)$, while the backward-induction outcome associated with the solution $((b, h), (c, e))$ is the play be with corresponding payoff $(3, 1)$.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 3.6.3 at the end of this chapter.

3.4 Relationship between backward induction and other solutions

If you have gone through the exercises for the previous three sections, you will have seen that in all those games the backward-induction solutions are also Nash equilibria. This is always true, as stated in the following theorem.

Theorem 3.4.1 Every backward-induction solution of a perfect-information game is a Nash equilibrium of the associated strategic form.

In some games the set of backward-induction solutions coincides with the set of Nash equilibria (see, for example, Exercise 3.9), but typically the set of Nash equilibria is larger than (is a proper superset of) the set of backward-induction solutions (for example the game of Figure 3.7 has two backward-induction solutions – shown in Figure 3.9 – but five Nash equilibria, shown in Figure 3.8).

Nash equilibria that are not backward-induction solutions often involve *incredible threats*. To see this, consider the following game.

An industry is currently a monopoly and the incumbent monopolist is making a profit of \$5 million. A potential entrant is considering whether or not to enter this industry.

- If she does not enter, she will make \$1 million in an alternative investment.
- If she does enter, then the incumbent can either fight entry with a price war whose outcome is that both firms make zero profits, or it can accommodate entry, by sharing the market with the entrant, in which case both firms make a profit of \$2 million.

This situation is illustrated in Figure 3.10 with the associated strategic form. Note that we are assuming that each player is selfish and greedy, that is, cares only about its own profit and prefers more money to less.

The backward-induction solution is $(in, accommodate)$ and it is also a Nash equilibrium. However, there is another Nash equilibrium, namely $(out, fight)$. The latter should be discarded as a “rational solution” because it involves an incredible threat on the part of the incumbent, namely that it will fight entry if the potential entrant enters.

- It is true that, if the potential entrant believes the incumbent’s threat, then she is better off staying out; however, she should ignore the incumbent’s threat because she should realize that – when faced with the *fait accompli* of entry – the incumbent would not want to carry out the threat.

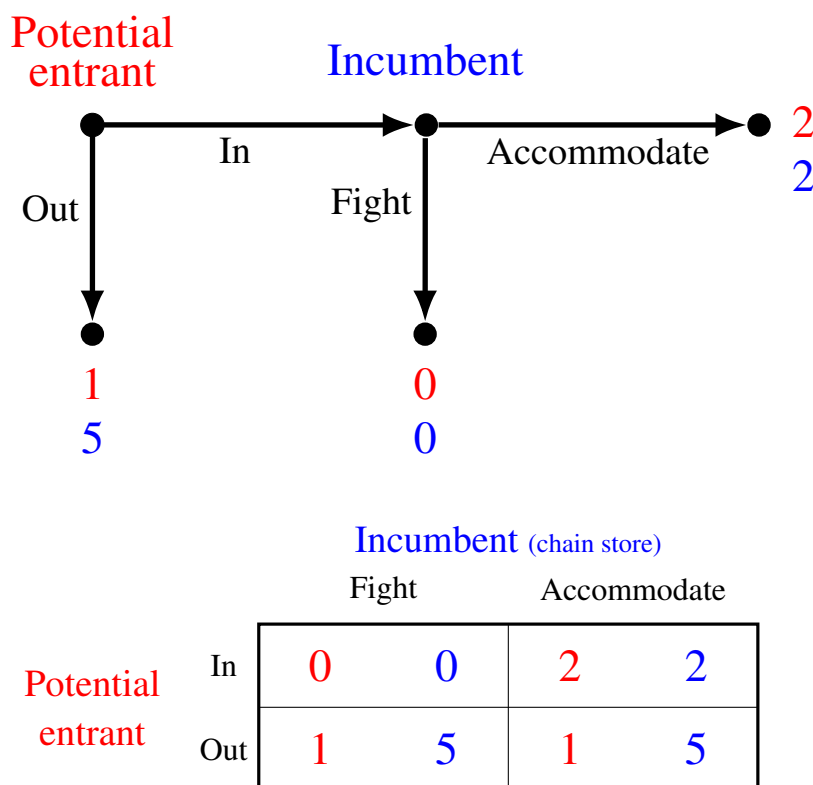


Figure 3.10: The entry game

Reinhard Selten (who shared the 1994 Nobel Memorial prize in economics with two other game theorists, John Harsanyi and John Nash) discussed a repeated version of the above entry game, which has become known as *Selten's Chain Store Game*. The story is as follows:

- A chain store is a monopolist in an industry. It owns stores in m different towns ($m \geq 2$).
- In each town the chain store makes \$5 million if left to enjoy its privileged position undisturbed.
- In each town there is a businesswoman who could enter the industry in that town, but earns \$1 million if she chooses not to enter; if she decides to enter, then the monopolist can either fight the entrant, leading to zero profits for both the chain store and the entrant in that town, or it can accommodate entry and share the market with the entrant, in which case both players make \$2 million in that town.

Thus, in each town the interaction between the incumbent monopolist and the potential entrant is as illustrated in Figure 3.10.

However, decisions are made sequentially, as follows:

At date t ($t = 1, \dots, m$) the businesswoman in town t decides whether or not to enter and if she enters then the chain store decides whether or not to fight in that town.

What happens in town t at date t becomes known to everybody. Thus, for example, the businesswoman in town 2 at date 2 knows what happened in town 1 at date 1 (either that there was no entry or that entry was met with a fight or that entry was accommodated).

Intuition suggests that in this game the threat by the incumbent to fight early entrants might be credible, for the following reason. The incumbent could tell Businesswoman 1 the following:

“It is true that, if you enter and I fight, I will make zero profits, while by accommodating your entry I would make \$2 million and thus it would seem that it cannot be in my interest to fight you. However, somebody else is watching us, namely Businesswoman 2. If she sees that I have fought your entry then she might fear that I would do the same with her and decide to stay out, in which case in town 2, I would make \$5 million, so that my total profits in towns 1 and 2 would be $(0 + 5) = \$5$ million. On the other hand, if I accommodate your entry, then she will be encouraged to enter herself and I will make \$2 million in each town, for a total profit of \$4 million. Hence, as you can see, it is indeed in my interest to fight you and thus you should stay out.”

Does the notion of backward induction capture this intuition? To check this, let us consider the case where $m = 2$, so that the extensive game is not too large to draw. It is shown in Figure 3.11, where at each terminal node the top number is the profit of the incumbent monopolist (it is the sum of the profits in the two towns), the middle number is the profit of Businesswoman 1 and the bottom number is the profit of Businesswoman 2. All profits are expressed in millions of dollars. We assume that all the players are selfish and greedy, so that we can take the profit of each player to be that player’s payoff. The backward-induction solution is unique and is shown by the double directed edges in Figure 3.11. The corresponding outcome is that both businesswomen will enter and the incumbent monopolist accommodates entry in both towns.

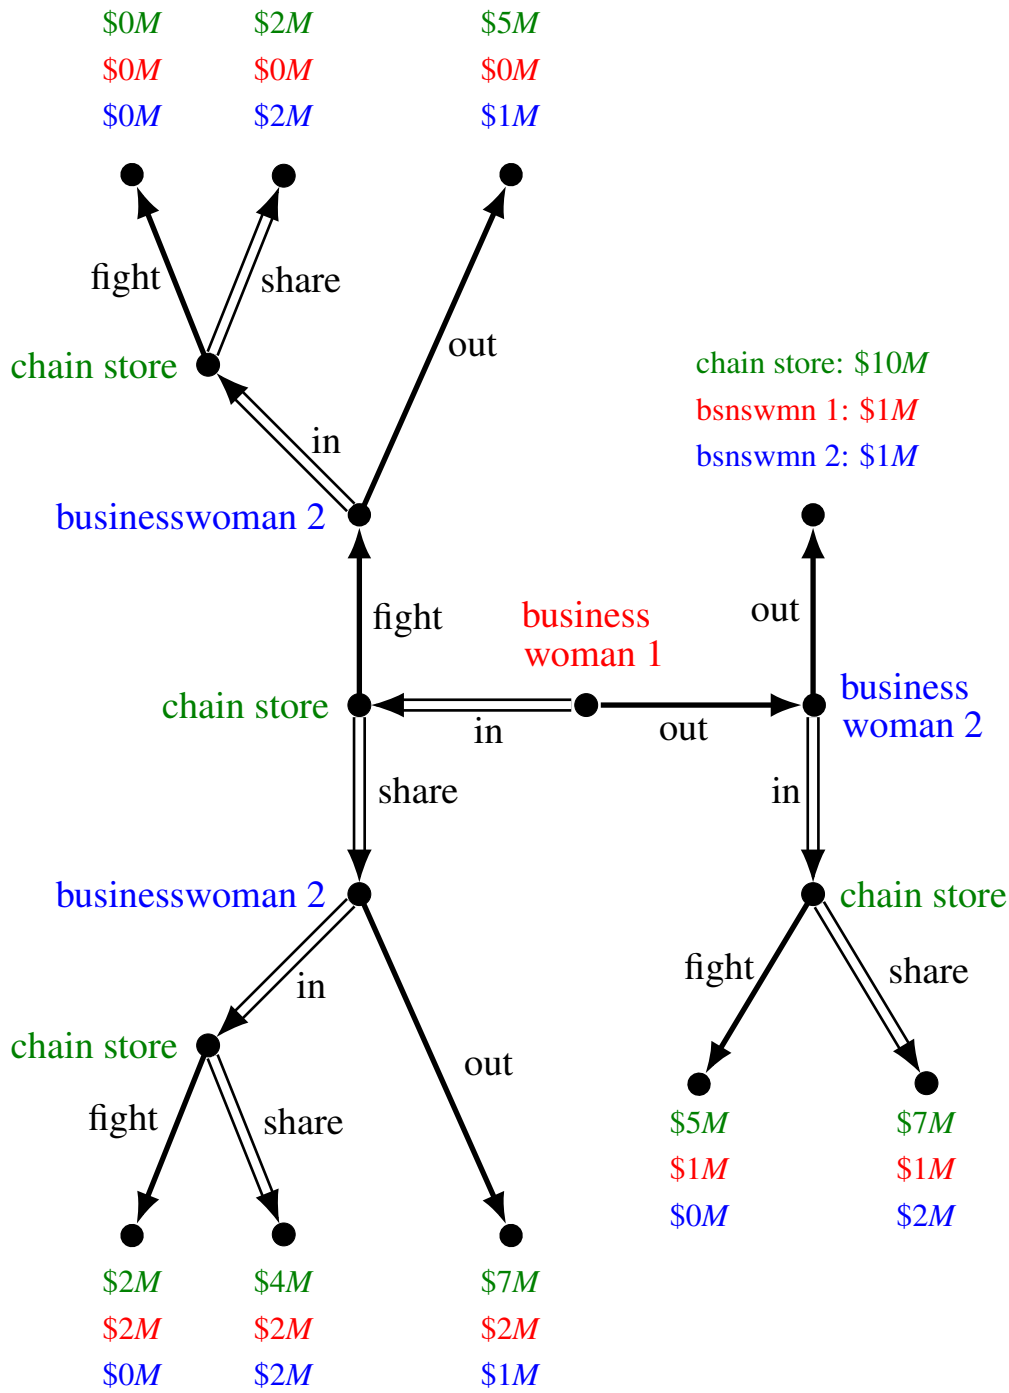


Figure 3.11: Selten's Chain-Store game

Thus the backward-induction solution does not capture the “reputation” argument outlined above. However, the backward-induction solution does seem to capture the notion of rational behavior in this game. Indeed, Businesswoman 1 should reply to the incumbent with the following counter-argument:

“Your reasoning is not valid. Whatever happens in town 1, it will be common knowledge between you and Businesswoman 2 that your interaction in town 2 will be the last; in particular, nobody else will be watching and thus there won’t be an issue of establishing a reputation in the eyes of another player. Hence in town 2 it will be in your interest to accommodate entry, since in essence you will be playing the one-shot entry game of Figure 3.10. Hence a rational Businesswoman 2 will decide to enter in town 2 *whatever happened in town 1*: what you do against me will have no influence on her decision. Thus your “reputation” argument does not apply and it will in fact be in your interest not to fight my entry: your choice will be between a profit of $\$(0 + 2) = \2 million, if you fight me, and a profit of $\$(2 + 2) = \4 million, if you don’t fight me. Hence I will enter and you will not fight me.”

In order to capture the reputation argument described above we need to allow for some uncertainty in the mind of some of the players, as we will show in a later chapter. In a perfect-information game uncertainty is ruled out by definition.

By Theorem 3.4.1 the notion of backward induction can be seen as a refinement of the notion of Nash equilibrium. Another solution concept that is related to backward induction is the iterated elimination of weakly dominated strategies. Indeed the backward-induction algorithm could be viewed as a step-wise procedure that eliminates dominated choices at decision nodes, and thus strategies that contain those choices. What is the relationship between the two notions? In general this is all that can be said: applying the iterated deletion of weakly dominated strategies to the strategic form associated with a perfect-information game leads to a set of strategy profiles that contains at least one backward-induction solution; however,

- (1) it may also contain strategy profiles that are not backward-induction solutions, and
- (2) it may fail to contain all the backward-induction solutions, as shown in Exercise 3.8.

3.5 Perfect-information games with two players

We conclude this chapter with a discussion of finite two-player extensive games with perfect information.

We will start with games that have only two outcomes, namely “Player 1 wins” (denoted by W_1) and “Player 2 wins” (denoted by W_2). We assume that Player 1 strictly prefers W_1 to W_2 and Player 2 strictly prefers W_2 to W_1 . Thus we can use utility functions with values 0 and 1 and associate with each terminal node either the payoff vector $(1, 0)$ (if the outcome is W_1) or the payoff vector $(0, 1)$ (if the outcome is W_2). We call these games *win-lose games*. An example of such a game is the following.

■ **Example 3.2** Two players take turns choosing a number from the set $\{1, 2, \dots, 10\}$, with Player 1 moving first. The first player who brings the sum of all the chosen numbers to 100 or more wins. ■

The following is one possible play of the game (the bold-face numbers are the ones chosen by Player 1 and the underlined numbers the ones chosen by Player 2):

10, 9, **9**, 10, **8**, 7, **4**, 10, **1**, 8, **3**, 3, **8**, 10.

In this play Player 2 wins: before her last move the sum is 90 and with her final choice of 10 she brings the total to 100. However, in this game Player 1 has a *winning strategy*, that is, a strategy that guarantees that he will win, no matter what numbers Player 2 chooses. To see this, we can use backward-induction reasoning. Drawing the tree is not a practical option, since the number of nodes is very large: one needs 10,000 nodes just to represent the first 4 moves! But we can imagine drawing the tree, placing ourselves towards the end of the tree and ask what partial sum represents a “losing position”, in the sense that whoever is choosing in that position cannot win, while the other player can then win with his subsequent choice. With some thought one can see that 89 is the largest losing position: whoever moves there can take the sum to any number in the set $\{90, 91, \dots, 99\}$, thus coming short of 100, while the other player can then take the sum to 100 with an appropriate choice. What is the largest losing position that precedes 89? The answer is 78: whoever moves at 78 must take the sum to a number in the set $\{79, 80, \dots, 88\}$ and then from there the other player can make sure to take the sum to 89 and then we know what happens from there! Repeating this reasoning we see that the losing positions are: 89, 78, 67, 56, 45, 34, 23, 12, 1. Since Player 1 moves first he can choose 1 and put Player 2 in the first losing position; then, whatever Player 2 chooses, Player 1 can put her in the next losing position, namely 12, etc. Recall that a strategy for Player 1 must specify what to do in every possible situation in which he might find himself. In his game Player 1’s winning strategy is as follows:

Start with the number 1. Then, at every turn, choose the number $(11 - n)$, where n is the number that was chosen by Player 2 in the immediately preceding turn.

Here is an example of a possible play of the game where Player 1 employs the winning strategy and does in fact win:

1, 9, **2**, 6, **5**, 7, **4**, 10, **1**, 8, **3**, 3, **8**, 9, **2**, 5, **6**, 1, **10**

We can now state a general result about this class of games.

Theorem 3.5.1 In every finite two-player, win-lose game with perfect information one of the two players has a winning strategy.

Although we will not give a detailed proof, the argument of the proof is rather simple. By applying the backward-induction algorithm we assign to every decision node either the payoff vector $(1, 0)$ or the payoff vector $(0, 1)$. Imagine applying the algorithm up to the point where the immediate successors of the root have been assigned a payoff vector. Two cases are possible.

Case 1: at least one of the immediate successors of the root has been assigned the payoff vector $(1, 0)$. In this case Player 1 is the one who has a winning strategy and his initial choice should be such that a node with payoff vector $(1, 0)$ is reached and then his future choices should also be such that only nodes with payoff vector $(1, 0)$ are reached.

Case 2: all the immediate successors of the root have been assigned the payoff vector $(0, 1)$. In this case it is Player 2 who has a winning strategy. An example of a game where it is Player 2 who has a winning strategy is given in Exercise 3.11.

We now turn to finite two-player games where there are three possible outcomes: “Player 1 wins” (W_1), “Player 2 wins” (W_2) and “Draw” (D). We assume that the rankings of the outcomes are as follows: $W_1 \succ_1 D \succ_1 W_2$ and $W_2 \succ_2 D \succ_2 W_1$.

Examples of such games are Tic-Tac-Toe (<http://en.wikipedia.org/wiki/Tic-tac-toe>), Draughts or Checkers (<http://en.wikipedia.org/wiki/Draughts>) and Chess (although there does not seem to be agreement as to whether the rules of Chess guarantee that every possible play of the game is finite). What can we say about such games? The answer is provided by the following theorem.

Theorem 3.5.2 Every finite two-player, perfect-information game with three outcomes: Player 1 wins (W_1), Player 2 wins (W_2) and Draw (D), and preferences $W_1 \succ_1 D \succ_1 W_2$ and $W_2 \succ_2 D \succ_2 W_1$, falls within one of the following three categories:

1. Player 1 has a strategy that guarantees outcome W_1 .
2. Player 2 has a strategy that guarantees outcome W_2 .
3. Player 1 has a strategy that guarantees that the outcome will be W_1 or D and Player 2 has a strategy that guarantees that the outcome will be W_2 or D , so that, if both players employ these strategies, the outcome will be D .

The logic of the proof is as follows. By applying the backward-induction algorithm we assign to every decision node either the payoff vector $(2, 0)$ (corresponding to outcome W_1) or the payoff vector $(0, 2)$ (corresponding to outcome W_2) or the payoff vector $(1, 1)$ (corresponding to outcome D). Imagine applying the algorithm up to the point where the immediate successors of the root have been assigned a payoff vector. Three cases are possible.

Case 1: at least one of the immediate successors of the root has been assigned the payoff vector $(2, 0)$; in this case Player 1 is the one who has a winning strategy.

Case 2: all the immediate successors of the root have been assigned the payoff vector $(0, 2)$; in this case it is Player 2 who has a winning strategy.

Case 3: there is at least one immediate successor of the root to which the payoff vector $(1, 1)$ has been assigned and all the other immediate successors of the root have been assigned either $(1, 1)$ or $(0, 2)$. In this case we fall within the third category of Theorem 3.5.2.

Both Tic-Tac-Toe and Checkers fall within the third category (http://en.wikipedia.org/wiki/Solved_game#Solved_games). As of the time of writing this book, it is not known to which category the game of Chess belongs.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 3.6.4 at the end of this chapter.

3.6 Exercises

3.6.1 Exercises for Section 3.1: Trees, frames and games

The answers to the following exercises are in Section 3.7 at the end of this chapter.

Exercise 3.1 How could they do that! They abducted Speedy, your favorite tortoise! They asked for \$1,000 in unmarked bills and threatened to kill Speedy if you don't pay. Call the tortoise-napper Mr. T. Let the possible outcomes be as follows:

- o_1 : you don't pay and speedy is released
- o_2 : you pay \$ 1,000 and speedy is released
- o_3 : you don't pay and speedy is killed
- o_4 : you pay \$ 1,000 and speedy is killed

You are attached to Speedy and would be willing to pay \$1,000 to get it back. However, you also like your money and you prefer not to pay, conditional on each of the two separate events "Speedy is released" and "Speedy is killed". Thus your ranking of the outcomes is $o_1 \succ_{you} o_2 \succ_{you} o_3 \succ_{you} o_4$. On the other hand, you are not quite sure of what Mr. T's ranking is.

- (a) Suppose first that Mr T has communicated that he wants you to go to Central Park tomorrow at 10:00 a.m. and leave the money in a garbage can; he also said that, two miles to the East and at the exact same time, he will decide whether or not to free Speedy in front of the police station and then go and collect his money in Central Park. What should you do?
- (b) Suppose that Mr T is not as dumb as in part (a) and instead gives you the following instructions: first you leave the money in a garbage can in Central Park and then he will go there to collect the money. He also told you that if you left the money there then he will free Speedy, otherwise he will kill it. Draw an extensive form or frame to represent this situation.
- (c) Now we want to construct a game based on the extensive form of part (b). For this we need Mr T's preferences. There are two types of criminals in Mr T's line of work: the professionals and the one-timers. Professionals are in the business for the long term and thus, besides being greedy, worry about reputation; they want it to be known that (1) every time they were paid they honored their promise to free the hostage and (2) their threats are to be taken seriously: every time they were *not* paid, the hostage was killed. The one-timers hit once and then they disappear; they don't try to establish a reputation and the only thing they worry about, besides money, is not to be caught: whether or not they get paid, they prefer to kill the hostage in order to eliminate any kind of evidence (DNA traces, fingerprints, etc.). Construct two games based on the extensive form of part (b) representing the two possible types of Mr T.

Exercise 3.2 A three-man board, composed of A , B , and C , has held hearings on a personnel case involving an officer of the company. This officer was scheduled for promotion but, prior to final action on his promotion, he made a decision that cost the company a good deal of money. The question is whether he should be (1) promoted anyway, (2) denied the promotion, or (3) fired. The board has discussed the matter at length and is unable to reach unanimous agreement. In the course of the discussion it has become clear to all three of them that their separate opinions are as follows:

- A considers the officer to have been a victim of bad luck, not bad judgment, and wants to go ahead and promote him but, failing that, would keep him rather than fire him.
- B considers the mistake serious enough to bar promotion altogether; he'd prefer to keep the officer, denying promotion, but would rather fire than promote him.
- C thinks the man ought to be fired but, in terms of personnel policy and morale, believes the man ought not to be kept unless he is promoted, i.e., that keeping an officer who has been declared unfit for promotion is even worse than promoting him.

	PROMOTE	KEEP	FIRE
A :	best	middle	worst
B :	worst	best	middle
C :	middle	worst	best

Assume that everyone's preferences among the three outcomes are fully evident as a result of the discussion. The three must proceed to a vote.

Consider the following voting procedure. First A proposes an action (either promote or keep or fire). Then it is B 's turn. If B accepts A 's proposal, then this becomes the final decision. If B disagrees with A 's proposal, then C makes the final decision (which may be *any of the three*: promote, keep or fire). Represent this situation as an extensive game with perfect information. (Use utility numbers from the set $\{1, 2, 3\}$.) ■

3.6.2 Exercises for Section 3.2: Backward induction

The answers to the following exercises are in Section 3.7 at the end of this chapter.

Exercise 3.3 Apply the backward-induction algorithm to the two games of Exercise 3.1 Part (c). ■

Exercise 3.4 Apply the backward-induction algorithm to the game of Exercise 3.2. ■

3.6.3 Exercises for Section 3.3: Strategies in perfect-information games

The answers to the following exercises are in Section 3.7 at the end of this chapter.

Exercise 3.5 Write the strategic form of the game of Figure 3.2, find all the Nash equilibria and verify that the backward-induction solution is a Nash equilibrium. ■

Exercise 3.6 Write the strategic form of the game of Figure 3.3, find all the Nash equilibria and verify that the backward-induction solutions are Nash equilibria. ■

Exercise 3.7 Consider the game of Exercise 3.2.

- (a) Write down all the strategies of Player *B*.
- (b) How many strategies does Player *C* have?

Exercise 3.8 Consider the perfect-information game shown in Figure 3.12.

- (a) Find the backward-induction solutions.
- (b) Write down all the strategies of Player 1.
- (c) Write down all the strategies of Player 2.
- (d) Write the strategic form associated with this game.
- (e) Does Player 1 have a dominant strategy?
- (f) Does Player 2 have a dominant strategy?
- (g) Is there a dominant-strategy equilibrium?
- (h) Does Player 1 have any dominated strategies?
- (i) Does Player 2 have any dominated strategies?
- (j) What do you get when you apply the iterative elimination of weakly dominated strategies?
- (k) What are the Nash equilibria?




Figure 3.12: The perfect-information game for Exercise 3.8

Exercise 3.9 Consider an industry where there are two firms, a large firm, Firm 1, and a small firm, Firm 2. The two firms produce identical products.

- Let x be the output of Firm 1 and y the output of Firm 2. Industry output is $Q = x + y$.
- The price P at which each unit of output can be sold is determined by the inverse demand function $P = 130 - 10Q$. For example, if Firm 1 produces 4 units and Firm 2 produces 2 units, then industry output is 6 and each unit is sold for $P = 130 - 60 = \$70$.
- For each firm the cost of producing q units of output is $C(q) = 10q + 62.5$.
- Each firm is only interested in its own profits.
- The profit of Firm 1 depends on both x and y and is given by

$$\Pi_1(x, y) = \underbrace{x [130 - 10(x + y)]}_{\text{revenue}} - \underbrace{(10x + 62.5)}_{\text{cost}}$$

and similarly the profit function of Firm 2 is given by

$$\Pi_2(x, y) = \underbrace{y [130 - 10(x + y)]}_{\text{revenue}} - \underbrace{(10y + 62.5)}_{\text{cost}}$$

- The two firms play the following sequential game. First Firm 1 chooses its own output x and commits to it; then Firm 2, after having observed Firm 1's output, chooses its own output y ; then the price is determined according to the demand function and the two firms collect their own profits. In what follows assume, for simplicity, that x can only be 6 or 6.5 units and y can only be 2.5 or 3 units.

- (a) Represent this situation as an extensive game with perfect information.
- (b) Solve the game using backward induction.
- (c) Write the strategic form associated with the perfect-information game.
- (d) Find the Nash equilibria of this game and verify that the backward-induction solutions are Nash equilibria.

■

Exercise 3.10 Consider the perfect-information game shown in Figure 3.13 where x is an integer.

- For every value of x find the backward induction solution(s).
- Write the corresponding strategic-form and find all the Nash equilibria.




Figure 3.13: A perfect-information game

3.6.4 Exercises for Section 3.5: Two-player games

The answers to the following exercises are in Section 3.7 at the end of this chapter.

Exercise 3.11 Consider the following perfect-information game. Player 1 starts by choosing a number from the set $\{1, 2, 3, 4, 5, 6, 7\}$, then Player 2 chooses a number from this set, then Player 1 again, followed by Player 2, etc. The first player who brings the cumulative sum of all the numbers chosen (up to and including the last one) to 48 or more wins. By Theorem 3.5.1 one of the two players has a winning strategy. Find out who that player is and fully describe the winning strategy.

Exercise 3.12 Consider Figure 3.14 and the following two-player, perfect-information game. A coin is placed in the cell marked ‘START’ (cell A1). Player 1 moves first and can move the coin one cell up (to A2) or one cell to the left (to B1) or one cell diagonally in the left-up direction (to B2). Then Player 2 moves, according to the same rules (e.g. if the coin is in cell B2 then the admissible moves are shown by the directed edges). The players alternate moving the coin. Black cells are not accessible (so that, for example, from A3 the coin can only be moved to A4 or B3 and from F3 it can only be moved to G4, as shown by the directed edge). The player who manages to place the coin in the cell marked ‘END’ wins.

- Represent this game by means of an extensive form with perfect information by drawing the initial part of the tree that covers the first two moves (the first move of Player 1 and the first move of Player 2).
- Suppose that the coin is currently in cell G4 and it is Player 1’s turn to move. Show that Player 1 has a strategy that allows her to win the game starting from cell G4. Describe the strategy in detail.
- Describe a play of the game (from cell A1) where Player 1 wins (describe it by means of the sequence of cells visited by the coin).
- Describe a play of the game (from cell A1) where Player 2 wins (describe it by means of the sequence of cells visited by the coin).
- Now go back to the beginning of the game. The coin is in cell A1 and player 1 has the first move. By Theorem 3.5.1 one of the two players has a winning strategy. Find out who that player is and fully describe the winning strategy.




Figure 3.14: The coin game

Exercise 3.13 — *** Challenging Question ***.

Two women, Anna and Bess, claim to be the legal owners of a diamond ring that - each claims - has great sentimental value. Neither of them can produce evidence of ownership and nobody else is staking a claim on the ring. Judge Sabio wants the ring to go to the legal owner, but he does not know which of the two women is in fact the legal owner. He decides to proceed as follows. First he announces a fine of $\$F > 0$ and then asks Anna and Bess to play the following game.

Move 1: Anna moves first. Either she gives up her claim to the ring (in which case Bess gets the ring, the game ends and nobody pays the fine) or she asserts her claim, in which case the game proceeds to Move 2.

Move 2: Bess either accepts Anna's claim (in which case Anna gets the ring, the game ends and nobody pays the fine) or challenges her claim. In the latter case, Bess must put in a bid, call it B , and Anna must pay the fine of $\$F$ to Sabio. The game goes on to Move 3.

Move 3: Anna now either matches Bess's bid (in which case Anna gets the ring, Anna pays $\$B$ to Sabio in addition to the fine that she already paid and Bess pays the fine of $\$F$ to Sabio) or chooses not to match (in which case Bess gets the ring and pays her bid of $\$B$ to Sabio and, furthermore, Sabio keeps the fine that Anna already paid).

Denote by C_A the monetary equivalent of getting the ring for Anna (that is, getting the ring is as good, in Anna's mind, as getting $\$C_A$) and C_B the monetary equivalent of getting the ring for Bess. Not getting the ring is considered by both to be as good as getting zero dollars.

- (a) Draw an extensive game with perfect information to represent the above situation, assuming that there are only two possible bids: B_1 and B_2 . Write the payoffs to Anna and Bess next to each terminal node.
- (b) Find the backward-induction solution of the game you drew in part (a) for the case where $B_1 > C_A > C_B > B_2 > F > 0$.

Now consider the general case where the bid B can be any non-negative number and assume that both Anna and Bess are very wealthy. Assume also that C_A, C_B and F are positive numbers and that C_A and C_B are common knowledge between Anna and Bess. We want to show that, at the backward-induction solution of the game, the ring always goes to the legal owner. Since we (like Sabio) don't know who the legal owner is, we must consider two cases.

Case 1: the legal owner is Anna. Let us assume that this implies that $C_A > C_B$.

Case 2: the legal owner is Bess. Let us assume that this implies that $C_B > C_A$.

- (c) Find the backward-induction solution for Case 1 and show that it implies that the ring goes to Anna.
- (d) Find the backward-induction solution for Case 2 and show that it implies that the ring goes to Bess.
- (e) How much money does Sabio make at the backward-induction solution? How much money do Ann and Bess end up paying at the backward-induction solution?

3.7 Solutions to exercises

Solution to Exercise 3.1.

- (a) For you it is a strictly dominant strategy to not pay and thus you should not pay.
- (b) The extensive form is shown in Figure 3.15.
- (c) For the professional, concern with reputation implies that $o_2 \succ_{MrT} o_4$ and $o_3 \succ_{MrT} o_1$. If we add the reasonable assumption that after all money is what they are after, then we can take the full ranking to be $o_2 \succ_{MrT} o_4 \succ_{MrT} o_3 \succ_{MrT} o_1$. Representing preferences with ordinal utility functions with values in the set $\{1, 2, 3, 4\}$, we have

outcome →	o_1	o_2	o_3	o_4
utility function ↓				
U_{you}	4	3	2	1
U_{MrT}	1	4	2	3

The corresponding game is obtained by replacing in Figure 3.15 o_1 with the payoff vector (4,1), o_3 with the payoff vector (2,2), etc.

For the one-timer, the ranking can be taken to be (although this is not the only possibility) $o_4 \succ_{MrT} o_2 \succ_{MrT} o_3 \succ_{MrT} o_1$, with corresponding utility representation:

outcome →	o_1	o_2	o_3	o_4
utility function ↓				
U_{you}	4	3	2	1
U_{MrT}	1	3	2	4

The corresponding extensive-form game is shown in Figure 3.16. □




Figure 3.15: The game-frame for Part (b) of Exercise 3.1




Figure 3.16: The game for Part (c) of Exercise 3.1 when Mr T is a one-timer

Solution to Exercise 3.2. The game is shown in Figure 3.17 ('P' stands for promote, 'K' for keep (without promoting), 'F' for fire). □

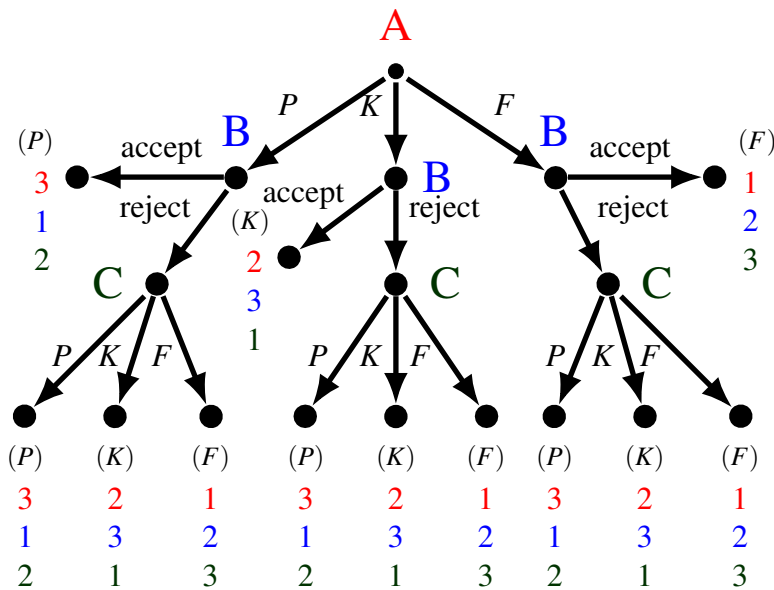


Figure 3.17: The game for Exercise 3.2

Solution to Exercise 3.3. The application of the backward-induction algorithm is shown by double edges in Figure 3.18 for the case of a professional Mr. T and in Figure 3.19 for the case of a one-timer Mr. T. Thus, against a professional you will pay and against a one-timer you would not pay. With the professional you would get Speedy back, with the one-timer you will hold a memorial service for Speedy. □

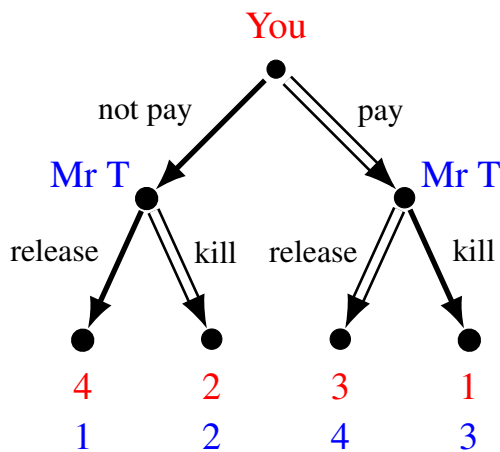


Figure 3.18: The game for Part (b) of Exercise 3.3




Figure 3.19: The game for Part (c) of Exercise 3.3

Solution to Exercise 3.4. The backward-induction algorithm yields two solutions, shown in Figures 3.20 and 3.21. The difference between the two solutions lies in what Player B would do if Player A proposed F. In both solutions the officer is kept without promotion. □

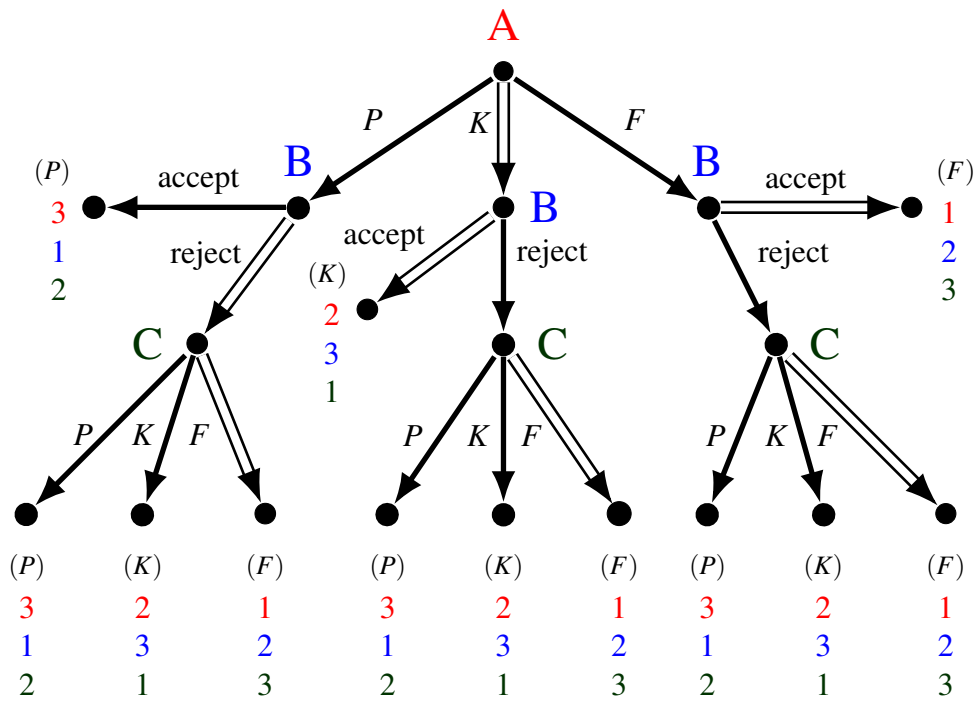


Figure 3.20: The first game for Exercise 3.4

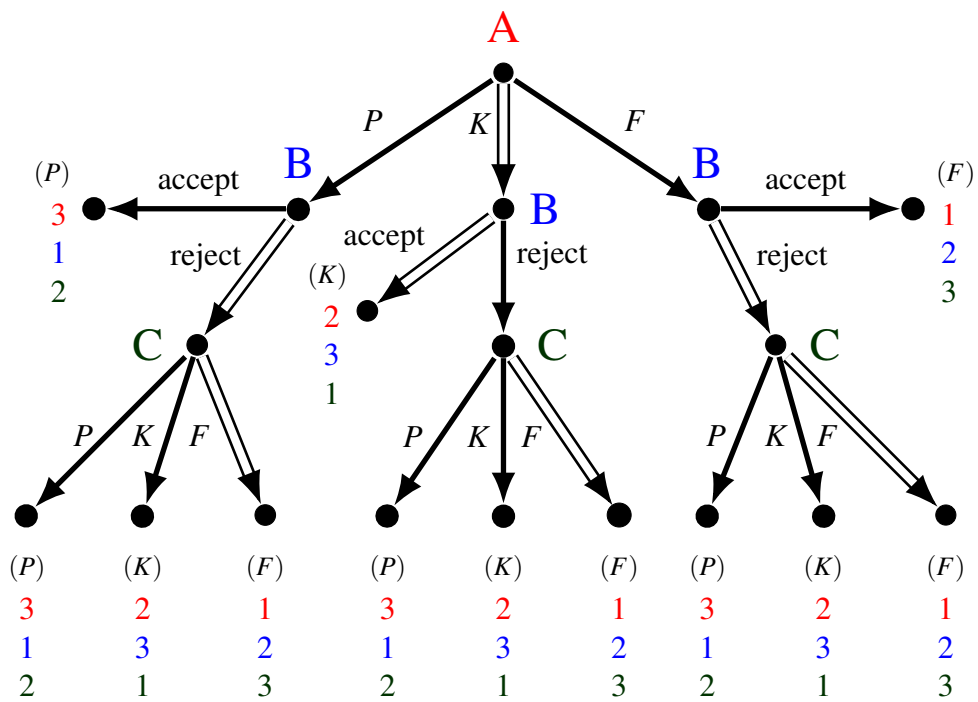


Figure 3.21: The second game for Exercise 3.4

Solution to Exercise 3.5. The game of Figure 3.2 is reproduced in Figure 3.22, with the unique backward-induction solution marked by double edges. The corresponding strategic form is shown in Figure 3.23 (for each of Player 2’s strategies, the first element in the pair is what Player 2 would do at her left node and the second element what she would do at her right node). The Nash equilibria are highlighted. One Nash equilibrium, namely $(Offer\ 50-50, (Accept, Reject))$, corresponds to the backward induction solution, while the other Nash equilibrium, namely $(Offer\ 70-30, (Reject, Reject))$ does not correspond to a backward-induction solution. \square




Figure 3.22: The extensive-form game for Exercise 3.5

		Player 2			
		(Accept,Accept)	(Accept,Reject)	(Reject,Accept)	(Reject,Reject)
Player 1	Offer 50-50	2 3	2 3	1 2	1 2
	Offer 70-30	3 1	1 2	3 1	1 2

Figure 3.23: The strategic-form game for Exercise 3.5

Solution to Exercise 3.6. The game of Figure 3.3 is reproduced in Figure 3.24 with the two backward-induction solutions marked by double edges. The corresponding strategic form is shown in Figure 3.25. The Nash equilibria are highlighted. The backward-induction solutions are $(a, (c, f), g)$ and $(b, (c, e), h)$ and both of them are Nash equilibria. There are three more Nash equilibria which are not backward-induction solutions, namely $(b, (d, f), g)$, $(a, (c, f), h)$ and $(b, (d, e), h)$. \square




Figure 3.24: The two backward-induction solutions for the game of Exercise 3.6

		Player 2											
		<i>ce</i>			<i>cf</i>			<i>de</i>			<i>df</i>		
Player 1	a	2	1	0	2	1	0	0	0	2	0	0	2
	b	3	1	0	1	2	1	3	1	0	1	2	1
		Player 3: <i>g</i>											
.....													
		Player 2											
		<i>ce</i>			<i>cf</i>			<i>de</i>			<i>df</i>		
Player 1	a	2	1	0	2	1	0	0	0	2	0	0	2
	b	3	1	0	0	0	1	3	1	0	0	0	1
		Player 3: <i>h</i>											

Figure 3.25: The strategic-form game for Exercise 3.6

Solution to Exercise 3.7. The game of Exercise 3.2 is reproduced in Figure 3.26.

- (a) All the possible strategies of Player B are shown in Figure 3.27.
- (b) Player C has three decision nodes and three choices at each of her nodes. Thus she has $3 \times 3 \times 3 = 27$ strategies. □

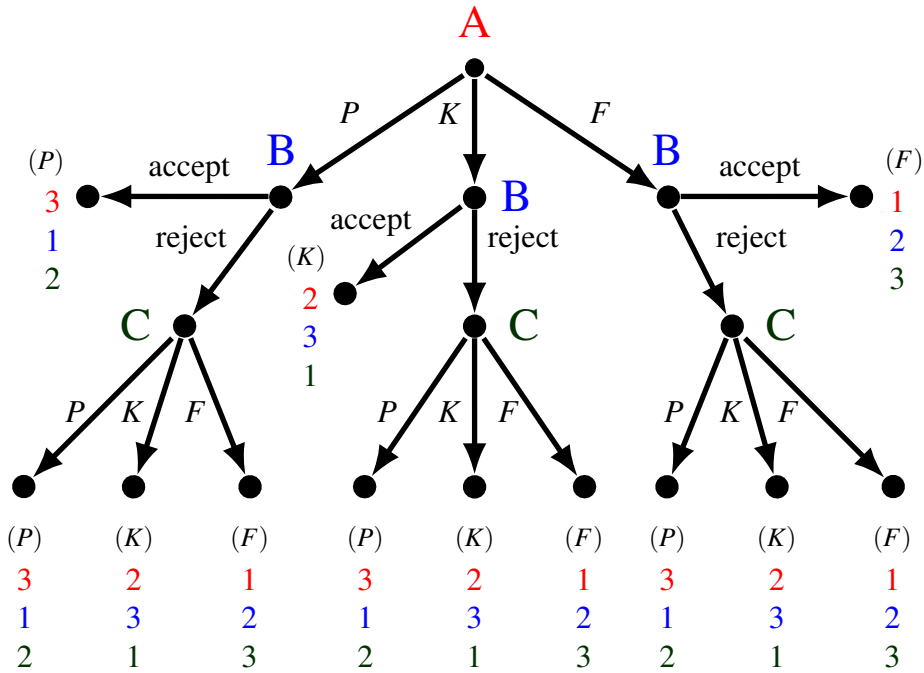


Figure 3.26: The extensive-form game for Exercise 3.7

	If A chooses <i>P</i>	If A chooses <i>K</i>	If A chooses <i>F</i>
1.	accept	accept	accept
2.	accept	accept	reject
3.	accept	reject	accept
4.	accept	reject	reject
5.	reject	accept	accept
6.	reject	accept	reject
7.	reject	reject	accept
8.	reject	reject	reject

Figure 3.27: The eight strategies of Player B

Solution to Exercise 3.8.

- (a) One backward-induction solution is the strategy profile $((L,W),(a,e))$ shown by double edges in Figure 3.28. The corresponding backward-induction outcome is the play La with associated payoff vector $(2,1)$. The other backward-induction solution is the strategy profile $((R,W),(a,d))$ shown in Figure 3.29. The corresponding backward-induction outcome is the play Rd with associated payoff vector $(3,2)$.




Figure 3.28: One backward-induction solution of the game of Part (a) of Exercise 3.8




Figure 3.29: A second backward-induction solution of the game of Part (a) of Exercise 3.8

- (b) Player 1 has four strategies: LW, LE, RW and RE .
 (c) Player 2 has six strategies: ac, ad, ae, bc, bd and be .

(d) The strategic form is shown in Figure 3.30.

		Player 2											
		<i>ac</i>	<i>ad</i>	<i>ae</i>	<i>bc</i>	<i>bd</i>	<i>be</i>						
Player 1	<i>LW</i>	2	1	2	1	2	1	4	0	4	0	4	0
	<i>LE</i>	2	1	2	1	2	1	4	0	4	0	4	0
	<i>RW</i>	2	0	3	2	1	2	2	0	3	2	1	2
	<i>RE</i>	2	0	3	2	0	3	2	0	3	2	0	3

Figure 3.30: The strategic-form game for Part (d) of Exercise 3.8

- (e) Player 1 does not have a dominant strategy.
- (f) For Player 2 *ae* is a weakly dominant strategy.
- (g) There is no dominant strategy equilibrium.
- (h) For Player 1 *RE* is weakly dominated by *RW* (and *LW* and *LE* are equivalent).
- (i) For Player 2 *ac* is weakly dominated by *ad* (and *ae*), *ad* is weakly dominated by *ae*, *bc* is (strictly or weakly) dominated by every other strategy, *bd* is weakly dominated by *be* (and by *ae* and *ad*), *be* is weakly dominated by *ae*. Thus the dominated strategies are: *ac, ad, bc, bd* and *be*.
- (j) The iterative elimination of weakly dominated strategies yields the following reduced game (in Step 1 eliminate *RE* for Player 1 and *ac, ad, bc, bd* and *be* for Player 2; in Step 2 eliminate *RW* for Player 1):

		Player 2	
		<i>ae</i>	<i>ae</i>
Player 1	<i>LW</i>	2, 1	2, 1
	<i>LE</i>	2, 1	2, 1

Thus we are left with one of the two backward-induction solutions, namely $((L, W), (a, e))$ but also with $((L, E), (a, e))$ which is not a backward-induction solution.

- (k) The Nash equilibria are highlighted in Figure 3.31.
There are five Nash equilibria: (LW, ac) , (LE, ac) , (RW, ad) , (LW, ae) and (LE, ae) . □

		Player 2											
		<i>ac</i>	<i>ad</i>	<i>ae</i>	<i>bc</i>	<i>bd</i>	<i>be</i>						
Player 1	<i>LW</i>	2	1	2	1	2	1	4	0	4	0	4	0
	<i>LE</i>	2	1	2	1	2	1	4	0	4	0	4	0
	<i>RW</i>	2	0	3	2	1	2	2	0	3	2	1	2
	<i>RE</i>	2	0	3	2	0	3	2	0	3	2	0	3

Figure 3.31: The highlighted cells are the Nash equilibria (for Part (k) of Exercise 3.8)

Solution to Exercise 3.9.

(a) The extensive game is shown in Figure 3.32.




Figure 3.32: The extensive-form game for Exercise 3.9

(b) There are two backward-induction solutions.

The first is the strategy profile shown in Figure 3.33. The corresponding backward-induction outcome is given by Firm 1 producing 6 units and Firm 2 producing 3 units with profits 117.5 for Firm 1 and 27.5 for Firm 2.

The other backward-induction solution is the strategy profile shown in Figure 3.34. The corresponding backward-induction outcome is given by Firm 1 producing 6.5 units and Firm 2 producing 2.5 units with profits 132.5 for Firm 1 and 12.5 for Firm 2.




Figure 3.33: One backward-induction solution of the game of Figure 3.32




Figure 3.34: A second backward-induction solution of the game of Figure 3.32

(c) The strategic form is shown in Figure 3.35.

		Firm 2							
		(2.5, 2.5)		(2.5, 3)		(3, 2.5)		(3, 3)	
Firm 1	6	147.5	25	147.5	25	117.5	27.5	117.5	27.5
	6.5	132.5	12.5	100	12.5	132.5	12.5	100	12.5

Figure 3.35: The strategic-form game for Part (d) of Exercise 3.9

(d) The Nash equilibria are highlighted in Figure 3.35. In this game the set of Nash equilibria coincides with the set of backward-induction solutions. □

Solution to Exercise 3.10. The game under consideration is shown in Figure 3.36, where x is an integer.




Figure 3.36: The extensive-form game for Part (a) of Exercise 3.10

- (a) The backward-induction strategy of Player 2 is the same, no matter what x is, namely (C, F) . Thus the backward induction solutions are as follows.
- If $x < 2$, there is only one: $(A, (C, F))$.
 - If $x = 2$ there are two: $(A, (C, F))$ and $(B, (C, F))$.
 - If $x > 2$, there is only one: $(B, (C, F))$.

- (b) The strategic form is shown in Figure 3.37. First note that $(A, (C, E))$ is a Nash equilibrium for every value of x . Now, depending on the value of x the other Nash equilibria are as follows:
- If $x < 1$, $(A, (C, F))$.
 - If $1 \leq x < 2$, $(A, (C, F))$ and $(B, (D, F))$.
 - If $x = 2$, $(A, (C, F))$, $(B, (C, F))$ and $(B, (D, F))$.
 - If $x > 2$, $(B, (C, F))$ and $(B, (D, F))$.

		Player 2							
		CE		CF		DE		DF	
Player 1	a	2	5	2	5	1	0	1	0
	b	1	0	x	2	1	0	x	2

Figure 3.37: The strategic-form game for Part (b) of Exercise 3.10

Solution to Exercise 3.11. Let us find the losing positions. If Player i , with his choice, can bring the sum to **40** then he can win (the other player with her next choice will take the sum to a number between 41 and 47 and then Player i can win with his next choice). Working backwards, the previous losing position is **32** (from here the player who has to move will take the sum to a number between 33 and 39 and after this the opponent can take it to 40). Reasoning backwards, the earlier losing positions are **24, 16, 8** and **0**. Thus Player 1 starts from a losing position and therefore it is Player 2 who has a winning strategy. The winning strategy is: at every turn, if Player 1's last choice was n then Player 2 should choose $(8 - n)$.

Solution to Exercise 3.12.

- (a) The initial part of the game is shown in Figure 3.38.




Figure 3.38: The initial part of the game of Part (a) of Exercise 3.12

(b) From G4 Player 1 should move the coin to H5. From there Player 2 has to move it to H6 and Player 1 to H7 and Player 2 to H8 and from there Player 1 wins by moving it to H9.

(c) $A1 \xrightarrow{1} B2 \xrightarrow{2} C3 \xrightarrow{1} D4 \xrightarrow{2} E5 \xrightarrow{1} F5 \xrightarrow{2} G6 \xrightarrow{1} H7 \xrightarrow{2} H8 \xrightarrow{1} H9$.

(d) $A1 \xrightarrow{1} B2 \xrightarrow{2} C3 \xrightarrow{1} D4 \xrightarrow{2} E5 \xrightarrow{1} F5 \xrightarrow{2} G6 \xrightarrow{1} G7 \xrightarrow{2} H7 \xrightarrow{1} H8 \xrightarrow{2} H9$.

(e) Using backward induction we can label each cell with a W (meaning that the player who has to move when the coin is there has a winning continuation strategy) or with an L (meaning that the player who has to move when the coin is there can be made to lose).

If all the cells that are accessible from a given cell are marked with a W then that cell must be marked with an L.

If from a cell there is an accessible cell marked with an L then that cell should be marked with a W. See Figure 3.39.

From the picture it is clear that it is Player 1 who has a winning strategy. The winning strategy of Player 1 is: move the coin to cell B1 and from then on, after every move of Player 2, move the coin to a cell marked L. \square

END	W	L	W	L				9				
W	W	W			W	W	L		8			
L	W	L	W	L			W			7		
W	W				W	L	W			6		
L	W	L	W			L	W	W			5	
W	W			W	L	W			L	4		
L			L	W	W	W	L	W			3	
		W	W			L	W	W	W			2
		L	W	L	W	W	L			START	1	
H	G	F	E	D	C	B	A					

Figure 3.39: Solution for the coin game

Solution to Exercise 3.13.

(a) The game is shown in Figure 3.40.

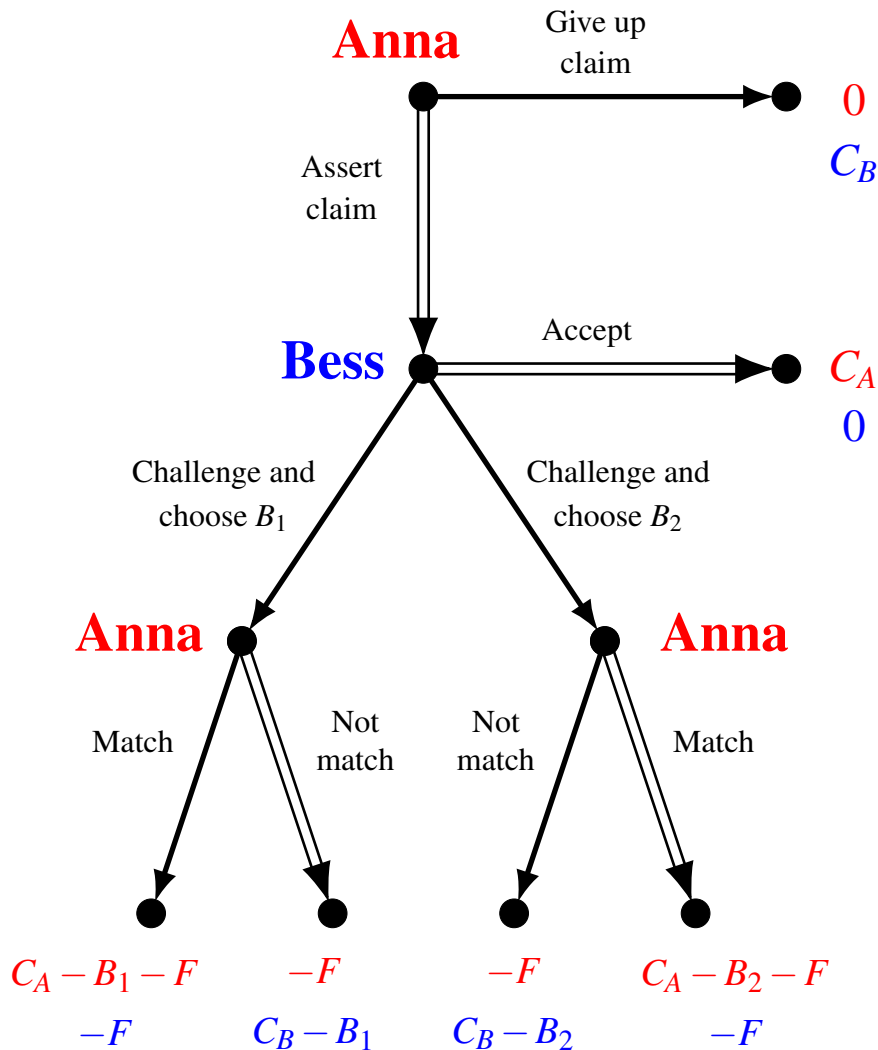


Figure 3.40: The extensive-form game for Part (a) of Exercise 3.13

(b) The backward-induction solution is marked by double arrows in Figure 3.40.

(c) The sequence of moves is shown in Figure 3.41.

Suppose that Anna is the legal owner and values the ring more than Bess does: $C_A > C_B$. At the last node Anna will choose “match” if $C_A > B$ and “don’t match” if $B > C_A$. In the first case Bess’s payoff will be $-F$, while in the second case it will be $C_B - B$, which is negative since $B > C_A$ and $C_A > C_B$. Thus in either case Bess’s payoff would be negative. Hence at her decision node Bess will choose “accept” (Bess can get the ring at this stage only if she bids more than the ring is worth to her). Anticipating this, Anna will assert her claim at the first decision node. Thus at the backward-induction solution the ring goes to Anna, the legal owner. The payoffs are C_A for Anna and 0 for Bess. **Note that no money changes hands.**

- (d) Suppose that Bess is the legal owner and values the ring more than Anna does: $C_B > C_A$. At the last node Anna will choose "match" if $C_A > B$ and "don't match" if $B > C_A$. In the first case Bess's payoff will be $-F$, while in the second case it will be $C_B - B$, which will be positive as long as $C_B > B$. Hence at her decision node Bess will choose to challenge and bid any amount B such that $C_B > B > C_A$. Anticipating this, at her first decision node Anna will give up (and get a payoff of 0), because if she asserted her claim then her final payoff would be negative. Thus at the backward-induction solution the ring goes to Bess, the legal owner. The payoffs are 0 for Anna and C_B for Bess. **Note that no money changes hands.**
- (e) As pointed out above, in both cases no money changes hands at the backward-induction solution. Thus Judge Sabio collects no money at all and both Ann and Bess pay nothing. \square

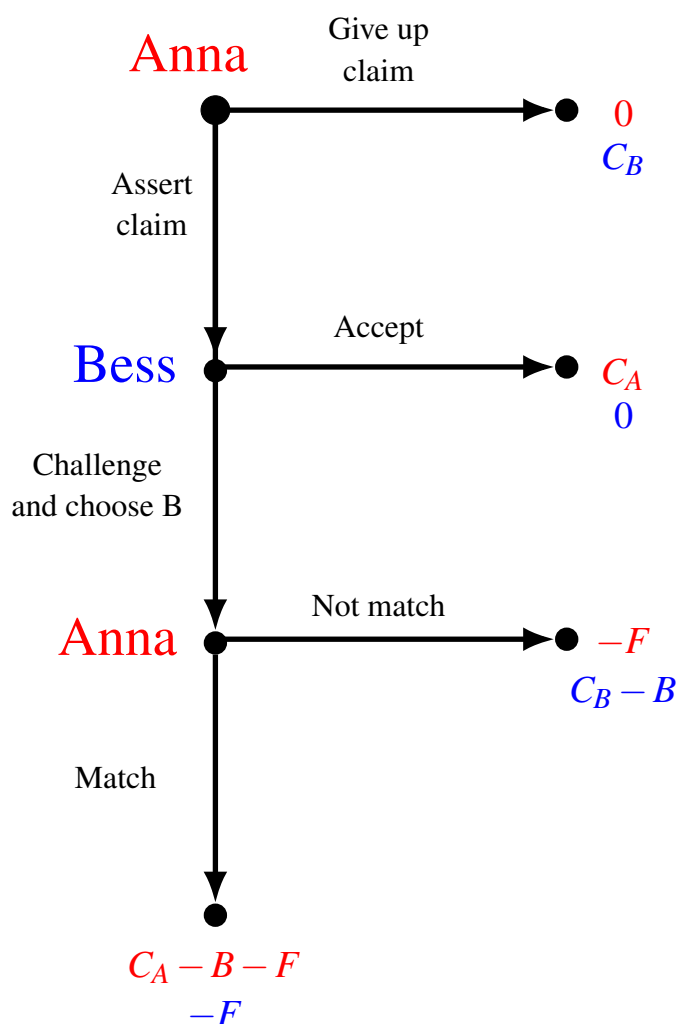


Figure 3.41: The extensive-form game for Part ((c)) of Exercise 3.13

4. General Dynamic Games

4.1 Imperfect Information

There are many situations where players have to make decisions with only partial information about previous moves by other players. Here is an example from my professional experience: in order to discourage copying and cheating in exams, I prepare two versions of the exam, print one version on white paper and the other on pink paper and distribute the exams in such a way that if a student gets, say, the white version then the students on his left and right have the pink version. For simplicity let us assume that there is only one question in the exam. What matters for my purpose is not that the question is indeed different in the two versions, but rather that the students *believe* that they are different and thus refrain from copying from their neighbors. The students, however, are not naïve and realize that I might be bluffing; indeed, introducing differences between the two versions of the exam involves extra effort on my part. Consider a student who finds himself in the embarrassing situation of not having studied for the final exam and is tempted to copy from his neighbor, whom he knows to be a very good student. Let us assume that, if he does not copy, then he turns in a blank exam; in this case, because of his earlier grades in the quarter, he will get a C; on the other hand, if he copies he will get an A if the two versions are identical but will be caught cheating and get an F if the two versions are slightly different. How can we represent such a situation?

Clearly this is a situation in which decisions are made sequentially: first the Professor decides whether to write identical versions (albeit printed on different-color paper) or different versions and then the Student chooses between copying and leaving the exam blank. We can easily represent this situation using a tree as we did with the case of perfect-information games, but the crucial element here is the fact that the Student *does not know* whether the two versions are identical or different. In order to represent this uncertainty (or lack of information) in the mind of the Student, we use the notion of *information set*. An information set for a player is a collection of decision nodes of that player and the

interpretation is that the player does not know at which of these nodes he is making his decision. Graphically, we represent an information set by enclosing the corresponding nodes in a rounded rectangle. Figure 4.1 represents the situation described above.

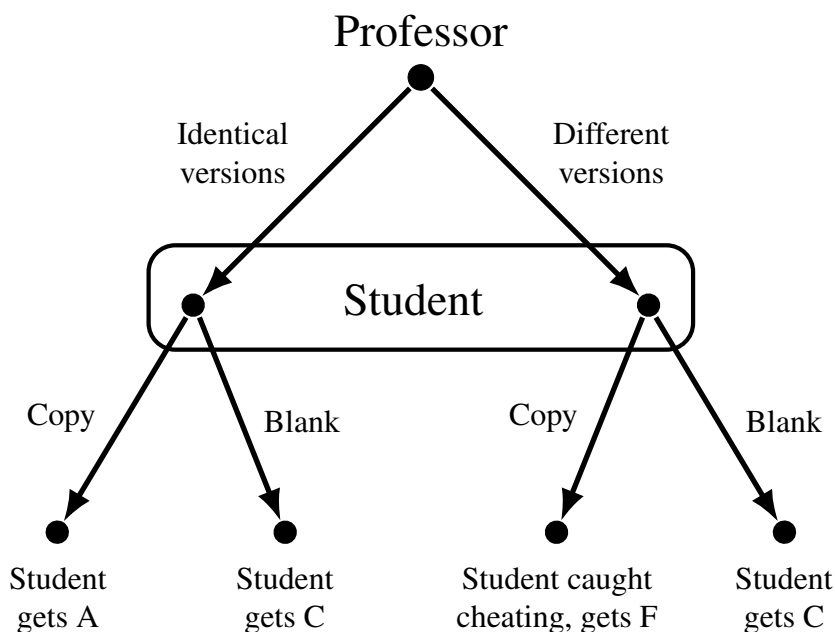


Figure 4.1: An extensive form, or frame, with imperfect information

As usual we need to distinguish between a game-frame and a game. Figure 4.1 depicts a game-frame: in order to obtain a game from it we need to add a ranking of the outcomes for each player. For the moment we shall ignore payoffs and focus on frames. A game-frame such as the one shown in Figure 4.1 is called an *extensive form (or frame) with imperfect information*: in this example it is the Student who has imperfect information, or uncertainty, about the earlier decision of the Professor.

Before we give the definition of extensive form, we need to introduce some additional terminology and notation. Given a directed tree and two nodes x and y we say that y is a *successor of x* , or x is a *predecessor of y* , if there is a sequence of directed edges from x to y .¹

A *partition* of a set H is a collection $\mathcal{H} = \{H_1, \dots, H_m\}$ ($m \geq 1$) of non-empty subsets of H such that:

- (1) any two elements of \mathcal{H} are disjoint (if $H_j, H_k \in \mathcal{H}$ with $j \neq k$ then $H_j \cap H_k = \emptyset$) and
- (2) the elements of \mathcal{H} cover H : $H_1 \cup \dots \cup H_m = H$.

¹If the sequence consists of a single directed edge then we say that y is an *immediate successor* of x or x is the *immediate predecessor* of y .

The definition of extensive form given below allows for perfect information as a special case. The first four items of Definition 4.1.1 (marked by the bullet symbol \bullet), coincide with Definition 3.1.1 in Chapter 3 (which covers the case of perfect information); what is new is the additional item marked by the symbol \star .

Definition 4.1.1 A *finite extensive form (or frame) with perfect recall* consists of the following items.

- \bullet A finite rooted directed tree.
- \bullet A set of players $I = \{1, \dots, n\}$ and a function that assigns one player to every decision node.
- \bullet A set of actions A and a function that assigns one action to every directed edge, satisfying the restriction that no two edges out of the same node are assigned the same action.
- \bullet A set of outcomes O and a function that assigns an outcome to every terminal node.
- \star For every player $i \in I$, a partition \mathcal{D}_i of the set D_i of decision nodes assigned to player i (thus \mathcal{D}_i is a collection of mutually disjoint subsets of D_i whose union is equal to D_i). Each element of \mathcal{D}_i is called an *information set of player i* .
The elements of \mathcal{D}_i satisfy the following restrictions:
 - (1) the actions available at any two nodes in the same information set must be the same (that is, for every $D \in \mathcal{D}_i$, if $x, y \in D$ then the outdegree of x is equal to the outdegree of y and the set of actions assigned to the directed edges out of x is equal to the set of actions assigned to the directed edges out of y),
 - (2) if x and y are two nodes in the same information set then it is not the case that one node is a predecessor of the other,
 - (3) each player has *perfect recall* in the sense that if node $x \in D \in \mathcal{D}_i$ is a predecessor of node $y \in D' \in \mathcal{D}_i$ (thus, by (2), $D \neq D'$), and a is the action assigned to the directed edge out of x in the sequence of edges leading from x to y , then for every node $z \in D'$ there is a predecessor $w \in D$ such that the action assigned to the directed edge out of w in the sequence of edges leading from w to z is that same action a .

The perfect-recall restriction says that if a player takes action a at an information set and later on has to move again, then at the later time she remembers that she took action a at that earlier information set (because every node she is uncertain about at the later time comes after taking action a at that information set). Perfect recall can be interpreted as requiring that *a player always remember what she knew in the past and what actions she herself took in the past*.

Figure 4.2 shows two examples of violation of perfect recall. In the frame shown in Panel (i) Player 1 first chooses between a and b and then chooses between c and d having forgotten his previous choice: he does not remember *what* he chose previously. In the frame shown in Panel (ii) when Player 2 has to choose between e and f she is uncertain whether this is the first time she moves (left node) or the second time (right node): she is uncertain *whether* she moved in the past.




Figure 4.2: Examples of violations of perfect recall

If every information set of every player consists of a single node, then the frame is said to be a *perfect-information frame*: it is easy to verify that, in this case, the last item of Definition 4.1.1 (marked by the symbol \star) is trivially satisfied and thus Definition 4.1.1 coincides with Definition 3.1.1 (Chapter 3). Otherwise (that is, if at least one player has at least one information set that consists of at least two nodes), the frame is said to have *imperfect information*. An example of an extensive frame with imperfect information is the one shown in Figure 4.1. We now give two more examples. In order to simplify the figures, when representing an extensive frame we enclose an information set in a rounded rectangle if and only if that information set contains at least two nodes.

■ **Example 4.1** There are three players, Ann, Bob and Carla. Initially, Ann and Bob are in the same room and Carla is outside the room. Ann moves first, chooses either a red card or a black card from a full deck of cards, shows it to Bob and puts it, face down, on the table. Now Carla enters the room and Bob makes a statement to Carla: he either says “Ann chose a Red card” or he says “Ann chose a Black card”; Bob could be lying or could be telling the truth. After hearing Bob’s statement Carla guesses the color of the card that was picked by Ann. The card is then turned and if Carla’s guess was correct then Ann and Bob give \$1 each to Carla, otherwise Carla gives \$1 to each of Ann and Bob. When drawing an extensive frame to represent this situation, it is important to be careful about what Carla knows, when she makes her guess, and what she is uncertain about. The extensive frame is shown in Figure 4.3. ■

- Carla’s top information set captures the situation she is in after hearing Bob say “Ann chose a black card” and not knowing if he is telling the truth (left node) or he is lying (right node).
- Carla’s bottom information set captures the alternative situation where she hears Bob say “Ann chose a red card” and does not know if he is lying (left node) or telling the truth (right node).
- In both situations Carla knows something, namely what Bob tells her, but lacks information about something else, namely what Ann chose.
- The fact that Bob knows the color of the card chosen by Ann is captured by giving Bob two information sets, each consisting of a single node: Bob’s left node represents the situation he is in when he sees that Ann picked a black card, while his right node represents the situation he is in when he sees that Ann picked a red card.

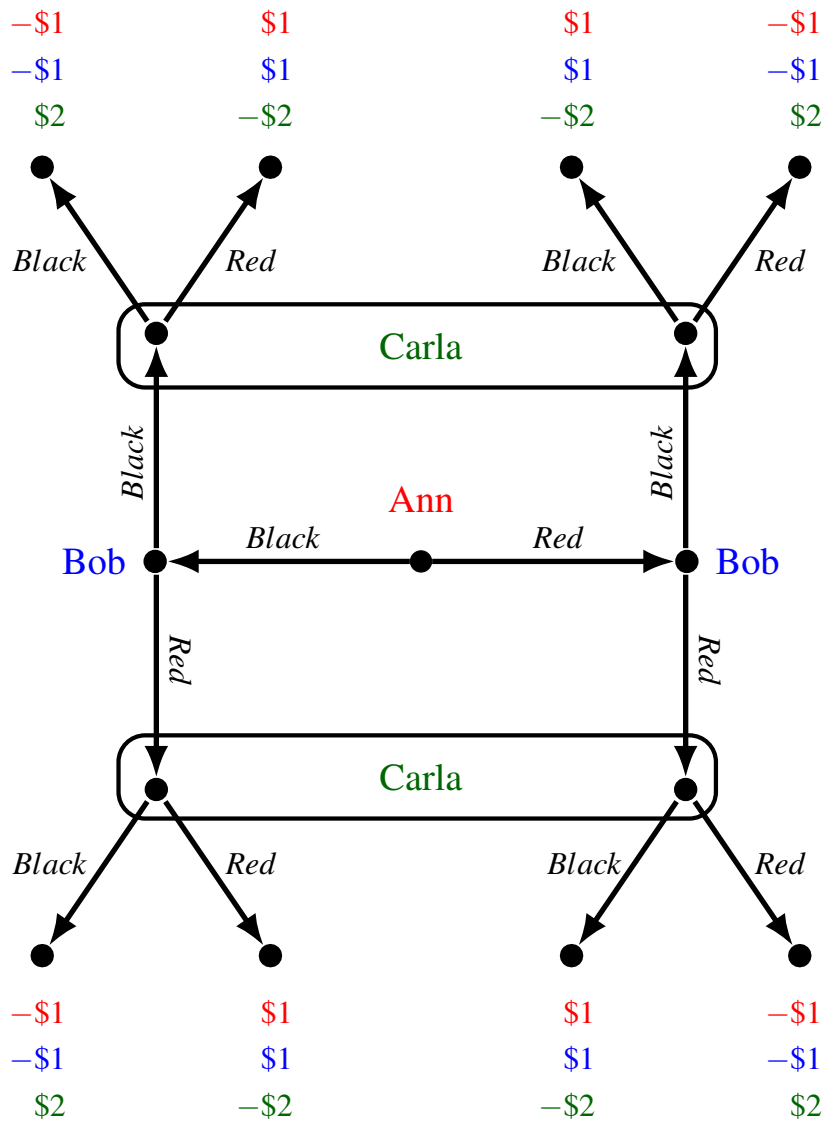


Figure 4.3: The extensive form, or frame, representing Example 4.1

■ **Example 4.2** Yvonne and Fran were both interviewed for the same job, but only one person can be hired. The employer told each candidate: “don’t call me, I will call you if I want to offer you the job”. He also told them that he desperately needs to fill the position and thus, if turned down by one candidate, he will automatically make the offer to the other candidate, without revealing whether he is making a first offer or a “recycled” offer. This situation is represented in the extensive frame shown in Figure 4.4. ■

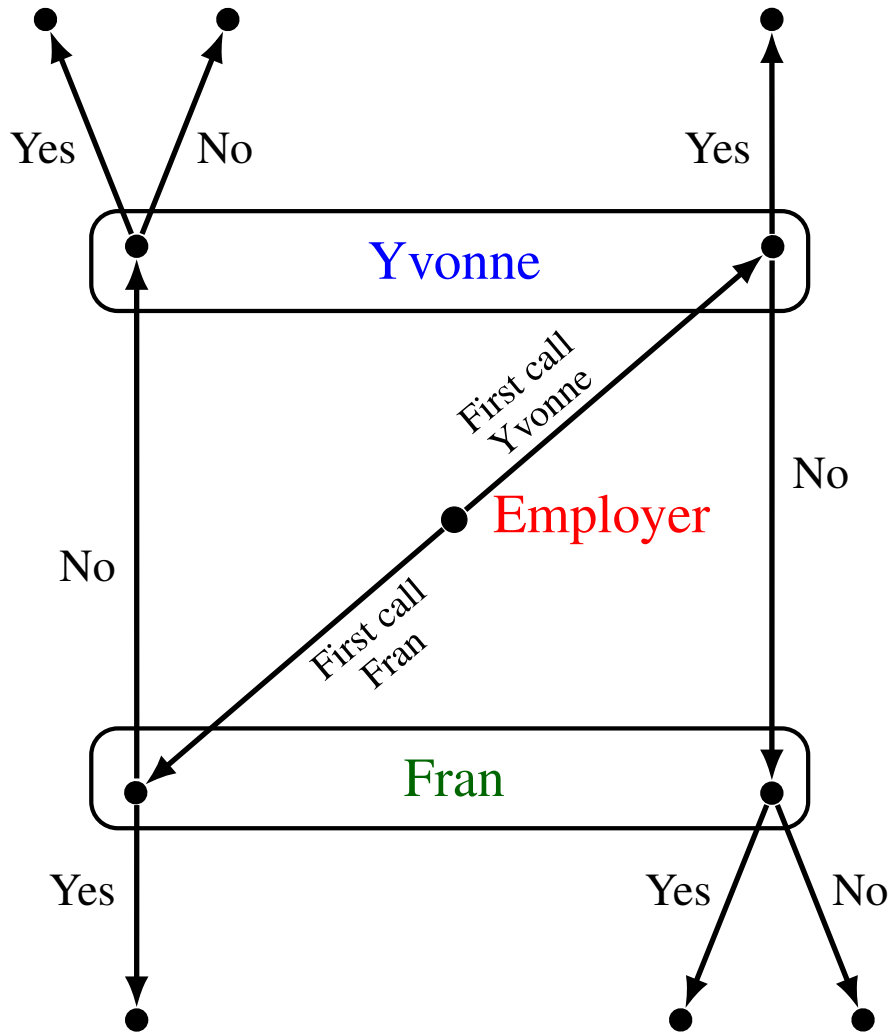


Figure 4.4: The extensive form, or frame, representing Example 4.2

As before, in order to obtain a *game* from an extensive *frame* all we need to do is add a ranking of the outcomes for each player. As usual, the best way to represent such rankings is by means of an ordinal utility function for each player and thus represent an extensive-form game by associating a vector of utilities with each terminal node. For instance, expanding on Example 4.2, suppose that the employer only cares about whether the position is filled or not, prefers filling the position to not filling it, but is indifferent between filling it with Yvonne or with Fran; thus we can assign a utility of 1 for the employer to every outcome where one of the two candidates accepts the offer and a utility of 0 to every other outcome.

Yvonne’s favorite outcome is to be hired if she was the recipient of the first call by the employer; her second best outcome is not to be hired and her worst outcome is to accept a recycled offer (in the latter case Fran would take pleasure telling Yvonne “You took that job?! It was offered to me but I turned it down. Who, in her right mind, would want that job? What’s wrong with you?!”). Thus for Yvonne we can use utilities of 2 (if she accepts a first offer), 1 (if she is not hired) and 0 (if she accepts a recycled offer). Finally, suppose that Fran has preferences similar (but symmetric) to Yvonne’s. Then the extensive frame of Figure 4.4 gives rise to the extensive game shown in Figure 4.5.

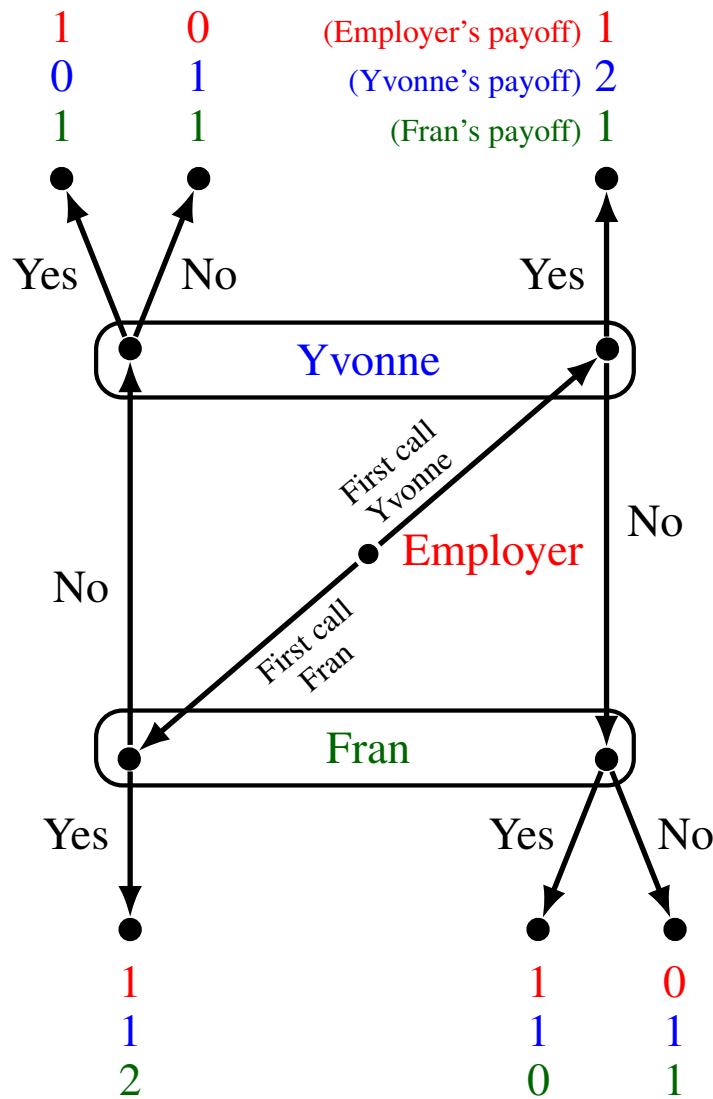


Figure 4.5: A game based on the extensive form of Figure 4.4

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 4.6.1 at the end of this chapter.

4.2 Strategies

The notion of strategy for general extensive games is the same as before: a strategy for Player i is a complete, contingent plan that covers all the possible situations Player i might find herself in. In the case of a perfect-information game a “possible situation” for a player is a decision node of that player; in the general case, where there may be imperfect information, a “possible situation” for a player is an *information set* of that player.

The following definition reduces to Definition 3.3.1 (Chapter 3) if the game is a perfect-information game (where each information set consists of a single node).

Definition 4.2.1 A *strategy* for a player in an extensive-form game is a list of choices, one for every information set of that player.

For example, in the game of Figure 4.5, Yvonne has only one information set and thus a strategy for her is what to do at that information set, namely either say Yes or say No. Yvonne cannot make the plan “if the employer calls me first I will say Yes and if he calls me second I will say No”, because when she receives the call she is not told if this is a first call or a recycled call and thus she cannot make her decision dependent on information she does not have.

As in the case of perfect-information games, the notion of strategy allows us to associate with every extensive-form game a strategic-form game. For example, the strategic form associated with the game of Figure 4.5 is shown in Figure 4.6 with the Nash equilibria highlighted.

		Yvonne					
		Yes			No		
Employer	Yvonne first	1	2	1	1	1	0
	Fran first	1	1	2	1	1	2
Fran: Yes							
.....							
		Yvonne					
		Yes			No		
Employer	Yvonne first	1	2	1	0	1	1
	Fran first	1	0	1	0	1	1
Fran: No							

Figure 4.6: The strategic form of the game of Figure 4.5 with the Nash equilibria highlighted

As another example, consider the extensive form of Figure 4.3 and view it as a game by assuming that each player is selfish and greedy (only cares about how much money he/she gets and prefers more money to less). Then the associated strategic form is shown in Figure 4.7, where Bob’s strategy (x,y) means “I say x if Ann chose a black card and I say y if Ann chose a red card”. Thus (R,B) means “if Ann chose a black card I say Red and if Ann chose a red card I say Black” (that is, Bob plans to lie in both cases). Similarly, Carla’s strategy (x,y) means “I guess x if Bob tells me Black and I guess y if Bob tells me Red”. Thus (B,R) means “if Bob tells me Black I guess Black and if Bob tells me Red I guess Red” (that is, Carla plans to repeat what Bob says).

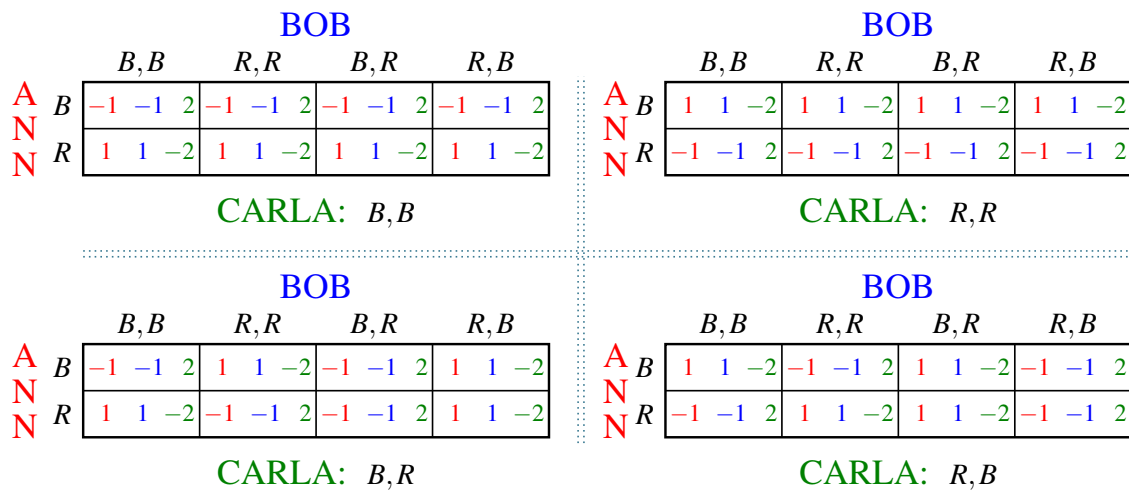


Figure 4.7: The strategic form of the game of Figure 4.3

In order to “solve” an extensive-form game we could simply construct the associated strategic-form game and look for the Nash equilibria. However, we saw in Chapter 3 that in the case of perfect-information games not all Nash equilibria of the associated strategic form can be considered “rational solutions” and we introduced the notion of backward induction to select the “reasonable” Nash equilibria. What we now need is a generalization of the notion of backward induction that can be applied to general extensive-form games. This generalization is called *subgame-perfect equilibrium*. First we need to define the notion of subgame.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 4.6.2 at the end of this chapter.

4.3 Subgames

Roughly speaking, a subgame of an extensive-form game is a portion of the game that could be a game in itself. What we need to be precise about is the meaning of “portion of the game”.

Definition 4.3.1 A *proper subgame* of an extensive-form game is obtained as follows:

1. Start from a decision node x , different from the root, whose information set consists of node x only and enclose in an oval node x itself and all its successors.
2. If the oval does not “cut” any information sets (that is, there is no information set S and two nodes $y, z \in S$ such that y is a successor of x while z is not) then what is included in the oval is a proper subgame, otherwise it is not.

The reason why we use the qualifier ‘proper’ is that one could start from the root, in which case one would end up taking the entire game and consider this as a (trivial) subgame (just like any set is a subset of itself; a proper subgame is analogous to a proper subset).

Consider, for example, the extensive-form game of Figure 4.8. There are three possible starting points for identifying a proper subgame: nodes x , y and z (the other nodes fail to satisfy condition (1) of Definition 4.3.1).

1. Starting from node x and including all of its successors, we do indeed obtain a proper subgame, which is the portion included in the oval on the left.
2. Starting from node y and including all of its successors we obtain the portion of the game that is included in the oval on the right; in this case, condition (2) of Definition 4.3.1 is violated, since we are cutting the top information set of Player 3; hence the portion of the game inside this oval is not a proper subgame.
3. Finally, starting from node z and including all of its successors, we do obtain a proper subgame, which is the portion included in the oval at the bottom.

Thus the game of Figure 4.8 has two proper subgames.

Definition 4.3.2 A proper subgame of an extensive-form game is called *minimal* if it does not strictly contain another proper subgame (that is, if there is no other proper subgame which is strictly contained in it).

For example, the game shown in Figure 4.9 on the following page has three proper subgames, one starting at node x , another at node y and the third at node z . The ones starting at nodes x and z are minimal subgames, while the one that starts at node y is not a minimal subgame, since it strictly contains the one that starts at node z .

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 4.6.3 at the end of this chapter.

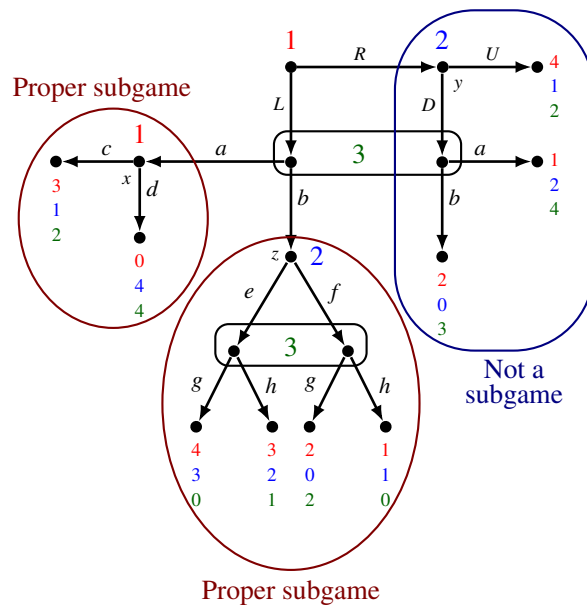


Figure 4.8: An extensive-form game with two proper subgames.




Figure 4.9: An extensive-form game with three proper subgames, two of which are minimal.

4.4 Subgame-perfect equilibrium

A subgame-perfect equilibrium of an extensive-form game is a Nash equilibrium of the entire game which remains an equilibrium in every proper subgame. Consider an extensive-form game and let s be a strategy profile for that game. Let G be a proper subgame. Then the *restriction of s to G* , denoted by $s|_G$, is that part of s which prescribes choices at every information set of G and only at those information sets.

For example, consider the extensive-form game of Figure 4.9 and the strategy profile

$$\left(\underbrace{(a, C)}_{1\text{'s strategy}}, \underbrace{(d, f, E)}_{2\text{'s strategy}}, \underbrace{(h, B)}_{3\text{'s strategy}} \right)$$

Let G be the subgame that starts at node y of Player 2. Then

$$s|_G = \left(\underbrace{C}_{1\text{'s strategy in } G}, \underbrace{(f, E)}_{2\text{'s strategy in } G}, \underbrace{B}_{3\text{'s strategy in } G} \right)$$

Definition 4.4.1 . Given an extensive-form game, let s be a strategy profile for the entire game. Then s is a *subgame-perfect equilibrium* if

1. s is a Nash equilibrium of the entire game and
2. for every proper subgame G , $s|_G$ (the restriction of s to G) is a Nash equilibrium of G .

For example, consider again the extensive-form game of Figure 4.9 and the strategy profile $s = ((a, C), (d, f, E), (h, B))$. Then s is a Nash equilibrium of the entire game: Player 1's payoff is 2 and if he were to switch to any strategy where he plays b his payoff would be 0; Player 2's payoff is 1 and if she were to switch to any strategy where she plays c her payoff would be 0; Player 3's payoff is 2 and if he were to switch to any strategy where he plays g his payoff would be 1. However, s is not a subgame-perfect equilibrium, because the restriction of s to the proper subgame that starts at node z of Player 1, namely (C, E) , is not a Nash equilibrium of that subgame: in that subgame, for Player 2 the unique best reply to C is F .

One way of finding the subgame-perfect equilibria of a given game is to first find the Nash equilibria and then, for each of them, check if it satisfies condition (2) of Definition 4.4.1. However, this is not a practical way to proceed. A quicker and easier way is to apply the following algorithm, which generalizes the backward-induction algorithm for games with perfect information (Definition 3.2.1, Chapter 3).

Definition 4.4.2 Given an extensive-form game, the *subgame-perfect equilibrium algorithm* is the following procedure.

1. Start with a minimal proper subgame and select a Nash equilibrium of it.
2. Delete the selected proper subgame and replace it with the payoff vector associated with the selected Nash equilibrium, making a note of the strategies that constitute the Nash equilibrium. This yields a smaller extensive-form game.
3. Repeat Steps 1 and 2 in the smaller game so obtained.

For example, let us apply the algorithm to the game of Figure 4.9. Begin with the proper subgame that starts at node x of Player 2, shown in Figure 4.10 with its associated strategic form, where the unique Nash equilibrium (d, h) is highlighted. Note that this is a game only between Players 2 and 3 and thus in Figure 4.10 we only show the payoffs of these two players.




Figure 4.10: A minimal proper subgame of the game of Figure 4.9 and its strategic form

Now we delete the proper subgame, thereby turning node x into a terminal node to which we attach the full payoff vector associated, in the original game, with the terminal node following history adh , namely $(2, 1, 2)$.

Hence we obtain the smaller game shown in Figure 4.11.




Figure 4.11: The reduced game after replacing a proper minimal subgame in the game of Figure 4.9

Now, in the reduced game of Figure 4.11 we select the only minimal proper subgame, namely the one that starts at the bottom decision node of Player 1. This subgame is shown in Figure 4.12 together with its associated strategic form. The unique Nash equilibrium of this subgame is (C, F) .

Then, in the reduced game of Figure 4.11, we replace the selected proper subgame with the payoff vector associated with the history $beACF$, namely $(1, 1, 1)$, thus obtaining the smaller game shown in Figure 4.13. The game of Figure 4.13 has a unique proper subgame, which has a unique Nash equilibrium, namely (f, A) .




Figure 4.12: A minimal proper subgame of the game of Figure 4.11 and its strategic form




Figure 4.13: The reduced game after replacing a proper minimal subgame in the game of Figure 4.11

Replacing the subgame with the payoff vector associated with the history bfA we get the smaller game shown in Figure 4.14.




Figure 4.14: The reduced game after replacing the subgame in the game of Figure 4.13

In the reduced game of Figure 4.14 the unique Nash equilibrium is b . Now patching together the choices selected during the application of the algorithm we get the following subgame-perfect equilibrium for the game of Figure 4.9: $((b, C), (d, f, F), (h, A))$.

As a second example, consider the game of Figure 4.15 (which reproduces Figure 4.8).




Figure 4.15: Copy of Figure 4.8

Begin with the subgame that starts at node x and replace it with the payoff vector $(3, 1, 2)$. Next replace the subgame that starts at node z with the payoff vector $(3, 2, 1)$ which corresponds to the Nash equilibrium (e, h) of that subgame, so that the game is reduced to the one shown in Figure 4.16, together with its strategic form.




Figure 4.16: The game of Figure 4.15 reduced after solving the proper subgames, together with the associated strategic form with the Nash equilibria highlighted

The reduced game of Figure 4.16 has two Nash equilibria: (L, D, a) and (R, U, b) . Thus the game of Figure 4.15 has two subgame-perfect equilibria:

$$\left(\underbrace{(L, c)}_{\text{Player 1}}, \underbrace{(D, e)}_{\text{Player 2}}, \underbrace{(a, h)}_{\text{Player 3}} \right) \quad \text{and} \quad \left(\underbrace{(R, c)}_{\text{Player 1}}, \underbrace{(U, e)}_{\text{Player 2}}, \underbrace{(b, h)}_{\text{Player 3}} \right).$$



- As shown in the last example, it is possible that – when applying the subgame-perfect equilibrium algorithm – one encounters a proper subgame or a reduced game that has several Nash equilibria. In this case one Nash equilibrium must be selected in order to continue the procedure and in the end one obtains one subgame-perfect equilibrium. One then has to repeat the procedure by selecting a different Nash equilibrium and thus obtain a different subgame-perfect equilibrium, and so on. This is similar to what happens with the backward-induction algorithm in perfect-information games.
- It is also possible that – when applying the subgame-perfect equilibrium algorithm – one encounters a proper subgame or a reduced game that has no Nash equilibria.² In such a case the game under consideration does not have any subgame-perfect equilibria.
- When applied to perfect-information games, the notion of subgame-perfect equilibrium coincides with the notion of backward-induction solution. Thus subgame-perfect equilibrium is a generalization of backward induction.
- For extensive-form games that have no proper subgames (for example, the game of Figure 4.3) the set of Nash equilibria coincides with the set of subgame-perfect equilibria. In general, however, the notion of subgame-perfect equilibrium is a refinement of the notion of Nash equilibrium.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 4.6.4 at the end of this chapter.

4.5 Games with chance moves

So far we have only considered games where the outcomes do not involve any uncertainty. As a way of introducing the topic discussed in Part II, in this section we consider games where uncertain, probabilistic events are incorporated in the extensive form.

We begin with an example: There are three cards, one black and two red. They are shuffled well and put face down on the table. Adele picks the top card, looks at it without showing it to Ben and then tells Ben either “the top card is black” or “the top card is red”: she could be telling the truth or she could be lying. Ben then has to guess the true color of the top card. If he guesses correctly he gets \$9 from Adele, otherwise he gives her \$9. How can we represent this situation?

Whether the top card is black or red is not the outcome of a player’s decision, but the outcome of a random event, namely the shuffling of the cards. In order to capture this random event we introduce a fictitious player called *Nature* or *Chance*. We assign a probability distribution to Nature’s “choices”. In this case, since one card is black and the other two are red, the probability that the top card is black is $\frac{1}{3}$ and the probability that the top card is red is $\frac{2}{3}$. Note that we don’t assign payoffs to Nature and thus the only ‘real’ players are Adele and Ben. The situation can be represented as shown in Figure 4.17, where the numbers associated with the terminal nodes are dollar amounts.

²We will see in Part II that, when payoffs are cardinal and one allows for mixed strategies, then every finite game has at least one Nash equilibrium in mixed strategies.

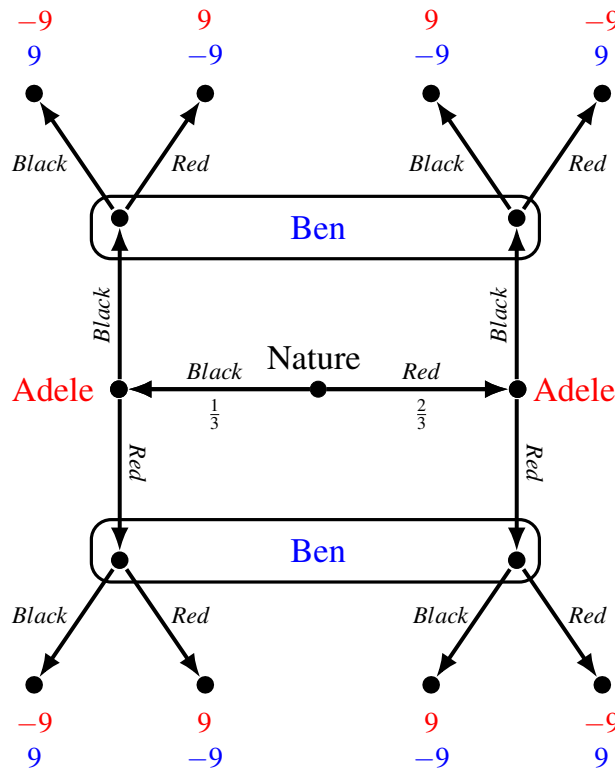


Figure 4.17: An extensive form with a chance move

Clearly, the notion of strategy is not affected by the presence of chance moves. In the game of Figure 4.17 Adele has four strategies and so does Ben. However, we do encounter a difficulty when we try to write the associated strategic form. For example, consider the following strategy profile: $((B, R), (B, B))$ where Adele’s strategy is to be truthful (say “Black” if she sees a black card and say “Red” if she sees a red card) and Ben’s strategy is to guess Black no matter what Adele says. What is the outcome in this case? It depends on what the true color of the top card is and thus the outcome is a probabilistic one:

$$\begin{pmatrix} \text{outcome} & \text{Adele gives \$9 to Ben} & \text{Ben gives \$9 to Adele} \\ \text{probability} & \frac{1}{3} & \frac{2}{3} \end{pmatrix}$$

We call such probabilistic outcomes *lotteries*. In order to convert the game-frame into a game we need to specify how the players rank probabilistic outcomes. Consider the case where Adele is selfish and greedy, in the sense that she only cares about her own wealth and she prefers more money to less. Then, from her point of view, the above probabilistic outcome reduces to the following monetary lottery $\begin{pmatrix} -\$9 & \$9 \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$.

If Ben is also selfish and greedy, then he views the same outcome as the lottery $\begin{pmatrix} \$9 & -\$9 \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$.

How do we convert a lottery into a payoff or utility? The general answer to this question will be provided in Chapter 5. Here we consider one possibility, which is particularly simple.

Definition 4.5.1 Given a lottery whose outcomes are sums of money

$$\begin{pmatrix} \$x_1 & \$x_2 & \dots & \$x_n \\ p_1 & p_2 & \dots & p_n \end{pmatrix}$$

(with $p_i \geq 0$, for all $i = 1, 2, \dots, n$, and $p_1 + p_2 + \dots + p_n = 1$) the *expected value* of the lottery is the following sum of money: $(x_1p_1 + x_2p_2 + \dots + x_np_n)$.

We call lotteries whose outcomes are sums of money, *money lotteries*.

For example, the expected value of the lottery

$$\begin{pmatrix} \$5 & \$15 & \$25 \\ \frac{1}{5} & \frac{2}{5} & \frac{2}{5} \end{pmatrix}$$

is

$$\$ \left[5 \left(\frac{1}{5} \right) + 15 \left(\frac{2}{5} \right) + 25 \left(\frac{2}{5} \right) \right] = \$ (1 + 6 + 10) = \$17$$

and the expected value of the lottery

$$\begin{pmatrix} -\$9 & \$9 \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$$

is \$3.

Definition 4.5.2 A player is defined to be *risk neutral* if she considers a money lottery to be just as good as its expected value. Hence a risk neutral person ranks money lotteries according to their expected value.^a

^aIt is important to stress that our focussing on the case of risk neutrality should *not* be taken to imply that a rational individual ought to be risk neutral nor that risk neutrality is empirically particularly relevant. At this stage we assume risk neutrality only because it yields a very simple type of preference over money lotteries and allows us to introduce the notion of backward induction without the heavy machinery of expected utility theory.

For example, consider the following money lotteries:

$$L_1 = \begin{pmatrix} \$5 & \$15 & \$25 \\ \frac{1}{5} & \frac{2}{5} & \frac{2}{5} \end{pmatrix}, L_2 = \begin{pmatrix} \$16 \\ 1 \end{pmatrix} \text{ and } L_3 = \begin{pmatrix} \$0 & \$32 & \$48 \\ \frac{5}{8} & \frac{1}{8} & \frac{1}{4} \end{pmatrix}.$$

The expected value of L_1 is \$17 and the expected value of both L_2 and L_3 is \$16. Thus a risk-neutral player would have the following ranking: $L_1 \succ L_2 \sim L_3$, that is, she would prefer L_1 to L_2 and be indifferent between L_2 and L_3 .

For a selfish and greedy player who is risk neutral we can take the expected value of a money lottery as the utility of that lottery. For example, if we make the assumption that, in the extensive form of Figure 4.17, Adele and Ben are selfish, greedy and risk neutral then we can associate a strategic-form game to it as shown in Figure 4.18. Note that inside each cell we have two numbers: the first is the utility (= expected value) of the underlying money lottery as perceived by Adele and the second number is the utility (= expected value) of the underlying money lottery as perceived by Ben.

The first element of Adele's strategy is what she says if she sees a black card and the second element is what she says if she sees a red card. The first element of Ben's strategy is what he guesses if Adele says "Red", the second element is what he guesses if Adele says "Black".

		Ben			
		<i>BB</i>	<i>BR</i>	<i>RB</i>	<i>RR</i>
A d e l e	<i>BB</i>	3 -3	-3 3	3 -3	-3 3
	<i>BR</i>	3 -3	9 -9	-9 9	-3 3
	<i>RB</i>	3 -3	-9 9	9 -9	-3 3
	<i>RR</i>	3 -3	3 -3	-3 3	-3 3

Figure 4.18: The strategic form of the game of Figure 4.17 when the two players are selfish, greedy and risk neutral.

We conclude this section with one more example.

■ **Example 4.3** There are three unmarked, opaque envelopes. One contains \$100, one contains \$200 and the third contains \$300. They are shuffled well and then one envelope is given to Player 1 and another is given to Player 2 (the third one remains on the table).

- Player 1 opens her envelope and checks its content without showing it to Player 2. Then she either says “pass” – in which case each player gets to keep his/her envelope – or she asks Player 2 to trade his envelope for hers.
- Player 2 is not allowed to see the content of his envelope and has to say either Yes or No. If he says No, then the two players get to keep their envelopes. If Player 2 says Yes, then they trade envelopes and the game ends. Each player is selfish, greedy and risk neutral.

This situation is represented by the extensive-form game shown in Figure 4.19, where (100, 200) means that Player 1 gets the envelope with \$100 and Player 2 gets the envelope with \$200, etc.; P stands for “pass” and T for “suggest a trade”; Y for “Yes” and N for “No”.

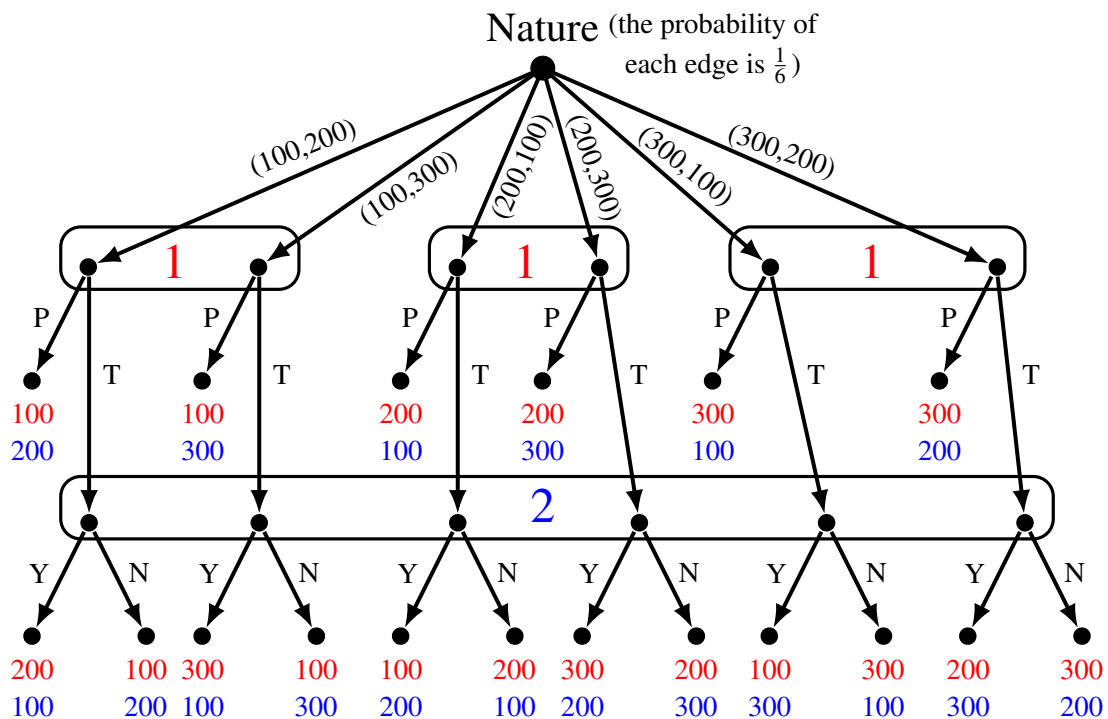


Figure 4.19: The extensive-form game of Example 4.3

In this game Player 1 has eight strategies. One possible strategy is: “if I get \$100 I will pass, if I get \$200 I will propose a trade, if I get \$300 I will pass”: we will use the shorthand PTP for this strategy. Similarly for the other strategies. Player 2 has only two strategies: Yes and No.

The strategic form associated with the game of Figure 4.19 is shown in Figure 4.20, where the Nash equilibria are highlighted.

		Player 2			
		Y		H	
Player 1	<i>PPP</i>	200	200	200	200
	<i>PPT</i>	150	250	200	200
	<i>PTP</i>	200	200	200	200
	<i>PTT</i>	150	250	200	200
	<i>TPP</i>	250	150	200	200
	<i>TPT</i>	200	200	200	200
	<i>TTP</i>	250	150	200	200
	<i>TTT</i>	200	200	200	200

Figure 4.20: The strategic form of the game of Figure 4.19

How did we get those payoffs? Consider, for example, the first cell. Given the strategies PPP and Y, the outcomes are:

(\$100, \$200) with probability $\frac{1}{6}$, (\$100, \$300) with probability $\frac{1}{6}$,
 (\$200, \$100) with probability $\frac{1}{6}$, (\$200, \$300) with probability $\frac{1}{6}$,
 (\$300, \$100) with probability $\frac{1}{6}$, (\$300, \$200) with probability $\frac{1}{6}$.

Being risk neutral, Player 1 views his corresponding money lottery as equivalent to getting its expected value $\$(100 + 100 + 200 + 200 + 300 + 300)(\frac{1}{6}) = \200 . Similarly for Player 2 and for the other cells.

Since the game of Figure 4.19 has no proper subgames, all the Nash equilibria are also subgame-perfect equilibria. Are some of the Nash equilibria more plausible than others? For Player 1 all the strategies are weakly dominated, except for TPP and TTP. Elimination of the weakly dominated strategies leads to a game where Y is strictly dominated for Player 2. Thus one could argue that (TPP, N) and (TTP, N) are the most plausible equilibria; in both of them Player 2 refuses to trade.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 4.6.5 at the end of this chapter.

4.6 Exercises

4.6.1 Exercises for Section 4.1: Imperfect information

The answers to the following exercises are in Section 4.7 at the end of this chapter.

Exercise 4.1 Amy and Bill simultaneously write a bid on a piece of paper. The bid can only be either 2 or 3. A referee then looks at the bids, announces the amount of the lowest bid (without revealing who submitted it) and invites Amy to either pass or double her initial bid.

- The outcome is determined by comparing Amy's final bid to Bill's bid: if one is greater than the other then the higher bidder gets the object and pays his/her own bid; if they are equal then Bill gets the object and pays his bid.

Represent this situation by means of two alternative extensive frames.

Note: (1) when there are simultaneous moves we have a choice as to which player we select as moving first: the important thing is that the second player does not know what the first player did;

(2) when representing, by means of information sets, what a player is uncertain about, we typically assume that a player is smart enough to deduce relevant information, even if that information is not explicitly given to him/her. ■

Exercise 4.2 Consider the following situation. An incumbent monopolist decides at date 1 whether to build a small plant or a large plant. At date 2 a potential entrant observes the plant built by the incumbent and decides whether or not to enter.

- If she does not enter then her profit is 0 while the incumbent's profit is \$25 million with a small plant and \$20 million with a large plant.
- If the potential entrant decides to enter, she pays a cost of entry equal to \$K million.
- At date 3 the two firms simultaneously decide whether to produce high output or low output.
- The profits of the firms are as shown in the following table, where 'L' means 'low output' and 'H' means 'high output' (these figure do not include the cost of entry for the entrant; thus you need to subtract that cost for the entrant); in each cell, the first number is the profit of the entrant (in millions of dollars) and the second is the profit of the incumbent.

		Incumbent				Incumbent	
		L	H			L	H
Entrant	L	10, 10	7, 7	Entrant	L	10, 7	5, 9
	H	7, 6	4, 3		H	7, 3	4, 5
If Incumbent has small plant				If Incumbent has large plant			

Draw an extensive-form game that represents this situation, assuming that each player is selfish and greedy (that is, cares only about its own profits and prefers more money to less). ■

4.6.2 Exercises for Section 4.2: Strategies

The answers to the following exercises are in Section 4.7 at the end of this chapter.

Exercise 4.3 Write the strategic-form game-frame of the extensive form of Exercise 4.1 (that is, instead of writing payoffs in each cell, you write the outcome). Verify that the strategic forms of the two possible versions of the extensive form are identical. ■

Exercise 4.4 Consider the extensive-form game of Exercise 4.2.

- (a) Write down in words one of the strategies of the potential entrant.
- (b) How many strategies does the potential entrant have?
- (c) Write the strategic-form game associated with the extensive-form game.
- (d) Find the Nash equilibria for the case where $K = 2$.

4.6.3 Exercises for Section 4.3: Subgames

The answers to the following exercises are in Section 4.7 at the end of this chapter.

Exercise 4.5 How many proper subgames does the extensive form of Figure 4.3 have?
namaste ■

Exercise 4.6 How many proper subgames does the extensive form of Figure 4.5 have?
namaste ■

Exercise 4.7 Consider the extensive game Figure 4.21.

- (a) How many proper subgames does the game have?
- (b) How many of those proper subgames are minimal?




Figure 4.21: The game of Exercise 4.7

Exercise 4.8 Consider the extensive game Figure 4.22.
 (a) How many proper subgames does the game have?
 (b) How many of those proper subgames are minimal?




Figure 4.22: The game of Exercise 4.8

4.6.4 Exercises for Section 4.4: Subgame-perfect equilibrium

The answers to the following exercises are in Section 4.7 at the end of this chapter.

Exercise 4.9 Find the Nash equilibria and the subgame-perfect equilibria of the game shown in Figure 4.23.




Figure 4.23: The game of Exercise 4.9

Exercise 4.10 Find the subgame-perfect equilibria of the game shown in Figure 4.24. ■




Figure 4.24: The game of Exercise 4.10

Exercise 4.11 Find the subgame-perfect equilibria of the game shown in Figure 4.25, assuming the following about the players' preferences. Both Amy and Bill are selfish and greedy, are interested in their own net gain, Amy values the object at \$5 and Bill at \$4.

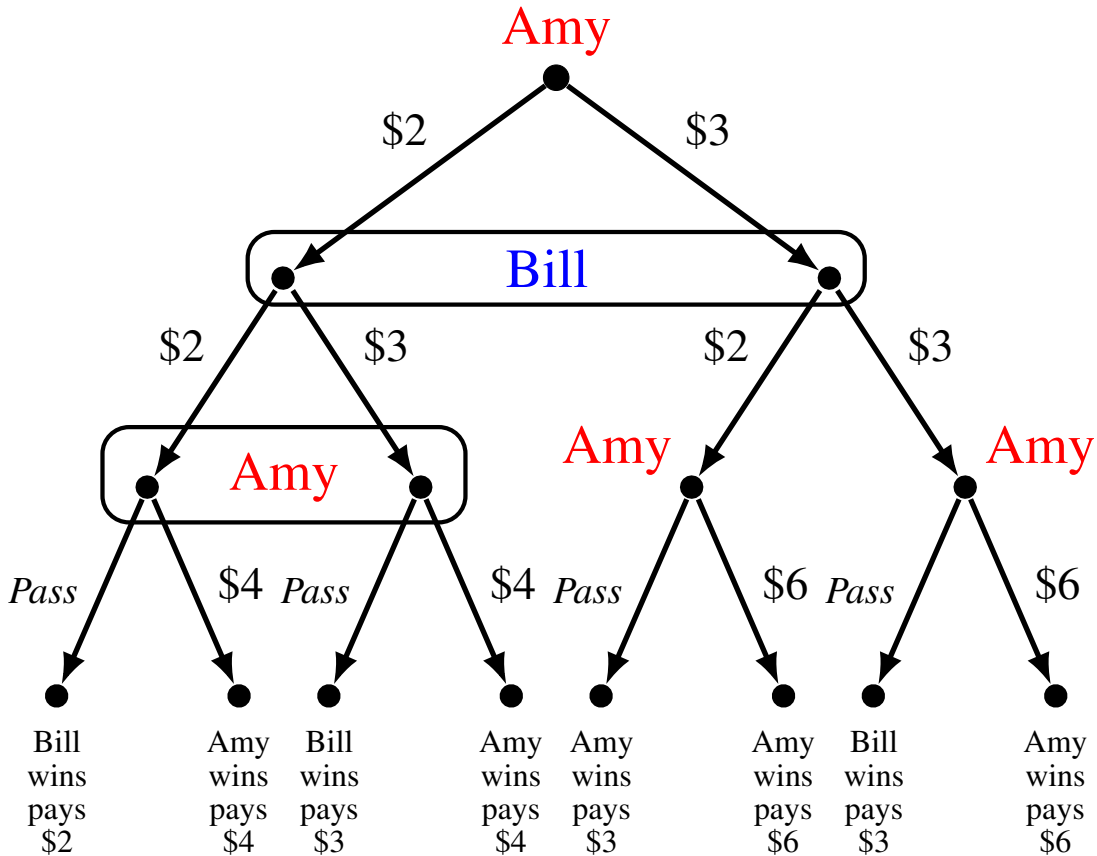


Figure 4.25: The game of Exercise 4.11

4.6.5 Exercises for Section 4.5: Games with chance moves

The answers to the following exercises are in Section 4.7 at the end of this chapter.

Exercise 4.12 Modify the game of Example 4.3 as follows: Player 2 is allowed to privately check the content of his envelope before he decides whether or not to accept Player 1's proposal.

- Represent this situation as an extensive-form game.
- List all the strategies of Player 1 and all the strategies of Player 2.

Exercise 4.13 Three players, Avinash, Brian and John, play the following game. Two cards, one red and the other black, are shuffled well and put face down on the table. Brian picks the top card, looks at it without showing it to the other players (Avinash and John) and puts it back face down on the table. Then Brian whispers either “Black” or “Red” in Avinash’s ear, making sure that John doesn’t hear. Avinash then tells John either “Black” or “Red”. Note that both players could be lying. Finally John announces either “Black” or “Red” and this exciting game ends.

The payoffs are as follows: if John’s final announcement matches the true color of the card Brian looked at, then Brian and Avinash give \$2 each to John. In every other case John gives \$2 each to Brian and Avinash.

- (a) Represent this situation as an extensive-form.
- (b) Write the corresponding strategic form assuming that the players are selfish, greedy and risk neutral. [At least try to fill in a few cells in at least one table.]

Exercise 4.14 Consider the following highly simplified version of Poker.

- There are three cards, marked A, B and C. A beats B and C, B beats C.
 - There are two players, Yvonne and Zoe. Each player contributes \$1 to the pot before the game starts. The cards are then shuffled and the top card is given, face down, to Yvonne and the second card (face down) to Zoe. Each player looks at, and only at, her own card: she does not see the card of the other player nor the remaining card.
 - Yvonne, the first player, may pass, or bet \$1. If she passes, the game ends, the cards are turned and the pot goes to the high-card holder (recall that A beats B and C, B beats C).
 - If Yvonne bets, then Zoe can fold, in which case the game ends and the pot goes to Yvonne, or Zoe can see by betting \$1, in which case the game ends, the cards are turned and the pot goes to the high-card holder. Both players are selfish, greedy and risk neutral.
- (a) Draw the extensive-form game.
 - (b) How many strategies does Yvonne have?
 - (c) How many strategies does Zoe have?
 - (d) Consider the following strategies. For Yvonne: If A pass, if B pass, if C bet. For Zoe: if Yvonne bets, I will fold no matter which card I get. Calculate the corresponding payoffs.
 - (e) Redo the same with the following strategies. For Yvonne: If A pass, if B pass, if C bet. For Zoe: see always (that is, no matter what card she gets).
 - (f) Now that you have understood how to calculate the payoffs, represent the entire game as a normal form game, assigning the rows to Yvonne and the columns to Zoe. [This might take you the entire night, so make sure you have a lot of coffee!]
 - (g) What strategies of Yvonne are weakly dominated? What strategies of Zoe are weakly dominated?
 - (h) What do you get when you apply the procedure of iterative elimination of weakly dominated strategies?

Exercise 4.15 — * Challenging Question ***.**

In an attempt to reduce the deficit, the government of Italy has decided to sell a 14th century palace near Rome. The palace is in disrepair and is not generating any revenue for the government. From now on we will call the government Player G. A Chinese millionaire has offered to purchase the palace for $\$p$. Alternatively, Player G can organize an auction among n interested parties ($n \geq 2$). The participants to the auction (we will call them players) have been randomly assigned labels $1, 2, \dots, n$. Player i is willing to pay up to $\$p_i$ for the palace, where $\$p_i$ is a positive integer. For the auction assume the following:

1. it is a simultaneous, sealed-bid second-price auction,
2. bids must be non-negative integers,
3. each player only cares about his own wealth,
4. the tie-breaking rule for the auction is that the palace is given to that player who has the lowest index (e.g. if the highest bid was submitted by Players 3, 7 and 12 then the palace is given to Player 3).

All of the above is commonly known among everybody involved, as is the fact that for every $i, j \in \{1, \dots, n\}$ with $i \neq j$, $p_i \neq p_j$.

We shall consider four different scenarios. In all scenarios you can assume that the p_i 's are common knowledge.

Scenario 1. Player G first decides whether to sell the palace to the Chinese millionaire or make a public and irrevocable decision to auction it.

- (a) Draw the extensive form of this game for the case where $n = 2$ and the only possible bids are $\$1$ and $\$2$. [List payoffs in the following order: first G then 1 then 2; don't forget that this is a *second-price* auction.]
- (b) For the general case where $n \geq 2$ and every positive integer is a possible bid, find a pure-strategy subgame-perfect equilibrium of this game. What are the players' payoffs at the equilibrium?

Scenario 2. Here we assume that $n = 2$, and $p_1 > p_2 + 1 > 2$.

- First Player G decides whether to sell the palace to the Chinese or make a public and irrevocable decision to auction it. In the latter case he first asks Player 2 to publicly announce whether or not he is going to participate in the auction.
- If Player 2 says Yes, then he has to pay $\$1$ to Player G as a participation fee, which is non-refundable. If he says No, then she is out of the game.
- After Player 2 has made his announcement (and paid his fee if he decided to participate), Player 1 is asked to make the same decision (participate and pay a non-refundable fee of $\$1$ to Player G or stay out); Player 1 knows Player 2's decision when he makes his own decision.

After both players have made their decisions, player G proceeds as follows:

- if both 1 and 2 said Yes, then he makes them play a simultaneous second-price auction,
 - if only one player said Yes, then he is asked to put an amount $\$x$ of his choice in an envelope (where x is a positive integer) and give it to Player G in exchange for the palace,
 - if both 1 and 2 said No, then G is no longer bound by his commitment to auction the palace and he sells it to the Chinese.
- (c) Draw the extensive form of this game for the case where the only possible bids are $\$1$ and $\$2$ and also $x \in \{1, 2\}$ [List payoffs in the following order: first G then 1 then 2; again, don't forget that this is a *second-price* auction.]
- (d) For the general case where all possible bids are allowed (subject to being positive integers) and x can be any positive integer, find a pure-strategy subgame-perfect equilibrium of this game. What are the players' payoffs at the equilibrium?

Scenario 3. Same as Scenario 2; the only difference is that if both Players 1 and 2 decide to participate in the auction then Player G gives to the loser the fraction a (with $0 < a < 1$) of the amount paid by the winner in the auction (note that player G still keeps 100% of the participation fees). This is publicly announced at the beginning and is an irrevocable commitment.

- (e) For the general case where all possible bids are allowed (subject to being positive integers) find a subgame-perfect equilibrium of this game. What are the players' payoff at the equilibrium?

Scenario 4. Player G tells the Chinese millionaire the following:

“First you (= the Chinese) say Yes or No; if you say No I will sell you the palace at the price that you offered me, namely $\$100$ (that is, we now assume that $p = 100$); if you say Yes then we play the following perfect information game. I start by choosing a number from the set $\{1, 2, 3\}$, then you (= the Chinese) choose a number from this set, then I choose again, followed by you, etc. The first player who brings the cumulative sum of all the numbers chosen (up to and including the last one) to 40 wins. If you win I will sell you the palace for $\$50$, while if I win I will sell you the palace for $\$200$.”

Thus there is no auction in this scenario. Assume that the Chinese would actually be willing to pay up to $\$300$ for the palace.

- (f) Find a pure-strategy subgame-perfect equilibrium of this game. ■

4.7 Solutions to exercises

Solution to Exercise 4.1. One possible extensive frame is shown in Figure 4.26, where Amy moves first. Note that we have only one non-trivial information set for Amy, while each of the other three consists of a single node. The reason is as follows: if Amy initially bids \$3 and Bill bids \$2 then the referee announces “the lowest bid was \$2”; this announcement does not directly reveal to Amy that Bill’s bid was \$2, but she can figure it out from her knowledge that her own bid was \$3; similarly, if the initial two bids are both \$3 then the referee announces “the lowest bid was \$3”, in which case Amy is able to figure out that Bill’s bid was also \$3. If we included those two nodes in the same information set for Amy, we would not show much faith in Amy’s reasoning ability!




Figure 4.26: One possible game-frame for Exercise 4.1

Another possible extensive frame is shown in Figure 4.27, where Bill moves first. □

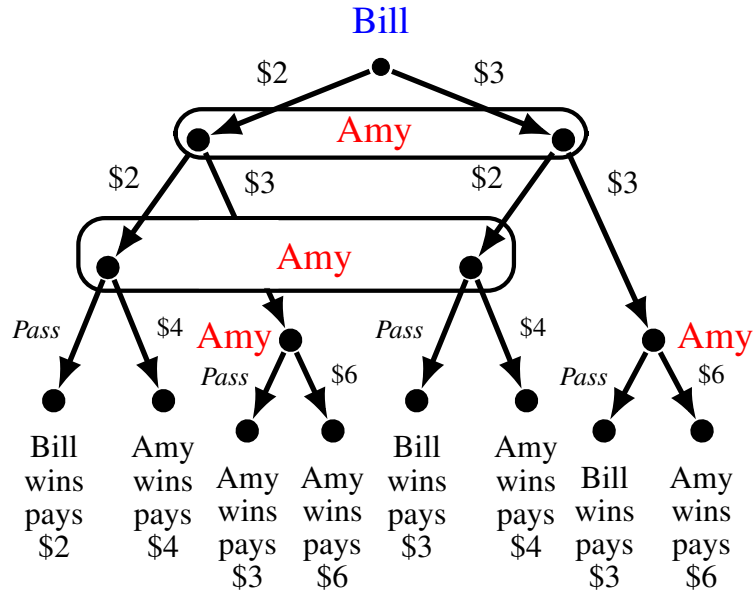


Figure 4.27: Another possible game-frame for Exercise 4.1

Solution to Exercise 4.2. The extensive form is shown in Figure 4.28 (since the players are selfish and greedy we can take a player's utility of an outcome to be the profit of that player at that outcome). □

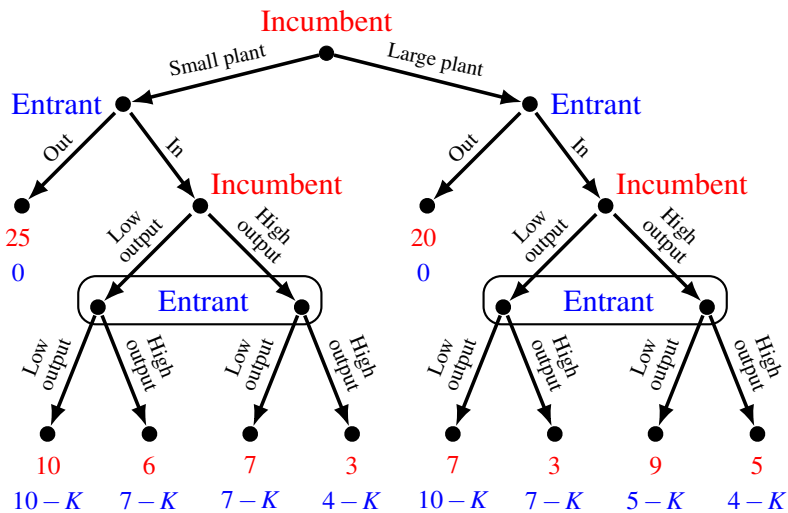


Figure 4.28: The extensive-form game for Exercise 4.2

Solution to Exercise 4.3. The strategic form is shown in Figure 4.29.

Amy's strategy (x, y, w, z) means: at the beginning I bid $\$x$, at the non-trivial information set on the left I choose y , at the singleton node in the middle I choose w and at the singleton node on the right I choose z . The numbers are bid amounts and P stands for "Pass". \square

		BILL	
		Bid \$2	Bid \$3
A M Y	$\$2, P, P, P$	Bill wins, pays \$2	Bill wins, pays \$3
	$\$2, P, P, \6	Bill wins, pays \$2	Bill wins, pays \$3
	$\$2, P, \$6, P$	Bill wins, pays \$2	Bill wins, pays \$3
	$\$2, P, \$6, \$6$	Bill wins, pays \$2	Bill wins, pays \$3
	$\$2, \$4, P, P$	Amy wins, pays \$4	Amy wins, pays \$4
	$\$2, \$4, P, \$6$	Amy wins, pays \$4	Amy wins, pays \$4
	$\$2, \$4, \$6, P$	Amy wins, pays \$4	Amy wins, pays \$4
	$\$2, \$4, \$6, \6	Amy wins, pays \$4	Amy wins, pays \$4
	$\$3, P, P, P$	Amy wins, pays \$3	Bill wins, pays \$3
	$\$3, P, P, \6	Amy wins, pays \$3	Amy wins, pays \$6
	$\$3, P, \$6, P$	Amy wins, pays \$6	Bill wins, pays \$3
	$\$3, P, \$6, \$6$	Amy wins, pays \$6	Amy wins, pays \$6
	$\$3, \$4, P, P$	Amy wins, pays \$3	Bill wins, pays \$3
	$\$3, \$4, P, \$6$	Amy wins, pays \$3	Amy wins, pays \$6
	$\$3, \$4, \$6, P$	Amy wins, pays \$6	Bill wins, pays \$3
	$\$3, \$4, \$6, \6	Amy wins, pays \$6	Amy wins, pays \$6

Figure 4.29: The strategic form for Exercise 4.3

Solution to Exercise 4.4.

- (a) The potential entrant has four information sets, hence a strategy has to specify what she would do in each of the four situations. A possible strategy is: "if the incumbent chooses a small plant I stay out, if the incumbent chooses a large plant I enter, if small plant and I entered then I choose low output, if large plant and I entered then I choose high output".
- (b) The potential entrant has $2^4 = 16$ strategies.
- (c) The strategic form is shown in Figure 4.30.
- (d) For the case where $K = 2$ the Nash equilibria are highlighted in Figure 4.30. \square

		INCUMBENT							
		SLL	SLH	SHL	SHH	LLL	LLH	LHL	LHH
E N T R A N T	OOLL	0, 25	0, 25	0, 25	0, 25	0, 20	0, 20	0, 20	0, 20
	OOLH	0, 25	0, 25	0, 25	0, 25	0, 20	0, 20	0, 20	0, 20
	OOHL	0, 25	0, 25	0, 25	0, 25	0, 20	0, 20	0, 20	0, 20
	OOHH	0, 25	0, 25	0, 25	0, 25	0, 20	0, 20	0, 20	0, 20
	OILL	0, 25	0, 25	0, 25	0, 25	10-K, 7	5-K, 9	10-K, 7	5-K, 9
	OILH	0, 25	0, 25	0, 25	0, 25	7-K, 3	4-K, 5	7-K, 3	4-K, 5
	OIHL	0, 25	0, 25	0, 25	0, 25	10-K, 7	5-K, 9	10-K, 7	5-K, 9
	OIHH	0, 25	0, 25	0, 25	0, 25	7-K, 3	4-K, 5	7-K, 3	4-K, 5
	IOLL	10-K, 10	10-K, 10	7-K, 7	7-K, 7	0, 20	0, 20	0, 20	0, 20
	IO LH	10-K, 10	10-K, 10	7-K, 7	7-K, 7	0, 20	0, 20	0, 20	0, 20
	IOHL	7-K, 6	7-K, 6	4-K, 3	4-K, 3	0, 20	0, 20	0, 20	0, 20
	IOHH	7-K, 6	7-K, 6	4-K, 3	4-K, 3	0, 20	0, 20	0, 20	0, 20
	IILL	10-K, 10	10-K, 10	7-K, 7	7-K, 7	10-K, 7	5-K, 9	10-K, 7	5-K, 9
	IILH	10-K, 10	10-K, 10	7-K, 7	7-K, 7	7-K, 3	4-K, 5	7-K, 3	4-K, 5
	IIHL	7-K, 6	7-K, 6	4-K, 3	4-K, 3	10-K, 7	5-K, 9	10-K, 7	5-K, 9
	IIHH	7-K, 6	7-K, 6	4-K, 3	4-K, 3	7-K, 3	4-K, 5	7-K, 3	4-K, 5

Figure 4.30: The strategic form for Exercise 4.4

Solution to Exercise 4.5. There are no proper subgames.

Solution to Exercise 4.6. There are no proper subgames.

Solution to Exercise 4.7.

- (a) Only one proper subgame: it starts at Player 2's node.
 (b) Since it is the only subgame, it is minimal. □

Solution to Exercise 4.8. The game under consideration is shown in Figure 4.31.




Figure 4.31: The game considered in Exercise 4.8

- (a) There are three proper subgames: one starting at node x , one starting at node y and one starting at the node of Player 3.
 (b) Two: the one starting at node x and the one starting at the decision node of Player 3. In a perfect-information game a minimal proper subgame is one that starts at a decision node followed only by terminal nodes. □

Solution to Exercise 4.9. The strategic form is shown in Figure 4.32.

The Nash equilibria are: $(A, (G, C), E)$ and $(B, (H, C), F)$.

The extensive-form game has two proper subgames. The one on the left has a unique Nash equilibrium, (G, E) , and the one on the right has a unique Nash equilibrium, C .

Hence the game reduces to the game shown in Figure 4.33. In that game A is the unique optimal choice. Hence there is only one subgame-perfect equilibrium, namely $(A, (G, C), E)$. □

		Player 2											
		<i>GC</i>			<i>GD</i>			<i>HC</i>			<i>HD</i>		
Player 1	<i>A</i>	2	2	4	2	2	4	2	0	1	2	0	1
	<i>B</i>	1	4	3	4	3	2	1	4	3	4	3	2
		Player 3: <i>E</i>											
.....													
		Player 2											
		<i>GC</i>			<i>GD</i>			<i>HC</i>			<i>HD</i>		
Player 1	<i>A</i>	4	5	3	4	5	3	0	6	0	0	6	0
	<i>B</i>	1	4	3	4	3	2	1	4	3	4	3	2
		Player 3: <i>F</i>											

Figure 4.32: The strategic form for Exercise 4.9




Figure 4.33: The reduced extensive-form game for Exercise 4.9

Solution to Exercise 4.10. Consider first the subgame that starts at Player 2's decision node following choice A of Player 1. The strategic-form of this game is shown in Figure 4.34 (where only the payoff of Players 2 and 3 are shown). The unique Nash equilibrium is (E, H) .

		Player 3					
		F		G		H	
Player 2	C	0	0	2	0	1	1
	D	1	0	1	2	2	1
	E	0	1	0	1	3	2

Figure 4.34: The strategic-form of subgame that starts at Player 2's decision node following choice A of Player 1

Now consider the subgame that starts at Player 2's decision node following choice B of Player 1. The strategic-form of this game is shown in Figure 4.35. This game has two Nash equilibria: (L, P) and (M, R) .

		Player 3			
		P		R	
Player 2	L	2	1	1	0
	M	0	0	1	1

Figure 4.35: The strategic-form of subgame that starts at Player 2's decision node following choice B of Player 1

Thus there are two subgame-perfect equilibria of the entire game:

1. Player 1's strategy: A ; Player 2's strategy: E if A and L if B ;
Player 3's strategy: H if A and P if B .
2. Player 1's strategy: B ; Player 2's strategy: E if A and M if B ;
Player 3's strategy: H if A and R if B .

□

Solution to Exercise 4.11. Given the players' preferences, we can assign the following utilities to the outcomes:

outcome	Amy's utility	Bill's utility
Amy wins and pays \$3	2	0
Amy wins and pays \$4	1	0
Amy wins and pays \$6	-1	0
Bill wins and pays \$2	0	2
Bill wins and pays \$3	0	1

In the extensive form there are only two proper subgames: they start at Amy's singleton information sets (the two nodes on the right).

- In both subgames Amy will choose to pass (since she values the object at \$5 and is thus not willing to pay \$6).
- Replacing Amy's left singleton node with payoffs of 2 for Amy and 0 for Bill, and Amy's right singleton node with payoffs of 0 for Amy and 1 for Bill, we get a reduced game whose associated strategic form is as follows (in Amy's strategy the first component is the initial bet and the second component is her choice at her information set following Bill's choices after Amy's bet of \$2):

		Bill	
		\$2	\$3
Amy	\$2, pass	0 , 2	0 , 1
	\$2, \$4	1 , 0	1 , 0
	\$3, pass	2 , 0	0 , 1
	\$3, \$4	2 , 0	0 , 1

This game has one Nash equilibrium: (((\$2,\$4),\$3)). Thus the initial extensive-form game has one subgame-perfect equilibrium, which is as follows: Amy's strategy is (\$2, \$4, pass, pass), Bill's strategy is \$3. The corresponding outcome is: Amy wins the auction and pays \$4. □

Solution to Exercise 4.12.

(a) The extensive-form game is shown in Figure 4.36.

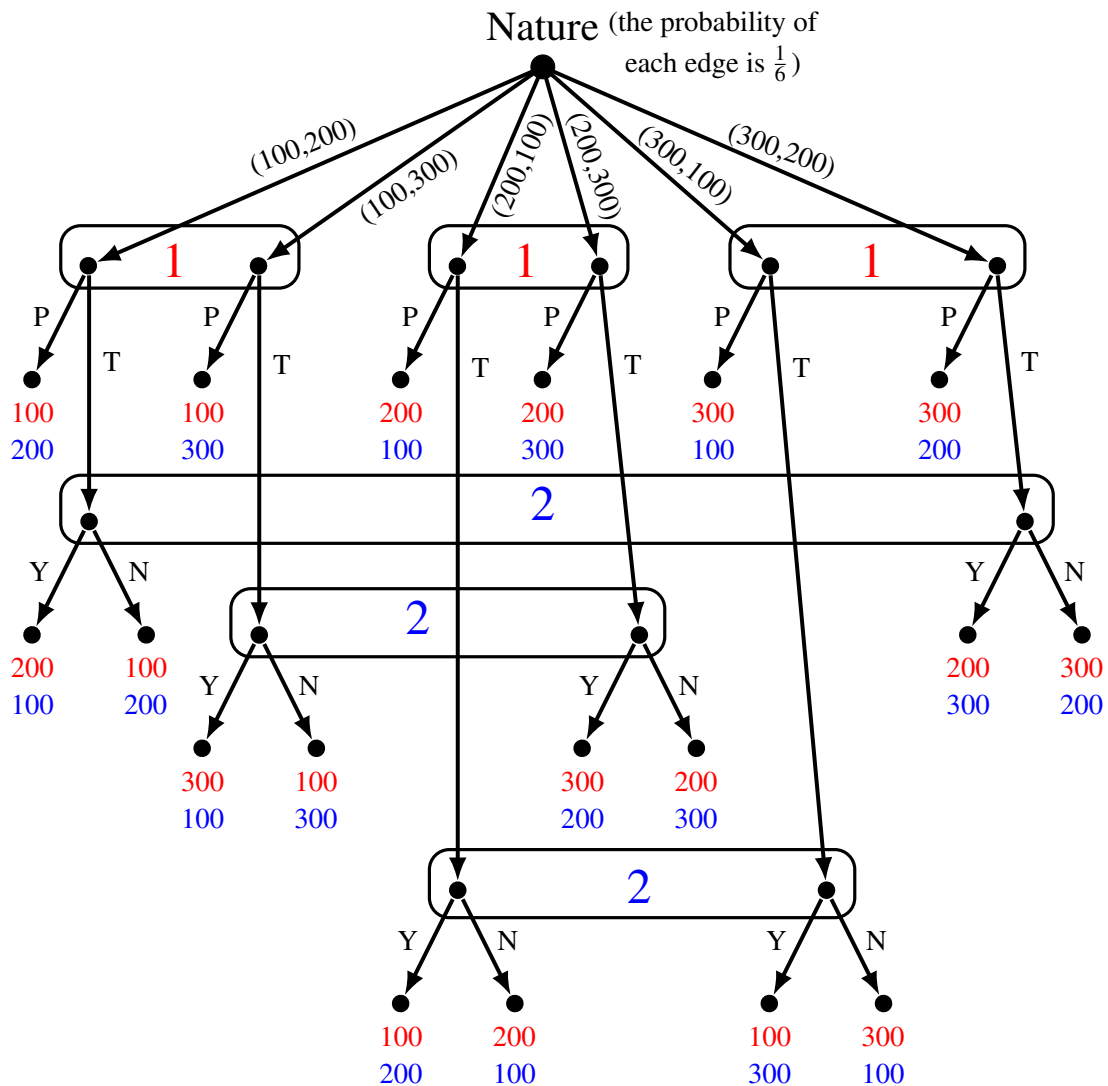


Figure 4.36: The extensive-form game for Exercise 4.12

(b) Player 1's strategies are the same as in Example 4.3. Player 2 now has 8 strategies. Each strategy has to specify how to reply to Player 1's proposal depending on the sum he (Player 2) has. Thus one possible strategy is: if I have \$100 I say No, if I have \$200 I say Yes and if I have \$300 I say No. \square

Solution to Exercise 4.13.

(a) The extensive-form game is shown in Figure 4.37

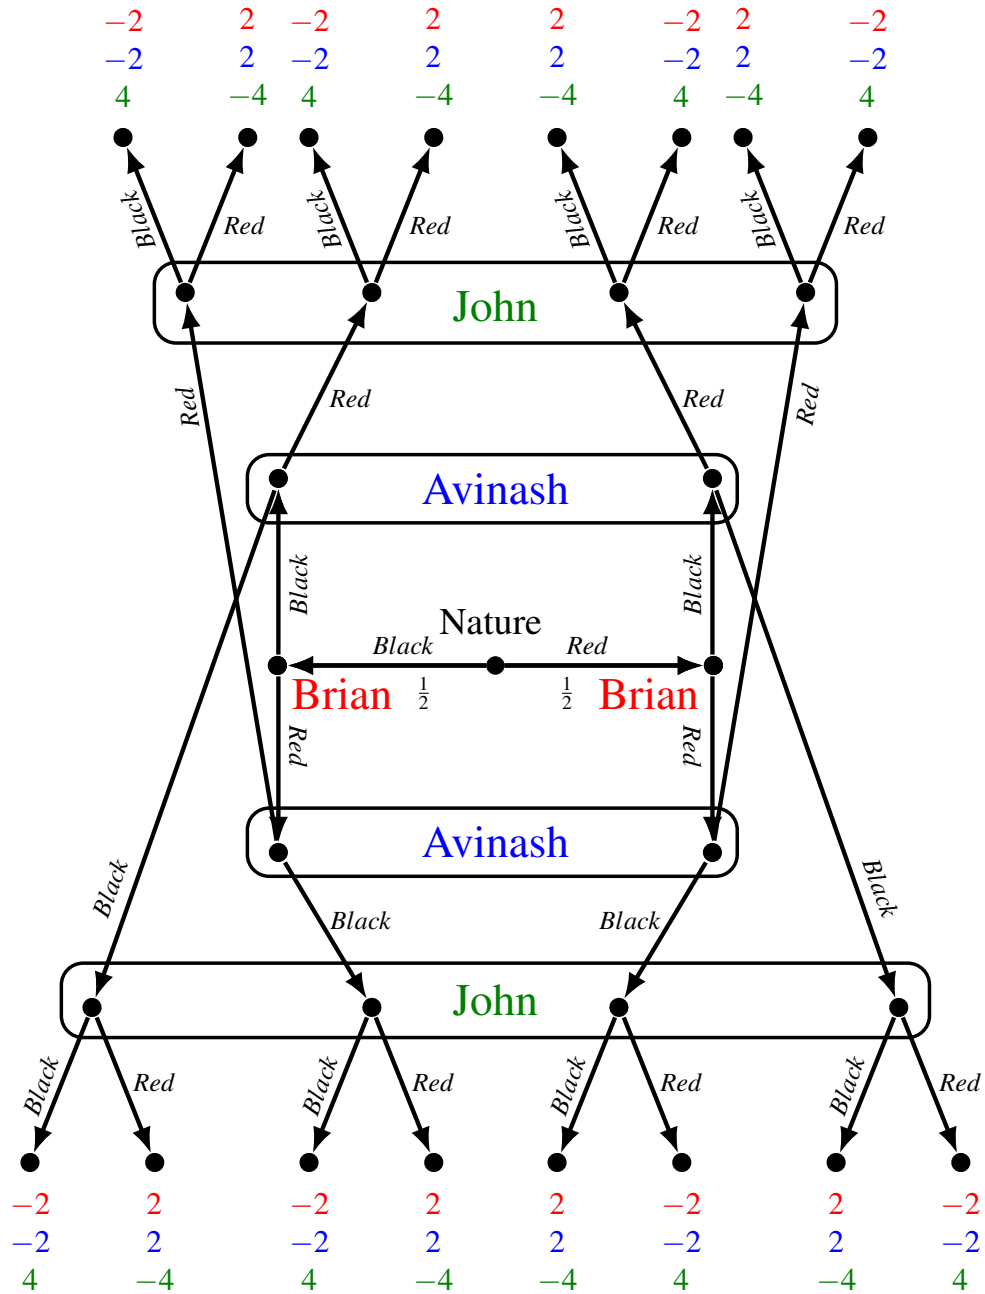


Figure 4.37: The extensive-form game for Part (a) of Exercise 4.13

- (b) Each player has two information sets, two choices at each information set, hence four strategies.

The strategic form is shown in Figure 4.38

(interpretation: for Avinash “if B, R, if R, B” means “if Brian tells me B then I say R and if Brian tells me R then I say B”; similarly for the other strategies and for the other players).

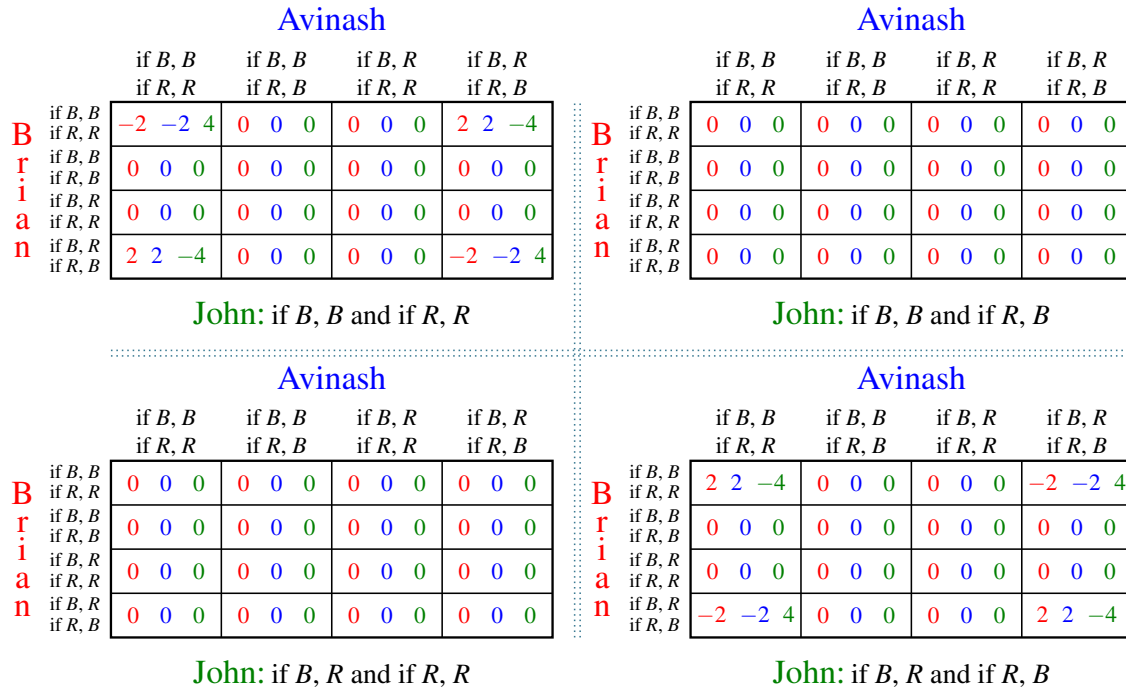


Figure 4.38: The strategic form for Part (b) of Exercise 4.13

How can we fill in the payoffs without spending more than 24 hours on this problem? There is a quick way of doing it. First of all, when John’s strategy is to guess Black, no matter what Avinash says, he has a 50% chance of being right and a 50% chance of being wrong. Thus his expected payoff is $\frac{1}{2}(4) + \frac{1}{2}(-4) = 0$ and the expected payoff of each of the other two players is $\frac{1}{2}(2) + \frac{1}{2}(-2) = 0$. This explains why the second table is filled with the same payoff vector, namely $(0, 0, 0)$. The same reasoning applies to the case where when John’s strategy is to guess Red, no matter what Avinash says (leading to the third table, filled with the same payoff vector $(0, 0, 0)$).

For the remaining strategies of John's, one can proceed as follows:
 Start with the two colors, B and R. Under B write T (for true) if Brian's strategy says "if B then B" and write F (for false) if Brian's strategy says "if B then R"; similarly, under R write T (for true) if Brian's strategy says "if R then R" and write F (for false) if Brian's strategy says "if R then B".

In the next row, in the B column rewrite what is in the previous row if Avinash's strategy says "if B then B" and change a T into an F or an F into a T if Avinash's strategy says "if B then R". Similarly for the R column. Now repeat the same for John (in the B column a T remains a T and an F remains an F if John's strategy is "if B then B", while a T is changed into an F and an F is changed into a T if John's strategy is "if B then R").

Now in each column the payoffs are $(-2, -2, 4)$ if the last row has a T and $(2, 2, -4)$ if the last row has an F. The payoffs are then given by $\frac{1}{2}$ the payoff in the left column plus $\frac{1}{2}$ the payoff in the right column. For example, for the cell in the second row, third column of the third table we have the calculations shown in Figure 4.39. \square

	B	R
Brian's strategy: if B, B and if R, B	T	F
Avinash's strategy: if B, R and if R, R	F	T
John's strategy: if B, R and if R, R	F	T
Payoffs:	$(2, 2, -4)$	$(-2, -2, 4)$
Expected payoffs:	$\frac{1}{2}(2, 2, -4) + \frac{1}{2}(-2, -2, 4) = (0, 0, 0)$	

Figure 4.39: The calculations for the expected payoffs

Solution to Exercise 4.14.

- (a) The extensive-form representation of the simplified poker game is shown in Figure 4.40 (the top number is Yvonne's net take in dollars and the bottom number is Zoe's net take).

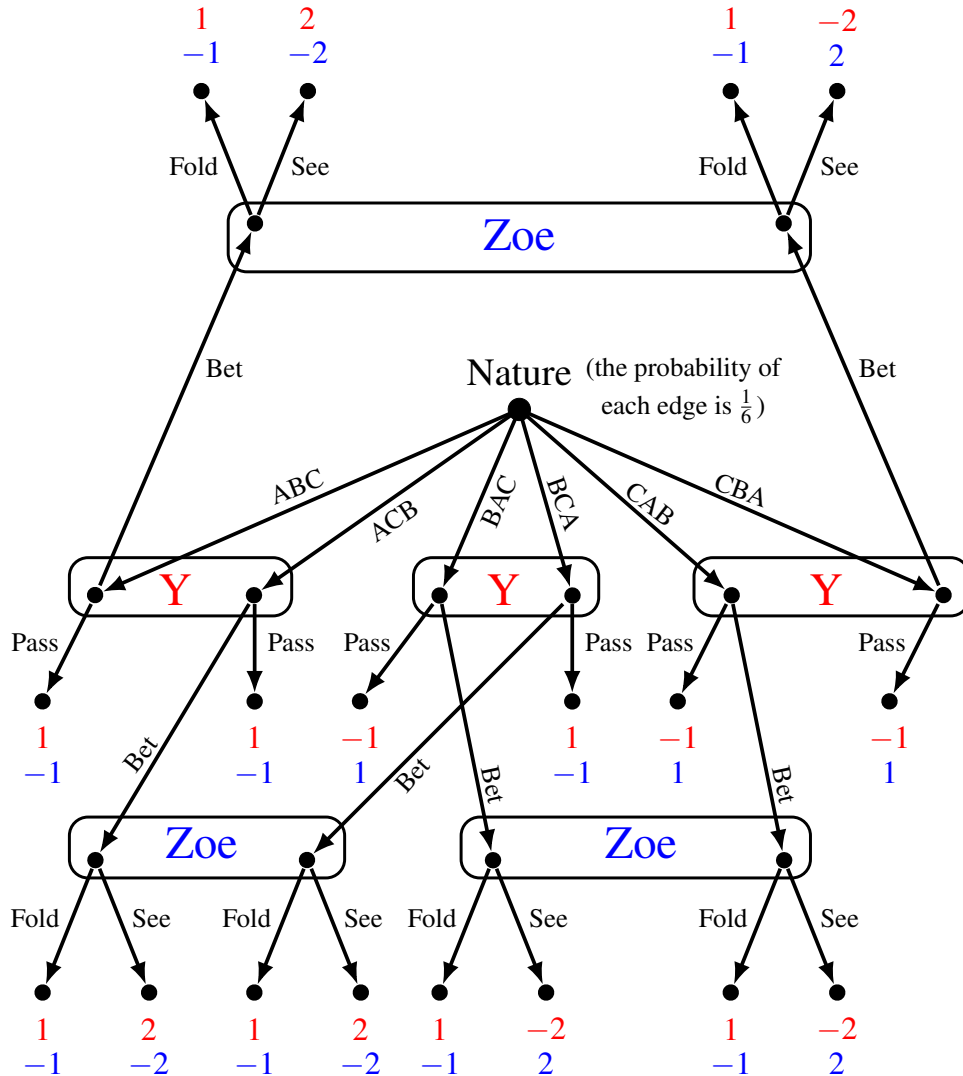


Figure 4.40: The extensive-form game for Exercise 4.14

- (b) Yvonne has eight strategies (three information sets, two choices at each information set, thus $2 \times 2 \times 2 = 8$ possible strategies).
- (c) Similarly, Zoe has eight strategies.

- (d) Yvonne uses the strategy “If A pass, if B pass, if C bet” and Zoe uses the strategy “If A fold, if B fold, if C fold”. The table below shows how to compute the expected net payoff for Yvonne. Zoe’s expected net payoff is the negative of that.

Top card is:	<i>A</i>	<i>A</i>	<i>B</i>	<i>B</i>	<i>C</i>	<i>C</i>
Second card is:	<i>B</i>	<i>C</i>	<i>A</i>	<i>C</i>	<i>A</i>	<i>B</i>
Probability:	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$
Yvonne’s action:	<i>pass</i>	<i>pass</i>	<i>pass</i>	<i>pass</i>	<i>bet</i>	<i>bet</i>
Zoe’s action:	--	--	--	--	<i>fold</i>	<i>fold</i>
Yvonne’s payoff:	1	1	-1	1	1	1

$$\text{Yvonne's expected payoff: } \frac{1}{6}(1 + 1 - 1 + 1 + 1 + 1) = \frac{4}{6}.$$

- (e) Yvonne uses the strategy “If A pass, if B pass, if C bet” and Zoe uses the strategy “see with any card”. The table below shows how to compute the expected net payoff for Yvonne. Zoe’s expected net payoff is the negative of that.

Top card is:	<i>A</i>	<i>A</i>	<i>B</i>	<i>B</i>	<i>C</i>	<i>C</i>
Second card is:	<i>B</i>	<i>C</i>	<i>A</i>	<i>C</i>	<i>A</i>	<i>B</i>
Probability:	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$
Yvonne’s action:	<i>pass</i>	<i>pass</i>	<i>pass</i>	<i>pass</i>	<i>bet</i>	<i>bet</i>
Zoe’s action:	--	--	--	--	<i>see</i>	<i>see</i>
Yvonne’s payoff:	1	1	-1	1	-2	-2

$$\text{Yvonne's expected payoff: } \frac{1}{6}(1 + 1 - 1 + 1 - 2 - 2) = -\frac{2}{6}.$$

- (f) The strategic form is shown in Figure 4.41.

- (g) Let \triangleright denote weak dominance, that is, $a \triangleright b$ means that a weakly dominates b .
 FOR YVONNE (row player): 3^{rd} row $\triangleright 1^{st}$ row, $6^{th} \triangleright 4^{th}$, $7^{th} \triangleright 4^{th}$, $7^{th} \triangleright 5^{th}$, $2^{nd} \triangleright 8^{th}$.
 FOR ZOE (column player): 1^{st} col $\triangleright 5^{th}$ col, $3^{rd} \triangleright 1^{st}$, $3^{rd} \triangleright 4^{th}$, $3^{rd} \triangleright 5^{th}$, $3^{rd} \triangleright 7^{th}$, $3^{rd} \triangleright 8^{th}$, $2^{nd} \triangleright 8^{th}$, $4^{th} \triangleright 5^{th}$, $4^{th} \triangleright 8^{th}$, $6^{th} \triangleright 2^{nd}$, $6^{th} \triangleright 4^{th}$, $6^{th} \triangleright 5^{th}$, $6^{th} \triangleright 7^{th}$, $6^{th} \triangleright 8^{th}$, $7^{th} \triangleright 4^{th}$, $7^{th} \triangleright 5^{th}$, $7^{th} \triangleright 8^{th}$.

		ZOE							
		If A fold, If B fold, If C fold	If A see, If B see, If C see	If A see, If B fold, If C fold	If A fold, If B see, If C fold	If A fold, If B fold, If C see	If A see, If B see, If C fold	If A see, If B fold, If C see	If A fold, If B see, If C see
Y V O N N O N E	If A pass, if B pass if C pass	0, 0	0, 0	0, 0	0, 0	0, 0	0, 0	0, 0	0, 0
	If A bet, if B bet if C bet	1, -1	0, 0	0, 0	4/6, -4/6	8/6, -8/6	-2/6, 2/6	2/6, -2/6	1, -1
	If A bet, if B pass if C pass	0, 0	2/6, -2/6	0, 0	1/6, -1/6	1/6, -1/6	1/6, -1/6	1/6, -1/6	2/6, -2/6
	If A pass, if B bet if C pass	2/6, -2/6	0, 0	-1/6, 1/6	2/6, -2/6	3/6, -3/6	-1/6, 1/6	0, 0	3/6, -3/6
	If A pass, if B pass if C bet	4/6, -4/6	-2/6, 2/6	1/6, -1/6	1/6, -1/6	4/6, -4/6	-2/6, 2/6	1/6, -1/6	1/6, -1/6
	If A bet, if B bet if C pass	2/6, -2/6	2/6, -2/6	-1/6, 1/6	3/6, -3/6	4/6, -4/6	0, 0	1/6, -1/6	5/6, -5/6
	If A bet, if B pass if C bet	4/6, -4/6	0, 0	1/6, -1/6	2/6, -2/6	5/6, -5/6	-1/6, 1/6	2/6, -2/6	3/6, -3/6
	If A Pass, if B or C, Bet,	1, -1	-2/6, 2/6	0, 0	3/6, -3/6	7/6, -7/6	-3/6, 3/6	1/6, -1/6	4/6, -4/6

Figure 4.41: The strategic form for Part (f) of Exercise 4.14

(h) Eliminating rows 1, 4, 5 and 8 and all columns except 3 and 6 we are left with the reduced game shown below:

		Zoe	
		See only if A	See only with A or B
Yvonne	Bet always	0, 0	$-\frac{2}{6}, \frac{2}{6}$
	Bet only if A	0, 0	$\frac{1}{6}, -\frac{1}{6}$
	Bet only if A or B	$-\frac{1}{6}, \frac{1}{6}$	0, 0
	Bet only if A or C	$\frac{1}{6}, -\frac{1}{6}$	$-\frac{1}{6}, \frac{1}{6}$

In this reduced game, the second row dominates the first and the third. Eliminating them we are led to the reduced game shown below, which is a remarkable simplification of the original strategic form:

		Zoe	
		See only if A	See only with A or B
Yvonne	Bet only if A	0, 0	$\frac{1}{6}, -\frac{1}{6}$
	Bet only if A or C	$\frac{1}{6}, -\frac{1}{6}$	$-\frac{1}{6}, \frac{1}{6}$

□

Solution to Exercise 4.15.

- (a) The extensive-form is shown in Figure 4.42.
- (b) In the auction subgame for every player it is a weakly dominant strategy to bid his own value. Thus a natural Nash equilibrium for this subgame is the dominant-strategy equilibrium (although there are other Nash equilibria and one could choose any one of the alternative Nash equilibria).

Let $p_j = \max\{p_1, \dots, p_n\}$ be the highest value and $p_k = \max\{p_1, \dots, p_n\} \setminus \{p_j\}$ be the second highest value. Then the auction, if it takes place, will be won by Player j and he will pay p_k . Hence there are three cases.

Case 1: $p > p_k$. In this case Player G will sell to the Chinese (and the strategy of Player i in the subgame is to bid p_i), G's payoff is p and the payoff of Player i is 0.

Case 2: $p < p_k$. In this case Player G announces the auction, the strategy of Player i in the subgame is to bid p_i , the winner is Player j and he pays p_k , so that the payoff of G is p_k , the payoff of player j is $p_j - p_k$ and the payoff of every other player is 0.

Case 3: $p = p_k$. In this case there are two subgame-perfect equilibria: one as in Case 1 and the other as in Case 2 and G is indifferent between the two.

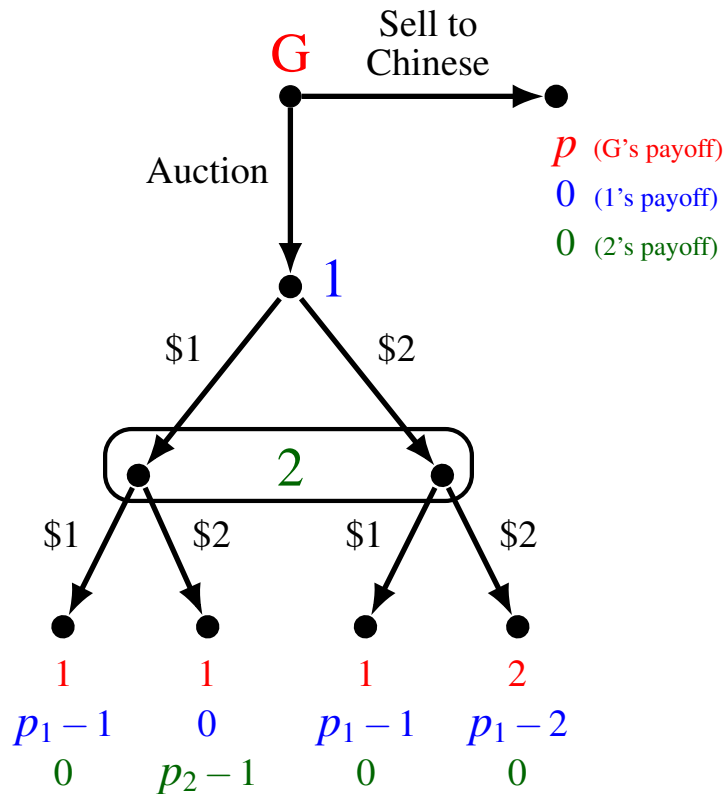


Figure 4.42: The extensive form for Part (a) of Exercise 4.15

(c) The extensive-form game is shown in Figure 4.43.




Figure 4.43: The extensive-form game for Part (c) of Exercise 4.15

(d) In the simultaneous subgame after both players have said Yes, the participation fee paid is a sunk cost and for every player bidding the true value is a weakly dominant strategy. Thus the outcome there is as follows: Player 1 bids p_1 , gets the palace by paying p_2 , G's payoff is $(p_2 + 2)$, 1's payoff is $(p_1 - p_2 - 1)$ and Player 2's payoff is -1 .

In the subgames where one player said No and the other said Yes the optimal choice is obviously $x = 1$, with payoffs of 2 for Player G, 0 for the player who said No and $p_i - 2$ for the player who said Yes.

Thus the game reduces to the one shown in Panel A of Figure 4.44.

By assumption, $p_1 > p_2 + 1 > 2$, so that $p_1 - p_2 - 1 > 0$ and $p_1 - 2 > 0$. Thus at the bottom node and at the left node Player 1 prefers Yes to No.

Thus the game reduces to the one shown in Panel B of Figure 4.44.

Hence Player 2 will say No. The subgame-perfect equilibrium is as follows: (1) if $p > 2$ then player G will sell to the Chinese (and the choices off the equilibrium path are as explained above) and the payoffs are $(p, 0, 0)$; (2) if $p < 2$ then G chooses to auction, 2 says No, 1 says Yes and then offers \$1 and the payoffs are $(2, p_1 - 2, 0)$ (and the choices off the equilibrium path are as explained above); (3) if $p = 2$ then there are two equilibria: one as in (1) and the other as in (2).

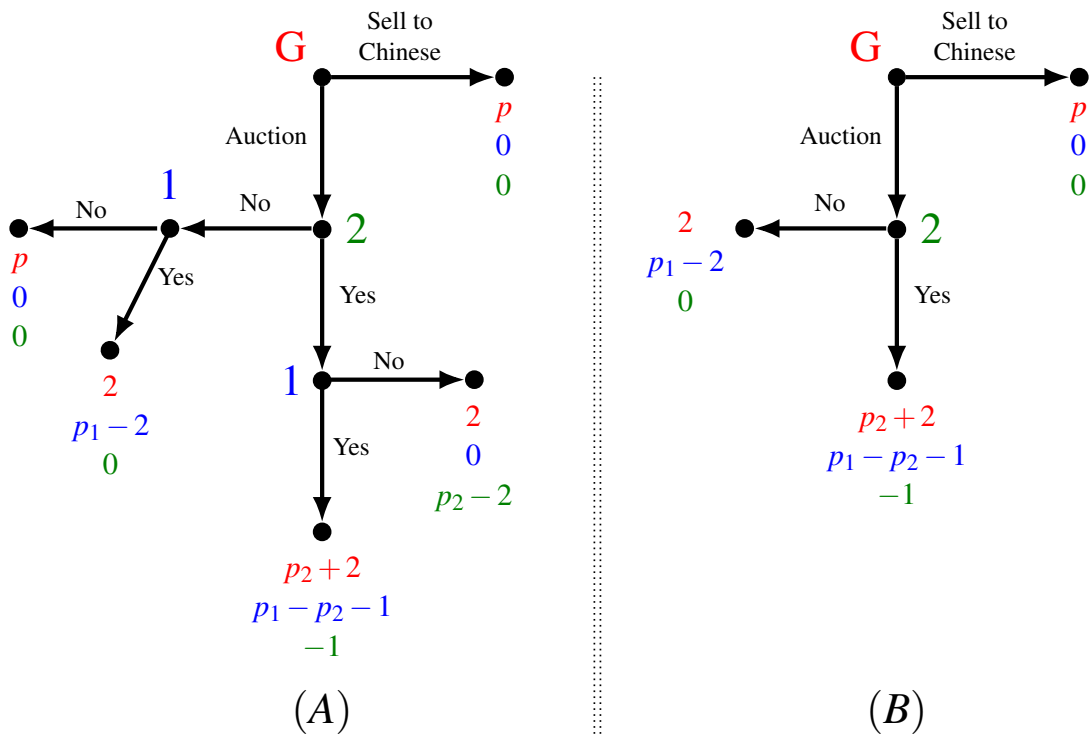



Figure 4.44: The extensive-form game for Part (d) of Exercise 4.15

(e) When the loser is given the fraction a of the amount paid by the winner (that is, the loser is given the fraction a of his own bid), **it is no longer true that bidding one's true value is a dominant strategy**. In fact, (p_1, p_2) is not even a Nash equilibrium any more. To see this, imagine that Player 1's true value is 10 and Player 2's true value is 6 and $a = 50\%$. Then if Player 1 bids 10 and 2 bids 6, Player 2 ends up losing the auction but being given \$3, while if he increased his bid to 8 then he would still lose the auction but receive \$4. This shows that there cannot be a Nash equilibrium where Player 2 bids less than Player 1. Now there are several Nash equilibria of the auction, for example, all pairs (b_1, b_2) with $b_1 = b_2 = b$ and $p_2 \leq b < p_1$ provided that $p_1 - b \geq a(b - 1)$, that is, $b \leq \frac{p_1 + a}{1 + a}$ (but there are more: for example all pairs (b_1, b_2) with $b_1 = b_2 = b$ and $b < p_2$ provided that $p_1 - b \geq a(b - 1)$ and $ab \geq p_2 - b$). Thus to find a subgame-perfect equilibrium of the game one first has to select a Nash equilibrium of the auction game and then apply backward induction to see if the players would want to say Yes or No to the auction, etc.

- (f) Let us start by considering the perfect-information game that is played if the Chinese says Yes. This is a game similar to the one discussed in Example 3.2 (Chapter 3, Section 3.5). We first determine the losing positions. Whoever has to move when the sum is 36 cannot win. Thus 36 is a losing position. Working backwards, the losing positions are 32, 28, 24, 20, 16, 12, 8, 4 and 0. Thus the first player (= player G) starts from a losing position: whatever his initial choice, he can be made to choose the second time when the sum is 4, and then 8, etc. Hence **the second player (= the Chinese) has a winning strategy**, which is as follows: if Player G just chose n , then choose $(4 - n)$. If the Chinese says Yes and then follows this strategy he can guarantee that he will buy the palace for \$50. Thus the subgame-perfect equilibrium of this game is: the Chinese says Yes and uses the winning strategy in the ensuing game, while for Player G we can pick any arbitrary choices (so that, in fact, there are many subgame-perfect equilibria, but they share the same outcome). \square



Games with Cardinal Payoffs

5	Expected Utility Theory	169
5.1	Money lotteries and attitudes to risk	
5.2	Expected utility: theorems	
5.3	Expected utility: the axioms	
5.4	Exercises	
5.5	Solutions to Exercises	
6	Strategic-form Games	193
6.1	Strategic-form games with cardinal payoffs	
6.2	Mixed strategies	
6.3	Computing the mixed-strategy Nash equilibria	
6.4	Strict dominance and rationalizability	
6.5	Exercises	
6.6	Solutions to exercises	
7	Extensive-form Games	227
7.1	Behavioral strategies in dynamic games	
7.2	Subgame-perfect equilibrium revisited	
7.3	Problems with the notion of subgame-perfect equilibrium	
7.4	Exercises	
7.5	Solutions to exercises	

5. Expected Utility Theory

5.1 Money lotteries and attitudes to risk

The introduction of chance moves gives rise to probabilistic outcomes, which we called lotteries. In Chapter 4 we restricted attention to lotteries whose outcomes are sums of money (money lotteries) and to one possible way of ranking such lotteries, based on the notion of risk neutrality. In this section we will continue to focus on money lotteries and define other possible attitudes to risk.¹

As before, we restrict attention to finite lotteries. Recall that a money lottery is a probability distribution of the form

$$\begin{pmatrix} \$x_1 & \$x_2 & \dots & \$x_n \\ p_1 & p_2 & \dots & p_n \end{pmatrix}$$

(with $0 \leq p_i \leq 1$, for all $i = 1, 2, \dots, n$, and $p_1 + p_2 + \dots + p_n = 1$) and that (Definition 4.2.2, Chapter 4) its expected value is the number $(x_1p_1 + x_2p_2 + \dots + x_np_n)$.

If L is a money lottery, we denote by $\mathbb{E}[L]$ the expected value of L . Thus, for example, if

$$L = \begin{pmatrix} \$30 & \$45 & \$90 \\ \frac{1}{3} & \frac{5}{9} & \frac{1}{9} \end{pmatrix} \quad \text{then} \quad \mathbb{E}[L] = \frac{1}{3}(30) + \frac{5}{9}(45) + \frac{1}{9}(90) = 45.$$

Recall also (Definition 4.2.3, Chapter 4) that a person is said to be *risk neutral* if she considers a money lottery to be just as good as its expected value for certain. For example, a risk-neutral person would consider getting \$45 with certainty to be just as good as playing lottery $L = \begin{pmatrix} \$30 & \$45 & \$90 \\ \frac{1}{3} & \frac{5}{9} & \frac{1}{9} \end{pmatrix}$.

We can now consider different attitudes to risk, besides risk neutrality.

¹In the next section we will consider more general lotteries, where the outcomes need not be sums of money, and introduce the theory of expected utility.

Definition 5.1.1 Let L be a money lottery and consider the choice between L and getting $\mathbb{E}[L]$ (the expected value of L) for certain. Then

- An individual who prefers $\mathbb{E}[L]$ for certain to L is said to be *risk averse*.
- An individual who is indifferent between $\mathbb{E}[L]$ for certain and L is said to be *risk neutral*.
- An individual who prefers L to $\mathbb{E}[L]$ for certain is said to be *risk loving*.

Note that if an individual is risk **neutral**, has transitive preferences over money lotteries and prefers more money to less, then we can tell how that individual ranks any two money lotteries. For example, how would a risk neutral individual rank the two lotteries

$$L_1 = \begin{pmatrix} \$30 & \$45 & \$90 \\ \frac{1}{3} & \frac{5}{9} & \frac{1}{9} \end{pmatrix} \text{ and } L_2 = \begin{pmatrix} \$5 & \$100 \\ \frac{3}{5} & \frac{2}{5} \end{pmatrix}?$$

Since $\mathbb{E}[L_1] = 45$ and the individual is risk neutral, $L_1 \sim \$45$; since $\mathbb{E}[L_2] = 43$ and the individual is risk neutral, $\$43 \sim L_2$; since the individual prefers more money to less, $\$45 \succ \43 ; thus, by transitivity, $L_1 \succ L_2$.

On the other hand, knowing that an individual is risk **averse**, has transitive preferences over money lotteries and prefers more money to less is not sufficient to predict how she will choose between two arbitrary money lotteries.

For example, as we will see later (see Exercise 5.11), it is possible that one risk-averse individual will prefer $L_3 = \begin{pmatrix} \$28 \\ 1 \end{pmatrix}$ (whose expected value is 28) to $L_4 = \begin{pmatrix} \$10 & \$50 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$ (whose expected value is 30), while another risk-averse individual will prefer L_4 to L_3 .

Similarly, knowing that an individual is risk **loving**, has transitive preferences over money lotteries and prefers more money to less is not sufficient to predict how she will choose between two arbitrary money lotteries.

R Note that “rationality” does not, and should not, dictate whether an individual should be risk neutral, risk averse or risk loving: an individual’s attitude to risk is merely a reflection of that individual’s preferences. It is a generally accepted principle that *de gustibus non est disputandum* (in matters of taste, there can be no disputes). According to this principle, there is no such thing as an irrational preference and thus there is no such thing as an irrational attitude to risk. From an empirical point of view, however, most people reveal through their choices (e.g. the decision to buy insurance) that they are risk averse, at least when the stakes are high.

As noted above, with the exception of risk-neutral individuals, even if we restrict attention to money lotteries we are not able to say much – in general – about how an individual would choose among lotteries. What we need is a theory of “rational” preferences over lotteries that (1) is general enough to cover lotteries whose outcomes are not necessarily sums of money and (2) is capable of accounting for different attitudes to risk in the case of money lotteries. One such theory is the theory of expected utility, to which we now turn.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 5.4.1 at the end of this chapter.

5.2 Expected utility: theorems

The theory of expected utility was developed by the founders of game theory, namely John von Neumann and Oskar Morgenstern, in their 1944 book *Theory of Games and Economic Behavior*. In a rather unconventional way, we shall first (in this section) state the main result of the theory (which we split into two theorems) and then (in the following section) explain the assumptions (or axioms) behind that result. The reader who is not interested in understanding the conceptual foundations of expected utility theory, but wants to understand what the theory says and how it can be used, can study this section and skip the next.

Let O be a set of *basic outcomes*. Note that a basic outcome need not be a sum of money: it could be the state of an individual's health, or whether the individual under consideration receives an award, or whether it will rain on the day of her planned outdoor party, etc.

Let $\mathcal{L}(O)$ be the set of *simple lotteries* (or probability distributions) over O . We will assume throughout that O is a finite set: $O = \{o_1, o_2, \dots, o_m\}$ ($m \geq 1$).

Thus, an element of $\mathcal{L}(O)$ is of the form $\begin{pmatrix} o_1 & o_2 & \dots & o_m \\ p_1 & p_2 & \dots & p_m \end{pmatrix}$ with $0 \leq p_i \leq 1$, for all $i = 1, 2, \dots, m$, and $p_1 + p_2 + \dots + p_m = 1$.

We will use the symbol L (with or without subscript) to denote an element of $\mathcal{L}(O)$, that is, a simple lottery. Lotteries are used to represent situations of uncertainty. For example, if $m = 4$ and the individual faces the lottery $L = \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ \frac{2}{5} & 0 & \frac{1}{5} & \frac{2}{5} \end{pmatrix}$ then she knows that, eventually, the outcome will be one and only one of o_1, o_2, o_3, o_4 , but does not know which one; furthermore, she is able to quantify her uncertainty by assigning probabilities to these outcomes.

We interpret these probabilities either as objectively obtained from relevant (past) data or as subjective estimates by the individual. For example, an individual who is considering whether or not to insure her bicycle against theft for the following 12 months knows that there are two relevant basic outcomes: either the bicycle will be stolen or it will not be stolen. Furthermore, she can look up data on past bicycle thefts in her area and use the proportion of bicycles that were stolen as an "objective" estimate of the probability that her bicycle will be stolen. Alternatively, she can use a more subjective estimate: for example she might use a lower probability of theft than suggested by the data because she knows herself to be very conscientious and – unlike other people – to always lock her bicycle when left unattended.

The assignment of *zero probability* to a particular basic outcome is taken to be an expression of *belief, not impossibility*: the individual is confident that the outcome will not arise, but she cannot rule out that outcome on logical grounds or by appealing to the laws of nature.

Among the elements of $\mathcal{L}(O)$ there are the degenerate lotteries that assign probability 1 to one basic outcome: for example, if $m = 4$ one degenerate lottery is $\begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ 0 & 0 & 1 & 0 \end{pmatrix}$.

To simplify the notation we will often denote degenerate lotteries as basic outcomes, that is, instead of writing $\begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ 0 & 0 & 1 & 0 \end{pmatrix}$ we will simply write o_3 .

Thus, in general, the degenerate lottery $\begin{pmatrix} o_1 & \dots & o_{i-1} & o_i & o_{i+1} & \dots & o_m \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{pmatrix}$ will be

denoted by o_i . As another simplification, we will often omit those outcomes that are assigned zero probability. For example, if $m = 4$, the lottery $\begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ \frac{1}{3} & 0 & \frac{2}{3} & 0 \end{pmatrix}$ will be written more simply as $\begin{pmatrix} o_1 & o_3 \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$.

In this chapter we shall call the individual under consideration the Decision-Maker, or *DM* for short. The theory of expected utility assumes that the *DM* has a complete and transitive ranking \succsim of the elements of $\mathcal{L}(O)$ (indeed, this is one of the axioms listed in the next section). As in Chapter 2, the interpretation of $L \succsim L'$ is that the *DM* considers L to be at least as good as L' . By completeness, given any two lotteries L and L' , either $L \succ L'$ (the *DM* prefers L to L') or $L' \succ L$ (the *DM* prefers L' to L) or $L \sim L'$ (the *DM* is indifferent between L and L'). Furthermore, by transitivity, for any three lotteries L_1, L_2 and L_3 , if $L_1 \succsim L_2$ and $L_2 \succsim L_3$, then $L_1 \succsim L_3$. Besides completeness and transitivity, a number of other “rationality” constraints are postulated on the ranking \succsim of the elements of $\mathcal{L}(O)$; these constraints are the so-called Expected Utility Axioms and are discussed in the next section.

Definition 5.2.1 A ranking \succsim of the elements of $\mathcal{L}(O)$ that satisfies the Expected Utility Axioms (listed in the next section) is called a *von Neumann-Morgenstern ranking*.

The following two theorems are the key results in the theory of expected utility.

Theorem 5.2.1 [von Neumann-Morgenstern, 1944]. Let $O = \{o_1, o_2, \dots, o_m\}$ be a set of basic outcomes and let $\mathcal{L}(O)$ be the set of simple lotteries over O . If \succsim is a von Neumann-Morgenstern ranking of the elements of $\mathcal{L}(O)$ then there exists a function $U : O \rightarrow \mathbb{R}$, called a *von Neumann-Morgenstern utility function* that assigns a number (called *utility*) to every basic outcome and is such that, for any two lotteries

$$L = \begin{pmatrix} o_1 & o_2 & \dots & o_m \\ p_1 & p_2 & \dots & p_m \end{pmatrix} \text{ and } L' = \begin{pmatrix} o_1 & o_2 & \dots & o_m \\ q_1 & q_2 & \dots & q_m \end{pmatrix},$$

$L \succ L'$ if and only if $\mathbb{E}[U(L)] > \mathbb{E}[U(L')]$, and

$L \sim L'$ if and only if $\mathbb{E}[U(L)] = \mathbb{E}[U(L')]$, where

$$U(L) = \begin{pmatrix} U(o_1) & U(o_2) & \dots & U(o_m) \\ p_1 & p_2 & \dots & p_m \end{pmatrix}, \quad U(L') = \begin{pmatrix} U(o_1) & U(o_2) & \dots & U(o_m) \\ q_1 & q_2 & \dots & q_m \end{pmatrix},$$

$\mathbb{E}[U(L)]$ is the expected value of the lottery $U(L)$ and $\mathbb{E}[U(L')]$ is the expected value of the lottery $U(L')$, that is,

$$\mathbb{E}[U(L)] = p_1 U(o_1) + p_2 U(o_2) + \dots + p_m U(o_m), \text{ and}$$

$$\mathbb{E}[U(L')] = q_1 U(o_1) + q_2 U(o_2) + \dots + q_m U(o_m).$$

$\mathbb{E}[U(L)]$ is called the *expected utility* of lottery L (and $\mathbb{E}[U(L')]$ the expected utility of lottery L'). We say that any function $U : O \rightarrow \mathbb{R}$ that satisfies the property that, for any two lotteries L and L' , $L \succsim L'$ if and only if $\mathbb{E}[U(L)] \geq \mathbb{E}[U(L')]$ represents the preferences (or ranking) \succsim .

Before we comment on Theorem 5.2.1 we give an example of how one can use it. Theorem 5.2.1 sometimes allows us to predict an individual's choice between two lotteries C and D if we know how that individual ranks two different lotteries A and B . For example, suppose that we observe that Susan is faced with the choice between lotteries A and B below and she says that she prefers A to B :

$$A = \begin{pmatrix} o_1 & o_2 & o_3 \\ 0 & 0.25 & 0.75 \end{pmatrix} \quad B = \begin{pmatrix} o_1 & o_2 & o_3 \\ 0.2 & 0 & 0.8 \end{pmatrix}$$

With this information we can predict which of the following two lotteries C and D she will choose, if she has von Neumann-Morgenstern preferences:

$$C = \begin{pmatrix} o_1 & o_2 & o_3 \\ 0.8 & 0 & 0.2 \end{pmatrix} \quad D = \begin{pmatrix} o_1 & o_2 & o_3 \\ 0 & 1 & 0 \end{pmatrix} = o_2.$$

Let U be a von Neumann-Morgenstern utility function whose existence is guaranteed by Theorem 5.2.1. Let $U(o_1) = a$, $U(o_2) = b$ and $U(o_3) = c$ (where a , b and c are numbers). Then, since Susan prefers A to B , the expected utility of A must be greater than the expected utility of B : $0.25b + 0.75c > 0.2a + 0.8c$. This inequality is equivalent to $0.25b > 0.2a + 0.05c$ or, dividing both sides by 0.25, $b > 0.8a + 0.2c$. It follows from this and Theorem 5.2.1 that Susan prefers D to C , because the expected utility of D is b and the expected utility of C is $0.8a + 0.2c$. Note that, in this example, we merely used the fact that a von Neumann-Morgenstern utility function *exists*, even though we do not know what the values of this function are.

Theorem 5.2.1 is an example of a “representation theorem” and is a generalization of a similar result for the case of the ranking of a finite set of basic outcomes O . It is not difficult to prove that if \succsim is a complete and transitive ranking of O then there exists a function $U : O \rightarrow \mathbb{R}$, called a utility function (see Chapter 2), such that, for any two basic outcomes $o, o' \in O$, $U(o) \geq U(o')$ if and only if $o \succsim o'$. Now, it is quite possible that an individual has a complete and transitive ranking of O , is fully aware of her ranking and yet she is not able to answer the question “what is your utility function?”, perhaps because she has never heard about utility functions. A utility function is a *tool* that we can use to represent her ranking, nothing more than that. The same applies to von Neumann-Morgenstern rankings: Theorem 5.2.1 tells us that if an individual has a von Neumann-Morgenstern ranking of the set of lotteries $\mathcal{L}(O)$ then there exists a von Neumann-Morgenstern utility function that we can use to represent her preferences, but it would not make sense for us to ask the individual “what is your von Neumann-Morgenstern utility function?” (indeed this was a question that could not even be conceived before von Neumann and Morgenstern stated and proved Theorem 5.2.1 in 1944!).

Theorem 5.2.1 tells us that a von Neumann-Morgenstern utility function exists; the next theorem can be used to actually construct such a function, by asking the individual to answer a few questions, formulated in a way that is fully comprehensible to her (without using the word ‘utility’). The theorem says that, although there are many utility functions that represent a given von Neumann-Morgenstern ranking, once you know one function you “know them all”, in the sense that there is a simple operation that transforms one function into the other.

Theorem 5.2.2 [von Neumann-Morgenstern, 1944].

Let \succsim be a von Neumann-Morgenstern ranking of the set of basic lotteries $\mathcal{L}(O)$, where $O = \{o_1, o_2, \dots, o_m\}$. Then the following are true.

- (A) If $U : O \rightarrow \mathbb{R}$ is a von Neumann-Morgenstern utility function that represents \succsim , then, for any two real numbers a and b , with $a > 0$, the function $V : O \rightarrow \mathbb{R}$ defined by $V(o_i) = aU(o_i) + b$ (for every $i = 1, \dots, m$) is also a von Neumann-Morgenstern utility function that represents \succsim .
- (B) If $U : O \rightarrow \mathbb{R}$ and $V : O \rightarrow \mathbb{R}$ are two von Neumann-Morgenstern utility functions that represent \succsim , then there exist two real numbers a and b , with $a > 0$, such that $V(o_i) = aU(o_i) + b$ (for every $i = 1, \dots, m$).

Proof. The proof of Part A of Theorem 5.2.2 is very simple. Let a and b be two numbers, with $a > 0$. The hypothesis is that $U : O \rightarrow \mathbb{R}$ is a von Neumann-Morgenstern utility function that represents \succsim , that is, that, for any two lotteries

$$L = \begin{pmatrix} o_1 & \dots & o_m \\ p_1 & \dots & p_m \end{pmatrix} \quad \text{and} \quad L' = \begin{pmatrix} o_1 & \dots & o_m \\ q_1 & \dots & q_m \end{pmatrix},$$

$$L \succsim L' \text{ if and only if } p_1U(o_1) + \dots + p_mU(o_m) \geq q_1U(o_1) + \dots + q_mU(o_m) \quad (5.1)$$

Multiplying both sides of the inequality in (5.1) by $a > 0$ and adding $(p_1 + \dots + p_m)b$ to the left-hand side and $(q_1 + \dots + q_m)b$ to the right-hand side we obtain

$$p_1[aU(o_1) + b] + \dots + p_m[aU(o_m) + b] \geq q_1[aU(o_1) + b] + \dots + q_m[aU(o_m) + b] \quad (5.2)$$

Defining $V(o_i) = aU(o_i) + b$, it follows from (5.1) and (5.2) that

$$L \succsim L' \text{ if and only if } p_1V(o_1) + \dots + p_mV(o_m) \geq q_1V(o_1) + \dots + q_mV(o_m),$$

that is, the function V is a von Neumann-Morgenstern utility function that represents the ranking \succsim . The proof of Part B will be given later, after introducing more notation and some observations. ■

Suppose that the *DM* has a von Neumann-Morgenstern ranking of the set of lotteries $\mathcal{L}(O)$. Since among the lotteries there are the degenerate ones that assign probability 1 to a single basic outcome, it follows that the *DM* has a complete and transitive ranking of the basic outcomes. We shall write o_{best} for a best basic outcome, that is, a basic outcome which is at least as good as any other basic outcome ($o_{best} \succsim o$, for every $o \in O$) and o_{worst} for a worst basic outcome, that is, a basic outcome such that every other outcome is at least as good as it ($o \succsim o_{worst}$, for every $o \in O$). Note that there may be several best outcomes (then the *DM* would be indifferent among them) and several worst outcomes; then o_{best} will denote an arbitrary best outcome and o_{worst} an arbitrary worst outcome. We shall assume throughout that the *DM* is not indifferent among all the outcomes, that is, we shall assume that $o_{best} \succ o_{worst}$.

We now show that, in virtue of Theorem 5.2.2, among the von Neumann-Morgenstern utility functions that represent a given von Neumann-Morgenstern ranking \succsim of $\mathcal{L}(O)$, there is one that assigns the value 1 to the best basic outcome(s) and the value 0 to the worst basic outcome(s). To see this, consider an arbitrary von Neumann-Morgenstern

utility function $F : O \rightarrow \mathbb{R}$ that represents \succsim and define $G : O \rightarrow \mathbb{R}$ as follows: for every $o \in O$, $G(o) = F(o) - F(o_{worst})$.

Then, by Theorem 5.2.2 (with $a = 1$ and $b = -F(o_{worst})$), G is also a utility function that represents \succsim and, by construction, $G(o_{worst}) = F(o_{worst}) - F(o_{worst}) = 0$; note also that, since $o_{best} \succ o_{worst}$, it follows that $G(o_{best}) > 0$.

Finally, define $U : O \rightarrow \mathbb{R}$ as follows: for every $o \in O$, $U(o) = \frac{G(o)}{G(o_{best})}$.

Then, by Theorem 5.2.2 (with $a = \frac{1}{G(o_{best})}$ and $b = 0$), U is a utility function that represents \succsim and, by construction, $U(o_{worst}) = 0$ and $U(o_{best}) = 1$.

For example, if there are six basic outcomes and the ranking of the basic outcomes is $o_3 \sim o_6 \succ o_1 \succ o_4 \succ o_2 \sim o_5$, then one can take as o_{best} either o_3 or o_6 and as o_{worst} either o_2 or o_5 ; furthermore, if F is given by

o_1	o_2	o_3	o_4	o_5	o_6
2	-2	8	0	-2	8

then G is the function

$$\begin{array}{cccccc}
 o_1 & o_2 & o_3 & o_4 & o_5 & o_6 \\
 4 & 0 & 10 & 2 & 0 & 10
 \end{array}
 \text{ and } U \text{ is the function }
 \begin{array}{cccccc}
 o_1 & o_2 & o_3 & o_4 & o_5 & o_6 \\
 0.4 & 0 & 1 & 0.2 & 0 & 1
 \end{array}
 .$$

Definition 5.2.2 Let $U : O \rightarrow \mathbb{R}$ be a utility function that represents a given von Neumann-Morgenstern ranking \succsim of the set of lotteries $\mathcal{L}(O)$. We say that U is *normalized* if $U(o_{worst}) = 0$ and $U(o_{best}) = 1$.

The transformations described above show how to normalize any given utility function. Armed with the notion of a normalized utility function we can now complete the proof of Theorem 5.2.2.

Proof of Part B of Theorem 5.2.2. Let $F : O \rightarrow \mathbb{R}$ and $G : O \rightarrow \mathbb{R}$ be two von Neumann-Morgenstern utility functions that represent a given von Neumann-Morgenstern ranking of $\mathcal{L}(O)$.

Let $U : O \rightarrow \mathbb{R}$ be the normalization of F and $V : O \rightarrow \mathbb{R}$ be the normalization of G . First we show that it must be that $U = V$, that is, $U(o) = V(o)$ for every $o \in O$.

Suppose, by contradiction, that there is an $\hat{o} \in O$ such that $U(\hat{o}) \neq V(\hat{o})$. Without loss of generality we can assume that $U(\hat{o}) > V(\hat{o})$.

Construct the following lottery: $L = \begin{pmatrix} o_{best} & o_{worst} \\ \hat{p} & 1 - \hat{p} \end{pmatrix}$ with $\hat{p} = U(\hat{o})$ (recall that U is normalized and thus takes on values in the interval from 0 to 1).

Then $\mathbb{E}[U(L)] = \mathbb{E}[V(L)] = U(\hat{o})$. Hence, according to U it must be that $\hat{o} \sim L$ (this follows from Theorem 5.2.1), while according to V it must be (again, by Theorem 5.2.1) that $L \succ \hat{o}$ (since $\mathbb{E}[V(L)] = U(\hat{o}) > V(\hat{o})$). Then U and V cannot be two representations of the same ranking. Now let $a_1 = \frac{1}{F(o_{best}) - F(o_{worst})}$ and $b_1 = -\frac{F(o_{worst})}{F(o_{best}) - F(o_{worst})}$.

Note that $a_1 > 0$. Then it is easy to verify that, for every $o \in O$, $U(o) = a_1 F(o) + b_1$.

Similarly let $a_2 = \frac{1}{G(o_{best}) - G(o_{worst})}$ and $b_2 = -\frac{G(o_{worst})}{G(o_{best}) - G(o_{worst})}$; again, $a_2 > 0$ and, for every $o \in O$, $V(o) = a_2 G(o) + b_2$. We can invert the latter transformation and obtain that, for every $o \in O$, $G(o) = \frac{V(o)}{a_2} - \frac{b_2}{a_2}$.

Thus, we can transform F into U , which – as proved above – is the same as V , and then transform V into G thus obtaining the following transformation of F into G : $G(o) = aF(o) + b$ where $a = \frac{a_1}{a_2} > 0$ and $b = \frac{b_1 - b_2}{a_2}$. ■

R Theorem 5.2.2 is often stated as follows: a utility function that represents a von Neumann-Morgenstern ranking \succsim of $\mathcal{L}(O)$ is *unique up to a positive affine transformation*.² Because of Theorem 5.2.2, a von Neumann-Morgenstern utility function is usually referred to as a *cardinal* utility function.

Theorem 5.2.1 guarantees the existence of a utility function that represents a given von Neumann-Morgenstern ranking \succsim of $\mathcal{L}(O)$ and Theorem 5.2.2 characterizes the set of such functions. Can one actually construct a utility function that represents a given ranking? The answer is affirmative: if there are m basic outcomes one can construct an individual's von Neumann-Morgenstern utility function by asking her at most $(m - 1)$ questions. The first question is “what is your ranking of the basic outcomes?”. Then we can construct the normalized utility function by first assigning the value 1 to the best outcome(s) and the value 0 to the worst outcome(s). This leaves us with at most $(m - 2)$ values to determine. For this we appeal to one of the axioms discussed in the next section, namely the Continuity Axiom, which says that, for every basic outcome o_i there is a probability $p_i \in [0, 1]$ such that the *DM* is indifferent between o_i for certain and the lottery that gives a best outcome with probability p_i and a worst outcome with probability $(1 - p_i)$: $o_i \sim \begin{pmatrix} o_{best} & o_{worst} \\ p_i & 1 - p_i \end{pmatrix}$. Thus, for each basic outcome o_i for which a utility has not been determined yet, we should ask the individual to tell us the value of p_i such that $o_i \sim \begin{pmatrix} o_{best} & o_{worst} \\ p_i & 1 - p_i \end{pmatrix}$; then we can set $U_i(o_i) = p_i$, because the expected utility of the lottery $\begin{pmatrix} o_{best} & o_{worst} \\ p_i & 1 - p_i \end{pmatrix}$ is $p_i U_i(o_{best}) + (1 - p_i) U_i(o_{worst}) = p_i(1) + (1 - p_i)0 = p_i$.

■ **Example 5.1** Suppose that there are five basic outcomes, that is, $O = \{o_1, o_2, o_3, o_4, o_5\}$ and the *DM*, who has von Neumann-Morgenstern preferences, tells us that her ranking of the basic outcomes is as follows: $o_2 \succ o_1 \sim o_5 \succ o_3 \sim o_4$.

- Then we can begin by assigning utility 1 to the best outcome o_2 and utility 0 to the worst outcomes o_3 and o_4 : $\begin{pmatrix} \text{outcome:} & o_1 & o_2 & o_3 & o_4 & o_5 \\ \text{utility:} & ? & 1 & 0 & 0 & ? \end{pmatrix}$.
- There is only one value left to be determined, namely the utility of o_1 (which is also the utility of o_5 , since $o_1 \sim o_5$).
- To find this value, we ask the *DM* to tell us what value of p makes her indifferent between the lottery $L = \begin{pmatrix} o_2 & o_3 \\ p & 1 - p \end{pmatrix}$ and outcome o_1 with certainty.
- Suppose that her answer is: 0.4. Then her normalized von Neumann-Morgenstern utility function is $\begin{pmatrix} \text{outcome:} & o_1 & o_2 & o_3 & o_4 & o_5 \\ \text{utility:} & 0.4 & 1 & 0 & 0 & 0.4 \end{pmatrix}$. Knowing this, we can predict her choice among any set of lotteries over these five basic outcomes.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 5.4.2 at the end of this chapter.

²An affine transformation is a function $f: \mathbb{R} \rightarrow \mathbb{R}$ of the form $f(x) = ax + b$ with $a, b \in \mathbb{R}$. The affine transformation is positive if $a > 0$.

5.3 Expected utility: the axioms

We can now turn to the list of rationality axioms proposed by von Neumann and Morgenstern. This section makes heavy use of mathematical notation and, as mentioned in the previous section, if the reader is not interested in understanding in what sense the theory of expected utility captures the notion of rationality, he/she can skip it without affecting his/her ability to understand the rest of this book.

Let $O = \{o_1, o_2, \dots, o_m\}$ be the set of basic outcomes and $\mathcal{L}(O)$ the set of simple lotteries, that is, the set of probability distributions over O . Let \succsim be a binary relation on $\mathcal{L}(O)$. We say that \succsim is a von Neumann-Morgenstern ranking of $\mathcal{L}(O)$ if it satisfies the following four axioms or properties.

Axiom 1 [Completeness and transitivity]. \succsim is complete (for every two lotteries L and L' either $L \succsim L'$ or $L' \succsim L$ or both) and transitive (for any three lotteries L_1, L_2 and L_3 , if $L_1 \succsim L_2$ and $L_2 \succsim L_3$ then $L_1 \succsim L_3$).

As noted in the previous section, Axiom 1 implies that there is a complete and transitive ranking of the basic outcomes. Recall that o_{best} denotes a best basic outcome and o_{worst} denotes a worst basic outcome and that we are assuming that $o_{best} \succ o_{worst}$, that is, that the DM is not indifferent among all the basic outcomes.

Axiom 2 [Monotonicity]. $\left(\begin{array}{cc} o_{best} & o_{worst} \\ p & 1-p \end{array} \right) \succsim \left(\begin{array}{cc} o_{best} & o_{worst} \\ q & 1-q \end{array} \right)$ if and only if $p \geq q$.

Axiom 3 [Continuity]. For every basic outcome o_i there is a $p_i \in [0, 1]$ such that $o_i \sim \left(\begin{array}{cc} o_{best} & o_{worst} \\ p_i & 1-p_i \end{array} \right)$.

Before we introduce the last axiom we need to define a compound lottery.

Definition 5.3.1 A *compound lottery* is a lottery of the form $\left(\begin{array}{cccc} x_1 & x_2 & \dots & x_r \\ p_1 & p_2 & \dots & p_r \end{array} \right)$ where each x_i is either an element of O or an element of $\mathcal{L}(O)$.

For example, let $m = 4$.

Then $L = \left(\begin{array}{cccc} o_1 & o_2 & o_3 & o_4 \\ \frac{2}{5} & 0 & \frac{1}{5} & \frac{2}{5} \end{array} \right)$ is a simple lottery (an element of $\mathcal{L}(O)$),

while $C = \left(\begin{array}{ccc} \left(\begin{array}{cccc} o_1 & o_2 & o_3 & o_4 \\ \frac{1}{3} & \frac{1}{6} & \frac{1}{3} & \frac{1}{6} \end{array} \right) & o_1 & \left(\begin{array}{cccc} o_1 & o_2 & o_3 & o_4 \\ \frac{1}{5} & 0 & \frac{1}{5} & \frac{3}{5} \end{array} \right) \\ \frac{1}{2} & \frac{1}{4} & \frac{1}{4} \end{array} \right)$ is a compound lottery.³

³With $r = 3$, $x_1 = \left(\begin{array}{cccc} o_1 & o_2 & o_3 & o_4 \\ \frac{1}{3} & \frac{1}{6} & \frac{1}{3} & \frac{1}{6} \end{array} \right)$, $x_2 = o_1$, $x_3 = \left(\begin{array}{cccc} o_1 & o_2 & o_3 & o_4 \\ \frac{1}{5} & 0 & \frac{1}{5} & \frac{3}{5} \end{array} \right)$, $p_1 = \frac{1}{2}$, $p_2 = \frac{1}{4}$ and $p_3 = \frac{1}{4}$.

$$\text{The compound lottery } C = \left(\begin{array}{c} \left(\begin{array}{cccc} o_1 & o_2 & o_3 & o_4 \\ \frac{1}{3} & \frac{1}{6} & \frac{1}{3} & \frac{1}{6} \end{array} \right) o_1 \left(\begin{array}{cccc} o_1 & o_2 & o_3 & o_4 \\ \frac{1}{5} & 0 & \frac{1}{5} & \frac{3}{5} \end{array} \right) \\ \frac{1}{2} \qquad \qquad \frac{1}{4} \qquad \qquad \frac{1}{4} \end{array} \right)$$

can be viewed graphically as a tree, as shown in Figure 5.1.




Figure 5.1: A compound lottery

Definition 5.3.2 Given a compound lottery $C = \begin{pmatrix} x_1 & x_2 & \dots & x_r \\ p_1 & p_2 & \dots & p_r \end{pmatrix}$ the corresponding simple lottery $L(C) = \begin{pmatrix} o_1 & o_2 & \dots & o_m \\ q_1 & q_2 & \dots & q_m \end{pmatrix}$ is defined as follows. First of all, for $i = 1, \dots, m$ and $j = 1, \dots, r$, define

$$o_i(x_j) = \begin{cases} 1 & \text{if } x_j = o_i \\ 0 & \text{if } x_j = o_k \text{ with } k \neq i \\ s_i & \text{if } x_j = \begin{pmatrix} o_1 & \dots & o_{i-1} & o_i & o_{i+1} & \dots & o_m \\ s_1 & \dots & s_{i-1} & s_i & s_{i+1} & \dots & s_m \end{pmatrix} \end{cases}$$

Then $q_i = \sum_{j=1}^r p_j o_i(x_j)$.

Continuing the above example where

$$C = \left(\begin{array}{c} \left(\begin{array}{cccc} o_1 & o_2 & o_3 & o_4 \\ \frac{1}{3} & \frac{1}{6} & \frac{1}{3} & \frac{1}{6} \end{array} \right) o_1 \left(\begin{array}{cccc} o_1 & o_2 & o_3 & o_4 \\ \frac{1}{5} & 0 & \frac{1}{5} & \frac{3}{5} \end{array} \right) \\ \frac{1}{2} \qquad \qquad \frac{1}{4} \qquad \qquad \frac{1}{4} \end{array} \right)$$

(see Figure 5.1) we have that $r = 3$, $x_1 = \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ \frac{1}{3} & \frac{1}{6} & \frac{1}{3} & \frac{1}{6} \end{pmatrix}$, $x_2 = o_1$ and

$x_3 = \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ \frac{1}{5} & 0 & \frac{1}{5} & \frac{3}{5} \end{pmatrix}$, so that

$$o_1(x_1) = \frac{1}{3}, \quad o_1(x_2) = 1, \quad \text{and} \quad o_1(x_3) = \frac{1}{5}$$

and thus $q_1 = \frac{1}{2} \left(\frac{1}{3} \right) + \frac{1}{4} (1) + \frac{1}{4} \left(\frac{1}{5} \right) = \frac{28}{60}$. Similarly, $q_2 = \frac{1}{2} \left(\frac{1}{6} \right) + \frac{1}{4} (0) + \frac{1}{4} (0) = \frac{1}{12} = \frac{5}{60}$,

$q_3 = \frac{1}{2} \left(\frac{1}{3}\right) + \frac{1}{4}(0) + \frac{1}{4} \left(\frac{1}{5}\right) = \frac{13}{60}$ and $q_4 = \frac{1}{2} \left(\frac{1}{6}\right) + \frac{1}{4}(0) + \frac{1}{4} \left(\frac{3}{5}\right) = \frac{14}{60}$. These numbers correspond to multiplying the probabilities along the edges of the tree of Figure 5.1 leading to an outcome, as shown in Figure 5.2 and then adding up the probabilities of each outcome, as shown in Figure 5.3. Thus, the simple lottery $L(C)$ that corresponds to C is $L(C) = \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ \frac{28}{60} & \frac{5}{60} & \frac{13}{60} & \frac{14}{60} \end{pmatrix}$, namely the lottery shown in Figure 5.3.




Figure 5.2: Simplification of Figure 5.1 obtained by merging paths into simple edges and associating with the simple edges the products of the probabilities along the path.




Figure 5.3: Simplification of Figure 5.2 obtained by adding, for each outcome, the probabilities of that outcome.

Axiom 4 [Independence or substitutability]. Consider an arbitrary basic outcome o_i and an arbitrary simple lottery $L = \begin{pmatrix} o_1 & \dots & o_{i-1} & o_i & o_{i+1} & \dots & o_m \\ p_1 & \dots & p_{i-1} & p_i & p_{i+1} & \dots & p_m \end{pmatrix}$. If \hat{L} is a simple lottery such that $o_i \sim \hat{L}$, then $L \sim M$ where M is the simple lottery corresponding to the compound lottery $C = \begin{pmatrix} o_1 & \dots & o_{i-1} & \hat{L} & o_{i+1} & \dots & o_m \\ p_1 & \dots & p_{i-1} & p_i & p_{i+1} & \dots & p_m \end{pmatrix}$ obtained by replacing o_i with \hat{L} in L .

We can now prove the first theorem of the previous section.

Proof of Theorem 5.2.1. To simplify the notation, throughout this proof we will assume that we have renumbered the basic outcomes in such a way that $o_{best} = o_1$ and $o_{worst} = o_m$. First of all, for every basic outcome o_i , let $u_i \in [0, 1]$ be such that $o_i \sim \begin{pmatrix} o_1 & o_m \\ u_i & 1 - u_i \end{pmatrix}$. The existence of such a value u_i is guaranteed by the Continuity Axiom (Axiom 3); clearly $u_1 = 1$ and $u_m = 0$. Now consider an arbitrary lottery

$$L_1 = \begin{pmatrix} o_1 & \dots & o_m \\ p_1 & \dots & p_m \end{pmatrix}.$$

First we show that

$$L_1 \sim \begin{pmatrix} o_1 & o_m \\ \sum_{i=1}^m p_i u_i & 1 - \sum_{i=1}^m p_i u_i \end{pmatrix} \quad (5.3)$$

This is done through a repeated application of the Independence Axiom (Axiom 4), as follows. Consider the compound lottery

$$\mathcal{C}_2 = \begin{pmatrix} o_1 & \begin{pmatrix} o_1 & o_m \\ u_2 & 1 - u_2 \end{pmatrix} & o_3 & \dots & o_m \\ p_1 & p_2 & p_3 & \dots & p_m \end{pmatrix}$$

obtained by replacing o_2 in lottery L_1 with the lottery $\begin{pmatrix} o_1 & o_m \\ u_2 & 1 - u_2 \end{pmatrix}$ that the DM considers to be just as good as o_2 . The simple lottery corresponding to \mathcal{C}_2 is

$$L_2 = \begin{pmatrix} o_1 & o_3 & \dots & o_{m-1} & o_m \\ p_1 + p_2 u_2 & p_3 & \dots & p_{m-1} & p_m + p_2(1 - u_2) \end{pmatrix}.$$

Note that o_2 is assigned probability 0 in L_2 and thus we have omitted it. By Axiom 4, $L_1 \sim L_2$. Now apply the same argument to L_2 : let

$$\mathcal{C}_3 = \begin{pmatrix} o_1 & \begin{pmatrix} o_1 & o_m \\ u_3 & 1 - u_3 \end{pmatrix} & \dots & o_m \\ p_1 + p_2 u_2 & p_3 & \dots & p_m + p_2(1 - u_2) \end{pmatrix}$$

whose corresponding simple lottery is

$$L_3 = \begin{pmatrix} o_1 & \dots & o_m \\ p_1 + p_2 u_2 + p_3 u_3 & \dots & p_m + p_2(1 - u_2) + p_3(1 - u_3) \end{pmatrix}.$$

Note, again, that o_3 is assigned probability zero in L_3 . By Axiom 4, $L_2 \sim L_3$; thus, by transitivity (since $L_1 \sim L_2$ and $L_2 \sim L_3$) we have that $L_1 \sim L_3$. Repeating this argument we get that $L_1 \sim L_{m-1}$, where

$$L_{m-1} = \begin{pmatrix} o_1 & o_m \\ p_1 + p_2 u_2 + \dots + p_{m-1} u_{m-1} & p_m + p_2(1 - u_2) + \dots + p_{m-1}(1 - u_{m-1}) \end{pmatrix}.$$

Since $u_1 = 1$ (so that $p_1 u_1 = p_1$) and $u_m = 0$ (so that $p_m u_m = 0$),

$$p_1 + p_2 u_2 + \dots + p_{m-1} u_{m-1} = \sum_{i=1}^m p_i u_i$$

and

$$\begin{aligned} p_2(1 - u_2) + \dots + p_{m-1}(1 - u_{m-1}) + p_m &= \sum_{i=2}^m p_i - \sum_{i=2}^{m-1} p_i u_i = p_1 + \sum_{i=2}^m p_i - \sum_{i=2}^{m-1} p_i u_i - p_1 \\ &= (\text{since } u_1=1 \text{ and } u_m=0) \sum_{i=1}^m p_i - \sum_{i=2}^{m-1} p_i u_i - p_1 u_1 - p_m u_m = \left(\text{since } \sum_{i=1}^m p_i=1 \right) 1 - \sum_{i=1}^m p_i u_i \end{aligned}$$

Thus, $L_{m-1} = \left(\begin{array}{cc} o_1 & o_m \\ \sum_{i=1}^m p_i u_i & 1 - \sum_{i=1}^m p_i u_i \end{array} \right)$, which proves (5.3). Now define the following utility function $U : \{o_1, \dots, o_m\} \rightarrow [0, 1]$: $U(o_i) = u_i$, where, as before, for every basic outcome o_i , $u_i \in [0, 1]$ is such that $o_i \sim \left(\begin{array}{cc} o_1 & o_m \\ u_i & 1 - u_i \end{array} \right)$. Consider two arbitrary lotteries

$L = \left(\begin{array}{ccc} o_1 & \dots & o_m \\ p_1 & \dots & p_m \end{array} \right)$ and $L' = \left(\begin{array}{ccc} o_1 & \dots & o_m \\ q_1 & \dots & q_m \end{array} \right)$. We want to show that $L \succsim L'$ if and only if $\mathbb{E}[U(L)] \geq \mathbb{E}[U(L')]$, that is, if and only if $\sum_{i=1}^m p_i u_i \geq \sum_{i=1}^m q_i u_i$. By (5.3), $L \sim M$, where

$$M = \left(\begin{array}{cc} o_1 & o_m \\ \sum_{i=1}^m p_i u_i & 1 - \sum_{i=1}^m p_i u_i \end{array} \right) \text{ and also } L' \sim M', \text{ where } M' = \left(\begin{array}{cc} o_1 & o_m \\ \sum_{i=1}^m q_i u_i & 1 - \sum_{i=1}^m q_i u_i \end{array} \right).$$

Thus, by transitivity of \succsim , $L \succsim L'$ if and only if $M \succsim M'$; by the Monotonicity Axiom (Axiom 2), $M \succsim M'$ if and only if $\sum_{i=1}^m p_i u_i \geq \sum_{i=1}^m q_i u_i$. \blacksquare

The following example, known as the *Allais paradox*, suggests that one should view expected utility theory as a “prescriptive” or “normative” theory (that is, as a theory about how rational people should choose) rather than as a descriptive theory (that is, as a theory about the actual behavior of individuals). In 1953 the French economist Maurice Allais published a paper regarding a survey he had conducted in 1952 concerning a hypothetical decision problem. Subjects “with good training in and knowledge of the theory of probability, so that they could be considered to behave rationally” were asked to rank the following pairs of lotteries:

$$A = \left(\begin{array}{cc} \$5 \text{ Million} & \$0 \\ \frac{89}{100} & \frac{11}{100} \end{array} \right) \text{ versus } B = \left(\begin{array}{cc} \$1 \text{ Million} & \$0 \\ \frac{90}{100} & \frac{10}{100} \end{array} \right)$$

and

$$C = \left(\begin{array}{ccc} \$5 \text{ Million} & \$1 \text{ Million} & \$0 \\ \frac{89}{100} & \frac{10}{100} & \frac{1}{100} \end{array} \right) \text{ versus } D = \left(\begin{array}{c} \$1 \text{ Million} \\ 1 \end{array} \right).$$

Most subjects reported the following ranking: $A \succ B$ and $D \succ C$. Such ranking violates the axioms of expected utility. To see this, let $O = \{o_1, o_2, o_3\}$ with $o_1 = \$5 \text{ Million}$, $o_2 = \$1 \text{ Million}$ and $o_3 = \$0$. Let us assume that the individual in question prefers more

money to less, so that $o_1 \succ o_2 \succ o_3$ and has a von Neumann-Morgenstern ranking of the lotteries over $\mathcal{L}(O)$. Let $u_2 \in (0, 1)$ be such that $D \sim \begin{pmatrix} \$5 \text{ Million} & \$0 \\ u_2 & 1 - u_2 \end{pmatrix}$ (the existence of such u_2 is guaranteed by the Continuity Axiom). Then, since $D \succ C$, by transitivity

$$\begin{pmatrix} \$5 \text{ Million} & \$0 \\ u_2 & 1 - u_2 \end{pmatrix} \succ C. \quad (5.4)$$

Let C' be the simple lottery corresponding to the compound lottery

$$\begin{pmatrix} \$5 \text{ Million} & \begin{pmatrix} \$5 \text{ Million} & \$0 \\ u_2 & 1 - u_2 \end{pmatrix} & \$0 \\ \frac{89}{100} & \frac{10}{100} & \frac{1}{100} \end{pmatrix}.$$

$$\text{Then } C' = \begin{pmatrix} \$5 \text{ Million} & \$0 \\ \frac{89}{100} + \frac{10}{100}u_2 & 1 - (\frac{89}{100} + \frac{10}{100}u_2) \end{pmatrix}.$$

By the Independence Axiom, $C \sim C'$ and thus, by (5.4) and transitivity,

$$\begin{pmatrix} \$5 \text{ Million} & \$0 \\ u_2 & 1 - u_2 \end{pmatrix} \succ \begin{pmatrix} \$5 \text{ Million} & \$0 \\ \frac{89}{100} + \frac{10}{100}u_2 & 1 - (\frac{89}{100} + \frac{10}{100}u_2) \end{pmatrix}.$$

Hence, by the Monotonicity Axiom, $u_2 > \frac{89}{100} + \frac{10}{100}u_2$, that is,

$$u_2 > \frac{89}{90}. \quad (5.5)$$

Let B' be the simple lottery corresponding to the following compound lottery, constructed from B by replacing the basic outcome '\$1 Million' with $\begin{pmatrix} \$5 \text{ Million} & \$0 \\ u_2 & 1 - u_2 \end{pmatrix}$:

$$\begin{pmatrix} \begin{pmatrix} \$5 \text{ Million} & \$0 \\ u_2 & 1 - u_2 \end{pmatrix} & \$0 \\ \frac{90}{100} & \frac{10}{100} \end{pmatrix}.$$

Then

$$B' = \begin{pmatrix} \$5 \text{ Million} & \$0 \\ \frac{90}{100}u_2 & 1 - \frac{90}{100}u_2 \end{pmatrix}.$$

By the Independence Axiom, $B \sim B'$; thus, since $A \succ B$, by transitivity, $A \succ B'$ and therefore, by the Monotonicity Axiom, $\frac{89}{100} > \frac{90}{100}u_2$, that is, $u_2 < \frac{89}{90}$, contradicting (5.5).

Thus, if one finds the expected utility axioms compelling as axioms of rationality, then one cannot consistently express a preference for A over B and also a preference for D over C .

Another well-known example is the *Ellsberg paradox*. Suppose that you are told that an urn contains 30 red balls and 60 more balls that are either blue or yellow. You don't know how many blue or how many yellow balls there are, but the number of blue balls plus the number of yellow ball equals 60 (they could be all blue or all yellow or any combination of the two). The balls are well mixed so that each individual ball is as likely to be drawn as any other. You are given a choice between the bets A and B , where

A = you get \$100 if you pick a red ball and nothing otherwise,
 B = you get \$100 if you pick a blue ball and nothing otherwise.

Many subjects in experiments state a strict preference for A over B : $A \succ B$. Consider now the following bets:

C = you get \$100 if you pick a red or yellow ball and nothing otherwise,
 D = you get \$100 if you pick a blue or yellow ball and nothing otherwise.

Do the axioms of expected utility constrain your ranking of C and D ? Many subjects in experiments state the following ranking: $A \succ B$ and $D \succ C$. All such people violate the axioms of expected utility. The fraction of red balls in the urn is $\frac{30}{90} = \frac{1}{3}$. Let p_2 be the fraction of blue balls and p_3 the fraction of yellow balls (either of these can be zero: all we know is that $p_2 + p_3 = \frac{60}{90} = \frac{2}{3}$). Then A, B, C and D can be viewed as the following lotteries:

$$A = \begin{pmatrix} \$100 & \$0 \\ \frac{1}{3} & p_2 + p_3 \end{pmatrix}, \quad B = \begin{pmatrix} \$100 & \$0 \\ p_2 & \frac{1}{3} + p_3 \end{pmatrix}$$

$$C = \begin{pmatrix} \$100 & \$0 \\ \frac{1}{3} + p_3 & p_2 \end{pmatrix}, \quad D = \begin{pmatrix} \$100 & \$0 \\ p_2 + p_3 = \frac{2}{3} & \frac{1}{3} \end{pmatrix}$$

Let U be the normalized von Neumann-Morgenstern utility function that represents the individual's ranking; then $U(\$100) = 1$ and $U(0) = 0$. Thus,

$$\mathbb{E}[U(A)] = \frac{1}{3}, \quad \mathbb{E}[U(B)] = p_2, \quad \mathbb{E}[U(C)] = \frac{1}{3} + p_3, \quad \text{and} \quad \mathbb{E}[U(D)] = p_2 + p_3 = \frac{2}{3}.$$

Hence, $A \succ B$ if and only if $\frac{1}{3} > p_2$, which implies that $p_3 > \frac{1}{3}$, so that $\mathbb{E}[U(C)] = \frac{1}{3} + p_3 > \mathbb{E}[U(D)] = \frac{2}{3}$ and thus $C \succ D$ (similarly, $B \succ A$ if and only if $\frac{1}{3} < p_2$, which implies that $\mathbb{E}[U(C)] < \mathbb{E}[U(D)]$ and thus $D \succ C$).

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 5.4.3 at the end of this chapter.

5.4 Exercises

The solutions to the following exercises are given in Section 5.5 at the end of this chapter.

5.4.1 Exercises for Section 5.1: Money lotteries and attitudes to risk

Exercise 5.1 What is the expected value of the following lottery?

$$\left(\begin{array}{cccc} 24 & 12 & 48 & 6 \\ \frac{1}{6} & \frac{2}{6} & \frac{1}{6} & \frac{2}{6} \end{array} \right)$$

Exercise 5.2 Consider the following lottery:

$$\left(\begin{array}{ccc} o_1 & o_2 & o_3 \\ \frac{1}{4} & \frac{1}{2} & \frac{1}{4} \end{array} \right)$$

where

- o_1 = you get an invitation to have dinner at the White House,
- o_2 = you get (for free) a puppy of your choice
- o_3 = you get \$600.

What is the expected value of this lottery?

Exercise 5.3 Consider the following money lottery

$$L = \left(\begin{array}{cccccc} \$10 & \$15 & \$18 & \$20 & \$25 & \$30 & \$36 \\ \frac{3}{12} & \frac{1}{12} & 0 & \frac{3}{12} & \frac{2}{12} & 0 & \frac{3}{12} \end{array} \right)$$

- (a) What is the expected value of the lottery?
- (b) Ann prefers more money to less and has transitive preferences. She says that, between getting \$20 for certain and playing the above lottery, she would prefer \$20 for certain. What is her attitude to risk?
- (c) Bob prefers more money to less and has transitive preferences. He says that, given the same choice as Ann, he would prefer playing the lottery. What is his attitude to risk?

Exercise 5.4 Sam has a debilitating illness and has been offered two mutually exclusive courses of action: (1) take some well-known drugs which have been tested for a long time and (2) take a new experimental drug. If he chooses (1) then for certain his pain will be reduced to a bearable level. If he chooses (2) then he has a 50% chance of being completely cured and a 50% chance of no benefits from the drug and possibly some harmful side effects. He chose (1). What is his attitude to risk?

5.4.2 Exercises for Section 5.2: Expected utility theory

Exercise 5.5 Ben is offered a choice between the following two money lotteries:

$A = \begin{pmatrix} \$4,000 & \$0 \\ 0.8 & 0.2 \end{pmatrix}$ and $B = \begin{pmatrix} \$3,000 \\ 1 \end{pmatrix}$. He says he strictly prefers B to A . Which of the following two lotteries, C and D , will Ben choose if he satisfies the axioms of expected utility and prefers more money to less?

$$C = \begin{pmatrix} \$4,000 & \$0 \\ 0.2 & 0.8 \end{pmatrix}, \quad D = \begin{pmatrix} \$3,000 & \$0 \\ 0.25 & 0.75 \end{pmatrix}. \quad \blacksquare$$

Exercise 5.6 There are three basic outcomes, o_1, o_2 and o_3 . Ann satisfies the axioms of expected utility and her preferences over lotteries involving these three outcomes can be represented by the following von Neumann-Morgenstern utility function:

$$V(o_2) = a > V(o_1) = b > V(o_3) = c. \text{ Normalize the utility function.} \quad \blacksquare$$

Exercise 5.7 Consider the following lotteries:

$$L_1 = \begin{pmatrix} \$3000 & \$500 \\ \frac{5}{6} & \frac{1}{6} \end{pmatrix}, \quad L_2 = \begin{pmatrix} \$3000 & \$500 \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix},$$

$$L_3 = \begin{pmatrix} \$3000 & \$2000 & \$1000 & \$500 \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \end{pmatrix}, \quad L_4 = \begin{pmatrix} \$2000 & \$1000 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}.$$

Jennifer says that she is indifferent between lottery L_1 and getting \$2,000 for certain. She is also indifferent between lottery L_2 and getting \$1,000 for certain. Finally, she says that between L_3 and L_4 she would chose L_3 .

Is she rational according to the theory of expected utility? [Assume that she prefers more money to less.] \blacksquare

Exercise 5.8 Consider the following basic outcomes:

- o_1 = a Summer internship at the White House,
- o_2 = a free one-week vacation in Europe,
- o_3 = \$800,
- o_4 = a free ticket to a concert.

Rachel says that her ranking of these outcomes is $o_1 \succ o_2 \succ o_3 \succ o_4$. She also says

that **(1)** she is indifferent between $\begin{pmatrix} o_2 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} o_1 & o_4 \\ \frac{4}{5} & \frac{1}{5} \end{pmatrix}$ and **(2)** she is indifferent

between $\begin{pmatrix} o_3 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} o_1 & o_4 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$. If she satisfies the axioms of expected utility theory,

which of the two lotteries $L_1 = \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ \frac{1}{8} & \frac{2}{8} & \frac{3}{8} & \frac{2}{8} \end{pmatrix}$ and $L_2 = \begin{pmatrix} o_1 & o_2 & o_3 \\ \frac{1}{5} & \frac{3}{5} & \frac{1}{5} \end{pmatrix}$ will she choose? \blacksquare

Exercise 5.9 Consider the following lotteries: $L_1 = \begin{pmatrix} \$30 & \$28 & \$24 & \$18 & \$8 \\ \frac{2}{10} & \frac{1}{10} & \frac{1}{10} & \frac{2}{10} & \frac{4}{10} \end{pmatrix}$

and $L_2 = \begin{pmatrix} \$30 & \$28 & \$8 \\ \frac{1}{10} & \frac{4}{10} & \frac{5}{10} \end{pmatrix}$.

- Which lottery would a risk neutral person choose?
- Paul's von Neumann-Morgenstern utility-of-money function is $U(m) = \ln(m)$, where \ln denotes the natural logarithm. Which lottery would Paul choose?

Exercise 5.10 There are five basic outcomes. Jane has a von Neumann-Morgenstern ranking of the set of lotteries over the set of basic outcomes that can be represented by

either of the following utility functions U and V : $\begin{pmatrix} o_1 & o_2 & o_3 & o_4 & o_5 \\ U : & 44 & 170 & -10 & 26 & 98 \\ V : & 32 & 95 & 5 & 23 & 59 \end{pmatrix}$.

- Show how to normalize each of U and V and verify that you get the same normalized utility function.
- Show how to transform U into V with a positive affine transformation of the form $x \mapsto ax + b$ with $a, b \in \mathbb{R}$ and $a > 0$.

Exercise 5.11 Consider the following lotteries: $L_3 = \begin{pmatrix} \$28 \\ 1 \end{pmatrix}$, $L_4 = \begin{pmatrix} \$10 & \$50 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$.

- Ann has the following von Neumann-Morgenstern utility function: $U_{Ann}(\$m) = \sqrt{m}$. How does she rank the two lotteries?
- Bob has the following von Neumann-Morgenstern utility function: $U_{Bob}(\$m) = 2m - \frac{m^4}{100^3}$. How does he rank the two lotteries?
- Verify that both Ann and Bob are risk averse, by determining what they would choose between lottery L_4 and its expected value for certain.

5.4.3 Exercises for Section 5.3: Expected utility axioms

Exercise 5.12 Let $O = \{o_1, o_2, o_3, o_4\}$. Find the simple lottery corresponding to the following compound lottery

$$\left(\begin{pmatrix} \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ \frac{2}{5} & \frac{1}{10} & \frac{3}{10} & \frac{1}{5} \end{pmatrix} & o_2 & \begin{pmatrix} o_1 & o_3 & o_4 \\ \frac{1}{5} & \frac{1}{5} & \frac{3}{5} \end{pmatrix} & \begin{pmatrix} o_2 & o_3 \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix} \\ \frac{1}{8} & \frac{1}{4} & \frac{1}{8} & \frac{1}{2} \end{pmatrix} \right)$$

Exercise 5.13 Let $O = \{o_1, o_2, o_3, o_4\}$. Suppose that the *DM* has a von Neumann-Morgenstern ranking of $\mathcal{L}(O)$ and states the following indifference:

$$o_1 \sim \begin{pmatrix} o_2 & o_4 \\ \frac{1}{4} & \frac{3}{4} \end{pmatrix} \text{ and } o_2 \sim \begin{pmatrix} o_3 & o_4 \\ \frac{3}{5} & \frac{2}{5} \end{pmatrix}.$$

Find a lottery that the *DM* considers just as good as $L = \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ \frac{1}{3} & \frac{2}{9} & \frac{1}{9} & \frac{1}{3} \end{pmatrix}$.

Do not add any information to what is given above (in particular, do not make any assumptions about which outcome is best and which is worst). ■

Exercise 5.14 — *** Challenging Question ***.

Would you be willing to pay more in order to reduce the probability of dying within the next hour from one sixth to zero or from four sixths to three sixths? Unfortunately, this is not a hypothetical question: you accidentally entered the office of a mad scientist and have been overpowered and tied to a chair. The mad scientist has put six glasses in front of you, numbered 1 to 6, and tells you that one of them contains a deadly poison and the other five contain a harmless liquid. He says that he is going to roll a die and make you drink from the glass whose number matches the number that shows from the rolling of the die. You beg to be exempted and he asks you “what is the largest amount of money that you would be willing to pay to replace the glass containing the poison with one containing a harmless liquid?”. Interpret this question as “what sum of money x makes you indifferent between (1) leaving the poison in whichever glass contains it and rolling the die, and (2) reducing your wealth by $\$x$ and rolling the die after the poison has been replaced by a harmless liquid”. Your answer is: $\$X$. Then he asks you “suppose that instead of one glass with poison there had been four glasses with poison (and two with a harmless liquid); what is the largest amount of money that you would be willing to pay to replace one glass with poison with a glass containing a harmless liquid (and thus roll the die with 3 glasses with poison and 3 with a harmless liquid)?”. Your answer is: $\$Y$. Show that if $X > Y$ then you do not satisfy the axioms of Expected Utility Theory. [Hint: think about what the basic outcomes are; assume that you do not care about how much money is left in your estate if you die and that, when alive, you prefer more money to less.] ■

5.5 Solutions to Exercises

Solution to Exercise 5.1 The expected value of the lottery $\begin{pmatrix} 24 & 12 & 48 & 6 \\ \frac{1}{6} & \frac{2}{6} & \frac{1}{6} & \frac{2}{6} \end{pmatrix}$ is

$$\frac{1}{6}(24) + \frac{2}{6}(12) + \frac{1}{6}(48) + \frac{2}{6}(6) = 18. \quad \square$$

Solution to Exercise 5.2 This was a trick question! There is no expected value because the basic outcomes are not numbers. □

Solution to Exercise 5.3

(a) The expected value of the lottery

$$L = \left(\begin{array}{ccccccc} \$10 & \$15 & \$18 & \$20 & \$25 & \$30 & \$36 \\ \frac{3}{12} & \frac{1}{12} & 0 & \frac{3}{12} & \frac{2}{12} & 0 & \frac{3}{12} \end{array} \right)$$

is $\mathbb{E}[L] = \frac{3}{12}(10) + \frac{1}{12}(15) + (0)(18) + \frac{3}{12}(20) + \frac{2}{12}(25) + (0)(30) + \frac{3}{12}(36) = \frac{263}{12} = \21.92

(b) Since Ann prefers more money to less, she prefers \$21.92 to \$20 ($\$21.92 \succ \20). She said that she prefers \$20 to lottery L ($\$20 \succ L$). Thus, since her preferences are transitive, she prefers \$21.92 to lottery L ($\$21.92 \succ L$). Hence, she is risk averse.

(c) The answer is: we cannot tell. First of all, since Bob prefers more money to less, he prefers \$21.92 to \$20 ($\$21.92 \succ \20). Bob could be risk neutral, because a risk neutral person would be indifferent between L and \$21.92 ($L \sim \21.92); since Bob prefers \$21.92 to \$20 and has transitive preferences, if risk neutral he would prefer L to \$20.

However, Bob could also be risk loving: a risk-loving person prefers L to \$21.92 ($L \succ \21.92) and we know that he prefers \$21.92 to \$20; thus, by transitivity, if risk loving, he would prefer L to \$20.

But Bob could also be risk averse: he could consistently prefer \$21.92 to L and L to \$20 (for example, he could consider L to be just as good as \$20.50). \square

Solution to Exercise 5.4 Just like Exercise 5.2, this was a trick question! Here the basic outcomes are not sums of money but states of health. Since the described choice is not one between money lotteries, the definitions of risk aversion/neutrality/love are not applicable. \square

Solution to Exercise 5.5 Since Ben prefers B to A , he must prefer D to C .

Proof. Let U be a von Neumann-Morgenstern utility function that represents Ben's preferences.

- Let $U(\$4,000) = a$, $U(\$3,000) = b$ and $U(\$0) = c$.
- Since Ben prefers more money to less, $a > b > c$.
- Then $\mathbb{E}[U(A)] = 0.8U(\$4,000) + 0.2U(\$0) = 0.8a + 0.2c$ and $\mathbb{E}[U(B)] = U(\$3,000) = b$.
- Since Ben prefers B to A , it must be that $b > 0.8a + 0.2c$.

Let us now compare C and D : $\mathbb{E}[U(C)] = 0.2a + 0.8c$ and $\mathbb{E}[U(D)] = 0.25b + 0.75c$.

- Since $b > 0.8a + 0.2c$, $0.25b > 0.25(0.8a + 0.2c) = 0.2a + 0.05c$ and thus, adding $0.75c$ to both sides, we get that $0.25b + 0.75c > 0.2a + 0.8c$, that is, $\mathbb{E}[U(D)] > \mathbb{E}[U(C)]$, so that $D \succ C$.

Note that the proof would have been somewhat easier if we had taken the normalized utility function, so that $a = 1$ and $c = 0$. \square

Solution to Exercise 5.6 Define the function U as follows:

$$U(x) = \frac{1}{a-c}V(x) - \frac{c}{a-c} = \frac{V(x)-c}{a-c} \text{ (note that, by hypothesis, } a > c \text{ and thus } \frac{1}{a-c} > 0).$$

Then U represents the same preferences as V .

$$\text{Then } U(o_2) = \frac{V(o_2)-c}{a-c} = \frac{a-c}{a-c} = 1, U(o_1) = \frac{V(o_1)-c}{a-c} = \frac{b-c}{a-c}, \text{ and } U(o_3) = \frac{V(o_3)-c}{a-c} = \frac{c-c}{a-c} = 0.$$

Note that, since $a > b > c$, $0 < \frac{b-c}{a-c} < 1$. \square

Solution to Exercise 5.7 We can take the set of basic outcomes to be

$\{\$3000, \$2000, \$1000, \$500\}$. Suppose that there is a von Neumann-Morgenstern utility function U that represents Jennifer's preferences. We can normalize it so that $U(\$3000) = 1$ and $U(\$500) = 0$.

- Since Jennifer is indifferent between L_1 and \$2000, $U(\$2000) = \frac{5}{6}$ (since the expected utility of L_1 is $\frac{5}{6}(1) + \frac{1}{6}(0) = \frac{5}{6}$).

- Since she is indifferent between L_2 and \$1000, $U(\$1000) = \frac{2}{3}$ (since the expected utility of L_2 is $\frac{2}{3}(1) + \frac{1}{3}(0) = \frac{2}{3}$).

$$\text{Thus, } \mathbb{E}[U(L_3)] = \frac{1}{4}(1) + \frac{1}{4}\left(\frac{5}{6}\right) + \frac{1}{4}\left(\frac{2}{3}\right) + \frac{1}{4}(0) = \frac{5}{8} \text{ and } \mathbb{E}[U(L_4)] = \frac{1}{2}\left(\frac{5}{6}\right) + \frac{1}{2}\left(\frac{2}{3}\right) = \frac{3}{4}.$$

Since $\frac{3}{4} > \frac{5}{8}$, Jennifer should prefer L_4 to L_3 . Hence, she is not rational according to the theory of expected utility. \square

Solution to Exercise 5.8 Normalize her utility function so that $U(o_1) = 1$ and $U(o_4) = 0$.

Since Rachel is indifferent between $\begin{pmatrix} o_2 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} o_1 & o_4 \\ \frac{4}{5} & \frac{1}{5} \end{pmatrix}$, we have that $U(o_2) = \frac{4}{5}$.

Similarly, since she is indifferent between $\begin{pmatrix} o_3 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} o_1 & o_4 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$, $U(o_3) = \frac{1}{2}$.

Then the expected utility of $L_1 = \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ \frac{1}{8} & \frac{2}{8} & \frac{3}{8} & \frac{2}{8} \end{pmatrix}$ is $\frac{1}{8}(1) + \frac{2}{8}\left(\frac{4}{5}\right) + \frac{3}{8}\left(\frac{1}{2}\right) + \frac{2}{8}(0) = \frac{41}{80} = 0.5125$,

while the expected utility of $L_2 = \begin{pmatrix} o_1 & o_2 & o_3 \\ \frac{1}{5} & \frac{2}{5} & \frac{1}{5} \end{pmatrix}$ is $\frac{1}{5}(1) + \frac{2}{5}\left(\frac{4}{5}\right) + \frac{1}{5}\left(\frac{1}{2}\right) = \frac{39}{50} = 0.78$.

Hence, she prefers L_2 to L_1 . \square

Solution to Exercise 5.9

(a) The expected value of L_1 is $\frac{2}{10}(30) + \frac{1}{10}(28) + \frac{1}{10}(24) + \frac{2}{10}(18) + \frac{4}{10}(8) = 18$ and the expected value of L_2 is $\frac{1}{10}(30) + \frac{4}{10}(28) + \frac{5}{10}8 = 18.2$.

Hence, a risk-neutral person would prefer L_2 to L_1 .

(b) The expected utility of L_1 is $\frac{1}{5}\ln(30) + \frac{1}{10}\ln(28) + \frac{1}{10}\ln(24) + \frac{1}{5}\ln(18) + \frac{2}{5}\ln(8) = 2.741$ while the expected utility of L_2 is $\frac{1}{10}\ln(30) + \frac{2}{5}\ln(28) + \frac{1}{2}\ln(8) = 2.713$.

Thus, Paul would choose L_1 (since he prefers L_1 to L_2). \square

Solution to Exercise 5.10

- (a) To normalize U first add 10 to each value and then divide by 180. Denote the normalization of U by \bar{U} .

Then

$$\bar{U}: \frac{o_1}{180} = 0.3 \quad \frac{o_2}{180} = 1 \quad \frac{o_3}{180} = 0 \quad \frac{o_4}{180} = 0.2 \quad \frac{o_5}{180} = 0.6$$

To normalize V first subtract 5 from each value and then divide by 90. Denote the normalization of V by \bar{V} .

Then

$$\bar{V}: \frac{o_1}{90} = 0.3 \quad \frac{o_2}{90} = 1 \quad \frac{o_3}{90} = 0 \quad \frac{o_4}{90} = 0.2 \quad \frac{o_5}{90} = 0.6$$

- (b) The transformation is of the form $V(o) = aU(o) + b$. To find the values of a and b plug in two sets of values and solve the system of equations $\begin{cases} 44a + b = 32 \\ 170a + b = 95 \end{cases}$.

The solution is $a = \frac{1}{2}$, $b = 10$. Thus, $V(o) = \frac{1}{2}U(o) + 10$. \square

Solution to Exercise 5.11

- (a) Ann prefers L_3 to L_4 ($L_3 \succ_{Ann} L_4$). In fact, $\mathbb{E}[U_{Ann}(L_3)] = \sqrt{28} = 5.2915$ while $\mathbb{E}[U_{Ann}(L_4)] = \frac{1}{2}\sqrt{10} + \frac{1}{2}\sqrt{50} = 5.1167$.

- (b) Bob prefers L_4 to L_3 ($L_4 \succ_{Bob} L_3$). In fact, $\mathbb{E}[U_{Bob}(L_3)] = 2(28) - \frac{28^4}{100^3} = 55.3853$ while $\mathbb{E}[U_{Bob}(L_4)] = \frac{1}{2}\left[2(10) - \frac{10^4}{100^3}\right] + \frac{1}{2}\left[2(50) - \frac{50^4}{100^3}\right] = 56.87$.

- (c) The expected value of lottery L_4 is $\frac{1}{2}10 + \frac{1}{2}50 = 30$; thus, a risk-averse person would strictly prefer \$30 with certainty to the lottery L_4 . We saw in part (a) that for Ann the expected utility of lottery L_4 is 5.1167; the utility of \$30 is $\sqrt{30} = 5.4772$. Thus, Ann would indeed choose \$30 for certain over the lottery L_4 . We saw in part (b) that for Bob the expected utility of lottery L_4 is 56.87; the utility of \$30 is $2(30) - \frac{30^4}{100^3} = 59.19$. Thus, Bob would indeed choose \$30 for certain over the lottery L_4 . \square

Solution to Exercise 5.12 The simple lottery is $\left(\begin{array}{cccc} o_1 & o_2 & o_3 & o_4 \\ \frac{18}{240} & \frac{103}{240} & \frac{95}{240} & \frac{24}{240} \end{array} \right)$. For example,

the probability of o_2 is computed as follows: $\frac{1}{8}\left(\frac{1}{10}\right) + \frac{1}{4}(1) + \frac{1}{8}(0) + \frac{1}{2}\left(\frac{1}{3}\right) = \frac{103}{240}$. \square

Solution to Exercise 5.13 Using the stated indifference, use lottery L to construct the

compound lottery $\left(\left(\begin{array}{cc} o_2 & o_4 \\ \frac{1}{4} & \frac{3}{4} \end{array} \right) \left(\begin{array}{cc} o_3 & o_4 \\ \frac{3}{5} & \frac{2}{5} \end{array} \right) o_3 \quad o_4 \right)$, whose corresponding simple lottery is $L' = \left(\begin{array}{cccc} o_1 & o_2 & o_3 & o_4 \\ 0 & \frac{1}{12} & \frac{11}{45} & \frac{121}{180} \end{array} \right)$. Then, by the Independence Axiom, $L \sim L'$. \square

Solution to Exercise 5.14 Let W be your initial wealth. The basic outcomes are:

1. you do not pay any money, do not die and live to enjoy your wealth W (denote this outcome by A_0),
2. you pay $\$Y$, do not die and live to enjoy your remaining wealth $W - Y$ (call this outcome A_Y),
3. you pay $\$X$, do not die and live to enjoy your remaining wealth $W - X$ (call this outcome A_X),
4. you die (call this outcome D); this could happen because (a) you do not pay any money, roll the die and drink the poison or (b) you pay $\$Y$, roll the die and drink the poison; we assume that you are indifferent between these two outcomes.

Since, by hypothesis, $X > Y$, your ranking of these outcomes must be $A_0 \succ A_Y \succ A_X \succ D$. If you satisfy the von Neumann-Morgenstern axioms, then your preferences can be represented by a von Neumann-Morgenstern utility function U defined on the set of basic outcomes. We can normalize your utility function by setting $U(A_0) = 1$ and $U(D) = 0$. Furthermore, it must be that

$$U(A_Y) > U(A_X). \quad (5.6)$$

The maximum amount $\$P$ that you are willing to pay is that amount that makes you indifferent between (1) rolling the die with the initial number of poisoned glasses and (2) giving up $\$P$ and rolling the die with one less poisoned glass.

Thus – based on your answers – you are indifferent between the two lotteries

$$\left(\begin{array}{cc} D & A_0 \\ \frac{1}{6} & \frac{5}{6} \end{array} \right) \text{ and } \left(\begin{array}{c} A_X \\ 1 \end{array} \right)$$

and you are indifferent between the two lotteries:

$$\left(\begin{array}{cc} D & A_0 \\ \frac{4}{6} & \frac{2}{6} \end{array} \right) \text{ and } \left(\begin{array}{cc} D & A_Y \\ \frac{3}{6} & \frac{3}{6} \end{array} \right).$$

Thus,

$$\underbrace{\frac{1}{6}U(D) + \frac{5}{6}U(A_0)}_{=\frac{1}{6}0 + \frac{5}{6}1 = \frac{5}{6}} = U(A_X) \quad \text{and} \quad \underbrace{\frac{4}{6}U(D) + \frac{2}{6}U(A_0)}_{=\frac{4}{6}0 + \frac{2}{6}1 = \frac{2}{6}} = \underbrace{\frac{3}{6}U(D) + \frac{3}{6}U(A_Y)}_{=\frac{3}{6}0 + \frac{3}{6}U(A_Y)}.$$

Hence, $U(A_X) = \frac{5}{6}$ and $U(A_Y) = \frac{2}{3} = \frac{4}{6}$, so that $U(A_X) > U(A_Y)$, contradicting (5.6). \square

6. Strategic-form Games

6.1 Strategic-form games with cardinal payoffs

At the end of Chapter 4 we discussed the possibility of incorporating random events in extensive-form games by means of chance moves. The introduction of chance moves gives rise to probabilistic outcomes and thus to the issue of how a player might rank such outcomes. Random events can also occur in strategic-form games, as shown in Figure 6.1, which represents the simple first-price auction of Example 6.1 below.

		Player 2		
		bid \$100	bid \$200	
P l a y e r 1	bid \$100	$\begin{pmatrix} o_1 & o_2 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$	o_3	o_1 : Player 1 wins and pays \$100
	bid \$200	o_4	$\begin{pmatrix} o_3 & o_4 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$	o_2 : Player 2 wins and pays \$100 o_3 : Player 2 wins and pays \$200 o_4 : Player 1 wins and pays \$200

Figure 6.1: A game-frame in strategic form representing Example 6.1

■ **Example 6.1** Two players simultaneously submit a bid for a painting. Only two bids are possible: \$100 and \$200. If one player bids \$200 and the other \$100 then the high bidder wins the painting and has to pay her own bid. If the two players bid the same amount then a fair coin is tossed and if the outcome is Heads the winner is Player 1 (who then has to pay her own bid) while if the outcome is Tails the winner is Player 2 (who then has to pay his own bid). ■

Suppose that Player 1 ranks the basic outcomes as follows: $o_1 \succ_1 o_4 \succ_1 o_2 \sim_1 o_3$, that is, she prefers winning to not winning; conditional on winning, she prefers to pay less and, conditional on not winning, she is indifferent as to how much Player 2 pays. Suppose also that Player 1 believes that Player 2 is going to submit a bid of \$100 (perhaps she has been informed of this by somebody spying on Player 2). What should we expect Player 1 to do? Knowing her ranking of the basic outcomes is of no help, because we need to know how she ranks the probabilistic outcome $\left(\begin{array}{cc} o_1 & o_2 \\ \frac{1}{2} & \frac{1}{2} \end{array} \right)$ relative to the basic outcome o_4 .

The theory of expected utility introduced in Chapter 5 provides one possible answer to the question of how players rank probabilistic outcomes. With the aid of expected utility theory we can now generalize the definition of strategic-form game. First we generalize the notion of game-frame in strategic form (Definition 2.1.1, Chapter 2) by allowing probabilistic outcomes, or lotteries, to be associated with strategy profiles. In the following definition, the bulleted items coincide with the first three items of Definition 2.1.1 (Chapter 2); the modified item is the last one, preceded by the symbol \star .

Definition 6.1.1 A *game-frame in strategic form* is a quadruple $\langle I, (S_i)_{i \in I}, O, f \rangle$ where:

- $I = \{1, \dots, n\}$ is a set of *players* ($n \geq 2$).
- For every Player $i \in I$, S_i is the set of *strategies* (or choices) of Player i . As before, we denote by $S = S_1 \times \dots \times S_n$ the set of *strategy profiles*.
- O is a set of *basic outcomes*.
- \star $f : S \rightarrow \mathcal{L}(O)$ is a function that associates with every strategy profile s a lottery over the set of basic outcomes O (as in Chapter 5, we denote by $\mathcal{L}(O)$ the set of lotteries, or probability distributions, over O).

If, for every $s \in S$, $f(s)$ is a degenerate lottery (that is, a basic outcome) then we are back to Definition 2.1.1 (Chapter 2).

From a game-frame one obtains a game by adding, for every player $i \in I$, a von Neumann-Morgenstern ranking \succsim_i of the elements of $\mathcal{L}(O)$. It is more convenient to represent such a ranking by means of a von Neumann-Morgenstern utility function $U_i : O \rightarrow \mathbb{R}$. We denote by $\mathbb{E}[U_i(f(s))]$ the expected utility of lottery $f(s) \in \mathcal{L}(O)$ for Player i . The following definition mirrors Definition 2.1.2 of Chapter 2.

Definition 6.1.2 A *game in strategic form with cardinal payoffs* is a quintuple $\langle I, (S_i)_{i \in I}, O, f, (\succsim_i)_{i \in I} \rangle$ where:

- $\langle I, (S_i)_{i \in I}, O, f \rangle$ is a game-frame in strategic form (Definition 6.1.1) and
- for every Player $i \in I$, \succsim_i is a von Neumann-Morgenstern ranking of the set of lotteries $\mathcal{L}(O)$.

If we represent each ranking \succsim_i by means of a von Neumann-Morgenstern utility function U_i and define $\pi_i : S \rightarrow \mathbb{R}$ by $\pi_i(s) = \mathbb{E}[U_i(f(s))]$, then $\langle I, (S_1, \dots, S_n), (\pi_1, \dots, \pi_n) \rangle$ is called a *reduced-form game in strategic form with cardinal payoffs* ('reduced-form' because some information is lost, namely the specification of the possible outcomes). The function $\pi_i : S \rightarrow \mathbb{R}$ is called the *von Neumann-Morgenstern payoff function of Player i* .

For example, consider the first-price auction of Example 6.1 whose game-frame in strategic form was shown in Figure 6.1. Let $O = \{o_1, o_2, o_3, o_4\}$ and suppose that Player 1 has a von Neumann-Morgenstern ranking of $\mathcal{L}(O)$ that is represented by the following von Neumann-Morgenstern utility function U_1 (note that the implied ordinal ranking of the basic outcomes is indeed $o_1 \succ_1 o_4 \succ_1 o_2 \sim_1 o_3$):

$$\begin{array}{l} \text{outcome: } o_1 \quad o_2 \quad o_3 \quad o_4 \\ U_1: \quad 4 \quad 1 \quad 1 \quad 2 \end{array}$$

Then, for Player 1, the expected utility of lottery $\left(\begin{array}{cc} o_1 & o_2 \\ \frac{1}{2} & \frac{1}{2} \end{array} \right)$ is 2.5 and the expected utility of lottery $\left(\begin{array}{cc} o_3 & o_4 \\ \frac{1}{2} & \frac{1}{2} \end{array} \right)$ is 1.5.

Suppose also that Player 2 has (somewhat spiteful) preferences represented by the following von Neumann-Morgenstern utility function U_2 :

$$\begin{array}{l} \text{outcome: } o_1 \quad o_2 \quad o_3 \quad o_4 \\ U_2: \quad 1 \quad 6 \quad 4 \quad 5 \end{array}$$

Thus, for Player 2, the expected utility of lottery $\left(\begin{array}{cc} o_1 & o_2 \\ \frac{1}{2} & \frac{1}{2} \end{array} \right)$ is 3.5 and the expected utility of lottery $\left(\begin{array}{cc} o_3 & o_4 \\ \frac{1}{2} & \frac{1}{2} \end{array} \right)$ is 4.5. Then we can represent the game in reduced form as shown in Figure 6.2.

		Player 2			
		\$100	\$200		
Player 1	\$100	2.5	3.5	1	4
	\$200	2	5	1.5	4.5

Figure 6.2: A cardinal game in reduced form based on the game-frame of 6.1

The game of Figure 6.2 does not have any Nash equilibria. However, we will show in the next section that if we extend the notion of strategy, by allowing players to choose randomly, then the game of Figure 6.2 does have a Nash equilibrium.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 6.5.1 at the end of this chapter.

6.2 Mixed strategies

Definition 6.2.1 Consider a game in strategic form with cardinal payoffs and recall that S_i denotes the set of strategies of Player i . From now on, we shall call S_i the set of *pure strategies* of Player i . We assume that S_i is a finite set (for every $i \in I$). A *mixed strategy* of Player i is a probability distribution over the set of pure strategies S_i . The set of mixed strategies of Player i is denoted by Σ_i .

R Since among the mixed strategies of Player i there are the degenerate strategies that assign probability 1 to a pure strategy, the set of mixed strategies includes the set of pure strategies (viewed as degenerate probability distributions).

For example, one possible mixed strategy for Player 1 in the game of Figure 6.2 is $\begin{pmatrix} \$100 & \$200 \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$. The traditional interpretation of a mixed strategy is in terms of objective randomization: the player, instead of choosing a pure strategy herself, delegates the choice to a random device.¹ For example, Player 1 choosing the mixed strategy $\begin{pmatrix} \$100 & \$200 \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$ is interpreted as a decision to let, say, a die determine whether she will bid \$100 or \$200: Player 1 will roll a die and if the outcome is 1 or 2 then she will bid \$100, while if the outcome is 3, 4, 5 or 6 then she will bid \$200. Suppose that Player 1 chooses this mixed strategy and Player 2 chooses the mixed strategy $\begin{pmatrix} \$100 & \$200 \\ \frac{3}{5} & \frac{2}{5} \end{pmatrix}$. Since the players rely on independent random devices, this pair of mixed strategies gives rise to the following probabilistic outcome:

$$\begin{pmatrix} \text{strategy profile} & (\$100, \$100) & (\$100, \$200) & (\$200, \$100) & (\$200, \$200) \\ \text{outcome} & \begin{pmatrix} o_1 & o_2 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} & o_3 & o_4 & \begin{pmatrix} o_3 & o_4 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} \\ \text{probability} & \frac{1}{3} \left(\frac{3}{5}\right) = \frac{3}{15} & \frac{1}{3} \left(\frac{2}{5}\right) = \frac{2}{15} & \frac{2}{3} \left(\frac{3}{5}\right) = \frac{6}{15} & \frac{2}{3} \left(\frac{2}{5}\right) = \frac{4}{15} \end{pmatrix}$$

If the two players have von Neumann-Morgenstern preferences, then – by the Compound Lottery Axiom (Chapter 5) – they will view the above as the following lottery:

$$\begin{pmatrix} \text{outcome} & o_1 & o_2 & o_3 & o_4 \\ \text{probability} & \frac{3}{30} & \frac{3}{30} & \frac{8}{30} & \frac{16}{30} \end{pmatrix}.$$

¹An alternative interpretation of mixed strategies in terms of beliefs will be discussed in Chapter 10 and in Part V (Chapters 14-16).

Using the von Neumann-Morgenstern utility functions postulated in the previous section, namely

$$\begin{array}{cccc} \text{outcome : } & o_1 & o_2 & o_3 & o_4 \\ U_1 : & 4 & 1 & 1 & 2 \end{array} \quad \text{and} \quad \begin{array}{cccc} \text{outcome : } & o_1 & o_2 & o_3 & o_4 \\ U_2 : & 1 & 6 & 4 & 5 \end{array}$$

the lottery $\left(\begin{array}{cccc} o_1 & o_2 & o_3 & o_4 \\ \frac{3}{30} & \frac{3}{30} & \frac{8}{30} & \frac{16}{30} \end{array} \right)$ has an expected utility of

$$\text{For Player 1: } \frac{3}{30}(4) + \frac{3}{30}(1) + \frac{8}{30}(1) + \frac{16}{30}(2) = \frac{55}{30}$$

$$\text{For Player 2: } \frac{3}{30}(1) + \frac{3}{30}(6) + \frac{8}{30}(4) + \frac{16}{30}(5) = \frac{133}{30}.$$

Thus we can define the payoffs of the two players from this mixed strategy profile by

$$\Pi_1 \left[\left(\begin{array}{cc} \$100 & \$200 \\ \frac{1}{3} & \frac{2}{3} \end{array} \right), \left(\begin{array}{cc} \$100 & \$200 \\ \frac{3}{5} & \frac{2}{5} \end{array} \right) \right] = \frac{55}{30}$$

$$\Pi_2 \left[\left(\begin{array}{cc} \$100 & \$200 \\ \frac{1}{3} & \frac{2}{3} \end{array} \right), \left(\begin{array}{cc} \$100 & \$200 \\ \frac{3}{5} & \frac{2}{5} \end{array} \right) \right] = \frac{133}{30}$$

Note that we can calculate these payoffs in a different – but equivalent – way by using the reduced-form game of Figure 6.2, as follows:

$$\left(\begin{array}{ccccc} \text{strategy profile} & (\$100, \$100) & (\$100, \$200) & (\$200, \$100) & (\$200, \$200) \\ \text{expected utilities} & (2.5, 3.5) & (1, 4) & (2, 5) & (1.5, 4.5) \\ \text{probability} & \frac{1}{3} \left(\frac{3}{5} \right) = \frac{3}{15} & \frac{1}{3} \left(\frac{2}{5} \right) = \frac{2}{15} & \frac{2}{3} \left(\frac{3}{5} \right) = \frac{6}{15} & \frac{2}{3} \left(\frac{2}{5} \right) = \frac{4}{15} \end{array} \right)$$

so that the expected payoff of Player 1 is

$$\frac{3}{15}(2.5) + \frac{2}{15}(1) + \frac{6}{15}(2) + \frac{4}{15}(1.5) = \frac{55}{30}$$

and the expected payoff of Player 2 is

$$\frac{3}{15}(3.5) + \frac{2}{15}(4) + \frac{6}{15}(5) + \frac{4}{15}(4.5) = \frac{133}{30}.$$

The previous example provides the rationale for the following definition. First some notation.

- Let $\sigma_i \in \Sigma_i$ be a mixed strategy of Player i ; then, for every pure strategy $s_i \in S_i$ of Player i , we denote by $\sigma_i(s_i)$ the probability that σ_i assigns to s_i .²
- Let Σ be the set of mixed-strategy profiles, that is, $\Sigma = \Sigma_1 \times \cdots \times \Sigma_n$.
- Consider a mixed-strategy profile $\sigma = (\sigma_1, \dots, \sigma_n) \in \Sigma$ and a pure-strategy profile $s = (s_1, \dots, s_n) \in S$; then we denote by $\sigma(s)$ the product of the probabilities $\sigma_i(s_i)$, that is, $\sigma(s) = \prod_{i=1}^n \sigma_i(s_i) = \sigma_1(s_1) \times \dots \times \sigma_n(s_n)$.³

Definition 6.2.2 Consider a reduced-form game in strategic form with cardinal payoffs $G = \langle I, (S_1, \dots, S_n), (\pi_1, \dots, \pi_n) \rangle$ (Definition 6.1.2), where, for every Player $i \in I$, the set of pure strategies S_i is finite. Then the *mixed-strategy extension of G* is the reduced-form game in strategic form $\langle I, (\Sigma_1, \dots, \Sigma_n), (\Pi_1, \dots, \Pi_n) \rangle$ where, for every Player $i \in I$,

- Σ_i is the set of mixed strategies of Player i in G (that is, Σ_i is the set of probability distributions over S_i).
- The payoff function $\Pi_i : \Sigma \rightarrow \mathbb{R}$ is defined by $\Pi_i(\sigma) = \sum_{s \in S} \sigma(s) \pi_i(s)$.^a

^aIn the above example, if $\sigma_1 = \begin{pmatrix} \$100 & \$200 \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$ and $\sigma_2 = \begin{pmatrix} \$100 & \$200 \\ \frac{3}{5} & \frac{2}{5} \end{pmatrix}$

$$\text{then } \Pi_1(\sigma_1, \sigma_2) = \frac{3}{15}(2.5) + \frac{2}{15}(1) + \frac{6}{15}(2) + \frac{4}{15}(1.5) = \frac{55}{30}.$$

Definition 6.2.3 Fix a reduced-form game in strategic form with cardinal payoffs $G = \langle I, (S_1, \dots, S_n), (\pi_1, \dots, \pi_n) \rangle$ (Definition 6.1.2), where, for every player $i \in I$, the set of pure strategies S_i is finite. A *Nash equilibrium in mixed-strategies of G* is a Nash equilibrium of the mixed-strategy extension of G .

²In the above example, if $\sigma_1 = \begin{pmatrix} \$100 & \$200 \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$ then $\sigma_1(\$200) = \frac{2}{3}$.

³In the above example, if $\sigma = (\sigma_1, \sigma_2)$ with $\sigma_1 = \begin{pmatrix} \$100 & \$200 \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$ and $\sigma_2 = \begin{pmatrix} \$100 & \$200 \\ \frac{3}{5} & \frac{2}{5} \end{pmatrix}$ then $\sigma_1(\$200) = \frac{2}{3}$, $\sigma_2(\$100) = \frac{3}{5}$ and thus $\sigma((\$200, \$100)) = \frac{2}{3} \left(\frac{3}{5} \right) = \frac{6}{15}$.

For example, consider the reduced-form game of Figure 6.3 (which reproduces Figure 6.2: with all the payoffs multiplied by 10; this corresponds to representing the preferences of the players with different utility functions that are obtained from the ones used above by multiplying them by 10).

Is $\sigma = (\sigma_1, \sigma_2)$ with $\sigma_1 = \begin{pmatrix} \$100 & \$200 \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$ and $\sigma_2 = \begin{pmatrix} \$100 & \$200 \\ \frac{3}{5} & \frac{2}{5} \end{pmatrix}$ a mixed-strategy Nash equilibrium of this game?

		Player 2			
		\$100	\$200		
Player 1	\$100	25	35	10	40
	\$200	20	50	15	45

Figure 6.3: The game of Figure 6.2 with the payoffs multiplied by 10

The payoff of Player 1 is

$$\Pi_1(\sigma_1, \sigma_2) = \frac{3}{15}(25) + \frac{2}{15}(10) + \frac{6}{15}(20) + \frac{4}{15}(15) = \frac{55}{3}.$$

If Player 1 switched from $\sigma_1 = \begin{pmatrix} \$100 & \$200 \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$ to $\hat{\sigma}_1 = \begin{pmatrix} \$100 & \$200 \\ 1 & 0 \end{pmatrix}$, that is, to the pure strategy \$100, then Player 1's payoff would be larger:

$$\Pi_1(\hat{\sigma}_1, \sigma_2) = \frac{3}{5}(25) + \frac{2}{5}(10) = 19.$$

Thus, since $19 > \frac{55}{3}$, it is not a Nash equilibrium.

John Nash (who shared the 1994 Nobel Memorial prize in economics with John Harsanyi and Reinhard Selten), proved the following theorem.

Theorem 6.2.1 — Nash, 1951. Every reduced-form game in strategic form with cardinal payoffs $\langle I, (S_1, \dots, S_n), (\pi_1, \dots, \pi_n) \rangle$ (Definition 6.1.2), where, for every Player $i \in I$, the set of pure strategies S_i is finite, has at least one Nash equilibrium in mixed-strategies.

We will not give the proof of this theorem, since it is rather complex (it requires the use of fixed-point theorems).

		Player 2	
		\$100	\$200
Player 1	\$100	25 35	10 40
	\$200	20 50	15 45

Going back to the game of Figure 6.3 reproduced above, let us verify that, on the other hand, $\sigma^* = (\sigma_1^*, \sigma_2^*)$ with $\sigma_1^* = \sigma_2^* = \begin{pmatrix} \$100 & \$200 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$ is a Nash equilibrium in mixed strategies. The payoff of Player 1 is

$$\Pi_1(\sigma_1^*, \sigma_2^*) = \frac{1}{4}(25) + \frac{1}{4}(10) + \frac{1}{4}(20) + \frac{1}{4}(15) = \frac{70}{4} = 17.5.$$

Could Player 1 obtain a larger payoff with some other mixed strategy $\sigma_1 = \begin{pmatrix} \$100 & \$200 \\ p & 1-p \end{pmatrix}$ for some $p \neq \frac{1}{2}$?

Fix an arbitrary $p \in [0, 1]$ and let us compute Player 1's payoff if she uses the strategy $\sigma_1 = \begin{pmatrix} \$100 & \$200 \\ p & 1-p \end{pmatrix}$ against Player 2's mixed strategy $\sigma_2^* = \begin{pmatrix} \$100 & \$200 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$:

$$\begin{aligned} \Pi_1 \left[\begin{pmatrix} \$100 & \$200 \\ p & 1-p \end{pmatrix}, \begin{pmatrix} \$100 & \$200 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} \right] &= \frac{1}{2}p25 + \frac{1}{2}p10 + \frac{1}{2}(1-p)20 + \frac{1}{2}(1-p)15 \\ &= p \left(\frac{1}{2}25 + \frac{1}{2}10 \right) + (1-p) \left(\frac{1}{2}20 + \frac{1}{2}15 \right) = \frac{35}{2} = 17.5. \end{aligned}$$

Thus if Player 2 uses the mixed strategy $\sigma_2^* = \begin{pmatrix} \$100 & \$200 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$, then *Player 1 gets the same payoff no matter what mixed strategy she employs.*

It follows that any mixed strategy of Player 1 is a best reply to $\sigma_2^* = \begin{pmatrix} \$100 & \$200 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$; in particular, $\sigma_1^* = \begin{pmatrix} \$100 & \$200 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$ is a best reply to $\sigma_2^* = \begin{pmatrix} \$100 & \$200 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$.

It is easy to verify that the same applies to Player 2: any mixed strategy of Player 2 is a best reply to Player 1's mixed strategy $\sigma_1^* = \begin{pmatrix} \$100 & \$200 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$. Hence $\sigma^* = (\sigma_1^*, \sigma_2^*)$ is indeed a Nash equilibrium in mixed strategies.

We will see in the next section that this "indifference" phenomenon is true in general.

- R** Since, among the mixed strategies of Player i there are the degenerate strategies that assign probability 1 to a pure strategy, every Nash equilibrium in pure strategies is also a Nash equilibrium in mixed strategies. That is, the set of mixed-strategy Nash equilibria includes the set of pure-strategy Nash equilibria.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 6.5.2 at the end of this chapter.

6.3 Computing the mixed-strategy Nash equilibria

How can we find the mixed-strategy equilibria of a given game? The first important observation is that if a pure strategy is strictly dominated by another pure strategy then it cannot be played with positive probability at a Nash equilibrium. Thus, for the purpose of finding Nash equilibria, one can delete all the strictly dominated strategies and focus on the resulting game. But then the same reasoning applies to the resulting game and one can delete all the strictly dominated strategies in that game, and so on. Thus we have the following observation.

- R** In order to find the mixed-strategy Nash equilibria of a game one can first apply the iterated deletion of strictly dominated strategies (IDSDS: Chapter 2) and then find the Nash equilibria of the resulting game (which can then be viewed as Nash equilibria of the original game where all the pure strategies that were deleted are assigned zero probability). Note, however, that – as we will see in Section 6.4 – one can perform more deletions than allowed by the IDSDS procedure.

For example, consider the game of Figure 6.4.

		Player 2					
		E		F		G	
Player 1	A	2	4	3	3	6	0
	B	4	0	2	4	4	2
	C	3	3	4	2	3	1
	D	3	6	1	1	2	6

Figure 6.4: A reduced-form game with cardinal payoffs

In this game there are no pure-strategy Nash equilibria; however, by Nash's theorem there will be at least one mixed-strategy equilibrium. To find it we can first note that, for Player 1, D is strictly dominated by B ; deleting D we get a smaller game where, for Player 2, G is strictly dominated by F . Deleting G we are left with a smaller game where A is strictly dominated by C . Deleting A we are left with the game shown in Figure 6.5.

		Player 2			
		<i>E</i>	<i>F</i>		
Player 1	<i>B</i>	4	0	2	4
	<i>C</i>	3	3	4	2

Figure 6.5: The result of applying the IDSDS procedure to the game of Figure 6.4

We will see that the game of Figure 6.5 has a unique Nash equilibrium in mixed strategies

given by $\left[\left(\begin{matrix} B & C \\ \frac{1}{5} & \frac{4}{5} \end{matrix} \right), \left(\begin{matrix} E & F \\ \frac{2}{3} & \frac{1}{3} \end{matrix} \right) \right]$.

Thus the game of Figure 6.4 has a unique Nash equilibrium in mixed strategies given by

$$\left[\left(\begin{matrix} A & B & C & D \\ 0 & \frac{1}{5} & \frac{4}{5} & 0 \end{matrix} \right), \left(\begin{matrix} E & F & G \\ \frac{2}{3} & \frac{1}{3} & 0 \end{matrix} \right) \right].$$

Once we have simplified the game by applying the IDSDS procedure, in order to find the mixed-strategy Nash equilibria we can use the following result.

First we recall some notation that was introduced in Chapter 2. Given a mixed-strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ and a Player i , we denote by σ_{-i} the profile of strategies of the players other than i and use (σ_i, σ_{-i}) as an alternative notation for σ ; furthermore, (τ_i, σ_{-i}) denotes the result of replacing σ_i with τ_i in σ , that is, $(\tau_i, \sigma_{-i}) = (\sigma_1, \dots, \sigma_{i-1}, \tau_i, \sigma_{i+1}, \dots, \sigma_n)$.

Theorem 6.3.1 Consider a reduced-form game in strategic form with cardinal payoffs.

- Suppose that $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*)$ is a Nash equilibrium in mixed strategies.
- Consider an arbitrary Player i .
- Let $\pi_i^* = \Pi_i(\sigma^*)$ be the payoff of Player i at this Nash equilibrium and let $s_{ij}, s_{ik} \in S_i$ be two pure strategies of Player i such that $\sigma_i^*(s_{ij}) > 0$ and $\sigma_i^*(s_{ik}) > 0$, that is, s_{ij} and s_{ik} are two pure strategies to which the mixed strategy σ_i^* of Player i assigns positive probability.
- Then $\Pi_i(s_{ij}, \sigma_{-i}^*) = \Pi_i(s_{ik}, \sigma_{-i}^*) = \pi_i^*$.

In other words, when the other players use the mixed-strategy profile σ_{-i}^* , Player i gets the same payoff no matter whether she plays the mixed strategy σ_i^* or the pure strategy s_{ij} or the pure strategy s_{ik} .

The details of the proof of Theorem 6.3.1 will be omitted, but the idea is simple: if s_{ij} and s_{ik} are two pure strategies to which the mixed strategy σ_i^* of Player i assigns positive probability and $\Pi_i((s_{ij}, \sigma_{-i}^*)) > \Pi_i((s_{ik}, \sigma_{-i}^*))$, then Player i can increase her payoff from $\pi_i^* = \Pi_i(\sigma^*)$ to a larger number by reducing the probability of s_{ik} to zero and adding that probability to $\sigma_i^*(s_{ij})$, that is, by switching from σ_i^* to the mixed strategy $\hat{\sigma}_i$ obtained as follows: $\hat{\sigma}_i(s_{ik}) = 0$, $\hat{\sigma}_i(s_{ij}) = \sigma_i^*(s_{ij}) + \sigma_i^*(s_{ik})$ and, for every other $s_i \in S_i$, $\hat{\sigma}_i(s_i) = \sigma_i^*(s_i)$. But this would contradict the hypothesis that $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*)$ is a Nash equilibrium.

Let us now go back to the game of Figure 6.5, which is reproduced in Figure 6.5, and see how we can use Theorem 6.3.1 to find the Nash equilibrium in mixed strategies.

		Player 2			
		E	F		
Player 1	B	4	0	2	4
	C	3	3	4	2

Figure 6.6: Copy of Figure 6.5

We want to find values of p and q , strictly between 0 and 1, such that

$$\left[\left(\begin{array}{cc} B & C \\ p & 1-p \end{array} \right), \left(\begin{array}{cc} E & F \\ q & 1-q \end{array} \right) \right]$$

is a Nash equilibrium.

By Theorem 6.3.1, if Player 1 played the pure strategy B against $\left(\begin{array}{cc} E & F \\ q & 1-q \end{array} \right)$ she should get the same payoff as if she were to play the pure strategy C .

The former would give her a payoff of $4q + 2(1 - q)$ and the latter a payoff of $3q + 4(1 - q)$.

Thus we need q to be such that $4q + 2(1 - q) = 3q + 4(1 - q)$, that is, $q = \frac{2}{3}$.

When $q = \frac{2}{3}$, both B and C give Player 1 a payoff of $\frac{10}{3}$ and thus any mixture of B and C would also give the same payoff of $\frac{10}{3}$.

In other words, Player 1 is indifferent among all her mixed strategies and thus any mixed strategy is a best response to $\left(\begin{array}{cc} E & F \\ \frac{2}{3} & \frac{1}{3} \end{array} \right)$.

Similar reasoning for Player 2 reveals that, by Theorem 6.3.1, we need p to be such that $0p + 3(1 - p) = 4p + 2(1 - p)$, that is, $p = \frac{1}{5}$.

Against $\left(\begin{array}{cc} B & C \\ \frac{1}{5} & \frac{4}{5} \end{array} \right)$ any mixed strategy of Player 2 gives him the same payoff of $\frac{12}{5}$; thus

any mixed strategy of Player 2 is a best reply to $\left(\begin{array}{cc} B & C \\ \frac{1}{5} & \frac{4}{5} \end{array} \right)$.

It follows that $\left[\left(\begin{array}{cc} B & C \\ \frac{1}{5} & \frac{4}{5} \end{array} \right), \left(\begin{array}{cc} E & F \\ \frac{2}{3} & \frac{1}{3} \end{array} \right) \right]$ is a Nash equilibrium.

- R** It follows from Theorem 6.3.1, and was illustrated in the above example, that at a mixed strategy Nash equilibrium where Player i plays two or more pure strategies with positive probability, Player i does not have an incentive to use that mixed strategy: she would get the same payoff if, instead of randomizing, she played one of the pure strategies in the support of her mixed strategy (that is, if she increased the probability of any pure strategy from a positive number to 1).⁴ *The only purpose of randomizing is to make the other player indifferent among two or more of his own pure strategies.*

		Player 2			
		D	E		
Player 1	A	3	0	0	2
	B	0	2	3	0
	C	2	0	2	1

Figure 6.7: A reduced-form game with cardinal payoffs

The “indifference” condition explained above provides a necessary, *but not sufficient*, condition for a mixed-strategy profile to be a Nash equilibrium. To see that the condition is not sufficient, consider the game of Figure 6.7 and the mixed-strategy profile $\left[\left(\begin{matrix} A & B & C \\ \frac{1}{2} & \frac{1}{2} & 0 \end{matrix} \right), \left(\begin{matrix} D & E \\ \frac{1}{2} & \frac{1}{2} \end{matrix} \right) \right]$. Given that Player 2 plays the mixed strategy $\left(\begin{matrix} D & E \\ \frac{1}{2} & \frac{1}{2} \end{matrix} \right)$, Player 1 is indifferent between the two pure strategies that are in the support of her own mixed strategy, namely A and B : the payoff from playing A is 1.5 and so is the payoff from playing B (and 1.5 is also the payoff associated with the mixed strategy under consideration). However, the profile $\left[\left(\begin{matrix} A & B & C \\ \frac{1}{2} & \frac{1}{2} & 0 \end{matrix} \right), \left(\begin{matrix} D & E \\ \frac{1}{2} & \frac{1}{2} \end{matrix} \right) \right]$ is not a Nash equilibrium, because Player 1 could get a payoff of 2 by switching to the pure strategy C .

We know from Theorem 6.2.1 that this game does have a mixed-strategy Nash equilibrium. How can we find it? Let us calculate the best response of Player 1 to every possible mixed strategy $\left(\begin{matrix} D & E \\ q & 1-q \end{matrix} \right)$ of Player 2 (with $q \in [0, 1]$).

For Player 1 the payoff from playing A against $\left(\begin{matrix} D & E \\ q & 1-q \end{matrix} \right)$ is $3q$, the payoff from playing B is $3 - 3q$ and the payoff from playing C is constant and equal to 2. These functions are shown in Figure 6.8.

⁴The support of a mixed strategy is the set of pure strategies that are assigned positive probability by that mixed strategy.




Figure 6.8: Player 1's payoff from each pure strategy against an arbitrary mixed strategy of Player 2

The upward-sloping line plots the function $A(q) = 3q$, the downward-sloping line plots the function $B(q) = 3 - 3q$ and the horizontal dashed line the function $C(q) = 2$.

The downward-sloping and horizontal lines intersect when $q = \frac{1}{3}$ and the upward-sloping and horizontal lines intersect when $q = \frac{2}{3}$.

The maximum payoff is given by the downward-sloping line up to $q = \frac{1}{3}$, then by the horizontal line up to $q = \frac{2}{3}$ and then by the upward-sloping line.

Thus the best reply function of Player 1 is as follows:

$$\text{Player 1's best reply} = \begin{cases} B & \text{if } 0 \leq q < \frac{1}{3} \\ \left(\begin{array}{cc} B & C \\ p & 1-p \end{array} \right) \text{ for any } p \in [0, 1] & \text{if } q = \frac{1}{3} \\ C & \text{if } \frac{1}{3} < q < \frac{2}{3} \\ \left(\begin{array}{cc} A & C \\ p & 1-p \end{array} \right) \text{ for any } p \in [0, 1] & \text{if } q = \frac{2}{3} \\ A & \text{if } \frac{2}{3} < q \leq 1 \end{cases}$$

Hence if there is a mixed-strategy Nash equilibrium it is either of the form

$$\left[\left(\begin{array}{ccc} A & B & C \\ 0 & p & 1-p \end{array} \right), \left(\begin{array}{cc} D & E \\ \frac{1}{3} & \frac{2}{3} \end{array} \right) \right] \text{ or of the form } \left[\left(\begin{array}{ccc} A & B & C \\ p & 0 & 1-p \end{array} \right), \left(\begin{array}{cc} D & E \\ \frac{2}{3} & \frac{1}{3} \end{array} \right) \right].$$

The latter cannot be a Nash equilibrium for any p , because when Player 1 plays B with probability 0, E strictly dominates D for Player 2 and thus Player 2's mixed strategy is not a best reply (E is the unique best reply). Thus the only candidate for a Nash equilibrium is of the form

$$\left[\left(\begin{array}{ccc} A & B & C \\ 0 & p & 1-p \end{array} \right), \left(\begin{array}{cc} D & E \\ \frac{1}{3} & \frac{2}{3} \end{array} \right) \right].$$

In this case, by Theorem 6.3.1, we need p to be such that Player 2 is indifferent between D and E : we need $2p = 1 - p$, that is, $p = \frac{1}{3}$. Hence the Nash equilibrium is

$$\left[\left(\begin{array}{ccc} A & B & C \\ 0 & \frac{1}{3} & \frac{2}{3} \end{array} \right), \left(\begin{array}{cc} D & E \\ \frac{1}{3} & \frac{2}{3} \end{array} \right) \right].$$

In games where the number of strategies or the number of players are larger than in the examples we have considered so far, finding the Nash equilibria involves lengthier calculations. However, computer programs have been developed that can be used to compute all the Nash equilibria of a finite game in a very short time.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 6.5.3 at the end of this chapter.

6.4 Strict dominance and rationalizability

We remarked in the previous section that a pure strategy that is strictly dominated by another pure strategy cannot be played with positive probability at a Nash equilibrium. Thus, when looking for a Nash equilibrium, one can first simplify the game by applying the IDSDS procedure (Chapter 2). When payoffs are cardinal (von Neumann-Morgenstern payoffs) it turns out that, *in a two-person game*, a pure strategy cannot be a best response to any mixed-strategy of the opponent not only when it is strictly dominated by another pure strategy but also when it is strictly dominated *by a mixed strategy*. To see this, consider the game of Figure 6.9.

		Player 2			
		<i>D</i>	<i>E</i>		
Player 1	<i>A</i>	0	1	4	0
	<i>B</i>	1	2	1	4
	<i>C</i>	2	0	0	1

Figure 6.9: A strategic-form game with cardinal payoffs

The pure strategy B of Player 1 is not strictly dominated by another pure strategy and yet it cannot be a best reply to any mixed strategy of Player 2.

To see this, consider an arbitrary mixed strategy $\begin{pmatrix} D & E \\ q & 1-q \end{pmatrix}$ of Player 2 with $q \in [0, 1]$.

If Player 1 plays B against it, she gets a payoff of 1; if, instead, she plays the mixed strategy $\begin{pmatrix} A & B & C \\ \frac{1}{3} & 0 & \frac{2}{3} \end{pmatrix}$ then her payoff is $\frac{1}{3}4(1-q) + \frac{2}{3}2q = \frac{4}{3} > 1$.

Theorem 6.4.1 — Pearce, 1984. Consider a two-player reduced-form game in strategic form with cardinal payoffs, an arbitrary Player i and a pure strategy s_i of Player i . Then there is no mixed-strategy of the opponent to which s_i is a best response, if and only if s_i is strictly dominated by a mixed strategy σ_i of Player i (that is, there is a $\sigma_i \in \Sigma_i$ such that $\Pi_i(\sigma_i, \sigma_j) > \Pi_i(s_i, \sigma_j)$, for every $\sigma_j \in \Sigma_j$).

Note that, since the set of mixed strategies includes the set of pure strategies, strict dominance by a mixed strategy includes as a sub-case strict dominance by a pure strategy.

When the number of players is 3 or more, the generalization of Theorem 6.4.1 raises some subtle issues: see Exercise 6.14. However, we can appeal to the intuition behind Theorem 6.4.1 (see the remark below) to refine the IDSDS procedure for general n -player games with cardinal payoffs as follows.

Definition 6.4.1 — Cardinal IDSDS. The *Cardinal Iterated Deletion of Strictly Dominated Strategies* is the following algorithm. Given a finite n -player ($n \geq 2$) strategic-form game with cardinal payoffs G , let G^1 be the game obtained by removing from G , for every Player i , those pure strategies of Player i (if any) that are strictly dominated in G by some *mixed* strategy of Player i ; let G^2 be the game obtained by removing from G^1 , for every Player i , those pure strategies of Player i (if any) that are strictly dominated in G^1 by some mixed strategy of Player i , and so on. Let G^∞ be the output of this procedure. Since the initial game G is finite, G^∞ will be obtained in a finite number of steps. For every Player i , the pure strategies of Player i in G^∞ are called her *rationalizable strategies*.

Figure 6.10 illustrates this procedure as applied to the game in Panel (i).




Figure 6.10: Application of the cardinal IDSDS procedure

In the first step, the pure strategy C of Player 1 is deleted, because it is strictly dominated by the mixed strategy $\begin{pmatrix} A & B \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$ thus yielding game G^1 shown in Panel (ii).

In the second step, the pure strategy F of Player 2 is deleted, because it is strictly dominated by the mixed strategy $\begin{pmatrix} D & E \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$ thus yielding game G^2 shown in Panel (iii).

In the third step, B is deleted because it is strictly dominated by A thus yielding game G^3 shown in the top part of Panel (iv).

In the final step, E is deleted because it is strictly dominated by D so that the final output is the strategy profile (A, D) .

Hence the only rationalizable strategies are A for Player 1 and D for Player 2.

Note that, in the game of Figure 6.10 Panel (i), since the only rationalizable strategy profile is (A, D) , it follows that (A, D) is also the unique Nash equilibrium.

As noted in Chapter 2 the significance of the output of the IDSDS procedure is as follows. Consider game G in Panel (i) of Figure 5.9. Since, for Player 1, C is strictly dominated, if Player 1 is rational she will not play C . Thus, if Player 2 believes that Player 1 is rational then he believes that Player 1 will not play C , that is, he restricts attention to game G^1 ; since, in G^1 , F is strictly dominated for Player 2, if Player 2 is rational he will not play F . It follows that if Player 1 believes that Player 2 is rational and that Player 2 believes that Player 1 is rational, then Player 1 restricts attention to game G^2 where rationality requires that Player 1 not play B , etc.

- R** Define a player to be rational if her chosen pure strategy is a best reply to her belief about what the opponent will do. In a two-player game a belief of Player 1 about what Player 2 will do can be expressed as a probability distribution over the set of pure strategies of Player 2; but this is the same object as a mixed strategy of Player 2. Thus, by Theorem 6.4.1, a rational Player 1 cannot choose a pure strategy that is strictly dominated by one of her own mixed strategies. The iterated reasoning outlined above can be captured by means of the notion of common knowledge of rationality. Indeed, it will be shown in Chapter 10 that if there is common knowledge of rationality then only rationalizable strategy profiles can be played. In a game with more than two players a belief of Player i about her opponents is no longer the same object as a mixed-strategy profile of the opponents, because a belief can allow for correlation in the behavior of the opponents, while the notion of mixed-strategy profile rules out such correlation (see Exercise 6.14).
- R** The iterated reasoning outlined above *requires that the von Neumann-Morgenstern preferences of both players be common knowledge between them*. For example, if Player 2 believes that Player 1 is rational but only knows her ordinal ranking of the outcomes, then Player 2 will not be able to deduce that it is irrational for Player 1 to play C and thus it cannot be irrational for him to play F . Expecting a player to know the von Neumann-Morgenstern preferences of another player is often (almost always?) very unrealistic! Thus one should be aware of the implicit assumptions that one makes (and one should question the assumptions made by others in their analyses).

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 6.5.4 at the end of this chapter.

6.5 Exercises

6.5.1 Exercises for Section 6.1: Strategic-form games with cardinal payoffs

The answers to the following exercises are in Section 6.6 at the end of this chapter.

Exercise 6.1 Consider the following game-frame in strategic form, where o_1, o_2, o_3 and o_4 are basic outcomes:

		Player 2	
		c	d
Player 1	a	o_1	o_2
	b	o_3	o_4

Both players satisfy the axioms of expected utility.

- The best outcome for Player 1 is o_3 ; she is indifferent between outcomes o_1 and o_4 and ranks them both as worst; she considers o_2 to be worse than o_3 and better than o_4 ; she is indifferent between o_2 with certainty and the lottery $\left(\begin{array}{cc} o_3 & o_1 \\ 0.25 & 0.75 \end{array} \right)$.
- The best outcome for Player 2 is o_4 , which he considers to be just as good as o_1 ; he considers o_2 to be worse than o_1 and better than o_3 ; he is indifferent between o_2 with certainty and the lottery $\left(\begin{array}{cc} o_1 & o_3 \\ 0.4 & 0.6 \end{array} \right)$.

Find the normalized von Neumann-Morgenstern utility functions for the two players and write the corresponding reduced-form game. ■

Exercise 6.2 Consider the game-frame shown in Figure 6.11, where o_1, \dots, o_4 are basic outcomes. Both players have von Neumann-Morgenstern rankings of the basic outcomes. The ranking of Player 1 can be represented by the following von Neumann-Morgenstern utility function:

$$\begin{array}{l} \text{outcome: } o_1 \quad o_2 \quad o_3 \quad o_4 \\ U_1: \quad 12 \quad 10 \quad 6 \quad 16 \end{array}$$

and the ranking of Player 2 can be represented by the following von Neumann-Morgenstern utility function:

$$\begin{array}{l} \text{outcome: } o_1 \quad o_2 \quad o_3 \quad o_4 \\ U_2: \quad 6 \quad 14 \quad 8 \quad 10 \end{array}$$

Write the corresponding reduced-form game. ■

		Player 2	
		<i>C</i>	<i>D</i>
Player 1	<i>A</i>	$\begin{pmatrix} o_1 & o_4 \\ \frac{1}{4} & \frac{3}{4} \end{pmatrix}$	$\begin{pmatrix} o_1 & o_2 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$
	<i>B</i>	o_3	$\begin{pmatrix} o_3 & o_4 \\ \frac{2}{5} & \frac{3}{5} \end{pmatrix}$

Figure 6.11: A game-frame in strategic form

6.5.2 Exercises for Section 6.2: Mixed strategies

The answers to the following exercises are in Section 6.6 at the end of this chapter.

		Player 2	
		<i>D</i>	<i>E</i>
Player 1	<i>A</i>	0 1	6 3
	<i>B</i>	4 4	2 0
	<i>C</i>	3 0	4 2

Figure 6.12: A strategic-form game with cardinal payoffs

Exercise 6.3 Consider the reduced-form game with cardinal payoffs shown in Figure 6.12.

(a) Calculate the players' payoffs from the mixed strategy profile

$$\left[\begin{pmatrix} A & B & C \\ \frac{1}{4} & \frac{3}{4} & 0 \end{pmatrix} \begin{pmatrix} D & E \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} \right].$$

(b) Is $\left[\begin{pmatrix} A & B & C \\ \frac{1}{4} & \frac{3}{4} & 0 \end{pmatrix} \begin{pmatrix} D & E \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} \right]$ a Nash equilibrium?

Exercise 6.4 Consider the following reduced-form game with cardinal payoffs:

		Player 2	
		<i>D</i>	<i>E</i>
Player 1	<i>A</i>	<i>2</i> , <i>3</i>	<i>8</i> , <i>5</i>
	<i>B</i>	<i>6</i> , <i>6</i>	<i>4</i> , <i>2</i>

Prove that $\left[\left(\begin{array}{cc} A & B \\ \frac{2}{3} & \frac{1}{3} \end{array} \right) \left(\begin{array}{cc} D & E \\ \frac{1}{2} & \frac{1}{2} \end{array} \right) \right]$ is a Nash equilibrium. ■

6.5.3 Exercises for Section 6.3: Computing the mixed-strategy Nash equilibria

The answers to the following exercises are in Section 6.6 at the end of this chapter.

Exercise 6.5 Consider again the game of Exercise 6.1.

- (a) Find the mixed-strategy Nash equilibrium.
- (b) Calculate the payoffs of both players at the Nash equilibrium.

Exercise 6.6 Find the Nash equilibrium of the game of Exercise 6.2. ■

Exercise 6.7 Find all the mixed-strategy Nash equilibria of the game of Exercise 6.4 and calculate the payoffs of both players at every Nash equilibrium. ■

Exercise 6.8 Find the mixed-strategy Nash equilibria of the following game:

		Player 2	
		<i>L</i>	<i>R</i>
Player 1	<i>T</i>	<i>1</i> , <i>4</i>	<i>4</i> , <i>3</i>
	<i>C</i>	<i>2</i> , <i>0</i>	<i>1</i> , <i>2</i>
	<i>B</i>	<i>1</i> , <i>5</i>	<i>0</i> , <i>6</i>

Exercise 6.9 Consider the following two-player game, where o_1, o_2, \dots, o_6 are basic outcomes.

		Player 2	
		<i>d</i>	<i>e</i>
Player 1	<i>a</i>	<i>o</i> ₁	<i>o</i> ₂
	<i>b</i>	<i>o</i> ₃	<i>o</i> ₄
	<i>c</i>	<i>o</i> ₅	<i>o</i> ₆

The players rank the outcomes as indicated below (as usual, if outcome o is above outcome o' then o is strictly preferred to o' and if o and o' are on the same row then the player is indifferent between the two):

$$\text{Player 1 : } \begin{pmatrix} o_1 \\ o_6 \\ o_4, o_2 \\ o_5 \\ o_3 \end{pmatrix} \quad \text{Player 2 : } \begin{pmatrix} o_3, o_4 \\ o_2 \\ o_1, o_5 \\ o_6 \end{pmatrix}$$

- (a) One player has a strategy that is strictly dominated. Identify the player and the strategy.

[Note: in order to answer the following questions, you can make your life a lot easier if you simplify the game on the basis of your answer to part (a).]

Player 1 satisfies the axioms of Expected Utility Theory and is indifferent between o_6 and the lottery $\left(\begin{smallmatrix} o_1 & o_5 \\ \frac{4}{5} & \frac{1}{5} \end{smallmatrix}\right)$ and is indifferent between o_2 and the lottery $\left(\begin{smallmatrix} o_6 & o_5 \\ \frac{1}{2} & \frac{1}{2} \end{smallmatrix}\right)$.

- (b) Suppose that Player 1 believes that Player 2 is going to play d with probability $\frac{1}{2}$ and e with probability $\frac{1}{2}$. Which strategy should he play?

Player 2 satisfies the axioms of Expected Utility Theory and is indifferent between o_5 and the lottery $\left(\begin{smallmatrix} o_2 & o_6 \\ \frac{1}{4} & \frac{3}{4} \end{smallmatrix}\right)$.

- (c) Suppose that Player 2 believes that Player 1 is going to play a with probability $\frac{1}{4}$ and c with probability $\frac{3}{4}$. Which strategy should she play?
- (d) Find all the (pure- and mixed-strategy) Nash equilibria of this game. ■

Exercise 6.10 Consider the following game (where the payoffs are von Neumann-Morgenstern payoffs):

		Player 2	
		C	D
Player 1	A	x , y	3 , 0
	B	6 , 2	0 , 4

- (a) Suppose that $x = 2$ and $y = 2$. Find the mixed-strategy Nash equilibrium and calculate the payoffs of both players at the Nash equilibrium.
- (b) For what values of x and y is $\left[\left(\begin{smallmatrix} A & B \\ \frac{1}{5} & \frac{4}{5} \end{smallmatrix}\right), \left(\begin{smallmatrix} C & D \\ \frac{3}{4} & \frac{1}{4} \end{smallmatrix}\right)\right]$ a Nash equilibrium? ■

Exercise 6.11 Find the mixed-strategy Nash equilibria of the game of Exercise 6.3. Calculate the payoffs of both players at every Nash equilibrium that you find. ■

6.5.4 Exercises for Section 6.4: Strict dominance and rationalizability

The answers to the following exercises are in Section 6.6 at the end of this chapter.

Exercise 6.12 In the following game, for each player, find all the rationalizable pure strategies (that is, apply the cardinal IDSDS procedure).

		Player 2		
		<i>L</i>	<i>M</i>	<i>R</i>
Player 1	<i>A</i>	3 , 5	2 , 0	2 , 2
	<i>B</i>	5 , 2	1 , 2	2 , 1
	<i>C</i>	9 , 0	1 , 5	3 , 2

Note: The next three exercises are more difficult than the previous ones.

Exercise 6.13 Is the following statement true or false? Either prove that it is true or give a counterexample.

- “Consider a two-player strategic-form game with cardinal payoffs.
- Let A and B be two pure strategies of Player 1.
 - Suppose that both A and B are rationalizable (that is, they survive the cardinal IDSDS procedure).
 - Then any mixed strategy that attaches positive probability to both A and B and zero to every other strategy is a best reply to some mixed strategy of Player 2.”

Exercise 6.14 Consider the three-player game shown in Figure 6.13, where only the payoffs of Player 1 are recorded.

- (a) Show that if Player 1 assigns probability $\frac{1}{2}$ to the event “Player 2 will play E and Player 3 will play G ” and probability $\frac{1}{2}$ to the event “Player 2 will play F and Player 3 will play H ”, then playing D is a best reply.

Next we want to show that there is no mixed-strategy profile

$$\sigma_{-1} = \left(\left(\begin{array}{cc} E & F \\ p & 1-p \end{array} \right), \left(\begin{array}{cc} G & H \\ q & 1-q \end{array} \right) \right)$$

of Players 2 and 3 against which D is a best reply for Player 1.

Define the following functions: $A(p, q) = \Pi_1(A, \sigma_{-1})$ (that is, $A(p, q)$ is Player 1’s expected payoff if she plays the pure strategy A against σ_{-1}), $B(p, q) = \Pi_1(B, \sigma_{-1})$, $C(p, q) = \Pi_1(C, \sigma_{-1})$ and $D(p, q) = \Pi_1(D, \sigma_{-1})$.

- (b) In the (p, q) plane (with $0 \leq p \leq 1$ and $0 \leq q \leq 1$) draw the curve corresponding to the equation $A(p, q) = D(p, q)$ and identify the region where $A(p, q) > D(p, q)$.
- (c) In the (p, q) plane draw the curve corresponding to the equation $C(p, q) = D(p, q)$ and identify the region where $C(p, q) > D(p, q)$.
- (d) In the (p, q) plane draw the two curves corresponding to the equation $B(p, q) = D(p, q)$ and identify the region where $B(p, q) > D(p, q)$.
- (e) Infer from parts (b)-(c) that there is no mixed-strategy profile of Players 2 and 3 against which D is a best reply for Player 1.

		Player 2				Player 2	
		E	F			E	F
Player 1	A	3	0		A	0	0
	B	0	3		B	3	0
	C	0	0		C	0	3
	D	2	0		D	0	2
		Player 3: G				Player 3: H	

Figure 6.13: A three-player game where only the payoffs of Player 1 are shown

Exercise 6.15 — *** **Challenging Question** ***. A team of n professional swimmers ($n \geq 2$) – from now on called players – are partying on the bank of the Sacramento river on a cold day in January. Suddenly a passerby shouts “Help! My dog fell into the water!” Each of the swimmers has to decide whether or not to jump into the icy cold water to rescue the dog. One rescuer is sufficient: the dog will be saved if at least one player jumps into the water; if nobody does, then the dog will die. Each player prefers somebody else to jump in, but each player prefers to jump in himself if nobody else does.

Let us formulate this as a game. The strategy set of each player $i = 1, \dots, n$ is $S_i = \{J, \neg J\}$, where J stands for ‘jump in’ and $\neg J$ for ‘not jump in’.

The possible basic outcomes can be expressed as subsets of the set $I = \{1, \dots, n\}$ of players: outcome $N \subseteq I$ is interpreted as ‘the players in the set N jump into the water’; if $N = \emptyset$ the dog dies, while if $N \neq \emptyset$ the dog is saved.

Player i has the following ordinal ranking of the outcomes:

- (1) $N \sim N'$, for every $N \neq \emptyset, N' \neq \emptyset$ with $i \notin N$ and $i \notin N'$,
- (2) $N \succ N'$ for every $N \neq \emptyset, N' \neq \emptyset$ with $i \notin N$ and $i \in N'$,
- (3) $\{i\} \succ \emptyset$.

- (a) Find all the pure-strategy Nash equilibria.
- (b) Suppose that each player i has the following von Neumann-Morgenstern payoff function (which is consistent with the above ordinal ranking):

$$\pi_i(N) = \begin{cases} v & \text{if } N \neq \emptyset \text{ and } i \notin N \\ v - c & \text{if } N \neq \emptyset \text{ and } i \in N \\ 0 & \text{if } N = \emptyset \end{cases} \quad \text{with } 0 < c < v.$$

Find the symmetric mixed-strategy Nash equilibrium (symmetric means that all the players use the same mixed strategy).

- (c) Assuming that the players behave according to the symmetric mixed-strategy Nash equilibrium of Part (b), is it better for the dog if n (the number of players) is large or if n is small? Calculate the probability that the dog is saved at the mixed-strategy Nash equilibrium as a function of n , for all possible values of c and v (subject to $0 < c < v$), and plot it for the case where $c = 10$ and $v = 12$.

6.6 Solutions to exercises

Solution to Exercise 6.1. The normalized von Neumann-Morgenstern utility functions are:

	outcome	U_1		outcome	U_2
Player 1:	o_3	1		o_1, o_4	1
	o_2	0.25	Player 2:	o_2	0.4
	o_1, o_4	0		o_3	0

The reduced-form game is shown in Figure 6.14. □

		Player 2			
		<i>c</i>		<i>d</i>	
Player 1	<i>a</i>	0	1	0.25	0.4
	<i>b</i>	1	0	0	1

Figure 6.14: The reduced-form game for Exercise 6.1

Solution to Exercise 6.2.

The expected utility of the lottery $\left(\begin{array}{cc} o_1 & o_4 \\ \frac{1}{4} & \frac{3}{4} \end{array} \right)$ is $\frac{1}{4}(12) + \frac{3}{4}(16) = 15$ for Player 1 and $\frac{1}{4}(6) + \frac{3}{4}(10) = 9$ for Player 2.

The expected utility of the lottery $\left(\begin{array}{cc} o_1 & o_2 \\ \frac{1}{2} & \frac{1}{2} \end{array} \right)$ is 11 for Player 1 and 10 for Player 2.

The expected utility of the lottery $\left(\begin{array}{cc} o_3 & o_4 \\ \frac{2}{5} & \frac{3}{5} \end{array} \right)$ is 12 for Player 1 and 9.2 for Player 2.

The reduced-form game is shown in Figure 6.15. □

		Player 2			
		<i>C</i>		<i>D</i>	
Player 1	<i>A</i>	15	9	11	10
	<i>B</i>	6	8	12	9.2

Figure 6.15: The reduced-form game for Exercise 6.2

Solution to Exercise 6.3.

(a) $\Pi_1 = \frac{1}{8}(0) + \frac{1}{8}(6) + \frac{3}{8}(4) + \frac{3}{8}(2) = 3$ and $\Pi_2 = \frac{1}{8}(1) + \frac{1}{8}(3) + \frac{3}{8}(4) + \frac{3}{8}(0) = 2$.

(b) No, because if Player 1 switched to the pure strategy *C* then her payoff would be $\frac{1}{2}(3) + \frac{1}{2}(4) = 3.5 > 3$. □

Solution to Exercise 6.4. Player 1's payoff is $\Pi_1 = \frac{2}{6}(2) + \frac{2}{6}(8) + \frac{1}{6}(6) + \frac{1}{6}(4) = 5$.

If Player 1 switches to any other mixed strategy $\begin{pmatrix} A & B \\ p & 1-p \end{pmatrix}$, while Player 2's strategy is kept fixed at $\begin{pmatrix} C & D \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$, then her payoff is $\Pi_1 = \frac{1}{2}(p)(2) + \frac{1}{2}(p)(8) + \frac{1}{2}(1-p)(6) + \frac{1}{2}(1-p)(4) = 5$.

Thus any mixed strategy of Player 1 is a best response to $\begin{pmatrix} C & D \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$.

Similarly, Player 2's payoff is $\Pi_2 = \frac{2}{6}3 + \frac{2}{6}(5) + \frac{1}{6}(6) + \frac{1}{6}(2) = 4$. If Player 2 switches to any other mixed strategy $\begin{pmatrix} C & D \\ q & 1-q \end{pmatrix}$, while Player 1's strategy is kept fixed at $\begin{pmatrix} A & B \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}$, then her payoff is $\Pi_2 = \frac{2}{3}(q)(3) + \frac{2}{3}(1-q)(5) + \frac{1}{3}(q)(6) + \frac{1}{3}(1-q)(2) = 4$.

Thus any mixed strategy of Player 2 is a best response to $\begin{pmatrix} A & B \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}$.

Hence $\begin{pmatrix} A & B \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}$ is a best reply to $\begin{pmatrix} C & D \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$ and $\begin{pmatrix} C & D \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$ is a best reply to $\begin{pmatrix} A & B \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}$,

that is, $\left[\begin{pmatrix} A & B \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}, \begin{pmatrix} C & D \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} \right]$ is a Nash equilibrium. \square

Solution to Exercise 6.5.

(a) We have to find the Nash equilibrium of the following game:

		Player 2	
		<i>c</i>	<i>d</i>
Player	<i>a</i>	0 , 1	0.25 , 0.4
1	<i>b</i>	1 , 0	0 , 1

To make calculations easier, let us multiply all the payoffs by 100 (that is, we re-scale the von Neumann-Morgenstern utility functions by a factor of 100):

		Player 2	
		<i>c</i>	<i>d</i>
Player	<i>a</i>	0 , 100	25 , 40
1	<i>b</i>	100 , 0	0 , 100

There are no pure-strategy Nash equilibria. To find the mixed-strategy Nash equilibrium, let p be the probability with which Player 1 chooses a and q be the probability with which Player 2 chooses c .

Then, for Player 1, the payoff from playing a against $\begin{pmatrix} c & d \\ q & 1-q \end{pmatrix}$ must be equal to the payoff from playing b against $\begin{pmatrix} c & d \\ q & 1-q \end{pmatrix}$. That is, it must be that $25(1-q) = 100q$, which yields $q = \frac{1}{5}$. Similarly, for Player 2, the payoff from playing c against $\begin{pmatrix} a & b \\ p & 1-p \end{pmatrix}$ must be equal to the payoff from playing d against $\begin{pmatrix} a & b \\ p & 1-p \end{pmatrix}$. This requires $100p = 40p + 100(1-p)$, that is, $p = \frac{5}{8}$. Thus the Nash equilibrium is $\left[\begin{pmatrix} a & b \\ \frac{5}{8} & \frac{3}{8} \end{pmatrix}, \begin{pmatrix} c & d \\ \frac{1}{5} & \frac{4}{5} \end{pmatrix} \right]$.

- (b) At the Nash equilibrium the payoffs are 20 for Player 1 and 62.5 for Player 2. (If you worked with the original payoffs, then the Nash equilibrium payoffs would be 0.2 for Player 1 and 0.625 for Player 2.) \square

Solution to Exercise 6.6. We have to find the Nash equilibria of the following game.

		Player 2	
		C	D
Player 1	A	15, 9	11, 10
	B	6, 8	12, 9.2

For Player 2 D is a strictly dominant strategy, thus at a Nash equilibrium Player 2 must play D with probability 1. For Player 1, the unique best reply to D is B . Thus the pure-strategy profile (B, D) is the only Nash equilibrium. \square

Solution to Exercise 6.7. We have to find the Nash equilibria of the following game.

		Player 2	
		D	E
Player 1	A	2, 3	8, 5
	B	6, 6	4, 2

(B, D) (with payoffs (6,6)) and (A, E) (with payoffs (8,5)) are both Nash equilibria. To see if there is also a mixed-strategy equilibrium we need to solve the following equations, where p is the probability of A and q is the probability of D : $2q + 8(1-q) = 6q + 4(1-q)$ and $3p + 6(1-p) = 5p + 2(1-p)$. The solution is $p = \frac{2}{3}$ and $q = \frac{1}{2}$ so that

$$\left[\begin{pmatrix} A & B \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}, \begin{pmatrix} D & E \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} \right]$$

is a Nash equilibrium. The payoffs at this Nash equilibrium are 5 for Player 1 and 4 for Player 2. \square

Solution to Exercise 6.8. Since B is strictly dominated (by C), it cannot be assigned positive probability at a Nash equilibrium. Let p be the probability of T and q the probability of L . Then p must be such that $4p + 0(1 - p) = 3p + 2(1 - p)$ and q must be such that $q + 4(1 - q) = 2q + (1 - q)$. Thus $p = \frac{2}{3}$ and $q = \frac{3}{4}$. Hence there is only one mixed-strategy equilibrium, namely

$$\left[\left(\begin{array}{ccc} T & C & B \\ \frac{2}{3} & \frac{1}{3} & 0 \end{array} \right), \left(\begin{array}{cc} L & R \\ \frac{3}{4} & \frac{1}{4} \end{array} \right) \right].$$

□

Solution to Exercise 6.9.

(a) Since Player 1 prefers o_5 to o_3 and prefers o_6 to o_4 , strategy b is strictly dominated by strategy c .

Thus, at a Nash equilibrium, Player 1 will not play b with positive probability and we can simplify the game to

		Player 2	
		d	e
Player 1	a	o_1	o_2
	c	o_5	o_6

Of the remaining outcomes, for Player 1 o_1 is the best outcome (we can assign utility 1 to it) and o_5 is the worst (we can assign utility 0 to it). Since he is indifferent between o_6 and the lottery $\left(\begin{array}{cc} o_1 & o_5 \\ \frac{4}{5} & \frac{1}{5} \end{array} \right)$, the utility of o_6 is $\frac{4}{5}$. Hence the expected utility of $\left(\begin{array}{cc} o_5 & o_6 \\ \frac{1}{2} & \frac{1}{2} \end{array} \right)$ is $\frac{1}{2}(0) + \frac{1}{2}(\frac{4}{5}) = \frac{2}{5}$ and thus the utility of o_2 is also $\frac{2}{5}$.

(b) If Player 2 plays d with probability $\frac{1}{2}$ and e with probability $\frac{1}{2}$, then for Player 1 playing a gives a payoff of $\frac{1}{2}(1) + \frac{1}{2}(\frac{2}{5}) = \frac{7}{10}$, while playing c gives a payoff of $\frac{1}{2}(0) + \frac{1}{2}(\frac{4}{5}) = \frac{4}{10}$. Hence he should play a .

If you did not follow the suggestion to simplify the analysis as was done above, then you can still reach the same answer, although in a lengthier way. You would still set $U(o_1) = 1$. Then the expected payoff from playing a is

$$\Pi_1(a) = \frac{1}{2}U(o_1) + \frac{1}{2}U(o_2) = \frac{1}{2} + \frac{1}{2}U(o_2) \quad (\star)$$

Since o_2 is as good as $\left(\begin{array}{cc} o_5 & o_6 \\ \frac{1}{2} & \frac{1}{2} \end{array} \right)$,

$$U(o_2) = \frac{1}{2}U(o_5) + \frac{1}{2}U(o_6). \quad (\diamond)$$

Since o_6 is as good as $\left(\begin{array}{cc} o_1 & o_5 \\ \frac{4}{5} & \frac{1}{5} \end{array} \right)$,

$$U(o_6) = \frac{4}{5} + \frac{1}{5}U(o_5). \quad (\dagger)$$

Replacing (\dagger) in (\diamond) we get $U(o_2) = \frac{2}{5} + \frac{3}{5}U(o_5)$ and replacing this expression in (\star) we get $\Pi_1(a) = \frac{7}{10} + \frac{3}{10}U(o_5)$. Similarly,

$$\Pi_1(c) = \frac{1}{2}U(o_5) + \frac{1}{2}U(o_6) = \frac{1}{2}U(o_5) + \frac{1}{2}\left(\frac{4}{5} + \frac{1}{5}U(o_5)\right) = \frac{4}{10} + \frac{6}{10}U(o_5)$$

Now, $\Pi_1(a) > \Pi_1(c)$ if and only if $\frac{7}{10} + \frac{3}{10}U(o_5) > \frac{4}{10} + \frac{6}{10}U(o_5)$ if and only if $3 > 3U(o_5)$ if and only if $U(o_5) < 1$, which is the case because o_5 is worse than o_1 and $U(o_1) = 1$. Similar steps would be taken to answer parts **(c)** and **(d)**.

(c) In the reduced game, for Player 2 o_2 is the best outcome (we can assign utility 1 to it) and o_6 is the worst (we can assign utility 0 to it). Thus, since she is indifferent between o_5 and the lottery $\left(\begin{array}{cc} o_2 & o_6 \\ \frac{1}{4} & \frac{3}{4} \end{array} \right)$, the utility of o_5 is $\frac{1}{4}$ and so is the utility of o_1 . Thus playing d gives an expected payoff of $\frac{1}{4}(\frac{1}{4}) + \frac{3}{4}(\frac{1}{4}) = \frac{1}{4}$ and playing e gives an expected utility of $\frac{1}{4}(1) + \frac{3}{4}(0) = \frac{1}{4}$. Thus she is indifferent between playing d and playing e (and any mixture of d and e).

(d) Using the calculations of parts **(b)** and **(c)** the game is as follows:

		Player 2	
		d	e
Player 1	a	1, $\frac{1}{4}$	$\frac{2}{5}$, 1
1	c	0, $\frac{1}{4}$	$\frac{4}{5}$, 0

There is no pure-strategy Nash equilibrium. At a mixed-strategy Nash equilibrium, each player must be indifferent between his/her two strategies. From part **(c)** we already know that Player 2 is indifferent if Player 1 plays a with probability $\frac{1}{4}$ and c with probability $\frac{3}{4}$. Now let q be the probability with which Player 2 plays d . Then we need $q + \frac{2}{5}(1 - q) = \frac{4}{5}(1 - q)$, hence $q = \frac{2}{7}$. Thus the Nash equilibrium is $\left(\begin{array}{ccc|cc} a & b & c & d & e \\ \frac{1}{4} & 0 & \frac{3}{4} & \frac{2}{7} & \frac{5}{7} \end{array} \right)$ which can be written more succinctly as $\left(\begin{array}{cc|cc} a & c & d & e \\ \frac{1}{4} & \frac{3}{4} & \frac{2}{7} & \frac{5}{7} \end{array} \right)$. □

Solution to Exercise 6.10.

- (a) Let p be the probability of A and q the probability of B . Player 1 must be indifferent between playing A and playing B : $2q + 3(1 - q) = 6q$; this gives $q = \frac{3}{7}$. Similarly, Player 2 must be indifferent between playing C and playing D : $2 = 4(1 - p)$; this gives $p = \frac{1}{2}$. Thus the Nash equilibrium is given by

$$\left[\left(\begin{matrix} A & B \\ \frac{1}{2} & \frac{1}{2} \end{matrix} \right), \left(\begin{matrix} C & D \\ \frac{3}{7} & \frac{4}{7} \end{matrix} \right) \right]$$

The equilibrium payoffs are $\frac{18}{7} = 2.57$ for Player 1 and 2 for Player 2.

- (b) Player 1 must be indifferent between playing A and playing B : $\frac{3}{4}(x) + \frac{1}{4}(3) = \frac{3}{4}(6)$. Thus $x = 5$. Similarly, Player 2 must be indifferent between playing C and playing D : $\frac{1}{5}(y) + \frac{4}{5}(2) = \frac{4}{5}(4)$. Thus $y = 8$. \square

Solution to Exercise 6.11. We have to find the Nash equilibria of the following game:

		Player 2	
		D	E
Player 1	A	0 , 1	6 , 3
	B	4 , 4	2 , 0
	C	3 , 0	4 , 2

There are two pure-strategy equilibria, namely (B, D) and (A, E) . To see if there is a mixed-strategy equilibrium we calculate the best response of Player 1 to every possible mixed strategy $\begin{pmatrix} D & E \\ q & 1 - q \end{pmatrix}$ of Player 2 (with $q \in [0, 1]$). For Player 1 the payoff from playing A against $\begin{pmatrix} D & E \\ q & 1 - q \end{pmatrix}$ is $6 - 6q$, the payoff from playing B is $4q + 2(1 - q) = 2 + 2q$ and the payoff from playing C is $3q + 4(1 - q) = 4 - q$. These functions are shown in Figure 6.16, where the downward-sloping line plots the function where $A(q) = 6 - 6q$, the upward-sloping line plots the function $B(q) = 2 + 2q$ and the dotted line the function $C(q) = 4 - q$.

It can be seen from Figure 6.16 that

$$\text{Player 1's best reply} = \begin{cases} A & \text{if } 0 \leq q < \frac{2}{5} \\ \begin{pmatrix} A & C \\ p & 1 - p \end{pmatrix} \text{ for any } p \in [0, 1] & \text{if } q = \frac{2}{5} \\ C & \text{if } \frac{2}{5} < q < \frac{2}{3} \\ \begin{pmatrix} B & C \\ p & 1 - p \end{pmatrix} \text{ for any } p \in [0, 1] & \text{if } q = \frac{2}{3} \\ B & \text{if } \frac{2}{3} < q \leq 1 \end{cases}$$




Figure 6.16: The best-reply diagram for Exercise 6.11

Thus if there is a mixed-strategy equilibrium it is either of the form

$$\left[\left(\begin{array}{cc} A & C \\ p & 1-p \end{array} \right), \left(\begin{array}{cc} D & E \\ \frac{2}{5} & \frac{3}{5} \end{array} \right) \right] \quad \text{or of the form} \quad \left[\left(\begin{array}{cc} B & C \\ p & 1-p \end{array} \right), \left(\begin{array}{cc} D & E \\ \frac{2}{3} & \frac{1}{3} \end{array} \right) \right].$$

In the first case, where Player 1 chooses B with probability zero, E strictly dominates D for Player 2 and thus $\left(\begin{array}{cc} D & E \\ \frac{2}{5} & \frac{3}{5} \end{array} \right)$ is not a best reply for Player 2, so that

$$\left[\left(\begin{array}{cc} A & C \\ p & 1-p \end{array} \right), \left(\begin{array}{cc} D & E \\ \frac{2}{5} & \frac{3}{5} \end{array} \right) \right] \quad \text{is not a Nash equilibrium for any } p.$$

In the second case we need $D(p) = E(p)$, that is, $4p = 2(1-p)$, which yields $p = \frac{1}{3}$.

Thus the mixed-strategy Nash equilibrium is $\left[\left(\begin{array}{cc} B & C \\ \frac{1}{3} & \frac{2}{3} \end{array} \right), \left(\begin{array}{cc} D & E \\ \frac{2}{3} & \frac{1}{3} \end{array} \right) \right]$ with payoffs of $\frac{10}{3}$ for Player 1 and $\frac{4}{3}$ for Player 2. \square

Solution to Exercise 6.12. For Player 1, B is strictly dominated by $\begin{pmatrix} A & C \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$; for Player 2, R is strictly dominated by $\begin{pmatrix} L & M \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$. Eliminating B and R we are left with

		Player 2	
		L	M
Player 1	A	3 , 5	2 , 0
	C	9 , 0	1 , 5

In this game no player has a strictly dominated strategy. Thus for Player 1 both A and C are rationalizable and for Player 2 both L and M are rationalizable. \square

Solution to Exercise 6.13. The statement is false. Consider, for example, the following game:

		Player 2	
		L	R
Player 1	A	3 , 1	0 , 0
	B	0 , 0	3 , 1
	C	2 , 1	2 , 1

Here both A and B are rationalizable (indeed, they are both part of a Nash equilibrium; note that the cardinal IDSDS procedure leaves the game unchanged). However, the mixture $\begin{pmatrix} A & B \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$ (which gives Player 1 a payoff of 1.5, no matter what Player 2 does) *cannot* be a best reply to any mixed strategy of Player 2, since it is strictly dominated by C . \square

Solution to Exercise 6.14.

- (a) If Player 1 assigns probability $\frac{1}{2}$ to the event "Player 2 will play E and Player 3 will play G " and probability $\frac{1}{2}$ to the event "Player 2 will play F and Player will play H ", then A gives Player 1 an expected payoff of 1.5, B an expected payoff of 0, C an expected payoff of 1.5 and D an expected payoff of 2. Thus D is a best reply to those beliefs. The functions are as follows: $A(p, q) = 3pq$, $B(p, q) = 3(1 - p)q + 3p(1 - q)$, $C(p, q) = 3(1 - p)(1 - q)$, $D(p, q) = 2pq + 2(1 - p)(1 - q)$.
- (b) $A(p, q) = D(p, q)$ at those points (p, q) such that $q = \frac{2-2p}{2-p}$. The set of such points is the continuous curve in the Figure 6.17. The region where $A(p, q) > D(p, q)$ is the region **above** the continuous curve.
- (c) $C(p, q) = D(p, q)$ at those points (p, q) such that $q = \frac{1-p}{1+p}$. The set of such points is the dotted curve in the diagram shown in Figure 6.17. The region where $C(p, q) > D(p, q)$ is the region **below** the dotted curve.
- (d) $B(p, q) = D(p, q)$ at those points (p, q) such that $q = \frac{2-5p}{5-10p}$ (for $p \neq \frac{1}{2}$). The set of such points is given by the two dashed curves in the diagram below. The region where $B(p, q) > D(p, q)$ is the region between the two dashed curves.

Thus

- in the region strictly above the continuous curve, A is better than D ,
- in the region strictly below the dotted curve, C is better than D and
- in the region on and between the continuous curve and the dotted curve, B is better than D .

Hence, at every point in the (p, q) square there is a pure strategy of Player 1 which is strictly better than D . It follows that there is no mixed-strategy σ_{-1} against which D is a best reply. \square




Figure 6.17: The diagram for Exercise 6.14

Solution to Exercise 6.15.

- (a) There are n pure-strategy Nash equilibria: at each equilibrium exactly one player jumps in.
- (b) Let p be the probability with which each player jumps into the water. Consider a Player i . The probability that none of the other players jump in is $(1-p)^{n-1}$ and thus the probability that somebody else jumps in is $[1 - (1-p)^{n-1}]$. Player i 's payoff if he jumps in is $v - c$, while his expected payoff if he does not jump in is $v [1 - (1-p)^{n-1}] + 0(1-p)^{n-1} = v [1 - (1-p)^{n-1}]$.

Thus we need $v - c = v [1 - (1-p)^{n-1}]$, that is, $p = 1 - \left(\frac{c}{v}\right)^{\frac{1}{n-1}}$, which is strictly between 0 and 1 because $c < v$.

- (c) At the Nash equilibrium the probability that nobody jumps in is $(1 - p)^n = \left(\frac{c}{v}\right)^{\frac{n}{n-1}}$; thus this is the probability that the dog dies.

Hence, the dog is rescued with the remaining probability $1 - \left(\frac{c}{v}\right)^{\frac{n}{n-1}}$.

This is a decreasing function of n . The larger the number of swimmers who are present, the more likely it is that the dog dies.

The plot of this function when $c = 10$ and $v = 12$ is shown in Figure 6.18. \square




Figure 6.18: The probability that the dog is saved as a function of the number of potential rescuers

7. Extensive-form Games

7.1 Behavioral strategies in dynamic games

The definition of dynamic (or extensive-form) game-frame with cardinal payoffs is just like the definition of extensive-form game-frame with ordinal payoffs (Definition 4.1.1, Chapter 4), the only difference being that we postulate von Neumann-Morgenstern preferences instead of merely ordinal preferences.

In Chapter 6 we generalized the notion of strategic-form game-frame by allowing for lotteries (rather than just simple outcomes) to be associated with strategy profiles. One can do the same with extensive-form game-frames. For example, Figure 7.1 shows an extensive-form game-frame where associated with each terminal node (denoted by z_i , $i = 1, 2, \dots, 5$) is either a basic outcome (denoted by o_j , $j = 1, 2, \dots, 5$) or a lottery (probability distribution) over the set of basic outcomes $\{o_1, o_2, o_3, o_4, o_5\}$.




Figure 7.1: An extensive-form frame with probabilistic outcomes. The z_i 's are terminal nodes and the o_i 's are basic outcomes

In Figure 7.1 $\{z_1, z_2, \dots, z_5\}$ is the set of terminal nodes and $\{o_1, o_2, \dots, o_5\}$ is the set of basic outcomes.

Associated with z_1 is the lottery $\begin{pmatrix} o_1 & o_2 \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}$, while the lottery associated with z_3 is

$\begin{pmatrix} o_1 & o_3 & o_4 \\ \frac{1}{5} & \frac{3}{5} & \frac{1}{5} \end{pmatrix}$, etc.

However, as we saw at the end of Chapter 4, in extensive forms one can explicitly represent random events by means of chance moves (also called moves of Nature). Thus an alternative representation of the extensive-form frame of Figure 7.1 is the extensive form shown in Figure 7.2.

We can continue to use the definition of extensive-form frame given in Chapter 4, but from now on we will allow for the possibility of chance moves.

The notion of strategy remains, of course, unchanged: a strategy for a player is a list of choices, one for every information set of that player (Definition 4.2.1, Chapter 4). For example, the set of strategies for Player 1 in the extensive frame of Figure 7.2 is $S_1 = \{(a, e), (a, f), (b, e), (b, f)\}$. Thus mixed strategies can easily be introduced also in extensive frames. For example, in the extensive frame of Figure 7.2, the set of mixed strategies for Player 1, denoted by Σ_1 , is the set of probability distributions over S_1 :

$$\Sigma_1 = \left\{ \begin{pmatrix} (a, e) & (a, f) & (b, e) & (b, f) \\ p & q & r & 1-p-q-r \end{pmatrix} : p, q, r \in [0, 1] \text{ and } p+q+r \leq 1 \right\}.$$




Figure 7.2: An alternative representation of the extensive frame of Figure 7.1. The terminal nodes have not been labeled. The o_i 's are basic outcomes.

However, it turns out that in extensive forms with perfect recall one can use simpler objects than mixed strategies, namely behavioral strategies.

Definition 7.1.1 A behavioral strategy for a player in an extensive form is a list of probability distributions, one for every information set of that player; each probability distribution is over the set of choices at the corresponding information set.

For example, the set of behavioral strategies for Player 1 in the extensive frame of Figure 7.2 is:

$$\left\{ \left(\begin{array}{cc|cc} a & b & e & f \\ p & 1-p & q & 1-q \end{array} \right) : p, q \in [0, 1] \right\}$$

A behavioral strategy is a simpler object than a mixed strategy: in this example, specifying a *behavioral strategy* for Player 1 requires specifying the values of two parameters (p and q), while specifying a *mixed strategy* requires specifying the values of three parameters (p , r and q). Can one then use behavioral strategies rather than mixed strategies? The answer is affirmative, as Theorem 7.1.1 below states.

First we illustrate with an example based on the extensive form of Figure 7.3 (the z_i 's are terminal nodes and the outcomes have been omitted).




Figure 7.3: An extensive frame with the outcomes omitted. The z_i 's are terminal nodes.

Consider the mixed-strategy profile $\sigma = (\sigma_1, \sigma_2)$ with

$$\sigma_1 = \begin{pmatrix} (a, e) & (a, f) & (b, e) & (b, f) \\ \frac{1}{12} & \frac{4}{12} & \frac{2}{12} & \frac{5}{12} \end{pmatrix} \quad \text{and} \quad \sigma_2 = \begin{pmatrix} c & d \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}.$$

We can compute the probability of reaching terminal node z_i , denoted by $P(z_i)$, as follows:

$$P(z_1) = \sigma_1((a, e)) \sigma_2(c) + \sigma_1((a, f)) \sigma_2(c) = \frac{1}{12} \left(\frac{1}{3}\right) + \frac{4}{12} \left(\frac{1}{3}\right) = \frac{5}{36}$$

$$P(z_2) = \sigma_1((a, e)) \sigma_2(d) + \sigma_1((a, f)) \sigma_2(d) = \frac{1}{12} \left(\frac{2}{3}\right) + \frac{4}{12} \left(\frac{2}{3}\right) = \frac{10}{36}$$

$$P(z_3) = \sigma_1((b, e)) \sigma_2(c) + \sigma_1((b, f)) \sigma_2(c) = \frac{2}{12} \left(\frac{1}{3}\right) + \frac{5}{12} \left(\frac{1}{3}\right) = \frac{7}{36}$$

$$P(z_4) = \sigma_1((b, e)) \sigma_2(d) = \frac{2}{12} \left(\frac{2}{3}\right) = \frac{4}{36}$$

$$P(z_5) = \sigma_1((b, f)) \sigma_2(d) = \frac{5}{12} \left(\frac{2}{3}\right) = \frac{10}{36}.$$

That is, the mixed-strategy profile $\sigma = (\sigma_1, \sigma_2)$ gives rise to the following probability distribution over terminal nodes:

$$\begin{pmatrix} z_1 & z_2 & z_3 & z_4 & z_5 \\ \frac{5}{36} & \frac{10}{36} & \frac{7}{36} & \frac{4}{36} & \frac{10}{36} \end{pmatrix}.$$

Now consider the following behavioral strategy of Player 1:

$$\left(\begin{array}{cc|cc} a & b & e & f \\ \frac{5}{12} & \frac{7}{12} & \frac{2}{7} & \frac{5}{7} \end{array} \right).$$

What probability distribution over the set of terminal nodes would it induce in conjunction with Player 2's mixed strategy $\sigma_2 = \begin{pmatrix} c & d \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$? The calculations are simple:¹

$$P(z_1) = P(a) \sigma_2(c) = \frac{5}{12} \left(\frac{1}{3}\right) = \frac{5}{36},$$

$$P(z_2) = P(a) \sigma_2(d) = \frac{5}{12} \left(\frac{2}{3}\right) = \frac{10}{36},$$

$$P(z_3) = P(b) \sigma_2(c) = \frac{7}{12} \left(\frac{1}{3}\right) = \frac{7}{36},$$

$$P(z_4) = P(b) \sigma_2(d) P(e) = \frac{7}{12} \left(\frac{2}{3}\right) \left(\frac{2}{7}\right) = \frac{4}{36},$$

$$P(z_5) = P(b) \sigma_2(d) P(f) = \frac{7}{12} \left(\frac{2}{3}\right) \left(\frac{5}{7}\right) = \frac{10}{36}.$$

Thus, against $\sigma_2 = \begin{pmatrix} c & d \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$, Player 1's behavioral strategy $\left(\begin{array}{cc|cc} a & b & e & f \\ \frac{5}{12} & \frac{7}{12} & \frac{2}{7} & \frac{5}{7} \end{array} \right)$ and her mixed strategy $\begin{pmatrix} (a,e) & (a,f) & (b,e) & (b,f) \\ \frac{1}{12} & \frac{4}{12} & \frac{2}{12} & \frac{5}{12} \end{pmatrix}$ are equivalent, in the sense that they give rise to the same probability distribution over terminal nodes, namely

$$\begin{pmatrix} z_1 & z_2 & z_3 & z_4 & z_5 \\ \frac{5}{36} & \frac{10}{36} & \frac{7}{36} & \frac{4}{36} & \frac{10}{36} \end{pmatrix}.$$

¹ $P(x)$ denotes the probability of choice x for Player 1, according to the given behavioral strategy

Theorem 7.1.1 — Kuhn, 1953. In extensive forms *with perfect recall*, behavioral strategies and mixed strategies are equivalent, in the sense that, for every mixed strategy there is a behavioral strategy that gives rise to the same probability distribution over terminal nodes.^a

^aA more precise statement is as follows. Consider an extensive form with perfect recall and a Player i . Let x_{-i} be an arbitrary profile of strategies of the players other than i , where, for every $j \neq i$, x_j is either a mixed or a behavioral strategy of Player j . Then, for every mixed strategy σ_i of Player i there is a behavioral strategy b_i of Player i such that (σ_i, x_{-i}) and (b_i, x_{-i}) give rise to the same probability distribution over the set of terminal nodes.

Without perfect recall, Theorem 7.1.1 does not hold. To see this, consider the one-player extensive form shown in Figure 7.4 and the mixed strategy

$$\left(\begin{array}{cccc} (a, c) & (a, d) & (b, c) & (b, d) \\ \frac{1}{2} & 0 & 0 & \frac{1}{2} \end{array} \right)$$

which induces the probability distribution $\left(\begin{array}{cccc} z_1 & z_2 & z_3 & z_4 \\ \frac{1}{2} & 0 & 0 & \frac{1}{2} \end{array} \right)$ on the set of terminal nodes.




Figure 7.4: A one-player extensive frame without perfect recall.

Consider an arbitrary behavioral strategy

$$\left(\begin{array}{cc|cc} a & b & c & d \\ p & 1-p & q & 1-q \end{array} \right),$$

whose corresponding probability distribution over the set of terminal nodes is

$$\left(\begin{array}{cccc} z_1 & z_2 & z_3 & z_4 \\ pq & p(1-q) & (1-p)q & (1-p)(1-q) \end{array} \right).$$

In order to have $P(z_2) = 0$ it must be that either $p = 0$ or $q = 1$.

If $p = 0$ then $P(z_1) = 0$ and if $q = 1$ then $P(z_4) = 0$.

Thus the probability distribution

$$\left(\begin{array}{cccc} z_1 & z_2 & z_3 & z_4 \\ \frac{1}{2} & 0 & 0 & \frac{1}{2} \end{array} \right)$$


cannot be achieved with a behavioral strategy.

Since the focus of this book is on extensive-form games with perfect recall, by appealing to Theorem 7.1.1, from now on we can restrict attention to behavioral strategies.

As usual, one goes from a frame to a game by adding preferences over outcomes. Let O be the set of basic outcomes (recall that with every terminal node is associated a basic outcome) and $\mathcal{L}(O)$ the set of lotteries (probability distributions) over O .

Definition 7.1.2 An *extensive-form game with cardinal payoffs* is an extensive frame (with, possibly, chance moves) together with a von Neumann-Morgenstern ranking \succsim_i of the set of lotteries $\mathcal{L}(O)$, for every Player i .

As usual, it is convenient to represent a von Neumann-Morgenstern ranking by means of a von Neumann-Morgenstern utility function and replace the outcomes with a vector of utilities, one for every player.



For example, consider the extensive form above, which reproduces Figure 7.2, where the set of basic outcomes is $O = \{o_1, o_2, o_3, o_4, o_5\}$ and suppose that Player 1 has a von Neumann-Morgenstern ranking of $\mathcal{L}(O)$ that is represented by the following von Neumann-Morgenstern utility function:

outcome:	o_1	o_2	o_3	o_4	o_5	
	$U_1:$	5	2	0	1	3

Suppose also that Player 2 has preferences represented by the von Neumann-Morgenstern utility function

outcome :	o_1	o_2	o_3	o_4	o_5	
	$U_2:$	3	6	4	5	0

Then from the extensive frame of Figure 7.2 we obtain the extensive-form game with cardinal payoffs shown in Figure 7.5.

Since the expected utility of lottery $\begin{pmatrix} o_1 & o_2 \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}$ is 4 for both players, and the expected utility of lottery $\begin{pmatrix} o_1 & o_3 & o_4 \\ \frac{1}{5} & \frac{3}{5} & \frac{1}{5} \end{pmatrix}$ is 1.2 for Player 1 and 4 for Player 2, we can simplify the game by replacing the first move of Nature with the payoff vector (4,4) and the second move of Nature with the payoff vector (1.2, 4). The simplified game is shown in Figure 7.6.




Figure 7.5: An extensive game based on the frame of Figure 7.2. The terminal nodes have not been labeled. The o_i 's are basic outcomes.




Figure 7.6: A simplified version of the game of Figure 7.5. The z_i 's are terminal nodes. Note that this is a game based on the frame of Figure 7.3.

Given an extensive game with cardinal payoffs, associated with every behavioral strategy profile is a lottery over basic outcomes and thus – using a von Neumann-Morgenstern utility function for each player – a payoff for each player. For example, the behavioral strategy profile

$$\left[\left(\begin{array}{cc|cc} a & b & e & f \\ \frac{5}{12} & \frac{7}{12} & \frac{2}{7} & \frac{5}{7} \end{array} \right), \left(\begin{array}{cc} c & d \\ \frac{1}{3} & \frac{2}{3} \end{array} \right) \right]$$

for the extensive game of Figure 7.5 gives rise to the lottery

$$\left(\begin{array}{ccccc} o_1 & o_2 & o_3 & o_4 & o_5 \\ \frac{71}{540} & \frac{25}{540} & \frac{213}{540} & \frac{81}{540} & \frac{150}{540} \end{array} \right)$$

(for instance, the probability of basic outcome o_1 is calculated as follows:

$$P(o_1) = P(a)P(c)\frac{2}{3} + P(b)P(c)\frac{1}{3} = \frac{5}{12}\frac{1}{3}\frac{2}{3} + \frac{7}{12}\frac{1}{3}\frac{1}{3} = \frac{71}{540}.$$

Using the utility function postulated above for Player 1, namely

$$\begin{array}{l} \text{outcome : } o_1 \quad o_2 \quad o_3 \quad o_4 \quad o_5 \\ U_1 : \quad \quad 5 \quad 2 \quad 0 \quad 1 \quad 3 \end{array}$$

we get a corresponding payoff for Player 1 equal to

$$\frac{71}{540}5 + \frac{25}{540}2 + \frac{213}{540}0 + \frac{81}{540}1 + \frac{150}{540}3 = \frac{936}{540} = 1.733.$$

An alternative way of computing this payoff is by using the simplified game of Figure 7.6 where the behavioral strategy profile

$$\left[\left(\begin{array}{cc|cc} a & b & e & f \\ \frac{5}{12} & \frac{7}{12} & \frac{2}{7} & \frac{5}{7} \end{array} \right), \left(\begin{array}{cc} c & d \\ \frac{1}{3} & \frac{2}{3} \end{array} \right) \right]$$

yields the probability distribution over terminal nodes

$$\left(\begin{array}{ccccc} z_1 & z_2 & z_3 & z_4 & z_5 \\ \frac{5}{36} & \frac{10}{36} & \frac{7}{36} & \frac{4}{36} & \frac{10}{36} \end{array} \right),$$

which, in turn, yields the probability distribution

$$\left(\begin{array}{ccccc} 4 & 0 & 1.2 & 1 & 3 \\ \frac{5}{36} & \frac{10}{36} & \frac{7}{36} & \frac{4}{36} & \frac{10}{36} \end{array} \right)$$

over utilities for Player 1. From the latter we get that the expected payoff for Player 1 is

$$\frac{5}{36}4 + \frac{10}{36}0 + \frac{7}{36}1.2 + \frac{4}{36}1 + \frac{10}{36}3 = \frac{936}{540} = 1.733.$$

The calculations for Player 2 are similar (see Exercise 7.3).

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 7.4.1 at the end of this chapter.

7.2 Subgame-perfect equilibrium revisited

The notion of subgame-perfect equilibrium was introduced in Chapter 4 (Definition 4.4.1) for extensive-form games with ordinal payoffs.

When payoffs are *ordinal*, a subgame-perfect equilibrium may fail to exist because either the entire game or a proper subgame does not have any Nash equilibria.

In the case of finite extensive-form games with *cardinal* payoffs, a subgame-perfect equilibrium always exists, because – by Nash’s theorem (Theorem 6.2.1, Chapter 6) – every finite game has at least one Nash equilibrium in mixed strategies.

Thus, in the case of cardinal payoffs, the subgame-perfect equilibrium algorithm (Definition 4.4.2, Chapter 4) never halts and the output of the algorithm is a subgame-perfect equilibrium.

We shall illustrate this with the extensive-form game shown in Figure 7.7.




Figure 7.7: An extensive-form game with cardinal payoffs.

Let us apply the subgame-perfect equilibrium algorithm to this game. We start with the proper subgame that begins at Player 2's decision node on the left, whose strategic form is shown in Figure 7.8. Note that this subgame has no pure-strategy Nash equilibria. Thus if payoffs were merely ordinal payoffs the algorithm would halt and we would conclude that the game of Figure 7.7 has no subgame-perfect equilibria. However, we will assume that payoffs are cardinal (that is, that they are von Neumann-Morgenstern utilities).

		Player 3			
		<i>C</i>		<i>D</i>	
Player 2	<i>A</i>	3	1	0	2
	<i>B</i>	0	3	1	2

Figure 7.8: The strategic form of the proper subgame on the left in the game of Figure 7.7.

To find the mixed-strategy Nash equilibrium of the game of Figure 7.8, let p be the probability of A and q the probability of C .

- Then we need q to be such that $3q = 1 - q$, that is, $q = \frac{1}{4}$,
- and p to be such that $p + 3(1 - p) = 2$, that is, $p = \frac{1}{2}$.

Thus the Nash equilibrium of this proper subgame is:

$$\left[\left(\begin{array}{cc} A & B \\ \frac{1}{2} & \frac{1}{2} \end{array} \right), \left(\begin{array}{cc} C & D \\ \frac{1}{4} & \frac{3}{4} \end{array} \right) \right],$$

yielding the following payoffs:

$$\text{for Player 1: } \frac{1}{2} \frac{1}{4} 1 + \frac{1}{2} \frac{3}{4} 2 + \frac{1}{2} \frac{1}{4} 2 + \frac{1}{2} \frac{3}{4} 0 = 1.125$$

$$\text{for Player 2: } \frac{1}{2} \frac{1}{4} 3 + \frac{1}{2} \frac{3}{4} 0 + \frac{1}{2} \frac{1}{4} 0 + \frac{1}{2} \frac{3}{4} 1 = 0.75$$

$$\text{for Player 3: } \frac{1}{2} \frac{1}{4} 1 + \frac{1}{2} \frac{3}{4} 2 + \frac{1}{2} \frac{1}{4} 3 + \frac{1}{2} \frac{3}{4} 2 = 2.$$

Thus we can simplify the game of Figure 7.7 as shown in Figure 7.9.




Figure 7.9: The game of Figure 7.7 after replacing the proper subgame on the left with the payoffs associated with its Nash equilibrium.

Now consider the proper subgame of the game of Figure 7.9 (the subgame that starts at Player 2’s node). Its strategic form is shown in Figure 7.10.

		Player 3	
		<i>G</i>	<i>H</i>
Player 2	<i>E</i>	0 3	1 2
	<i>F</i>	2 1	0 3

Figure 7.10: The strategic form of the proper subgame of the game of Figure 7.9.

Again, there is no pure-strategy Nash equilibrium. To find the mixed-strategy equilibrium let p be the probability of E and q the probability of G .

- Then we need q to be such that $1 - q = 2q$, that is, $q = \frac{1}{3}$,
- and p to be such that $3p + 1 - p = 2p + 3(1 - p)$, that is, $p = \frac{2}{3}$.
- Hence the Nash equilibrium is $\left[\left(\begin{matrix} E & F \\ \frac{2}{3} & \frac{1}{3} \end{matrix} \right), \left(\begin{matrix} G & H \\ \frac{1}{3} & \frac{2}{3} \end{matrix} \right) \right]$ yielding the following payoffs:

for Player 1: $\frac{2}{3} \left(\frac{1}{3} \right) (2) + \frac{2}{3} \left(\frac{2}{3} \right) (0) + \frac{1}{3} \left(\frac{1}{3} \right) (1) + \frac{1}{3} \left(\frac{2}{3} \right) (2) = 1$.

for Player 2: $\frac{2}{3} \left(\frac{1}{3} \right) (0) + \frac{2}{3} \left(\frac{2}{3} \right) (1) + \frac{1}{3} \left(\frac{1}{3} \right) (2) + \frac{1}{3} \left(\frac{2}{3} \right) (0) = 0.67$.

for Player 3: $\frac{2}{3} \left(\frac{1}{3} \right) (3) + \frac{2}{3} \left(\frac{2}{3} \right) (2) + \frac{1}{3} \left(\frac{1}{3} \right) (1) + \frac{1}{3} \left(\frac{2}{3} \right) (3) = 2.33$.

Thus we can simplify the game of Figure 7.9 as shown in Figure 7.11, where the optimal choice for Player 1 is L .

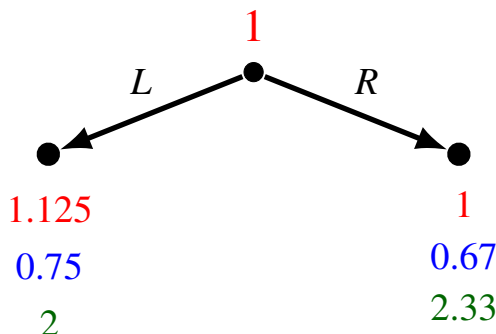


Figure 7.11: The game of Figure 7.9 after replacing the proper subgame with the payoffs associated with the Nash equilibrium.

Hence the subgame-perfect equilibrium of the game of Figure 7.7 (expressed in terms of behavioral strategies) is:

$$\left[\left(\begin{array}{cc} L & R \\ 1 & 0 \end{array} \right), \left(\begin{array}{cc|cc} A & B & E & F \\ \frac{1}{2} & \frac{1}{2} & \frac{2}{3} & \frac{1}{3} \end{array} \right), \left(\begin{array}{cc|cc} C & D & G & H \\ \frac{1}{4} & \frac{3}{4} & \frac{1}{3} & \frac{2}{3} \end{array} \right) \right]$$

We conclude this section with the following theorem, which is a corollary of Theorem 6.2.1 (Chapter 6).

Theorem 7.2.1 Every finite extensive-form game with cardinal payoffs has at least one subgame-perfect equilibrium in mixed strategies.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 7.4.2 at the end of this chapter.

7.3 Problems with the notion of subgame-perfect equilibrium

The notion of subgame-perfect equilibrium is a refinement of Nash equilibrium. As explained in Chapter 3, in the context of perfect-information games, the notion of subgame-perfect equilibrium eliminates some “unreasonable” Nash equilibria that involve incredible threats. However, not every subgame-perfect equilibrium can be viewed as a “rational solution”. To see this, consider the extensive-form game shown in Figure 7.12. This game has no proper subgames and thus the set of subgame-perfect equilibria coincides with the set of Nash equilibria. The pure-strategy Nash equilibria of this game are (a, f, c) , (a, e, c) , (b, e, c) and (b, f, d) . It can be argued that neither (a, f, c) nor (b, f, d) can be considered “rational solutions”.

Consider first the Nash equilibrium (a, f, c) . Player 2's plan to play f is rational only in the very limited sense that, given that Player 1 plays a , what Player 2 plans to do is irrelevant because it cannot affect anybody's payoff; thus f is as good as e . However, if we take Player 2's strategy as a "serious" plan specifying what Player 2 would actually do if she had to move, then – given that Player 3 plays c – e would give Player 2 a payoff of 2, while f would only give a payoff of 1. Thus e seems to be a better strategy than f , if Player 2 takes the contingency "seriously".




Figure 7.12: An extensive-form game showing the insufficiency of the notion of subgame-perfect equilibrium.

Consider now the Nash equilibrium (b, f, d) and focus on Player 3. As before, Player 3's plan to play d is rational only in the very limited sense that, given that Player 1 plays a and Player 2 plays f , what Player 3 plans to do is irrelevant, so that c is as good as d . However, if Player 3 did find himself having to play, it would not be rational for him to play d , since d is a strictly dominated choice: no matter whether he is making his choice at the left node or at the right node of his information set, c gives him a higher payoff than d . How can it be then that d can be part of a Nash equilibrium? The answer is that d is strictly dominated *conditional on Player 3's information set being reached* but not as a plan formulated before the play of the game starts. In other words, d is strictly dominated *as a choice but not as a strategy*.

The notion of subgame-perfect equilibrium is not strong enough to eliminate "unreasonable" Nash equilibria such as (a, f, c) and (b, f, d) in the game of Figure 7.12. In order to do that we will need a stronger notion. This issue is postponed to Part IV (Chapters 11-13).

7.4 Exercises

7.4.1 Exercises for section 7.1: Behavioral strategies in dynamic games

The answers to the following exercises are in Section 7.5 at the end of this chapter.

Exercise 7.1 What properties must an extensive-form frame satisfy in order for it to be the case that, for a given player, the set of mixed strategies coincides with the set of behavioral strategies? [Assume that there are at least two choices at every information set.] ■

Exercise 7.2 Suppose that, in a given extensive-form frame, Player 1 has four information sets: at one of them she has two choices and at each of the other three she has three choices.

- (a) How many parameters are needed to specify a mixed strategy of Player 1?
- (b) How many parameters are needed to specify a behavioral strategy of Player 1?

Exercise 7.3 From the behavioral strategy profile

$$\left[\left(\begin{array}{cc|cc} a & b & e & f \\ \frac{5}{12} & \frac{7}{12} & \frac{2}{7} & \frac{5}{7} \end{array} \right), \left(\begin{array}{cc} c & d \\ \frac{1}{3} & \frac{2}{3} \end{array} \right) \right]$$

calculate the payoff of Player 2 in two ways:

- (1) using the game of Figure 7.5 and
- (2) using the simplified game of Figure 7.6

Exercise 7.4 Consider the extensive form of Figure 7.13, where o_1, \dots, o_5 are basic outcomes. Player 1's ranking of O is

$$o_1 \succ_1 o_5 \succ_1 o_4 \succ_1 o_2 \sim_1 o_3;$$

furthermore, she is indifferent between o_5 and the lottery

$$\left(\begin{array}{ccc} o_1 & o_2 & o_3 \\ \frac{6}{8} & \frac{1}{8} & \frac{1}{8} \end{array} \right)$$

and is also indifferent between o_4 and the lottery

$$\left(\begin{array}{cc} o_2 & o_5 \\ \frac{2}{3} & \frac{1}{3} \end{array} \right).$$

Player 2's ranking of O is

$$o_1 \sim_2 o_2 \sim_2 o_4 \succ_2 o_3 \succ_2 o_5;$$

furthermore, he is indifferent between o_3 and the lottery

$$\left(\begin{array}{ccc} o_1 & o_2 & o_5 \\ \frac{1}{10} & \frac{1}{10} & \frac{8}{10} \end{array} \right).$$

Finally, Player 3's ranking of O is

$$o_2 \succ_3 o_4 \succ_3 o_3 \sim_3 o_5 \succ_3 o_1;$$

furthermore, she is indifferent between o_4 and the lottery

$$\left(\begin{array}{ccc} o_1 & o_2 & o_3 \\ \frac{1}{4} & \frac{1}{2} & \frac{1}{4} \end{array} \right)$$

and is also indifferent between o_3 and the lottery

$$\left(\begin{array}{cc} o_1 & o_2 \\ \frac{3}{5} & \frac{2}{5} \end{array} \right),$$

Write the corresponding extensive-form game. ■




Figure 7.13: The game for Exercise 7.4

7.4.2 Exercises for section 7.2: Subgame-perfect equilibrium revisited

The answers to the following exercises are in Section 7.5 at the end of this chapter.

Exercise 7.5 Consider the extensive-form game with cardinal payoffs shown in Figure 7.14 .

- Write the corresponding strategic-form game and find all the pure-strategy Nash equilibria.
- Find the subgame-perfect equilibrium.




Figure 7.14: The game for Exercise 7.5

Exercise 7.6 Consider the extensive form shown in Figure 7.15 (where the basic outcomes are denoted by x_j instead of o_j , $j = 1, \dots, 10$). All the players satisfy the axioms of expected utility. They rank the outcomes as indicated below (as usual, if outcome w is above outcome y then w is strictly preferred to y , and if w and y are written next to each other then the player is indifferent between the two):

$$\text{Player 1: } \begin{pmatrix} x_7, x_9 \\ x_1, x_2, x_4, x_5 \\ x_{10} \\ x_3, x_6, x_8 \end{pmatrix} \quad \text{Player 2: } \begin{pmatrix} x_1, x_3 \\ x_4, x_5 \\ x_2, x_7, x_8 \\ x_6 \\ x_9 \end{pmatrix} \quad \text{Player 3: } \begin{pmatrix} x_2, x_7 \\ x_8 \\ x_1, x_4, x_9 \\ x_3, x_5, x_6 \end{pmatrix}$$

Furthermore, Player 2 is indifferent between x_4 and the lottery $\begin{pmatrix} x_1 & x_2 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$

and Player 3 is indifferent between x_1 and the lottery $\begin{pmatrix} x_2 & x_5 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$.

Although the above information is not sufficient to determine the von Neumann-Morgenstern utility functions of the players, it is sufficient to compute the subgame-perfect equilibrium. [Hint: apply the IDSDS procedure to the subgame.] Find the subgame-perfect equilibrium. ■




Figure 7.15: The game for Exercise 7.6

Exercise 7.7 Consider the extensive-form game shown in Figure 7.16.

- Write the corresponding strategic-form game.
- Find all the pure-strategy Nash equilibria.
- Find the mixed-strategy subgame-perfect equilibrium.




Figure 7.16: The game for Exercise 7.7

Exercise 7.8 Consider the extensive-form game shown in Figure 7.17.

- (a) Find all the pure-strategy Nash equilibria. Which ones are also subgame perfect?
- (b) (This is a more challenging question) Prove that there is no mixed-strategy Nash equilibrium where Player 1 plays M with probability strictly between 0 and 1.




Figure 7.17: The game for Exercise 7.8

Exercise 7.9 — * Challenging Question ***.**

You have to go to a part of town where many people have been mugged recently. You consider whether you should leave your wallet at home or carry it with you. Of the four possible outcomes, your most preferred one is having your wallet with you and not being mugged. Being mugged is a very unpleasant experience, so your second favorite alternative is not carrying your wallet and not being mugged (although not having any money with you can be very inconvenient). If, sadly enough, your destiny is to be mugged, then you prefer to have your wallet with you (possibly with not too much money in it!) because you don't want to have to deal with a frustrated mugger.

A typical potential mugger's favorite outcome is the one where you have your wallet with you and he mugs you. His least preferred outcome is the one where he attempts to mug you and you don't have your wallet with you (he risks being caught for nothing). He is indifferent to whether or not you are carrying your wallet if he decides not to mug you. Denote the possible outcomes as shown in Figure 7.18.

- (a) What is the ordinal ranking of the outcomes for each player?

Suppose that both players have von Neumann-Morgenstern utility functions. You are indifferent between the following lotteries:

$$L_1 = \begin{pmatrix} z_1 & z_2 & z_3 & z_4 \\ \frac{3}{20} & \frac{14}{20} & \frac{3}{20} & 0 \end{pmatrix} \quad \text{and} \quad L_2 = \begin{pmatrix} z_1 & z_2 & z_3 & z_4 \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} \end{pmatrix};$$

furthermore, you are indifferent between

$$L_3 = \begin{pmatrix} z_1 & z_2 & z_3 & z_4 \\ 0 & \frac{2}{3} & \frac{1}{3} & 0 \end{pmatrix} \quad \text{and} \quad L_4 = \begin{pmatrix} z_1 & z_2 & z_3 & z_4 \\ \frac{1}{2} & \frac{1}{2} & 0 & 0 \end{pmatrix}.$$

The potential mugger is indifferent between the two lotteries

$$L_5 = \begin{pmatrix} z_1 & z_2 & z_3 & z_4 \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \end{pmatrix} \quad \text{and} \quad L_6 = \begin{pmatrix} z_1 & z_2 & z_3 & z_4 \\ \frac{8}{128} & \frac{67}{128} & \frac{16}{128} & \frac{37}{128} \end{pmatrix}.$$

- (b) For each player find the normalized von Neumann-Morgenstern utility function.

You have to decide whether or not to leave your wallet at home. Suppose that, if you leave your wallet at home, with probability p (with $0 < p < 1$) the potential mugger will notice that your pockets are empty and with probability $(1 - p)$ he will not notice; in the latter case he will be uncertain as to whether you have your wallet with you or you don't. He will be in the same state of uncertainty if you did take your wallet with you.

- (c) Represent this situation as an extensive game with imperfect information.
 (d) Write the corresponding normal form.
 (e) Find all the subgame-perfect equilibria (including the mixed-strategy ones, if any). (Hint: your answer should distinguish between different values of p).

		Potential mugger	
		Not mug	Mug
You	Leave wallet at home	z_1	z_2
	Take wallet with you	z_3	z_4

Figure 7.18: The outcomes for Exercise 7.9

7.5 Solutions to exercises

Solution to Exercise 7.1 It must be the case that the player under consideration has only one information set. □

Solution to Exercise 7.2

- (a) 53. The number of pure strategies is $2 \times 3 \times 3 \times 3 = 54$ and thus 53 probabilities are needed to specify a mixed strategy.
- (b) 7: one probability for the information set where she has two choices and two probabilities for each of the other three information sets. □

Solution to Exercise 7.3

1. The induced probability distribution on basic outcomes is

$$\left(\begin{array}{ccccc} o_1 & o_2 & o_3 & o_4 & o_5 \\ \frac{71}{540} & \frac{25}{540} & \frac{213}{540} & \frac{81}{540} & \frac{150}{540} \end{array} \right).$$

Thus Player 2's expected utility is

$$\frac{71}{540}3 + \frac{25}{540}6 + \frac{213}{540}4 + \frac{81}{540}5 + \frac{150}{540}0 = \frac{1620}{540} = 3.$$

2. The induced probability distribution on terminal nodes is

$$\left(\begin{array}{ccccc} z_1 & z_2 & z_3 & z_4 & z_5 \\ \frac{5}{36} & \frac{10}{36} & \frac{7}{36} & \frac{4}{36} & \frac{10}{36} \end{array} \right).$$

Thus Player 2's expected payoff is

$$\frac{5}{36}4 + \frac{10}{36}4 + \frac{7}{36}4 + \frac{4}{36}5 + \frac{10}{36}0 = \frac{108}{36} = 3.$$

Not surprisingly, the same number as in Part 1. □

Solution to Exercise 7.4 The normalized von Neumann-Morgenstern utility functions are

$$\begin{pmatrix} & o_1 & o_2 & o_3 & o_4 & o_5 \\ U_1 & 1 & 0 & 0 & 0.25 & 0.75 \\ U_2 & 1 & 1 & 0.2 & 1 & 0 \\ U_3 & 0 & 1 & 0.4 & 0.6 & 0.4 \end{pmatrix}$$

Thus the extensive-form game is shown in Figure 7.19.

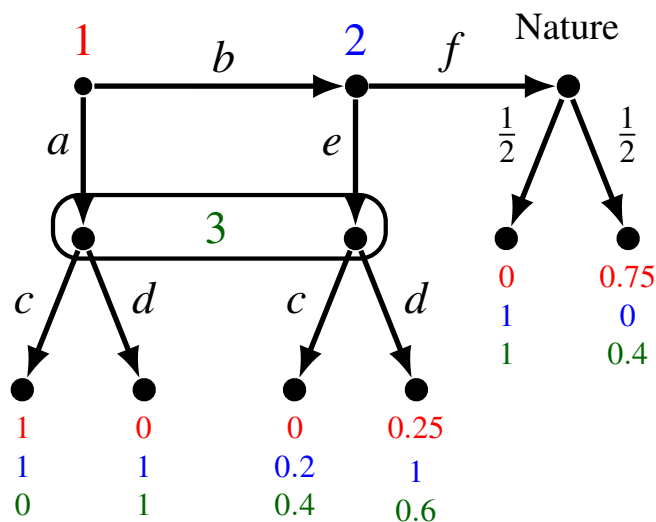


Figure 7.19: The game for Exercise 7.4

Or, in a simplified form obtained by removing the move of Nature, as shown in Figure 7.20. \square




Figure 7.20: The simplified game for Exercise 7.4

Solution to Exercise 7.5

- (a) The strategic form is shown in Figure 7.21. The pure-strategy Nash equilibria are (UL, r) and (UR, r) .

		Player 2			
		<i>l</i>		<i>r</i>	
Player 1	<i>UL</i>	1	5	1	5
	<i>UR</i>	1	5	1	5
	<i>DL</i>	3	2	0	3
	<i>DR</i>	2	4	1	1

Figure 7.21: The strategic form for Part (a) of Exercise 7.5

- (b) The strategic form of the proper subgame that starts at Player 2’s node is as follows:

		Player 2			
		<i>l</i>	<i>r</i>		
Player 1	<i>L</i>	3	2	0	3
	<i>R</i>	2	4	1	1

This game has a unique mixed-strategy Nash equilibrium given by

$$\left[\left(\begin{matrix} L & R \\ \frac{3}{4} & \frac{1}{4} \end{matrix} \right), \left(\begin{matrix} l & r \\ \frac{1}{2} & \frac{1}{2} \end{matrix} \right) \right], \text{ yielding Player 1 an expected payoff of 1.5.}$$

Thus the unique subgame-perfect equilibrium, expressed as a behavioral-strategy profile, is

$$\left[\left(\begin{matrix} U & D & | & L & R \\ 0 & 1 & | & \frac{3}{4} & \frac{1}{4} \end{matrix} \right), \left(\begin{matrix} l & r \\ \frac{1}{2} & \frac{1}{2} \end{matrix} \right) \right]$$

or, expressed as a mixed-strategy profile,

$$\left[\left(\begin{matrix} UL & UR & DL & DR \\ 0 & 0 & \frac{3}{4} & \frac{1}{4} \end{matrix} \right), \left(\begin{matrix} l & r \\ \frac{1}{2} & \frac{1}{2} \end{matrix} \right) \right].$$

□

Solution to Exercise 7.6 There is only one proper subgame starting from Player 2's node; its strategic-form frame is as follows:

		Player 3		
		<i>d</i>	<i>e</i>	<i>f</i>
Player 2	<i>a</i>	x_1	x_2	x_3
	<i>b</i>	x_4	x_5	x_6
	<i>c</i>	x_7	x_8	x_9

For Player 2 strategy *c* is strictly dominated by strategy *b* (she prefers x_4 to x_7 , and x_5 to x_8 and x_6 to x_9) and for Player 3 strategy *f* is strictly dominated by strategy *d* (she prefers x_1 to x_3 , and x_4 to x_6 and x_7 to x_9). Thus we can simplify the game as follows:

		Player 3	
		<i>d</i>	<i>e</i>
Player 2	<i>a</i>	x_1	x_2
	<i>b</i>	x_4	x_5

Restricted to these outcomes the payers' rankings are:

$$\text{Player 2 : } \begin{pmatrix} x_1 \\ x_4, x_5 \\ x_2 \end{pmatrix} \qquad \text{Player 3 : } \begin{pmatrix} x_2 \\ x_1, x_4 \\ x_5 \end{pmatrix}.$$

Let U be Player 2's von Neumann-Morgenstern utility function. We can set $U(x_1) = 1$ and $U(x_2) = 0$. Thus, since she is indifferent between x_4 and x_5 and also between x_4 and the lottery $\begin{pmatrix} x_1 & x_2 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$, $U(x_4) = U(x_5) = \frac{1}{2}$.

Let V be Player 3's von Neumann-Morgenstern utility function. We can set $V(x_2) = 1$ and $V(x_5) = 0$. Thus, since she is indifferent between x_1 and x_4 and also between x_1 and the lottery $\begin{pmatrix} x_2 & x_5 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$, $V(x_1) = V(x_4) = \frac{1}{2}$.

Hence the above game-frame becomes the following game:

		Player 3	
		<i>d</i>	<i>e</i>
Player 2	<i>a</i>	$1, \frac{1}{2}$	$0, 1$
	<i>b</i>	$\frac{1}{2}, \frac{1}{2}$	$\frac{1}{2}, 0$

There is no pure-strategy Nash equilibrium. Let p be the probability of *a* and q the probability of *d*. Then for a Nash equilibrium we need $q = \frac{1}{2}$ and $p = \frac{1}{2}$.

Hence in the subgame the outcome will be $\begin{pmatrix} x_1 & x_2 & x_4 & x_5 \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \end{pmatrix}$.

Since all of these outcomes are better than x_{10} for Player 1, Player 1 will play d . Thus the subgame-perfect equilibrium is

$$\left[\left(\begin{matrix} d & u \\ 1 & 0 \end{matrix} \right), \left(\begin{matrix} a & b & c \\ \frac{1}{2} & \frac{1}{2} & 0 \end{matrix} \right), \left(\begin{matrix} d & e & f \\ \frac{1}{2} & \frac{1}{2} & 0 \end{matrix} \right) \right].$$

□

Solution to Exercise 7.7

(a) The strategic form is shown in Figure 7.22.

		Player 2							
		AC		AD		BC		BD	
Player 1	LEG	2	0	2	0	0	2	0	2
	LEH	2	0	2	0	0	2	0	2
	LFG	0	6	0	6	4	1	4	1
	LFH	0	6	0	6	4	1	4	1
	REG	1	4	4	3	1	4	4	3
	REH	2	0	1	2	2	0	1	2
	RFG	1	4	4	3	1	4	4	3
	RFH	2	0	1	2	2	0	1	2

Figure 7.22: The strategic form for Exercise 7.7

(b) There are no pure-strategy Nash equilibria.

(c) First let us solve the subgame on the left, whose strategic form is as follows:

		Player 2	
		A	B
Player 1	E	2, 0	0, 2
	F	0, 6	4, 1

There is no pure-strategy Nash equilibrium. Let us find the mixed-strategy equilibrium. Let p be the probability assigned to E and q the probability assigned to A .

- Then p must be the solution to $6(1 - p) = 2p + (1 - p)$ and q must be the solution to $2q = 4(1 - q)$.

- Thus $p = \frac{5}{7}$ and $q = \frac{2}{3}$.

- The expected payoff of Player 1 is $\frac{4}{3} = 1.33$, while the expected payoff of player 2 is $\frac{12}{7} = 1.714$.

Next we solve the subgame on the right, whose strategic form is as follows:

		Player 2	
		C	D
Player 1	G	1, 4	4, 3
	H	2, 0	1, 2

There is no pure-strategy Nash equilibrium. Let us find the mixed-strategy equilibrium.

- Let p be the probability assigned to G and q the probability assigned to C .
- Then p must be the solution to $4p = 3p + 2(1 - p)$ and q must be the solution to $q + 4(1 - q) = 2q + (1 - q)$.
- Thus $p = \frac{2}{3}$ and $q = \frac{3}{4}$.
- The expected payoff of Player 1 is $\frac{7}{4} = 1.75$. Thus the game reduces to the one shown in Figure 7.23, where the optimal choice is R . Hence the subgame-perfect equilibrium is:

$$\left[\left(\begin{array}{cc|cc|cc} L & R & E & F & G & H \\ \hline 0 & 1 & \frac{5}{7} & \frac{2}{7} & \frac{2}{3} & \frac{1}{3} \end{array} \right), \left(\begin{array}{cc|cc} A & B & C & D \\ \hline \frac{2}{3} & \frac{1}{3} & \frac{3}{4} & \frac{1}{4} \end{array} \right) \right] \quad \square$$




Figure 7.23: The reduced game after eliminating the proper subgames

Solution to Exercise 7.8

(a) The strategic form is shown in Figure 7.24.

		Player 2					
		l			r		
Player 1	R	1	0	0	1	0	0
	M	0	1	1	2	2	1
	L	1	1	0	0	0	0
		Player 3: a					

		Player 2					
		l			r		
Player 1	R	1	0	0	1	0	0
	M	0	1	0	0	0	0
	L	1	1	0	0	0	0
		Player 3: b					

Figure 7.24: The strategic form for Exercise 7.8

The pure-strategy Nash equilibria are highlighted: (R, l, a) , (M, r, a) , (L, l, a) , (R, l, b) , (R, r, b) and (L, l, b) . They are all subgame perfect because there are no proper subgames.

- (b) Since, for Player 3, a strictly dominates b , conditional on his information set being reached, he will have to play a if his information set is reached with positive probability. Now, Player 3's information set is reached with positive probability if and only if Player 1 plays M with positive probability. Thus when $P(M) > 0$ the game essentially reduces to the one shown in Figure 7.25.

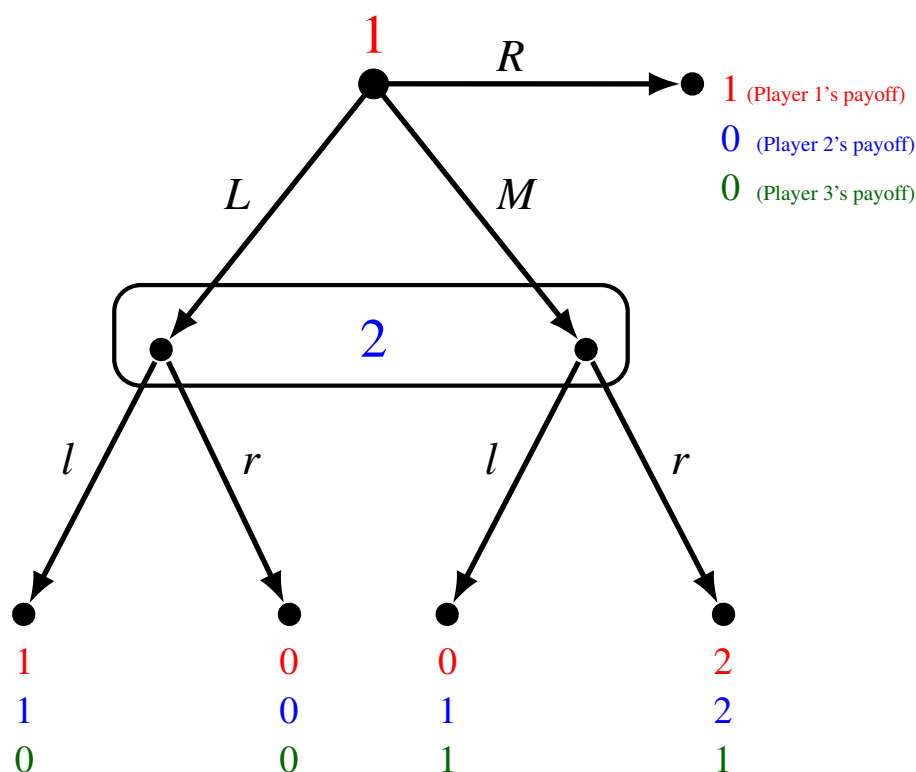


Figure 7.25: The extensive-form game for Part (b) of Exercise 7.8

Now, in order for Player 1 to be willing to assign positive probability to M he must expect a payoff of at least 1 (otherwise R would be better) and the only way he can expect a payoff of at least 1 is if Player 2 plays r with probability at least $\frac{1}{2}$.

- If Player 2 plays r with probability greater than $\frac{1}{2}$, then M gives Player 1 a higher payoff than both L and R and thus he will choose M with probability 1, in which case Player 2 will choose r with probability 1 (and Player 3 will choose a with probability 1) and so we get the pure strategy equilibrium (M, r, a) .
- If Player 2 plays r with probability exactly $\frac{1}{2}$ then Player 1 is indifferent between M and R (and can mix between the two), but finds L inferior and must give it probability 0. But then Player 2's best reply to a mixed strategy of Player 1 that assigns positive probability to M and R and zero probability to L is to play r with probability 1 (if his information set is reached it can only be reached at node x_2).
- Thus there cannot be a mixed-strategy equilibrium where Player 1 assigns to M probability p with $0 < p < 1$: it must be either $Pr(M) = 0$ or $Pr(M) = 1$. \square

Solution to Exercise 7.9

(a) The rankings are as follows:

$$\text{You: } \begin{pmatrix} \text{best} & z_3 \\ & z_1 \\ & z_4 \\ \text{worst} & z_2 \end{pmatrix}, \quad \text{Potential Mugger: } \begin{pmatrix} \text{best} & z_4 \\ & z_1, z_3 \\ \text{worst} & z_2 \end{pmatrix}$$

(b) Let U be your utility function. Let $U(z_3) = 1, U(z_1) = a, U(z_4) = b$ and $U(z_2) = 0$, with $0 < b < a < 1$. The expected utilities are as follows: $EU(L_1) = \frac{3}{20}a + \frac{3}{20}$, $EU(L_2) = \frac{1}{2}b$, $EU(L_3) = \frac{1}{3}$ and $EU(L_4) = \frac{1}{2}a$.

From $EU(L_3) = E(L_4)$ we get that $a = \frac{2}{3}$.

Substituting this into the equation $EU(L_1) = EU(L_2)$ gives $b = \frac{1}{2}$.

Thus $U(z_3) = 1$, $U(z_1) = \frac{2}{3}$, $U(z_4) = \frac{1}{2}$ and $U(z_2) = 0$.

Let V be the mugger's utility function. Let $V(z_4) = 1, V(z_1) = V(z_3) = c$ and $V(z_2) = 0$ with $0 < c < 1$. The expected utilities are as follows: $EV(L_5) = \frac{1}{4}(2c + 1)$ and $EV(L_6) = \frac{1}{128}(24c + 37)$.

Solving $EV(L_5) = EV(L_6)$ gives $c = \frac{1}{8}$.

Thus, $V(z_4) = 1, V(z_1) = V(z_3) = \frac{1}{8}$ and $V(z_2) = 0$.

(c) The extensive game is shown in Figure:7.26.

(d) The strategic form is shown in Figure 7.27 (for the mugger's strategy the first item refers to the left node, the second item to the information set on the right).

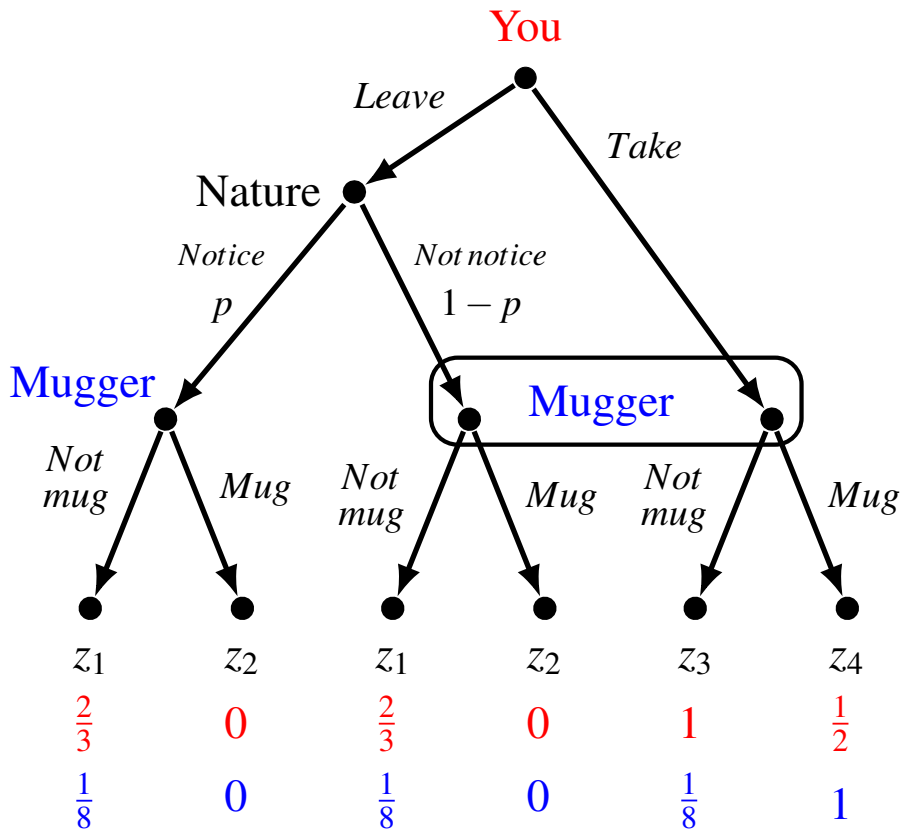


Figure 7.26: The extensive game for Exercise 7.9

Potential Mugger

		<i>NN</i>		<i>NM</i>		<i>MN</i>		<i>MM</i>	
You	<i>L</i>	$\frac{2}{3}$	$\frac{1}{8}$	$\frac{2}{3}p$	$\frac{1}{8}p$	$\frac{2}{3}(1-p)$	$\frac{1}{8}(1-p)$	0	0
	<i>T</i>	1	$\frac{1}{8}$	$\frac{1}{2}$	1	1	$\frac{1}{8}$	$\frac{1}{2}$	1

Figure 7.27: The strategic form for the game of Figure 7.26

- (e) At a subgame-perfect equilibrium the mugger will choose not to mug when he notices your empty pockets. Thus the normal form can be simplified as shown in Figure 7.28.

		Potential Mugger			
		NN		NM	
You	L	$\frac{2}{3}$	$\frac{1}{8}$	$\frac{2}{3}p$	$\frac{1}{8}p$
	T	1	$\frac{1}{8}$	$\frac{1}{2}$	1

Figure 7.28: The reduced game for Exercise 7.9

Thus,

- If $p < \frac{3}{4}$ then Take is a strictly dominant strategy for you and therefore there is a unique subgame-perfect equilibrium given by (Take, (Not mug, Mug)).
- If $p = \frac{3}{4}$ then there is a continuum of equilibria where the Mugger chooses (Not mug, Mug) with probability 1 and you choose L with probability q and T with probability $(1 - q)$ for any q with $0 \leq q \leq \frac{28}{29}$, obtained from the following condition about the Potential Mugger:

$$\underbrace{\frac{3}{32}q + 1 - q}_{\text{expected payoff from playing NM}} \geq \underbrace{\frac{1}{8}}_{\text{payoff from playing NN}}$$


- If $p > \frac{3}{4}$ then there is no pure-strategy subgame-perfect equilibrium. Let q be the probability that you choose L and r the probability that the mugger chooses NN. Then the unique mixed strategy equilibrium is given by the solution to:

$$\frac{2}{3}r + \frac{2}{3}p(1 - r) = r + \frac{1}{2}(1 - r) \quad \text{and} \quad \frac{1}{8} = \frac{1}{8}pq + (1 - q)$$

which is $q = \frac{7}{8-p}$ and $r = \frac{4p-3}{4p-1}$. Thus the unique subgame-perfect equilibrium is:

$$\left(\begin{array}{cc|cccc} L & T & NN & NM & MN & MM \\ \frac{7}{8-p} & \frac{1-p}{8-p} & \frac{4p-3}{4p-1} & \frac{2}{4p-1} & 0 & 0 \end{array} \right).$$

□



Knowledge, Common Knowledge, Beliefs

8	Common Knowledge	261
8.1	Individual knowledge	
8.2	Interactive knowledge	
8.3	Common knowledge	
8.4	Exercises	
8.5	Solutions to Exercises	
9	Adding Beliefs to Knowledge	295
9.1	Sets and probability: brief review	
9.2	Probabilistic beliefs	
9.3	Conditional probability and Bayes' rule	
9.4	Changing beliefs in response to new information	
9.5	Harsanyi consistency of beliefs or like-mindedness	
9.6	Agreeing to disagree	
9.7	Proof of the Agreement Theorem	
9.8	Exercises	
9.9	Solutions to Exercises	
10	Rationality	347
10.1	Models of strategic-form games	
10.2	Common knowledge of rationality in strategic-form games	
10.3	Common knowledge of rationality in extensive-form games	
10.4	Proofs of Theorems	
10.5	Exercises	
10.6	Solutions to Exercises	

8. Common Knowledge

8.1 Individual knowledge

In extensive-form games with imperfect information we use information sets to represent what the players know about past choices (when it is their turn to move). An information set of Player i is a collection of nodes in the tree where it is Player i 's turn to move and the interpretation is that Player i knows that she is making her choice at one of those nodes, but she does not know which of these nodes has actually been reached. In this chapter we extend the notion of information set to more general settings.

We start with an example. After listening to her patient's symptoms, a doctor reaches the conclusion that there are only five possible causes: (1) a bacterial infection, (2) a viral infection, (3) an allergic reaction to a drug, (4) an allergic reaction to food and (5) environmental factors. The doctor decides to do a lab test. If the lab test turns out to be positive then the doctor will be able to rule out causes (3)-(5), while a negative lab test will be an indication that causes (1) and (2) can be ruled out. To represent the doctor's possible states of information and knowledge we can use five states: a, b, c, d , and e . Each state represents a possible cause, as shown in Figure 8.1.




Figure 8.1: The information provided by the lab test.

We can partition the set $\{a, b, c, d, e\}$ into two sets: the set $\{a, b\}$, representing the state of knowledge of the doctor if she is informed that the lab test is positive, and the set $\{c, d, e\}$, representing the state of knowledge of the doctor if she is informed that the lab test is negative.

Consider the proposition “the cause of the patient’s symptoms is either an infection or environmental factors”. We can think of this proposition as the set of states $\{a, b, e\}$ where the proposition is in fact true; furthermore, we can ask the question “after receiving the result of the lab test, at which state would the doctor know the proposition represented by the set $\{a, b, e\}$?”

- If the true state is a , then – after viewing the result of the lab test – the doctor will think that it is possible that the state is either a or b and thus know that the cause of the patient’s symptoms is an infection (hence she will also know the weaker proposition that the cause is either an infection or environmental factors); the same is true if the true state is b .
- On the other hand, if the true state is e then the doctor will consider c , d and e as possibilities and thus will not be able to claim that she knows that the cause of the patient’s symptoms is either an infection or environmental factors.

Hence the answer to the question “after receiving the result of the lab test, at which state would the doctor know the proposition represented by the set $\{a, b, e\}$?” is “at states a and b only”.

We can now turn to the general definitions.

Definition 8.1.1 Let W be a finite set of *states*, where each state is to be understood as a complete specification of the relevant *facts* about the world. An *information partition* is a partition \mathcal{I} of W (that is, a collection of subsets of W that (1) are pairwise disjoint and (2) whose union covers the entire set W); the elements of the partition are called *information sets*. For every $w \in W$ we denote by $I(w)$ the information set that contains state w .

In the example of the doctor, $W = \{a, b, c, d, e\}$ and $\mathcal{I} = \{\{a, b\}, \{c, d, e\}\}$; furthermore, $I(a) = I(b) = \{a, b\}$ and $I(c) = I(d) = I(e) = \{c, d, e\}$.

Definition 8.1.2 Let W be a set of states. We will call the subsets of W events. Let \mathcal{I} be a partition of W , E an event (thus $E \subseteq W$) and $w \in W$ a state. We say that *at w the agent knows E* if and only if the information set to which w belongs is contained in E , that is, if and only if $I(w) \subseteq E$.

In the example of the doctor, where $W = \{a, b, c, d, e\}$ and $\mathcal{I} = \{\{a, b\}, \{c, d, e\}\}$, let $E = \{a, b, d, e\}$; then at a and b the doctor knows E because $I(a) = I(b) = \{a, b\} \subseteq E$, but at d it is not true that the doctor knows E because $I(d) = \{c, d, e\} \not\subseteq E$ (since $c \in I(d)$ but $c \notin E$) and, for the same reason, also at c and e it is not true that the doctor knows E .

Note that it is possible that there is no state where the agent knows a given event. In the doctor example, if we consider event $F = \{a, c\}$ then there is no state where the doctor knows F .

Definition 8.1.3 Using Definition 8.1.2, we can define a *knowledge operator* K on events that, given as input any event E , produces as output the event KE defined as the set of states at which the agent knows E .

Let 2^W denote the set of events, that is the set of subsets of W .^a Then the knowledge operator is the function $K : 2^W \rightarrow 2^W$ defined as follows: for every $E \subseteq W$, $KE = \{w \in W : I(w) \subseteq E\}$.

^aIf W contains n elements, then there are 2^n subsets of W , hence the notation 2^W . For example, if $W = \{a, b, c\}$ then $2^W = \{\emptyset, \{a\}, \{b\}, \{c\}, \{a, b\}, \{a, c\}, \{b, c\}, \{a, b, c\}\}$. (Recall that \emptyset denotes the empty set.)

In the example of the doctor, where $W = \{a, b, c, d, e\}$ and $\mathcal{I} = \{\{a, b\}, \{c, d, e\}\}$, let $E = \{a, b, d, e\}$ and $F = \{a, c\}$; then $KE = \{a, b\}$ and $KF = \emptyset$.

Given an event $G \subseteq W$ we denote by $\neg G$ the complement of G , that is, the set of states that are not in G . For example, if $W = \{a, b, c, d, e\}$ and $G = \{a, b, d\}$ then $\neg G = \{c, e\}$. Thus while KG is the event that the agent knows G , $\neg KG$ is the event that the agent does not know G .

Note the important difference between event $\neg KG$ (the agent does not know G) and event $K\neg G$ (the agent knows that it is not the case that G):

- if $w \in \neg KG$ then at state w the agent does not know G but she might not know $\neg G$ either, that is, it may be that she considers G possible ($I(w) \cap G \neq \emptyset$) and she also considers $\neg G$ possible ($I(w) \cap \neg G \neq \emptyset$).¹
- On the other hand, if $w \in K\neg G$ then at state w the agent knows that G is not true, because every state that she considers possible is in $\neg G$ ($I(w) \subseteq \neg G$).

Thus $K\neg G \subseteq \neg KG$ but the converse inclusion does not hold.

In the example of the doctor, where $W = \{a, b, c, d, e\}$ and $\mathcal{I} = \{\{a, b\}, \{c, d, e\}\}$, again let $E = \{a, b, d, e\}$ (so that $\neg E = \{c\}$) and $F = \{a, c\}$ (so that $\neg F = \{b, d, e\}$); then $KE = \{a, b\}$, $\neg KE = \{c, d, e\}$, $K\neg E = \emptyset$, $KF = \emptyset$, $\neg KF = W$ and $K\neg F = \emptyset$. Note the interesting fact that, since we can apply the knowledge operator to any event, we can also compute the event that the agent knows that she knows E (that is, the event $K(KE)$, which we will denote more simply as KKE) and the event that the agent knows that she does not know E (that is, the event $K\neg KE$).

Continuing the example of the doctor, where $W = \{a, b, c, d, e\}$, $\mathcal{I} = \{\{a, b\}, \{c, d, e\}\}$ and $E = \{a, b, d, e\}$, KKE is the set of states where the agent knows event $KE = \{a, b\}$; thus $KKE = \{a, b\}$. Furthermore, since $\neg KE = \{c, d, e\}$, $K\neg KE = \{c, d, e\}$. As noted in the following remark, this is not a coincidence.

R The knowledge operator $K : 2^W \rightarrow 2^W$ satisfies the following properties (which you are asked to prove in Exercise 8.5): for every event $E \subseteq W$,

- $KE \subseteq E$
- $KKE = KE$
- $K\neg KE = \neg KE$.

¹In the example of the doctor, where $W = \{a, b, c, d, e\}$ and $\mathcal{I} = \{\{a, b\}, \{c, d, e\}\}$, if $F = \{a, c\}$ (so that $\neg F = \{b, d, e\}$) then $KF = \emptyset$ and $K\neg F = \emptyset$; for instance, if the true state is a then the doctor considers F possible (because her information set is $I(a) = \{a, b\}$ and $I(a) \cap F = \{a\} \neq \emptyset$) but she also considers $\neg F$ possible (because $I(a) \cap \neg F = \{b\} \neq \emptyset$).

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 8.4.1 at the end of this chapter.

8.2 Interactive knowledge

We can now extend our analysis to the case of several agents and talk about not only what an individual knows about relevant facts but also about what she knows about what other individuals know and what they know about what she knows, etc. There is an entertaining episode of the TV series *Friends* in which Phoebe and Rachel reason about whether Chandler and Monica know that they (= Phoebe and Rachel) know that Chandler and Monica are having an affair: see <https://www.youtube.com/watch?v=LUN2YN0b0i8> (or search for the string ‘Friends-They Don’t Know That We Know They Know We Know’).

Again we start with a set of states W , where each state represents a complete description of the relevant *facts*. Let there be n individuals. To represent the possible states of mind of each individual we use an information partition: \mathcal{S}_i denotes the partition of individual $i \in 1, \dots, n$. As before, we call the subsets of W events. Using Definition 8.1.3 we can define a knowledge operator for every individual.

Let K_i be the knowledge operator of individual i ; thus, for every event $E \subseteq W$, $K_i E = \{w \in W : I_i(w) \subseteq E\}$. Now consider an event E and an individual, say Individual 1; since $K_1 E$ is an event (it is the set of states where Individual 1 knows event E), we can compute the event $K_2 K_1 E$, which is the event that Individual 2 knows event $K_1 E$, that is, the event that 2 knows that 1 knows E .

But there is no need to stop there: we can also compute the event $K_3 K_2 K_1 E$ (the event that 3 knows that 2 knows that 1 knows E) and the event $K_1 K_3 K_2 K_1 E$ (the event that 1 knows that 3 knows that 2 knows that 1 knows E), etc. A few examples will be useful.

We begin with an abstract example, without interpreting the states in terms of specific facts. Let the set of states be $W = \{a, b, c, d, e, f, g, h\}$. There are three individuals, Ann, Bob and Carol, with the information partitions shown in Figure 8.2.




Figure 8.2: The information partitions of three individuals.

Consider the event $E = \{a, b, c, f, g\}$. Let us compute the following events:

1. $K_{Ann}E$ (Ann knows E),
2. $K_{Bob}E$ (Bob knows E),
3. $K_{Carol}E$ (Carol knows E),
4. $K_{Carol}K_{Ann}E$ (Carol knows that Ann knows E),
5. $K_{Bob}K_{Carol}K_{Ann}E$ (Bob knows that Carol knows that Ann knows E),
6. $K_{Ann}\neg K_{Bob}K_{Carol}E$ (Ann knows that Bob does **not** know that Carol knows E).

All we need to do is apply Definition 8.1.3. First of all,

$$1. K_{Ann}E = \{a, b, c\}$$

(for example, $b \in K_{Ann}E$ because $I_{Ann}(b) = \{b, c\}$ and $\{b, c\} \subseteq E$,

while $f \notin K_{Ann}E$ because $I_{Ann}(f) = \{e, f, g\}$ and $\{e, f, g\}$ is not a subset of E).

Similarly,

$$2. K_{Bob}E = \{a, b, f\} \quad \text{and} \quad 3. K_{Carol}E = \{b, c, f, g\}.$$

To compute $K_{Carol}K_{Ann}E$ we need to find the set of states where Carol knows event $\{a, b, c\}$, since $K_{Ann}E = \{a, b, c\}$. Thus,

$$4. K_{Carol} \underbrace{K_{Ann}E}_{=\{a,b,c\}} = \{b, c\}.$$

Hence

$$5. K_{Bob} \underbrace{K_{Carol}K_{Ann}E}_{=\{b,c\}} = \emptyset,$$

that is, there is no state where Bob knows that Carol knows that Ann knows E .

To compute $K_{Ann}\neg K_{Bob}K_{Carol}E$ first we start with $K_{Carol}E$, which we have already computed: $K_{Carol}E = \{b, c, f, g\}$; then we compute $K_{Bob} \underbrace{K_{Carol}E}_{=\{b,c,f,g\}}$, which is $\{f\}$: $K_{Bob}K_{Carol}E =$

$\{f\}$; then we take the complement of this: $\neg K_{Bob}K_{Carol}E = \{a, b, c, d, e, g, h\}$ and finally we compute K_{Ann} of this event:

$$6. K_{Ann} \underbrace{\neg K_{Bob}K_{Carol}E}_{=\{a,b,c,d,e,g,h\}} = \{a, b, c, d, h\}.$$

Thus, for example, at state a it is true that Ann knows that Bob does not know that Carol knows E , while at state e this is not true.

Next we discuss a concrete example. The professor in a game theory class calls three students to the front of the room, shows them a large transparent box that contains many hats, some red and some white (there are no other colors) and tells them the following:

“I will blindfold you and put a hat on each of you, then I will remove the box from the room and, after that, I will remove your blindfolds, so that each of you can see the color(s) of the hats worn by the other two students, but you will not be able to see the color of your own hat. Then I will ask you questions starting with Student 1 then Student 2 then Student 3 then back to Student 1 and so on.”

After having placed the hats and removed the blindfolds, the professor asks Student 1 “Do you know the color of your own hat?” She replies “No.” Then he asks Student 2 the same question: “Do you know the color of your own hat?” Student 2 says “No.” Then he asks Student 3 and she says “No.” Then he asks the same question again to Student 1 and again the answer is “No,” and so on. After asking the same question over and over and always hearing the answer “No” he gets tired and tells them “I’ll give you a piece of information: **I did not pick three white hats.**” He then resumes the questioning: first he asks Student 1 “Do you know the color of your own hat?” She replies “No.” Then he asks Student 2 the same question: “Do you know the color of your own hat?” Student 2 says “No.” Then he asks Student 3 and she says “Yes I do!” What color hat does she have? What color hats do Students 1 and 2 have?

To answer these questions, we begin by defining the set of possible states. We can think of a state as a triple (x_1, x_2, x_3) , where $x_i \in \{R, W\}$ is the color of the hat of Student i (R means Red and W means White). Thus, for example, (R, W, R) is the state where Students 1 and 3 have a red hat, while Student 2 has a white hat. The possible states of information of the three students *before the professor announces that he did not pick three white hats* are represented by the information partitions shown in Figure 8.3, where we have connected with a line states that are in the same information set.




Figure 8.3: The initial information partitions for the Red/White hat example.

Whatever the state (that is, whatever hats the professor picks), each student is only uncertain about the color of her own hat: she can see, and thus knows, the colors of the hats of the other two students. Thus each information set contains only two elements.

Consider a particular state, say the state where all three hats are red: (R, R, R) . At that state, obviously, each student knows that not all hats are white: he/she can actually see two red hats. Furthermore, each student knows that every other student knows that not all hats are white.²

Take, for example, Student 1. She sees that the hats of the other two students are red and thus she reasons that Student 2 also sees that the hat of Student 3 is red and hence Student 2 knows that not all hats are white (similarly, she reasons that Student 3 knows that not all hats are white). But does Student 1 know that Student 2 knows that Student 3 knows that not all hats are white? The answer is No. Seeing two red hats, Student 1 must consider it possible that her own hat is white, in which case Student 2 would, like Student 1, see a red hat on Student 3 but (unlike Student 1) would also see a white hat on Student 1; thus Student 2 would have to consider the possibility that her own hat is white in which case, putting herself in the shoes of Student 3, would reason that Student 3 would see two white hats and consider it possible that his own hat was also white, that is, consider it possible that all hats were white. We can see this more clearly by using the information partitions and the associated knowledge operators. To simplify matters, let us assign names to the states and rewrite the information partitions using these names, as shown in Figure 8.4.

(W,W,W)	(W,W,R)	(W,R,W)	(W,R,R)	(R,R,R)	(R,R,W)	(R,W,R)	(R,W,W)
<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>	<i>g</i>	<i>h</i>

Student 1's partition: $\{\{a, h\}, \{b, g\}, \{c, f\}, \{d, e\}\}$

Student 2's partition: $\{\{a, c\}, \{b, d\}, \{e, g\}, \{f, h\}\}$

Student 3's partition: $\{\{a, b\}, \{c, d\}, \{e, f\}, \{g, h\}\}$

Figure 8.4: Assignment of names to the states in the Red/White hat example.

The proposition “not all hats are white” corresponds to event $E = \{b, c, d, e, f, g, h\}$ (the set of all states, excluding only state a).

²It is unfortunate that many people would use, incorrectly, the expression “all hats are not white” to mean that “it is not the case that all hats are white”. The latter expression is equivalent to “at least one hat is red (possibly one, possibly two, possibly all three)”, while the former is equivalent to “every hat is not white”, that is, in this context, “every hat is red”. [It is also unfortunate that when asked “how are you?” most people answer “I am good”, instead of “I am well” or “I am fine”!]

Using Definition 8.1.3 we get the following events (recall that $E = \{b, c, d, e, f, g, h\}$ represents the proposition “not all hats are white”):

$$\begin{aligned}
 E &= \{b, c, d, e, f, g, h\}, \\
 K_1E &= \{b, c, d, e, f, g\}, \\
 K_2E &= \{b, d, e, f, g, h\}, \\
 K_3E &= \{c, d, e, f, g, h\} \\
 K_1K_2E &= K_2K_1E = \{b, d, e, g\}, \\
 K_1K_3E &= K_3K_1E = \{c, d, e, f\}, \\
 K_2K_3E &= K_3K_2E = \{e, f, g, h\}
 \end{aligned}$$

Note that the intersection of all these events is the singleton set $\{e\}$. Thus at state e (where all the hats are red), and only at state e , everybody knows that not all hats are white and everybody knows that everybody knows that not all hats are white. Proceeding one step further, we have that

$$K_1K_2K_3E = K_1K_3K_2E = K_2K_1K_3E = K_2K_3K_1E = K_3K_1K_2E = K_3K_2K_1E = \emptyset.$$

Thus there is no state (not even e) where a student knows that another student knows that the third student knows that not all hats are white.

Let us now continue with the formal analysis of the story of the three hats. At some stage the professor makes the public announcement “I did not choose three white hats” (that is, he announces event $E = \{b, c, d, e, f, g, h\}$). This announcement makes it commonly known that state a is to be ruled out. Thus, after the announcement, the information partitions are reduced to the ones shown in Figure 8.5, obtained from Figure 8.3 by deleting the state (W, W, W) .

Note that, at this stage, if the true state is (R, W, W) Student 1 knows that her hat is red




Figure 8.5: The reduced information partitions for the Red/White hat example after the announcement that not all hats are white: state (W, W, W) is removed from the partitions of Figure 8.3.

(she sees two white hats and, having been informed that the professor did not choose three white hats, she can deduce that hers must be red). Similarly, if the true state is (W, R, W) , Student 2 knows that her own hat is red and, if the true state is (W, W, R) , Student 3 knows that her own hat is red. According to the story, after announcing that not all hats are white, the professor first asks Student 1 if she knows the color of her hat and she answers “No.” From this answer everybody can deduce that the state is not (R, W, W) and thus this state can be deleted and the information partitions reduce to the ones shown in Figure 8.6.




Figure 8.6: The reduced information partitions for the Red/White hat example after Student 1 says that she does not know the color of her own hat: state (R, W, W) is removed from the partitions of Figure 8.5.

Now, according to the story, Student 2 is asked whether she knows the color of her hat. Let us see what the possibilities are.

1. Student 2 will answer **Yes** if state is either (W, R, W) (in fact, in this case, she knew even before hearing Student 1’s answer) or (R, R, W) (in this case, before hearing Student 1’s answer, she thought the state might be either (R, R, W) or (R, W, W) but then, after hearing Student 1’s answer, she was able to eliminate (R, W, W) as a possibility). In either of these two states Student 2’s hat is red and thus she knows that her hat is red. Furthermore, upon hearing Student 2 say Yes, Student 3 learns that the state is either (W, R, W) or (R, R, W) and in both of these states his hat is white, thus he acquires the knowledge that his own hat is white.

2. In each of the remaining states, namely (W, W, R) , (W, R, R) , (R, R, R) and (R, W, R) , Student 2 will answer **No**. Then everybody learns that the state is neither (W, R, W) nor (R, R, W) and thus the information partitions reduce to the ones shown in Figure 8.7 (obtained from Figure 8.6 by removing states (W, R, W) and (R, R, W)). Note that each of the remaining states is now a singleton information set for Student 3 and thus, upon hearing Student 2 say No, he learns what the state is: in particular he learns that his own hat is red (at each of these states Student 3's hat is red). In the original story, Students 1 and 2 answered No and Student 3 answered Yes and thus we fall in this second case.




Figure 8.7: The reduced information partitions for the Red/White hat example after Student 2 also says that she does not know the color of her own hat: states (W, R, W) and (R, R, W) are removed from the partitions of Figure 8.6.

Now consider a blindfolded witness, who cannot see the color of anybody's hat, but knows that they are either red or white and hears the professor's announcement and subsequent questions, as well as the answers by the three students. What would the blindfolded witness learn about the true state?

The initial information set of the witness would consist of the set of all states; then,
 (1) after the announcement of the professor, he would be able to eliminate state (W, W, W) ,
 (2) after the negative answer of Student 1 he would be able to eliminate state (R, W, W) and
 (3) after the negative answer of Student 2 he would be able to eliminate states (W, R, W) and (R, R, W) .

The affirmative answer of Student 3 is not informative, because in each of these states Student 3 knows that the color of her own hat is red. Thus at the end of the exchange the witness's information set is $\{(W, W, R), (W, R, R), (R, R, R), (R, W, R)\}$. Hence all the witness knows is that the hat of Student 3 is red (on the other hand, Student 3 knows the color of all the hats, because she has figured out the color of her own hat and can see the hats of the other two students).

Let us now focus on state (R, R, R) where all hats are red. Initially, before the professor makes the announcement, no student is ever able to figure out the color of her own hat, no matter how many times the students are asked. However, as we saw, once the professor announces that not all hats are white, then after Students 1 and 2 reply negatively to the question whether they know the color of their own hat, Student 3 is able to deduce that her hat is red. Thus the professor's announcement provides crucial information. This, however, seems puzzling, because the professor merely tells the students *what they already knew*: each student, seeing two red hats, knows that not all hats are white (furthermore, as we saw above, each student also knows that every other student knows this). So how can giving the students a piece of information that they already possess make any difference? The answer is that the professor's *public* announcement makes it a matter of *common knowledge* that not all hats are white. Indeed, we saw above that – at the beginning – if the true state is (R, R, R) , although everybody knows that not all hats are white and also everybody knows that everybody knows that not all hats are white, it is not the case that Student 1 knows that Student 2 knows that Student 3 knows that not all hats are white. Thus it is not common knowledge that not all hats are white. The notion of common knowledge is discussed in the next section.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 8.4.2 at the end of this chapter.

8.3 Common knowledge

Common knowledge is the strongest form of interactive knowledge: an event E is common knowledge if everybody knows E and everybody knows that everybody knows E and everybody knows that everybody knows that everybody knows E , and so on. For example, in the case of two individuals, we say that at state w event E is common knowledge if

$$w \in K_1 E \cap K_2 E \cap K_1 K_2 E \cap K_2 K_1 E \cap K_1 K_2 K_1 E \cap K_2 K_1 K_2 E \cap \dots$$

We denote by CKE the event that (that is, the set of states where) event E is common knowledge. Thus, in the case of two individuals,

$$CKE = K_1 E \cap K_2 E \cap K_1 K_2 E \cap K_2 K_1 E \cap K_1 K_2 K_1 E \cap K_2 K_1 K_2 E \cap \dots$$

Given the definition of common knowledge, it may seem impossible to check if an event is common knowledge, because it requires checking an infinite number of conditions. We will see that, on the contrary, it is very easy to determine if an event is common knowledge at any state. We begin with an example.

■ **Example 8.1** Abby proposes the following to Bruno and Caroline.

“Tomorrow I will put you in two separate rooms, so that there will be no possibility of communication between you. I will then pick randomly an even number from the set $\{2,4,6\}$. Let n be that number. Then I will write the number $n - 1$ on a piece of paper and the number $n + 1$ on another piece of paper, shuffle the two pieces of paper and hand one to Bruno and the other to Caroline. For example, if I happen to pick the number 6, then I will write 5 on a piece of paper and 7 on another piece of paper, shuffle and give one piece of paper to each of you. After seeing the number handed to you, each of you will then write a pair of numbers on your piece of paper and return it to me. If

1. you write the same pair of numbers and
2. at least one of the two numbers is equal to the number that was actually given to Bruno then I will give \$1,000 to each of you, otherwise each of you will give me \$1,000.”

■

Should Bruno and Caroline accept to play this game? They can agree today on how they should act tomorrow under various contingencies, bearing in mind that they will be unable to communicate with each other tomorrow.

We will see that Bruno and Caroline should indeed accept, because they have a strategy that guarantees that they will each get \$1,000 from Abby. The first step is to represent the set of possible states and the information partitions. We will describe a state by a triple abc , where a is the number picked by Abby, b is the number given to Bruno and c is the number given to Caroline. Bruno only observes b and Caroline only observes c . Thus the information partitions are as shown in Figure 8.8.




Figure 8.8: The information partitions for Example 8.1.

Let us see if the following would be a successful strategy for Bruno and Caroline: “if Bruno gets a 1 or a 3 we will write the pair (1,3) and if Bruno gets a 5 or a 7 we will write the pair (5,7).”

Consider the event “Bruno gets a 1 or a 3”: call this event E . Then $E = \{213, 231, 435\}$. If this event occurs (e.g. because the actual state is 213), will it be common knowledge between Bruno and Caroline that it occurred? It is straightforward to check that (B stands for Bruno and C for Caroline) $K_B E = E$, $K_C E = \{231\}$ and thus $K_B K_C E = \emptyset$. Hence, while Bruno, of course, will know if he gets a 1 or a 3, Caroline might know (if the state is 231) or she might not know (if the state is 213 or 435), but Bruno will never know that Caroline knows. Thus event E is far from being common knowledge, if it occurs. It is easy to check that the same is true of the event “Bruno gets a 5 or a 7.” In order to be successful, a coordination strategy must be based on events that, when they occur, are commonly known.

Let us now consider an alternative strategy, namely

$$\left\{ \begin{array}{l} \text{if Bruno gets a 1 or a 5 let us write the pair (1,5), and} \\ \text{if Bruno gets a 3 or a 7 let us write the pair (3,7)} \end{array} \right. \quad (\star)$$

Let F be the event “Bruno gets a 1 or a 5”, that is, $F = \{213, 453, 657\}$. Then $K_B F = F$ and $K_C F = F$, so that $K_B \underbrace{K_C F}_{=F} = F$, $K_C \underbrace{K_B F}_{=F} = F$, $K_B \underbrace{K_C K_B F}_{=F} = F$, $K_C \underbrace{K_B K_C F}_{=F} = F$ and so on. Hence $CKF = F$, that is, if event F occurs then it is common knowledge between Bruno and Caroline that it occurred.³ Hence strategy (\star) will be a successful coordination strategy, since the conditioning events, when they occur, are common knowledge between Bruno and Caroline.

In the above example we were able to show directly that an event was common knowledge at a given state (we showed that $CKF = F$, that is, that, for every $w \in F$, $w \in CKF$). We now show a faster method for computing, for every event E , the event CKE . The crucial step is to derive from the individuals’ information partitions a new partition of the set of states which we call the *common knowledge partition*.

Definition 8.3.1 Consider a set of states W and n partitions $\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_n$ of W . As usual, if w is a state, we denote by $I_i(w)$ the element (information set) of the partition \mathcal{I}_i that contains w . Given two states $w, w' \in W$, we say that

- w' is *reachable from w in one step* if there exists an $i \in \{1, \dots, n\}$ such that $w' \in I_i(w)$.
- w' is *reachable from w in two steps* if there exists a state $x \in W$ such that x is reachable from w in one step and w' is reachable from x in one step.^a
- In general, w' is *reachable from w in m steps* ($m \geq 1$) if there is a sequence of states $\langle w_1, w_2, \dots, w_m \rangle$ such that (1) $w_1 = w$, (2) $w_m = w'$ and (3) for every $k = 2, \dots, m$, w_k is reachable from w_{k-1} in one step. Finally, we say that w' is *reachable from w* if, for some $m \geq 1$, w' is reachable from w in m steps.

^aThus w' is *reachable from w in two steps* if there exist $x \in W$ and $i, j \in \{1, \dots, n\}$ such that $x \in I_i(w)$ and $w' \in I_j(x)$.

³Similarly, letting $G = \{231, 435, 675\}$ be the event “Bruno gets a 3 or a 7” we have that $CKG = G$.

Definition 8.3.2 Consider a set of states W and n partitions $\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_n$ of W . Given a state w , the *common knowledge information set that contains w* , denoted by $I_{CK}(w)$, is the set of states reachable from w . The *common knowledge information partition* is the collection of common knowledge information sets.

Example 8.1 continued. Let us go back to Example 8.1, where there are two individuals, Bruno and Caroline, with the information partitions shown in Figure 8.9:




Figure 8.9: Copy of Figure 8.8.

Applying Definition 8.3.1 we have that $I_{CK}(213) = I_{CK}(453) = I_{CK}(657) = \{213, 453, 657\}$ ⁴ and $I_{CK}(231) = I_{CK}(435) = I_{CK}(675) = \{231, 435, 675\}$. Thus the common knowledge partition is as shown in Figure 8.10.




Figure 8.10: The common knowledge partition for the information partitions of Figure 8.9.

⁴In fact, 213 is reachable from itself in one step (through either Bruno or Caroline), 453 is reachable in one step from 213 (through Caroline) and 657 is reachable in two steps from 213 (the first step – to 453 – through Caroline and the second step – from 453 – through Bruno).

As a second example, consider the information partitions shown in Figure 8.11 (which reproduces Figure 8.2). The corresponding common knowledge partition is shown in Figure 8.12.




Figure 8.11: Copy of Figure 8.2.



Figure 8.12: The common knowledge partition for the information partitions of Figure 8.11.

The following theorem states that, determining whether an event E is common knowledge at a state w , is equivalent to determining whether an individual whose information partition coincided with the common knowledge partition would know E at w .

Theorem 8.3.1 At state $w \in W$, event $E \subseteq W$ is common knowledge (that is, $w \in CKE$) if and only if $I_{CK}(w) \subseteq E$. Thus we can define the common knowledge operator $CK : 2^W \rightarrow 2^W$ as follows: $CKE = \{w \in W : I_{CK}(w) \subseteq E\}$.

Example 8.1 continued. Let us go back to Example 8.1 about Bruno and Caroline. Let F be the event that Bruno gets a 1 or a 5: $F = \{213, 453, 657\}$. Then, using Theorem 8.3.1, $CKF = F$ because $I_{CK}(213) = I_{CK}(453) = I_{CK}(657) = \{213, 453, 657\}$; thus – confirming what we found in Example 8.1 – at any state where Bruno gets a 1 or a 5 it is common knowledge between Bruno and Caroline that Bruno got a 1 or a 5.

Now let H be the event “Bruno did not get a 5”, that is, $H = \{213, 231, 435, 675\}$. Then, using Theorem 8.3.1 we have that $CKH = \{231, 435, 675\}$. Thus while at state 231 Bruno does not get a 5 and this is common knowledge between Bruno and Caroline, at state 213 Bruno does not get a 5 but this is not common knowledge between Bruno and Caroline (in fact $213 \notin K_{Caroline}H = \{231, 435, 675\}$).

As a last example, consider again the information partitions of Figure 8.11, whose common knowledge partition was shown in Figure 8.12 and is reproduced in Figure 8.13.



Figure 8.13: Copy of Figure 8.12.

Let $E = \{a, b, c, d, e, f\}$. Then $CKE = \{a, b, c, d\}$. Let $F = \{a, b, f, g, h\}$. Then $CKF = \emptyset$.

The smallest event that is common knowledge at state a is $I_{CK}(a) = \{a, b, c, d\}$ and the smallest event that is common knowledge at state g is $I_{CK}(g) = \{e, f, g, h\}$.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 8.4.3 at the end of this chapter.

8.4 Exercises

8.4.1 Exercises for Section 8.1: Individual knowledge

The answers to the following exercises are in Section 8.5 at the end of this chapter.

Exercise 8.1

You are in a completely dark room, where you cannot see anything. You open a drawer that you know to contain individual socks, all of the same size but of different colors. You know that 5 are blue and 7 are white.

- (a) First use your intuition to answer the following question. What is the smallest number of socks you need to take out in order to know (that is, to be absolutely certain) that you have a matching pair (i.e. either a pair of blue socks or a pair of white socks)?
- (b) Now represent this situation using states and information sets. Do this to
 1. represent the situation after you have taken out one sock,
 2. represent the situation after you have taken out two socks and
 3. represent the situation after you have taken out three socks.

Remember that a state should encode a complete description of the relevant aspects of the situation; in particular, the state should tell us how many socks you have taken out and the color of each sock that you have taken out (thus the set of states changes over time as you take out more socks). ■

Exercise 8.2

Consider the situation described in Exercise 8.1: the room is dark and you have taken out three socks. Consider the following alternative scenarios.

- (a) Somebody tells you the color of the third sock (but you still don't know the color of the other two socks). Represent your state of knowledge by means of an information set.
- (b) Somebody tells you the color of the matching pair (but you don't know what socks you picked). Represent your state of knowledge by means of an information set.

Exercise 8.3

Let the set of states be $W = \{a, b, c, d, e, f, g, h, k, m, n\}$ and the information partition of an individual be $\{\{a, b, c\}, \{d\}, \{e, f, g, h\}, \{k, m\}, \{n\}\}$. Consider the following event: $E = \{a, b, d, k, n\}$.

- (a) Does the individual know E at state a ?
- (b) Does the individual know E at state c ?
- (c) Does the individual know E at state d ?
- (d) Does the individual know E at state h ?
- (e) Does the individual know E at state k ?
- (f) Let KE denote the event that the individual knows E (that is, the set of states where the individual knows E). What is KE ?

For the next question, recall that, given an event F , we denote by $\neg F$ the complement of F , that is, the set of states that are not in F .

- (g) Once again, let $E = \{a, b, d, k, n\}$. What is the event $\neg KE$, that is the event that the individual does not know E ? What is the event $K\neg KE$, that is, the event that the individual knows that she does not know E ?

Exercise 8.4

The famous pirate Sandokan has captured you and put you in front of three numbered chests containing coins. Chest 1 is labeled “gold,” Chest 2 is labeled “bronze,” and Chest 3 is labeled “gold or silver.” One chest contains gold coins only, another contains silver coins only, and the third bronze coins only.

- (a) Represent the set of possible states in the case where the labels might or might not be correct (a state must describe the label and content of each box).
- (b) Let E be the event “Chest 1 is mislabeled” (that is, what the label says is false). What states are in event E ?
- (c) Let F be the event “Chest 2 is mislabeled”. What states are in event F ?
- (d) What is the event “both Chests 1 and 2 are mislabeled”?
- (e) Suppose now that Sandokan tells you that **all** the chests are falsely labeled, that is, what the label says is false (for example, if it says “gold” then you can be sure that the chest does not contain gold coins). If you correctly announce the content of all the chests you will be given a total of \$1,000. If you make a mistake (e.g. state that a chest contains, say, gold coins while in fact it contains bronze coins) then you don’t get any money at all. You can open any number of chests you like in order to inspect the content. However, the first time you open a chest, you have to pay \$500, the second time \$300, the third time \$100.
 1. What is the set of possible states (assuming that Sandokan told you the truth)?
 2. What is the maximum amount of money you can be absolutely certain to make?

Exercise 8.5

Prove each of the following properties of the knowledge operator. The proofs are straightforward applications of Definition 8.1.2.

- **Truth:** $KE \subseteq E$, that is, if at a state w one knows E , then E is indeed true at w .
- **Consistency:** $KE \cap K\neg E = \emptyset$, that is, one never simultaneously knows E and also $\neg E$ ($\neg E$ denotes the complement of E).
- **Positive introspection:** $KE \subseteq KKE$, that is, if one knows E then one knows that one knows E (one is fully aware of what one knows). [Note that it follows from this and Truth that $KE = KKE$, because from Truth we get that $KKE \subseteq KE$.]
- **Negative Introspection:** $\neg KE \subseteq K\neg KE$, that is, if one does not know E , then one knows that one does not know E (one is fully aware of what one does not know). [Note that it follows from this and Truth that $\neg KE = K\neg KE$, because from Truth we get that $K\neg KE \subseteq \neg KE$.]

- **Monotonicity:** If $E \subseteq F$, then $KE \subseteq KF$, that is, if E implies F then if one knows E then one knows F .
- **Conjunction:** $KE \cap KF = K(E \cap F)$, that is, if one knows E and one knows F , then one knows E and F , and *vice versa*.

8.4.2 Exercises for Section 8.2: Interactive knowledge

The answers to the following exercises are in Section 8.5 at the end of this chapter.




Figure 8.14: The information partitions for Exercise 8.6.

Exercise 8.6

Let the set of states be $W = \{a, b, c, d, e, f, g, h\}$. There are three individuals with the information partitions shown in Figure 8.14.

Consider the event $E = \{a, b, c, f, g\}$. Find the following events.

- $K_{Ann}E$ (the event that Ann knows E).
- $K_{Bob}E$,
- $K_{Carol}E$,
- $K_{Carol}K_{Ann}E$ (the event that Carol knows that Ann knows E),
- $K_{Bob}K_{Carol}K_{Ann}E$,
- $K_{Ann} \neg K_{Bob}K_{Carol}E$ (the event that Ann knows that Bob does not know that Carol knows E).

Exercise 8.7

Dan is at the airport. He calls his office and leaves a voice message that says: “My flight was cancelled and I am waiting to see if they can re-route me through Los Angeles or San Francisco. I will call one of you at home tonight at 8:00 sharp to let that person know whether I am in San Francisco or Los Angeles.” Dan’s office staff, consisting of Ann, Barb and Carol, were out for lunch. When they come back they listen to the message together. They leave the office at 5:00 pm and each goes to her home.

- (a) Using information partitions, represent the possible states of knowledge of Ann, Barb and Carol concerning Dan’s whereabouts at 8:15 pm, after Dan’s call (there has been no communication among Ann, Barb and Carol after they left the office).
- (b) Let E be the event that Dan calls either Barb or Carol. What states are in E ?
- (c) For the event E of Part (b), find $K_A E$ (Ann knows E), $K_B K_A E$ (Barb knows that Ann knows E) and $\neg K_C E$ (Carol does not know E or it is not the case that Carol knows E).
- (d) For the event E of Part (b), find a state x where all of the following are true:
 - (1) at x Ann knows E ,
 - (2) at x Barb knows that Ann knows E ,
 - (3) at x it is not the case that Carol knows E .

Exercise 8.8

A set of lights is controlled by two switches, each of which can be in either the Up position or in the Down position. One switch is in room number 1, where Ann is; the other switch is in room number 2, where Bob is. The lights are in room number 3, where Carla is.




There are two lights: one red and one green. The red light is on if the two switches are in different positions (one up and the other down: it doesn’t matter which is up and which is down), while the green light is on if the two switches are in the same position (both up or both down).

All this is common knowledge among Ann, Bob and Carla.

- (a) Represent the possible states (you need to specify the position of each switch and which light is on).
- (b) Represent the possible states of information of Ann, Bob and Carla by means of information partitions.
- (c) Let G be the event “the green light is on”. Find the events G , $K_A G$ (Ann knows G), $K_B G$ (Bob knows G), $K_C G$ (Carla knows G).
- (d) Let L be the event “either the green light is on or the red light is on”. Find the events L , $K_A L$ (Ann knows L), $K_B L$ (Bob knows L), $K_C L$ (Carla knows L).

8.4.3 Exercises for Section 8.3: Common knowledge

The answers to the following exercises are in Section 8.5 at the end of this chapter.



Exercise 8.9

Consider again the information partitions of Exercise 8.6, reproduced above.

- Find the common knowledge partition,
- Let $E = \{a, b, c, f, g\}$.
Find the event CKE , that is, the event that E is common knowledge.
- Let $F = \{a, b, c, d, e, g\}$.
Find CKF , that is, the event that F is common knowledge.

Exercise 8.10

In Exercise 8.7,

- Find the common knowledge partition,
- Find the event CKE (where E is the event that Dan calls either Barb or Carol).

Exercise 8.11

In Exercise 8.8,

- Find the common knowledge partition,
- Find the event CKG (where G is the event “the green light is on”),
- Find the event CKL (where L is the event “either the green light is on or the red light is on”).

Exercise 8.12

The set of states is $W = \{a, b, c, d, e, f, g, h\}$. There are four individuals with the following information partitions:

Individual 1: $\{\{a, b\}, \{c\}, \{d\}, \{e, f\}, \{g\}, \{h\}\}$

Individual 2: $\{\{a\}, \{b, c\}, \{d, e\}, \{f\}, \{g\}, \{h\}\}$

Individual 3: $\{\{a, c\}, \{b\}, \{d\}, \{e\}, \{g\}, \{f, h\}\}$

Individual 4: $\{\{a\}, \{b, c\}, \{d, e\}, \{f, g\}, \{h\}\}$

(a) Let $E = \{a, c, d, e\}$.

Find the following events: K_1E, K_2E, K_3E, K_4E and $K_1K_2\neg K_3E$

(Recall that \neg denotes the complement of a set, that is, $\neg F$ is the set of all states that are not in F).

(b) Find the common knowledge partition.

(c) At what states is event $E = \{a, c, d, e\}$ common knowledge?

(d) Let $F = \{a, b, c, d, g, h\}$. Find the event that F is common knowledge.




Figure 8.15: The situation described in Exercise 8.13.

Exercise 8.13

Amy, Beth and Carla are now in Room 1 as shown in Figure 8.15. They are asked to proceed, one at a time, to Room 3 through Room 2. In Room 2 there are two large containers, one with many red hats and the other with many white hats. They have to choose one hat, put it on their head and then go and sit in the chair in Room 3 that has their name on it.

Amy goes first, then Beth then Carla. The chairs are turned with the back to the door. Thus a person entering Room 3 can see whomever is already seated there but cannot be seen by them. Beth and Carla don't know how many hats there were in each box.

(a) Use information partitions to represent the possible states of knowledge of Amy, Beth and Carla after they are seated in Room 3.

(b) Suppose that Amy chose a white hat, Beth a red hat and Carla a white hat. Find the smallest event that is common knowledge among them. Give also a verbal description of this event.

(c) Repeat Parts (a) and (b) for the modified setting where there is a mirror in Room 3 that allows Amy to see the hat of Carla (but not that of Beth).

Exercise 8.14 — ***Challenging Question.***

Francis and his three daughters Elise, Sophia and Justine are in the same room. Francis gives a sealed envelope to Elise and tells her (in a loud voice, so that everybody can hear)

“In this envelope I put a sum of money; I don’t remember how much I put in it, but I know for sure that it was either \$4 or \$8 or \$12. Now, my dear Elise, go to your room by yourself and open it. Divide the money into two equal amounts, put the two sums in two different envelopes, seal them and give one to Sophia and one to Justine.”

Unbeknownst to her sisters, Elise likes one of them more than the other and decides to disobey her father: after dividing the sum into two equal parts, she takes \$1 from one envelope and puts it in the other envelope. She then gives the envelope with more money to her favorite sister and the envelope with the smaller amount to the other sister. Sophia and Justine go to their respective rooms and privately open their envelopes, to discover, to their surprise, an odd number of dollars. So they realize that Elise did not follow their father’s instructions. Neither Sophia nor Justine suspect that Elise kept some money for herself; in fact, it is common knowledge between them that Elise simply rearranged the money, without taking any for herself. Of course, neither Sophia nor Justine know in principle how much money Elise took from one envelope (although in some cases they might be able to figure it out). Thus it is **not** common knowledge between Sophia and Justine that Elise took only \$1 from one of the two envelopes. Your answers should reflect this.

- (a) Use states and information partitions to represent the possible states of knowledge of Sophia and Justine.
- (b) Let E be the event that Sophia is Elise’s favorite sister. Find the events $K_S E$ (the event that Sophia knows it), $K_J E$ (the event that Justine knows it), $K_S K_J E$ and $K_J K_S E$.
- (c) Find the common knowledge partition.
- (d) Is there a state at which it is common knowledge between Sophia and Justine that Elise’s favorite sister is Sophia?
- (e) The night before, Francis was looking through his digital crystal ball and saw what Elise was planning to do. However the software was not working properly (the screen kept freezing) and he could not tell whether Elise was trying to favor Sophia or Justine. He knew he couldn’t stop Elise and wondered how much money he should put in the envelope to make sure that the mistreated sister (whether it be Sophia or Justine) would not know that she had been mistreated. How much money should he put in the envelope? ■

8.5 Solutions to Exercises

Solution to Exercise 8.1

- (a) The answer is that you need to take out three socks.
 (b) Describe a state by the number of socks you have taken out and the color of each sock that you have taken out. After you have taken out only one sock:

1	1
<i>B</i>	<i>W</i>

(the information set captures the fact that you cannot see the color of the sock because it is dark). After you have taken out two socks:

2	2	2	2
<i>B</i>	<i>B</i>	<i>W</i>	<i>W</i>
<i>B</i>	<i>W</i>	<i>B</i>	<i>W</i>

After you have taken out three socks:

3	3	3	3	3	3	3	3
<i>B</i>	<i>B</i>	<i>B</i>	<i>B</i>	<i>W</i>	<i>W</i>	<i>W</i>	<i>W</i>
<i>B</i>	<i>B</i>	<i>W</i>	<i>W</i>	<i>B</i>	<i>B</i>	<i>W</i>	<i>W</i>
<i>B</i>	<i>W</i>	<i>B</i>	<i>W</i>	<i>B</i>	<i>W</i>	<i>B</i>	<i>W</i>

Now at every state there are (at least) two socks of the same color, thus you know that you have a matching pair, even though you don't know what state you are in and hence you don't know the color of the matching pair. An alternative (and equivalent) way to proceed would be to describe the state as a quadruple of numbers as follows:

number of blue socks in drawer	x_1
number of white socks in drawer	x_2
number of blue socks in your hand	x_3
number of white socks in your hand	x_4

Clearly, it must always be that $x_1 + x_3 = 5$ and $x_2 + x_4 = 7$.

The initial state is $x_1 = 5, x_2 = 7, x_3 = x_4 = 0$.

Taking one sock from the drawer will modify the state as shown in Figure 8.16.

Now taking a second sock will modify the state as shown in Figure 8.17.

Now taking a third sock will modify the state as shown in Figure 8.18.

Now you know that you have a matching pair because in every possible state you have at least 2 socks of the same color. Thus the answer is indeed that you need to take out three socks. \square

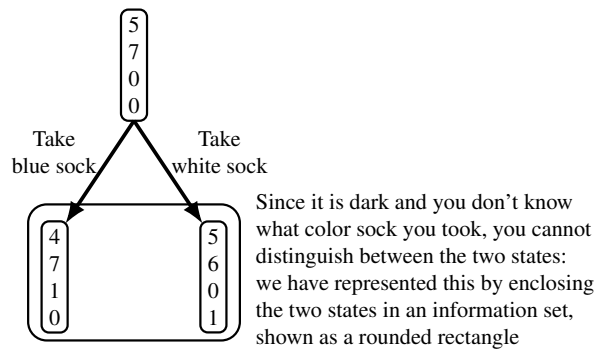


Figure 8.16: The possible states after taking one sock.

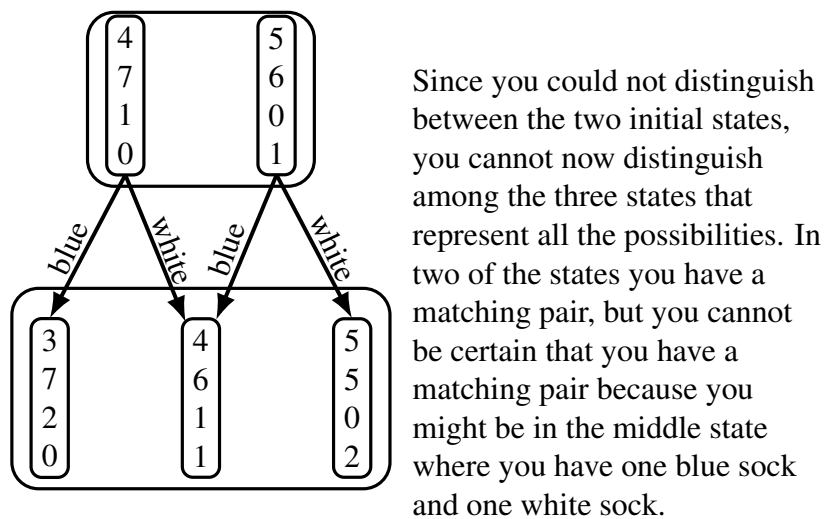


Figure 8.17: The possible states after taking two socks.




Figure 8.18: The possible states after taking three socks.

Solution to Exercise 8.2

(a)

3	3	3	3
B	B	W	W
B	W	B	W
B	B	B	B

3	3	3	3
B	B	W	W
B	W	B	W
W	W	W	W

(b)

3	3	3	3
B	B	W	B
B	W	B	B
B	B	B	W

3	3	3	3
W	B	W	W
W	W	B	W
B	W	W	W

You have a blue matching pair
You have a white matching pair

□

Solution to Exercise 8.3

- (a) No, because the information set containing a is $\{a, b, c\}$ which is not contained in $E = \{a, b, d, k, n\}$.
- (b) No, because the information set containing c is $\{a, b, c\}$ which is not contained in $E = \{a, b, d, k, n\}$.
- (c) Yes, because the information set containing d is $\{d\}$ which is contained in $E = \{a, b, d, k, n\}$.
- (d) No, because the information set containing h is $\{e, f, g, h\}$ which is not contained in $E = \{a, b, d, k, n\}$.
- (e) No, because the information set containing k is $\{k, m\}$ which is not contained in $E = \{a, b, d, k, n\}$.
- (f) $KE = \{d, n\}$.
- (g) $\neg KE = \{a, b, c, e, f, g, h, k, m\}$. $K\neg KE = \{a, b, c, e, f, g, h, k, m\} = \neg KE$. □

Solution to Exercise 8.4

- (a) Describe a state by a triple $((x_1, y_1), (x_2, y_2), (x_3, y_3))$ where x_i is the label on box number i ($i = 1, 2, 3$) and y_i is the content of box number i .

Then the set of possible states is

$$\begin{aligned}
 z_1 &= ((G, B), (B, G), (G \text{ or } S, S)), & z_2 &= ((G, B), (B, S), (G \text{ or } S, G)) \\
 z_3 &= ((G, G), (B, B), (G \text{ or } S, S)), & z_4 &= ((G, G), (B, S), (G \text{ or } S, B)) \\
 z_5 &= ((G, S), (B, B), (G \text{ or } S, G)), & z_6 &= ((G, S), (B, G), (G \text{ or } S, B)).
 \end{aligned}$$

Thus, for example, state z_4 is one where box number 1 (which is labeled ‘gold’) in fact contains gold coins, box number 2 (which is labeled ‘bronze’) as a matter of fact contains silver coins, and box number 3 (which is labeled ‘gold or silver’) as a matter of fact contains bronze coins. More simply, we could write a state as a triple (y_1, y_2, y_3) denoting the content of each box (since we are told what the labels are). In this simpler notation the states are:

$$z_1 = (B, G, S), z_2 = (B, S, G), z_3 = (G, B, S)$$

$$z_4 = (G, S, B), z_5 = (S, B, G), z_6 = (S, G, B)$$

- (b) $E = \{(B, G, S), (B, S, G), (S, B, G), (S, G, B)\}$ (recall that the label says “gold”).
- (c) $F = \{(B, G, S), (B, S, G), (G, S, B), (S, G, B)\}$ (recall that the label says “bronze”).
- (d) $E \cap F = \{(B, G, S), (B, S, G), (S, G, B)\}$.
- (e) Of all the states listed above, only state z_6 is such that all the labels are false. Hence Sandokan’s statement reduces the set of states to only one: $z_6 = (S, G, B)$. This is because the label “gold or silver” must be on the chest containing the bronze coins. Hence we are only left with gold and silver. Then silver must be in the chest labeled “gold”. Hence gold must be in the chest labeled “bronze”. Thus, by looking at the labels you can correctly guess the content without opening any chests: you know that the true state is (S, G, B) . So you will get \$1,000 (you are not guessing at random, you are deducing by reasoning and you don’t need to open any chests). \square

Solution to Exercise 8.5

- **Truth:** $KE \subseteq E$. Consider an arbitrary $w \in KE$. We have to show that $w \in E$. Since $w \in KE$, by Definition 8.1.2, $I(w) \subseteq E$. Since $w \in I(w)$, it follows that $w \in E$.
- **Consistency:** $KE \cap K\neg E = \emptyset$. Suppose that $w \in KE \cap K\neg E$ for some w and some E . Then, by Definition 8.1.2, $I(w) \subseteq E$ (because $w \in KE$) and $I(w) \subseteq \neg E$ (because $w \in \neg E$) and thus $I(w) \subseteq E \cap \neg E$. Since $E \cap \neg E = \emptyset$, this implies that $I(w) = \emptyset$, which is not true because $w \in I(w)$.
- **Positive introspection:** $KE \subseteq KKE$. Consider an arbitrary $w \in KE$. We need to show that $w \in KKE$, that is, that $I(w) \subseteq KE$ which means that $w' \in KE$ for every $w' \in I(w)$. Since $w \in KE$, by Definition 8.1.2, $I(w) \subseteq E$. Consider an arbitrary $w' \in I(w)$. By definition of partition, $I(w') = I(w)$. Thus $I(w') \subseteq E$, that is, by Definition 8.1.2, $w' \in KE$.
- **Negative introspection:** $\neg KE \subseteq K\neg KE$. Consider an arbitrary $w \in \neg KE$. We need to show that $w \in K\neg KE$, that is, that $I(w) \subseteq \neg KE$. By Definition 8.1.2, since $I(w) \subseteq \neg KE$, $I(w) \cap \neg E \neq \emptyset$. Consider an arbitrary $w' \in I(w)$; then, since (by definition of partition) $I(w') = I(w)$, $I(w') \cap \neg E \neq \emptyset$ so that $w' \in \neg KE$. Thus we have shown that, for every $w' \in I(w)$, $w' \in \neg KE$, that is, $I(w) \subseteq \neg KE$ which, by Definition 8.1.2, yields $w \in K\neg KE$.

- **Monotonicity:** if $E \subseteq F$, then $KE \subseteq KF$. Consider an arbitrary $w \in KE$. We need to show that $w \in KF$. Since $w \in KE$, by Definition 8.1.2, $I(w) \subseteq E$. Hence, since, by hypothesis, $E \subseteq F, I(w) \subseteq F$, that is, by Definition 8.1.2, $w \in KF$.
- **Conjunction:** $KE \cap KF = K(E \cap F)$. Let $w \in KE \cap KF$. Then $w \in KE$ and $w \in KF$; by Definition 8.1.2, the former implies that $I(w) \subseteq E$ and the latter that $I(w) \subseteq F$, so that $I(w) \subseteq E \cap F$ and hence, by Definition 8.1.2, $w \in K(E \cap F)$. Conversely, suppose that $w \in K(E \cap F)$. Then, by Definition 8.1.2, $I(w) \subseteq E \cap F$ and thus $I(w) \subseteq E$ and $I(w) \subseteq F$ so that, by Definition 8.1.2, $w \in KE$ and $w \in KF$; hence $w \in KE \cap KF$. □

Solution to Exercise 8.6

- (a) $K_{Ann}E = \{a, b, c\}$,
- (b) $K_{Bob}E = \{a, b, f\}$,
- (c) $K_{Carol}E = \{b, c, f, g\}$,
- (d) $K_{Carol}K_{Ann}E = K_{Carol}\{a, b, c\} = \{b, c\}$,
- (e) $K_{Bob}K_{Carol}K_{Ann}E = K_{Bob}\{b, c\} = \emptyset$,
- (f) $K_{Ann}\neg K_{Bob}K_{Carol}E = K_{Ann}\neg K_{Bob}\{b, c, f, g\} = K_{Ann}\neg\{f\} = K_{Ann}\{a, b, c, d, e, g, h\} = \{a, b, c, d, h\}$. □

Solution to Exercise 8.7

- (a) We can represent a state as a pair (x, y) where x is the city where Dan is (SF for San Francisco or LA for Los Angeles) and y is the person that Dan called (A for Ann, B for Barb and C for Carol). The information partitions are shown in Figure 8.19.



Figure 8.19: The information partitions for Exercise 8.7.

- (b) $E = \{(SF, B), (LA, B), (SF, C), (LA, C)\}$
- (c) $K_{Ann}E = E, K_{Barb}K_{Ann}E = \{(SF, B), (LA, B)\}$ and $K_{Carol}E = \{(SF, C), (LA, C)\}$ so that $\neg K_{Carol}E = \{(SF, A), (LA, A), (SF, B), (LA, B)\}$.
- (d) We want a state x such that $x \in K_{Ann}E$, $x \in K_{Barb}K_{Ann}E$ and $x \in \neg K_{Carol}E$ (that is, $x \notin K_{Carol}E$). There are only two such states: (SF, B) and (LA, B) . Thus either $x = (SF, B)$ or $x = (LA, B)$. □

Solution to Exercise 8.8

- (a) We can represent a state as a triple $\begin{pmatrix} x \\ y \\ z \end{pmatrix}$ where x is the position of the switch in Room 1 (Up or Down), y is the position of the switch in Room 2 and z is the light which is on in Room 3 (Green or Red).
- (b) The information partitions are shown in Figure 8.20

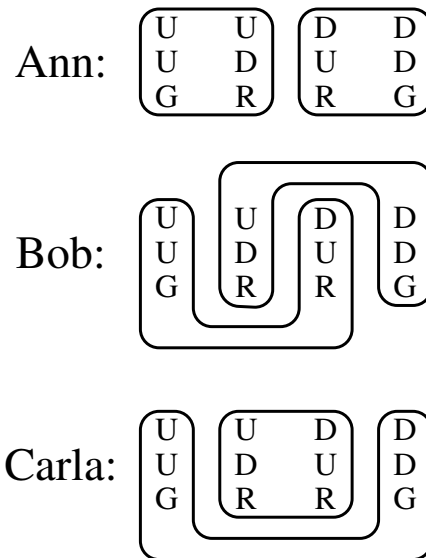


Figure 8.20: The information partitions for Exercise 8.8.

(c) $G = \left\{ \begin{array}{c|c} U & D \\ \hline U & D \\ G & G \end{array} \right\}, \quad K_A G = \emptyset, \quad K_B G = \emptyset, \quad K_C G = \left\{ \begin{array}{c|c} U & D \\ \hline U & D \\ G & G \end{array} \right\}.$

(d) L is the set of all states. Hence $K_A L = K_B L = K_C L = L$. □

Solution to Exercise 8.9

- (a) The common knowledge partition is shown in Figure 8.21.



Figure 8.21: The common knowledge partition for Exercise 8.9.

- (b) $CKE = \emptyset$ (where $E = \{a, b, c, f, g\}$).
- (c) $CKF = \{a, b, c, d\}$ (where $F = \{a, b, c, d, e, g\}$). □

Solution to Exercise 8.10

- (a) The common knowledge partition is the trivial partition shown in Figure 8.22.

$$\boxed{(\text{SF}, \text{A}) \quad (\text{LA}, \text{A}) \quad (\text{SF}, \text{B}) \quad (\text{LA}, \text{B}) \quad (\text{SF}, \text{C}) \quad (\text{LA}, \text{C})}$$

Figure 8.22: The common knowledge partition for Exercise 8.10.

- (b)
- $CKE = \emptyset$

□

Solution to Exercise 8.11

- (a) The common knowledge partition is the trivial partition shown below:

$$\boxed{\begin{array}{cccc} U & U & U & U \\ U & D & U & D \\ G & R & R & G \end{array}}$$

- (b)
- $CKG = \emptyset$
- .

- (c)
- $CKL = L$
- .

□

Solution to Exercise 8.12

- (a) Let
- $E = \{a, c, d, e\}$
- . Then:
- $K_1E = \{c, d\}$
- ,
- $K_2E = \{a, d, e\}$
- ,
- $K_3E = \{a, c, d, e\}$
- ,
- $K_4E = \{a, d, e\}$
- and
- $K_1K_2\neg K_3E = \{g, h\}$
- [in fact,
- $\neg K_3E = \{b, f, g, h\}$
- , so that
- $K_2\neg K_3E = \{f, g, h\}$
- and
- $K_1K_2\neg K_3E = \{g, h\}$
-].

- (b) The common knowledge partition is
- $\{\{a, b, c\}, \{d, e, f, g, h\}\}$
- .

- (c) At no state is event
- $E = \{a, c, d, e\}$
- common knowledge:
- $CKE = \emptyset$
- .

- (d)
- $CKF = \{a, b, c\}$
- (where
- $F = \{a, b, c, d, g, h\}$
-).

□

Solution to Exercise 8.13

- (a) Represent a state as a triple of letters, where the top letter denotes the color of Amy's hat, the second letter the color of Beth's hat and the bottom letter the color of Carla's hat. Each of them knows which hat she chose; furthermore, Beth can see Amy's hat and Carla can see everything. Thus the partitions are as shown in Figure 8.23.

- (b) When the true state is
- (W, R, W)
- , the smallest event that is common knowledge among them is the first information set of the common knowledge partition, that is, the fact that Amy has a white hat.

- (c) In the modified setting the information partitions are as shown in Figure 8.24. The common knowledge partition is the same as before, hence the smallest event that is common knowledge among all three of them is that Amy has a white hat. □




Figure 8.23: The information partitions for Part (a) of Exercise 8.13.




Figure 8.24: The information partitions for Part (c) of Exercise 8.13.

Solution to Exercise 8.14

- (a) Describe a state by three numbers where the top number is the amount of money that Francis put in the envelope, the middle number is the amount given to Sophia and the bottom number is the amount given to Justine. As a matter of fact, each sister can only find either \$1 or \$3 or \$5 or \$7 in her own envelope. Thus the objectively possible states are:

$$\begin{pmatrix} 4 \\ 1 \\ 3 \end{pmatrix}, \begin{pmatrix} 4 \\ 3 \\ 1 \end{pmatrix}, \begin{pmatrix} 8 \\ 3 \\ 5 \end{pmatrix}, \begin{pmatrix} 8 \\ 5 \\ 3 \end{pmatrix}, \begin{pmatrix} 12 \\ 5 \\ 7 \end{pmatrix} \text{ and } \begin{pmatrix} 12 \\ 7 \\ 5 \end{pmatrix}.$$

However, besides these, there are also subjectively possible states, namely

$$\begin{pmatrix} 8 \\ 1 \\ 7 \end{pmatrix}, \begin{pmatrix} 8 \\ 7 \\ 1 \end{pmatrix}, \begin{pmatrix} 12 \\ 1 \\ 11 \end{pmatrix}, \begin{pmatrix} 12 \\ 11 \\ 1 \end{pmatrix}, \begin{pmatrix} 12 \\ 3 \\ 9 \end{pmatrix} \text{ and } \begin{pmatrix} 12 \\ 9 \\ 3 \end{pmatrix}.$$

(because Sophia and Justine don't know that Elise only transferred \$1 from one envelope to the other).

The information partitions are as follows:

SOPHIA:

$$\begin{array}{ccc} \boxed{\begin{pmatrix} 4 \\ 1 \\ 3 \end{pmatrix} \begin{pmatrix} 8 \\ 1 \\ 7 \end{pmatrix} \begin{pmatrix} 12 \\ 1 \\ 11 \end{pmatrix}} & \boxed{\begin{pmatrix} 4 \\ 3 \\ 1 \end{pmatrix} \begin{pmatrix} 8 \\ 3 \\ 5 \end{pmatrix} \begin{pmatrix} 12 \\ 3 \\ 9 \end{pmatrix}} & \boxed{\begin{pmatrix} 8 \\ 5 \\ 3 \end{pmatrix} \begin{pmatrix} 12 \\ 5 \\ 7 \end{pmatrix}} \\ & \boxed{\begin{pmatrix} 8 \\ 7 \\ 1 \end{pmatrix} \begin{pmatrix} 12 \\ 7 \\ 5 \end{pmatrix}} & \boxed{\begin{pmatrix} 12 \\ 9 \\ 3 \end{pmatrix}} & \boxed{\begin{pmatrix} 12 \\ 11 \\ 1 \end{pmatrix}} \end{array}$$

JUSTINE:

$$\begin{array}{ccc} \boxed{\begin{pmatrix} 4 \\ 1 \\ 3 \end{pmatrix} \begin{pmatrix} 8 \\ 5 \\ 3 \end{pmatrix} \begin{pmatrix} 12 \\ 9 \\ 3 \end{pmatrix}} & \boxed{\begin{pmatrix} 4 \\ 3 \\ 1 \end{pmatrix} \begin{pmatrix} 8 \\ 7 \\ 1 \end{pmatrix} \begin{pmatrix} 12 \\ 11 \\ 1 \end{pmatrix}} & \boxed{\begin{pmatrix} 8 \\ 3 \\ 5 \end{pmatrix} \begin{pmatrix} 12 \\ 7 \\ 5 \end{pmatrix}} \\ & \boxed{\begin{pmatrix} 8 \\ 1 \\ 7 \end{pmatrix} \begin{pmatrix} 12 \\ 5 \\ 7 \end{pmatrix}} & \boxed{\begin{pmatrix} 12 \\ 3 \\ 9 \end{pmatrix}} & \boxed{\begin{pmatrix} 12 \\ 1 \\ 11 \end{pmatrix}} \end{array}$$

(b) The event that Sophia is Elise's favorite sister is

$$E = \left\{ \begin{pmatrix} 4 \\ 3 \\ 1 \end{pmatrix}, \begin{pmatrix} 8 \\ 5 \\ 3 \end{pmatrix}, \begin{pmatrix} 8 \\ 7 \\ 1 \end{pmatrix}, \begin{pmatrix} 12 \\ 7 \\ 5 \end{pmatrix}, \begin{pmatrix} 12 \\ 9 \\ 3 \end{pmatrix}, \begin{pmatrix} 12 \\ 11 \\ 1 \end{pmatrix} \right\}$$

$$K_S E = \left\{ \begin{pmatrix} 8 \\ 7 \\ 1 \end{pmatrix}, \begin{pmatrix} 12 \\ 7 \\ 5 \end{pmatrix}, \begin{pmatrix} 12 \\ 9 \\ 3 \end{pmatrix}, \begin{pmatrix} 12 \\ 11 \\ 1 \end{pmatrix} \right\}$$

$$K_J E = \left\{ \begin{pmatrix} 4 \\ 3 \\ 1 \end{pmatrix}, \begin{pmatrix} 8 \\ 7 \\ 1 \end{pmatrix}, \begin{pmatrix} 12 \\ 11 \\ 1 \end{pmatrix} \right\}$$

$$K_S K_J E = \left\{ \begin{pmatrix} 12 \\ 11 \\ 1 \end{pmatrix} \right\} \quad \text{and} \quad K_J K_S E = \emptyset.$$

(c) The common knowledge partition is:

$$\left[\begin{pmatrix} 4 \\ 1 \\ 3 \end{pmatrix}, \begin{pmatrix} 8 \\ 1 \\ 7 \end{pmatrix}, \begin{pmatrix} 12 \\ 1 \\ 11 \end{pmatrix}, \begin{pmatrix} 8 \\ 5 \\ 3 \end{pmatrix}, \begin{pmatrix} 12 \\ 9 \\ 3 \end{pmatrix}, \begin{pmatrix} 12 \\ 5 \\ 7 \end{pmatrix} \right]$$

$$\left[\begin{pmatrix} 4 \\ 3 \\ 1 \end{pmatrix}, \begin{pmatrix} 8 \\ 3 \\ 5 \end{pmatrix}, \begin{pmatrix} 12 \\ 3 \\ 9 \end{pmatrix}, \begin{pmatrix} 8 \\ 7 \\ 1 \end{pmatrix}, \begin{pmatrix} 12 \\ 11 \\ 1 \end{pmatrix}, \begin{pmatrix} 12 \\ 7 \\ 5 \end{pmatrix} \right]$$

(d) At no state. In fact, $CKE = \emptyset$.

(e) Francis should put either \$8 or \$12 in the envelope. If he were to put \$4, then one sister would end up with \$1 and know that she was mistreated. In no other case does a mistreated sister know that she got less money than the other. \square

9. Adding Beliefs to Knowledge

9.1 Sets and probability: brief review

We begin with a very brief review of definitions and concepts from set theory and probability theory.

9.1.1 Sets

We will focus on finite sets, that is, sets that have a finite number of elements. Let U be a finite set. The set of subsets of U is denoted by 2^U . The reason for this notation is that if U contains n elements then there are 2^n subsets of U . For example, if $U = \{a, b, c\}$ then the set of subsets of U is the following collection of $2^3 = 8$ sets:

$$2^U = \{\emptyset, \{a\}, \{b\}, \{c\}, \{a, b\}, \{a, c\}, \{b, c\}, \{a, b, c\}\},$$

where \emptyset denotes the empty set, that is, a set with no elements.

The following notation is used to denote membership in a set and to denote that one set is contained in another: $x \in A$ means that x is an element of the set A (capital letters are used to denote sets and lower-case letters to denote elements) and $A \subseteq B$ means that A is a subset of B , that is, every element of A is also an element of B . Note that $A \subseteq B$ allows for the possibility that $A = B$.

Next we review operations that can be performed on sets.

- Let $A \in 2^U$. The *complement of A in U* , denoted by $\neg A$, is the set of elements of U that are not in A . When the “universe of discourse” U is clear from the context, one simply refers to $\neg A$ as the *complement of A* .

For example, if $U = \{a, b, c, d, e, f\}$ and $A = \{b, d, f\}$ then $\neg A = \{a, c, e\}$.

Note that $\neg U = \emptyset$ and $\neg \emptyset = U$.

- Let $A, B \in 2^U$. The *intersection of A and B*, denoted by $A \cap B$, is the set of elements that belong to both A and B.
For example, if $A = \{b, d, f\}$ and $B = \{a, b, d, e\}$ then $A \cap B = \{b, d\}$.
If $A \cap B = \emptyset$ we say that A and B are *disjoint*.
- Let $A, B \in 2^U$. The *union of A and B*, denoted by $A \cup B$, is the set of elements that belong to either A or B (or both).
For example, if $A = \{b, d, f\}$ and $B = \{a, b, d, e\}$ then $A \cup B = \{a, b, d, e, f\}$.

The above operations on sets are illustrated in Figure 9.1.




Figure 9.1: Operations on sets.

We denote by $A \setminus B$ the set of elements of A that are not in B. Thus, $A \setminus B = A \cap \neg B$. For example, if $A = \{b, d, f\}$ and $B = \{a, b, d, e\}$ then $A \setminus B = \{f\}$ and $B \setminus A = \{a, e\}$.

The following are known as *De Morgan's Laws*:

- $\neg(A \cup B) = \neg A \cap \neg B$
- $\neg(A \cap B) = \neg A \cup \neg B$

Let us verify De Morgan's Laws in the following example:

$U = \{a, b, c, d, e, f, g, h, i, j, k\}$, $A = \{b, d, f, g, h, i\}$ and $B = \{a, b, f, i, k\}$.

Then $\neg A = \{a, c, e, j, k\}$, $\neg B = \{c, d, e, g, h, j\}$, $A \cup B = \{a, b, d, f, g, h, i, k\}$

so that $\neg(A \cup B) = \{c, e, j\} = \neg A \cap \neg B$;

furthermore, $A \cap B = \{b, f, i\}$ so that $\neg(A \cap B) = \{a, c, d, e, g, h, j, k\} = \neg A \cup \neg B$.

9.1.2 Probability

In probability theory the “universal set” U is called the *sample space* and the subsets of U are called *events*. A *probability measure on U* is a function $P : 2^U \rightarrow [0, 1]$ that assigns to every event $E \in 2^U$ a number greater than or equal to 0 and less than or equal to 1, as shown in Figure 9.2, with the following restrictions:

1. $P(U) = 1$.
2. For every two events $E, F \in 2^U$, if $E \cap F = \emptyset$ then $P(E \cup F) = P(E) + P(F)$.




Figure 9.2: A probability measure.

From the above two properties one can obtain the following properties (the reader might want to try to prove them using Properties 1 and 2 above):

- $P(\neg E) = 1 - P(E)$, for every event E (this follows from the fact that E and $\neg E$ are disjoint and their union is equal to U).
- $P(\emptyset) = 0$ (this follows from the previous line and the fact that $\emptyset = \neg U$).
- For every two events $E, F \in 2^U$, $P(E \cup F) = P(E) + P(F) - P(E \cap F)$ (see Exercise 9.5).
- For every two events $E, F \in 2^U$, if $E \subseteq F$ then $P(E) \leq P(F)$.
- If $E_1, E_2, \dots, E_m \in 2^U$ ($m \geq 2$) is a collection of mutually disjoint sets (that is, for every $i, j = 1, \dots, m$ with $i \neq j$, $E_i \cap E_j = \emptyset$) then $P(E_1 \cup E_2 \cup \dots \cup E_m) = P(E_1) + P(E_2) + \dots + P(E_m)$.

When the set U is finite, a *probability distribution p on U* is a function that assigns to each element $z \in U$ a number $p(z)$, with $0 \leq p(z) \leq 1$, and such that $\sum_{z \in U} p(z) = 1$.

Given a probability distribution $p : U \rightarrow [0, 1]$ on U one can obtain a probability measure $P : 2^U \rightarrow [0, 1]$ by defining, for every event A , $P(A) = \sum_{z \in A} p(z)$.

Conversely, given a probability measure $P : 2^U \rightarrow [0, 1]$, one can obtain from it a probability distribution $p : U \rightarrow [0, 1]$ by defining, for every $z \in U$, $p(z) = P(\{z\})$. Thus, the two notions are equivalent.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 9.8.1 at the end of this chapter.

9.2 Probabilistic beliefs

An information set contains all the states that an individual considers possible, that is, the states that the individual cannot rule out, given her information. However, of all the states that are possible, the individual might consider some to be more likely than others and might even dismiss some states as “extremely unlikely” or “implausible”.

For example, suppose that there are only three students in a class: Ann, Bob and Carla. The professor tells them that in the last exam one of them got 95 points (out of 100), another 78 and the third 54. We can think of a state as a triple (a, b, c) , where a is Ann’s score, b is Bob’s score and c is Carla’s score. Then, based on the information given by the professor, Ann must consider all of the following states as possible: $(95, 78, 54)$, $(95, 54, 78)$, $(78, 95, 54)$, $(78, 54, 95)$, $(54, 95, 78)$ and $(54, 78, 95)$.

Suppose, however, that in all the previous exams Ann and Bob always obtained a higher score than Carla and often Ann outperformed Bob. Then Ann might consider states $(95, 78, 54)$ and $(78, 95, 54)$ much more likely than $(78, 54, 95)$ and $(54, 78, 95)$.

To represent such judgments of relative likelihood we add to an information set a probability distribution over the states in the information set. The probability distribution expresses the individual’s *beliefs*, while the information set represents what the individual *knows*. In this example, Ann’s beliefs could be represented by the following probability distribution:

$$\begin{array}{cccccc} (95, 78, 54) & (95, 54, 78) & (78, 95, 54) & (54, 95, 78) & (78, 54, 95) & (54, 78, 95) & (\blacklozenge) \\ \frac{9}{16} & \frac{4}{16} & \frac{2}{16} & \frac{1}{16} & 0 & 0 & \end{array}$$

According to these beliefs, Ann considers it very likely that she got the highest score, is willing to dismiss the possibility that Carla received the highest score as extremely unlikely (she assigns probability 0 to the two states where Carla’s score is 95) and considers it much more likely that she, rather than Bob, received the highest score.

Recall that, given a probability distribution $p : U \rightarrow [0, 1]$ on a set U (U can be thought of as an information set for the individual under consideration) and an event $E \subseteq U$, the probability of event E , denoted by $P(E)$, is defined as the sum of the probabilities of the elements of E :

$$P(E) = \sum_{x \in E} p(x).$$

For instance, continuing the above example, the proposition “Ann received the highest score” corresponds to event

$$E = \{(95, 78, 54), (95, 54, 78)\}$$

and – according to Ann’s beliefs – the probability of that event is

$$P(E) = p(95, 78, 54) + p(95, 54, 78) = \frac{9}{16} + \frac{4}{16} = \frac{13}{16} = 81.25\%.$$

On the other hand, the proposition “Bob’s score is higher than Carla’s score” corresponds to the event

$$F = \{(95, 78, 54), (78, 95, 54), (54, 95, 78)\}$$

and – according to Ann’s beliefs – the probability of that event is

$$P(F) = p(95, 78, 54) + p(78, 95, 54) + p(54, 95, 78) = \frac{9}{16} + \frac{2}{16} + \frac{1}{16} = \frac{12}{16} = 75\%.$$

Definition 9.2.1 We say that an individual is *certain* of an event E if she attaches probability 1 to E (that is, if $P(E) = 1$).

In the above example, Ann is certain of event

$$G = \{(95, 78, 54), (95, 54, 78), (78, 95, 54), (54, 95, 78)\},$$

corresponding to the proposition “Carla did not get the highest score”. She is also certain of every event H such that $G \subseteq H$.

R Note the important difference between knowledge and certainty: if at a state x the individual knows an event E (that is, $x \in KE$) then, at that state, E is indeed true ($x \in E$),¹ that is, it is never the case that an individual knows something which is false; on the other hand, an individual can be certain of something which is false, that is, it is possible that the individual assigns probability 1 to an event E even though the true state is not in E .²

Continuing the example of the exam scores, suppose that – before distributing the exams – the professor says “I was surprised to see that, this time, Ann did not get the highest score”. This new announcement by the professor informs the students that the true state is neither $(95, 78, 54)$ nor $(95, 54, 78)$. Thus we can view the effect of the new piece of information as shrinking Ann’s information set from

$$\{(95, 78, 54), (95, 54, 78), (78, 95, 54), (54, 95, 78), (78, 54, 95), (54, 78, 95)\}$$

to

$$\{(78, 95, 54), (54, 95, 78), (78, 54, 95), (54, 78, 95)\}.$$

How should Ann revise her beliefs in response to the new information?

The answer cannot be that we simply drop states $(95, 78, 54)$ and $(95, 54, 78)$ from the probability distribution (\blacklozenge) because the result would be

<i>state</i>	$(78, 95, 54)$	$(54, 95, 78)$	$(78, 54, 95)$	$(54, 78, 95)$
<i>probability</i>	$\frac{2}{16}$	$\frac{1}{16}$	0	0

which is not a probability distribution, since the probabilities do not add up to 1 (they add up to $\frac{3}{16}$). The topic of belief revision is addressed in Section 9.4.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 9.8.2 at the end of this chapter.

¹It was proved in Exercise 8.5 (Chapter 8) that, for every event E , $KE \subseteq E$.

²In the above example, if the true state is $(78, 54, 95)$ then it does not belong to event G , representing the proposition “Carla did not get the highest score” and yet Ann assigns probability 1 to G , that is, she is certain of G .

9.3 Conditional probability and Bayes' rule

9.3.1 Conditional probability

Let $A, B \subseteq U$ be two events (where U is the universal set or sample space) and P a probability measure on U .

If $P(B) > 0$, the *conditional probability of A given B* , denoted by $P(A|B)$, is defined as follows:

$$P(A|B) = \frac{P(A \cap B)}{P(B)} \quad (9.1)$$

For example, if $P(A \cap B) = 0.2$ and $P(B) = 0.6$ then $P(A|B) = \frac{0.2}{0.6} = \frac{1}{3}$.

One way to visualize conditional probability is to think of U as a geometric shape of area 1 (e.g. a square with each side equal to 1 unit of measurement).

For a subset A of the unit square, $P(A)$ is the area of A .

If B is a non-empty subset of the square then $A \cap B$ is that part of A that lies in B and $P(A|B)$ is the area of $A \cap B$ relative to the area of B , that is, as a fraction of the area of B . This is illustrated in Figure 9.3.




Figure 9.3: Geometric interpretation of the conditional probability $P(A|B)$.

Next we derive from the conditional probability formula (9.1) three versions of what is known as *Bayes' rule*.

Let E and F be two events such that $P(E) > 0$ and $P(F) > 0$. Then, using the conditional probability formula (9.1) we get

$$P(E|F) = \frac{P(E \cap F)}{P(F)} \quad (9.2)$$

and

$$P(F|E) = \frac{P(E \cap F)}{P(E)}. \quad (9.3)$$

From (9.3) we get that

$$P(E \cap F) = P(F|E)P(E) \quad (9.4)$$

and replacing (9.4) in (9.2) we get

$$\text{Bayes' formula version 1 : } \boxed{P(E|F) = \frac{P(F|E)P(E)}{P(F)}} \quad (9.5)$$

As an illustration of how one can use (9.5), consider the following example. You are a doctor examining a middle-aged man who complains of lower back pain. You know that 25% of men in the age group of your patient suffer from lower back pain.

There are various causes of lower back pain; one of them is chronic inflammation of the kidneys. This is not a very common disease: it affects only 4% of men in the age group that you are considering. Among those who suffer from chronic inflammation of the kidneys, 85% complain of lower back pain.

What is the probability that your patient has chronic inflammation of the kidneys? Let I denote inflammation of the kidneys and L denote lower back pain.

The information you have is that $P(I) = \frac{4}{100}$, $P(L) = \frac{25}{100}$ and $P(L|I) = \frac{85}{100}$. Thus, using (9.5), we get that

$$P(I|L) = \frac{P(L|I)P(I)}{P(L)} = \frac{\frac{85}{100} \left(\frac{4}{100} \right)}{\frac{25}{100}} = 0.136 = 13.6\%.$$

We now derive a second version of Bayes' formula. According to Bayes' rule (9.5), $P(E|F) = \frac{P(F|E)P(E)}{P(F)}$.

From set theory we have that, given any two sets A and B , $A = (A \cap B) \cup (A \cap \neg B)$ and the two sets $A \cap B$ and $A \cap \neg B$ are disjoint. Thus, $P(A) = P(A \cap B) + P(A \cap \neg B)$.

Hence, in the denominator of Bayes' formula (9.5) we can replace $P(F)$ with $P(F \cap E) + P(F \cap \neg E)$.

Then, using the formula for conditional probability we get that $P(F \cap E) = P(F|E)P(E)$ and $P(F \cap \neg E) = P(F|\neg E)P(\neg E)$. Thus, $P(F) = P(F|E)P(E) + P(F|\neg E)P(\neg E)$.

Replacing this in Bayes' formula (9.5) we get

$$\text{Bayes' formula version 2 : } P(E|F) = \frac{P(F|E)P(E)}{P(F|E)P(E) + P(F|\neg E)P(\neg E)} \quad (9.6)$$

As an illustration of how one can use (9.6), consider the following example.

Enrollment in a Decision Making class is as follows: 60% economics majors (E), 40% other majors ($\neg E$). In the past, 80% of the economics majors passed and 65% of the other majors passed.

A student tells you that she passed the class. What is the probability that she is an economics major? Let A stand for "pass the class". Then, using (9.6),

$$P(E|A) = \frac{P(A|E)P(E)}{P(A|E)P(E) + P(A|\neg E)P(\neg E)} = \frac{\frac{80}{100} \left(\frac{60}{100} \right)}{\frac{80}{100} \left(\frac{60}{100} \right) + \frac{65}{100} \left(\frac{40}{100} \right)} = \frac{24}{37} = 64.86\%.$$

One more example: 0.3763% of the population (that is, approximately 4 in 100,000 individuals) is infected with the HIV virus.

Let H be the event "a randomly selected individual has the HIV virus".

Then $P(H) = 0.003763$ and $P(\neg H) = 0.996237$.

A blood test can be used to detect the virus. The blood test has a true positive rate (sensitivity) of 99.97% and a true negative rate (specificity) of 98.5%. Thus, (letting '+' denote a positive blood test and '-' a negative blood test) $P(+|H) = 0.9997$, $P(-|H) = 0.0003$, $P(+|\neg H) = 0.015$ and $P(-|\neg H) = 0.985$.

Now suppose that you pick an individual at random, administer the blood test and it turns out to be positive. What is the probability that the individual has the HIV virus? That is, what is $P(H|+)$? Using (9.6),

$$\begin{aligned} P(H|+) &= \frac{P(+|H)P(H)}{P(+|H)P(H) + P(+|\neg H)P(\neg H)} \\ &= \frac{0.9997(0.003763)}{0.9997(0.003763) + 0.015(0.996237)} = 0.201 = 20.1\%. \end{aligned}$$

A generalization of (9.6) is as follows: If $\{E_1, \dots, E_n\}$ is a partition of the sample space U ,³ then, for every event F , $P(F) = P(F|E_1)P(E_1) + \dots + P(F|E_n)P(E_n)$.

Thus, using (9.5) we obtain that, for every $i = 1, \dots, n$,

$$\text{Bayes' formula version 3 : } P(E_i|F) = \frac{P(F|E_i)P(E_i)}{P(F|E_1)P(E_1) + \dots + P(F|E_n)P(E_n)} \quad (9.7)$$

³ That is, the sets E_1, \dots, E_n (1) cover the set U (in the sense that $E_1 \cup \dots \cup E_n = U$) and (2) are pairwise disjoint (in the sense that, for all $i, j = 1, \dots, n$ with $i \neq j$, $E_i \cap E_j = \emptyset$).

Example: enrollment in a class is restricted to the following majors: economics (E), statistics (S) and math (M). Current enrollment is: 40% E , 35% S and 25% M . Let A be the event “pass the class”. According to past data, $P(A|E) = 60\%$, $P(A|S) = 50\%$ and $P(A|M) = 75\%$.

A student from this class tells you that she received a passing grade. What is the probability that she is an economics major? Using (9.7),

$$\begin{aligned}
 P(E|A) &= \frac{P(A|E)P(E)}{P(A|E)P(E) + P(A|S)P(S) + P(A|M)P(M)} \\
 &= \frac{\frac{60}{100} \left(\frac{40}{100} \right)}{\frac{60}{100} \left(\frac{40}{100} \right) + \frac{50}{100} \left(\frac{35}{100} \right) + \frac{75}{100} \left(\frac{25}{100} \right)} = \frac{96}{241} = 39.83\%.
 \end{aligned}$$

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 9.8.3 at the end of this chapter.

9.4 Changing beliefs in response to new information

The issue of how to “rationally” modify one’s initial beliefs – expressed as a probability measure P on a set U – after receiving an item of information (represented by a subset F of U) has been studied extensively by philosophers and logicians. Two different situations may arise:

- In one case, the item of information F was not ruled out by the initial beliefs, in the sense that event F was assigned positive probability ($P(F) > 0$). Information might still be somewhat surprising, in case $P(F)$ is small (close to zero), but it is not completely unexpected. We call this case *belief updating*.
- The other case is where the item of information was initially dismissed, in the sense that it was assigned zero probability ($P(F) = 0$). In this case the information received is completely surprising or completely unexpected. We call this case *belief revision*.

We shall first address the issue of belief updating.

9.4.1 Belief updating

It is generally agreed that the rational way to update one’s beliefs is by conditioning the initial probability measure on the information received, that is, by using the conditional probability formula (9.1).

Definition 9.4.1 We use the expression *belief updating* or *Bayesian updating* to refer to the modification of initial beliefs (expressed by an initial probability distribution P) obtained by applying the conditional probability rule; this assumes that the belief change is prompted by the arrival of new information, represented by an event F such that $P(F) > 0$.

Thus, when receiving a piece of information $F \subseteq U$ such that $P(F) > 0$, one would change one's initial probability measure P into a new probability measure P_{new} by

- reducing the probability of every state in $\neg F$ (the complement of F) to zero (this captures the notion that the information represented by F is trusted to be correct), and
- setting $P_{new}(s) = P(s|F)$ for every state $s \in F$.

Thus, for every state $s \in U$,

$$P_{new}(s) = P(s|F) = \begin{cases} 0 & \text{if } s \notin F \\ \frac{P(s)}{P(F)} & \text{if } s \in F \end{cases} \quad (9.8)$$

(recall the assumption that $P(F) > 0$).

Thus, for every event $E \subseteq U$, $P_{new}(E) = \sum_{s \in E} P_{new}(s) = \sum_{s \in E} P(s|F) = P(E|F)$.

As an illustration, let us go back to the example of Section 9.2 concerning three students: Ann, Bob and Carla. The professor tells them that in the last exam one of them got 95 points (out of 100), another 78 and the third 54. We represented a state as a triple (a, b, c) , where a is Ann's score, b is Bob's score and c is Carla's score. The initial information given by the professor is thus represented by the set

$$U = \{(95, 78, 54), (95, 54, 78), (78, 95, 54), (78, 54, 95), (54, 95, 78), (54, 78, 95)\}$$

Based on the results of previous exams, Ann forms the following probabilistic beliefs:

$$\begin{array}{cccccc} (95, 78, 54) & (95, 54, 78) & (78, 95, 54) & (54, 95, 78) & (78, 54, 95) & (54, 78, 95) \\ \frac{9}{16} & \frac{4}{16} & \frac{2}{16} & \frac{1}{16} & 0 & 0 \end{array}$$

We then supposed that – before distributing the exams – the professor made the additional remark “I was surprised to see that, this time, Ann did not get the highest score”. This new announcement by the professor informs the students that the true state is neither $(95, 78, 54)$ nor $(95, 54, 78)$.

Thus the new piece of information is represented by the event

$$F = \{(78, 95, 54), (54, 95, 78), (78, 54, 95), (54, 78, 95)\}.$$

How should Ann revise her beliefs in response to this new piece of information?

Conditioning Ann's initial beliefs on the event F yields the following updated beliefs:

$$\begin{array}{cccccc} (95, 78, 54) & (95, 54, 78) & (78, 95, 54) & (54, 95, 78) & (78, 54, 95) & (54, 78, 95) \\ 0 & 0 & \frac{2}{3} & \frac{1}{3} & 0 & 0 \end{array}$$

These updated beliefs can be represented more succinctly as a probability distribution on the set F (that is, by dropping the states that belong to the complement of F , which are zero-probability states in the updated beliefs):

$$\begin{array}{cccc} (78, 95, 54) & (54, 95, 78) & (78, 54, 95) & (54, 78, 95) \\ \frac{2}{3} & \frac{1}{3} & 0 & 0 \end{array}$$

9.4.2 Belief revision

[Note: the material of this section will not be needed until Chapter 13. It is presented here for completeness on the topic of beliefs.]

How should a rational individual revise her beliefs when receiving information that is completely surprising, that is, when informed of an event E to which her initial beliefs assigned zero probability ($P(E) = 0$)?

As we will see in Part IV, belief revision is very important in dynamic (or extensive-form) games. In such games a player may find herself at an information set that, according to her initial beliefs, had zero probability of being reached and thus will have to form new beliefs reflecting the unexpected information.

The best known theory of rational belief revision is the so-called *AGM theory*, which takes its name from its originators: Alchourrón (a legal scholar), Gärdenfors (a philosopher) and Makinson (a computer scientist); their pioneering contribution was published in 1985. Just like the theory of expected utility (Chapter 5), the AGM theory is an axiomatic theory: it provides a list of “rationality” axioms for belief revision and provides a representation theorem.⁴

Although the AGM theory was developed within the language of propositional logic, it can be restated in terms of a set of states and a collection of possible items of information represented as events. We first introduce the non-probabilistic version of the theory and then add graded beliefs, that is, probabilities.

Let U be a finite set of states and $\mathcal{E} \subseteq 2^U$ a collection of events (subsets of U) representing possible items of information; we assume that $U \in \mathcal{E}$ and $\emptyset \notin \mathcal{E}$. To represent initial beliefs and revised beliefs we introduce a function $f : \mathcal{E} \rightarrow 2^U$, which we call a belief revision function.

Definition 9.4.2 Let U be a finite set of states and \mathcal{E} a collection of events such that $U \in \mathcal{E}$ and $\emptyset \notin \mathcal{E}$. A *belief revision function* is a function $f : \mathcal{E} \rightarrow 2^U$ that satisfies the following properties: for every $E \in \mathcal{E}$, (1) $f(E) \subseteq E$ and (2) $f(E) \neq \emptyset$.

The interpretation of a belief revision function is as follows. First of all, $f(U)$ represents the initial beliefs, namely the set of states that the individual initially considers possible.⁵ Secondly, for every $E \in \mathcal{E}$, $f(E)$ is the set of states that the individual would consider possible if informed that the true state belongs to E ; thus $f(E)$ represents the individual’s *revised beliefs* after receiving information E .⁶

⁴We will not list and discuss the axioms here. The interested reader can consult <http://plato.stanford.edu/entries/formal-belief/> or, for a discussion which is closer to the approach followed in this section, Bonanno (2009).

⁵The universal set U can be thought of as representing minimum information: all states are possible. If the initial beliefs were to be expressed probabilistically, by means of a probability distribution P over U , then $f(U)$ would be the support of P , that is, the set of states to which P assigns positive probability. Thus, $f(U)$ would be the smallest event of which the individual would initially be certain (that is, to which she assigns probability 1): she would initially be certain of (assign probability 1 to) any event F such that $f(U) \subseteq F$.

⁶If the revised beliefs after receiving information E were to be expressed probabilistically, by means of a probability distribution P_E over U , then $f(E)$ would be the support of P_E , that is, the set of states to which P_E assigns positive probability. Thus, $f(E)$ would be the smallest event of which the individual would be certain *after having been informed that E* : according to her revised beliefs she would be certain of any event F such that $f(E) \subseteq F$. [Note that, since – by assumption – $f(E) \subseteq E$, the individual is assumed to be certain of the information received (e.g. because she trusts the source of the information).]

One of the implications of the AGM axioms for belief revision is the following condition, which is known as Arrow's Axiom (proposed by the Nobel laureate Ken Arrow in the context of rational choice, rather than rational belief revision):

$$\text{if } E, F \in \mathcal{E}, E \subseteq F \text{ and } E \cap f(F) \neq \emptyset \text{ then } f(E) = E \cap f(F).$$

Arrow's Axiom says that if information E implies information F ($E \subseteq F$) and there are states in E that would be considered possible upon receiving information F ($E \cap f(F) \neq \emptyset$), then the states that the individual would consider possible if informed that E are precisely those that belong to both E and $f(F)$ ($f(E) = E \cap f(F)$).

Although necessary for a belief revision policy that satisfies the AGM axioms, Arrow's Axiom is not sufficient. Before stating the necessary and sufficient conditions for rational belief revision, we remind the reader of the notion of a complete and transitive relation on a set U (Chapter 2, Section 2.1).

- ◇ In Chapter 2 the relation was denoted by \succsim and was interpreted in terms of *preference*: $o_1 \succsim o_2$ was interpreted as “the individual considers outcome o_1 to be at least as good as outcome o_2 ”.
- ◇ In the present context the interpretation is in terms of “plausibility”: $s \succsim s'$ means that the individual considers state s to be *at least as plausible* as state s' ; $s \succ s'$ means that s is considered to be *more plausible* than s' and $s \sim s'$ means that s is considered to be *just as plausible* as s' .

Definition 9.4.3 A *plausibility order* on a set of states U is a binary relation \succsim on U that is complete (for every two states s_1 and s_2 , either $s_1 \succsim s_2$ or $s_2 \succsim s_1$, or both) and transitive if $s_1 \succsim s_2$ and $s_2 \succsim s_3$ then $s_1 \succsim s_3$). We define $s_1 \succ s_2$ as “ $s_1 \succsim s_2$ and $s_2 \not\succeq s_1$ ” and we define $s_1 \sim s_2$ as “ $s_1 \succsim s_2$ and $s_2 \succsim s_1$ ”.

The following theorem is based on a result by Adam Grove.⁷

Theorem 9.4.1 Let U be a finite set of states, \mathcal{E} a collection of events (representing possible items of information), with $U \in \mathcal{E}$ and $\emptyset \notin \mathcal{E}$, and $f : \mathcal{E} \rightarrow 2^U$ a belief revision function (Definition 9.4.2). Then the belief revision policy represented by the function f is compatible with the AGM axioms of belief revision if and only if there exists a plausibility order \succsim on U that *rationalizes* f in the sense that, for every $E \in \mathcal{E}$, $f(E)$ is the set of most plausible states in E : $f(E) = \{s \in E : s \succsim s' \text{ for every } s' \in E\}$.

⁷Adam Grove, Two modelings for theory change, *Journal of Philosophical Logic*, 1988, Vol. 17, pages 157-170. That result was proved within the context of propositional logic. The version given here is proved in Giacomo Bonanno, Rational choice and AGM belief revision, *Artificial Intelligence*, 2009, Vol. 88, pages 221-241.

Definition 9.4.4 A belief revision function $f : \mathcal{E} \rightarrow 2^U$ which is rationalized by a plausibility order is called an *AGM belief revision function*.

R An AGM belief revision function satisfies Arrow's Axiom. The converse is not true: it is possible for a belief revision function $f : \mathcal{E} \rightarrow 2^U$ to satisfy Arrow's Axiom and yet fail to be rationalized by a plausibility order.

Within the context of probabilistic beliefs, let P be the probability distribution on a finite set of states U that represents the initial beliefs and P_E be the probability distribution representing the updated beliefs after receiving information E such that $P(E) > 0$ (thus the information is not surprising).

The *support* of a probability distribution P , denoted by $Supp(P)$, is the set of states to which P assigns positive probability: $Supp(P) = \{s \in U : P(s) > 0\}$.

The rule for *updating* beliefs upon receiving information E (Definition 9.4.1) implies the following:

$$\text{if } E \cap Supp(P) \neq \emptyset \text{ (that is, } P(E) > 0) \text{ then } Supp(P_E) = E \cap Supp(P). \quad (9.9)$$

We call this the *qualitative belief updating rule* or qualitative Bayes' rule. It is easy to check that the qualitative belief updating rule is implied by Arrow's Axiom (see Exercise 9.22). Thus, by the above remark, an AGM belief revision function has incorporated in it the qualitative belief updating rule. In other words, *belief updating is included in the notion of AGM belief revision*.

A belief revision function, however, goes beyond belief updating because it also encodes new beliefs after receipt of surprising information (that is, after being informed of an event E such that $P(E) = 0$).

What is the probabilistic version of AGM belief revision? It turns out that in order to obtain probabilistic beliefs we only need to make a simple addition to an AGM belief revision function $f : \mathcal{E} \rightarrow 2^U$.

- Let P_0 be any full-support probability distribution on U (that is, P_0 is such that $P_0(s) > 0$, for every $s \in U$).
- Then, for every $E \in \mathcal{E}$, let P_E be the probability distribution obtained by conditioning P_0 on $f(E)$ (note: on $f(E)$, not on E):

$$P_E(s) = P_0(s|f(E)) = \begin{cases} \frac{P_0(s)}{\sum_{s' \in f(E)} P_0(s')} & \text{if } s \in f(E) \\ 0 & \text{if } s \notin f(E) \end{cases}$$

- Then P_U gives the initial probabilistic beliefs and, for every other $E \in \mathcal{E}$, P_E gives the revised probabilistic beliefs after receiving information E .
- The collection $\{P_E\}_{E \in \mathcal{E}}$ of probability distributions on U so obtained gives the individual's *probabilistic belief revision policy* (while the function $f : \mathcal{E} \rightarrow 2^U$ gives the individual's *qualitative belief revision policy*).

Definition 9.4.5 Let U be a finite set of states and \mathcal{E} a collection of events such that $U \in \mathcal{E}$ and $\emptyset \notin \mathcal{E}$. A *probabilistic belief revision policy* is a collection $\{P_E\}_{E \in \mathcal{E}}$ of probability distributions on U such that, for every $E \in \mathcal{E}$, $\text{Supp}(P_E) \subseteq E$. P_U represents the initial beliefs and, for every other $E \in \mathcal{E}$, P_E represents the revised beliefs after receiving information E . The collection $\{P_E\}_{E \in \mathcal{E}}$ is called an *AGM probabilistic belief revision policy* if it satisfies the following properties:

1. there exists a plausibility order \succsim on U such that, for every $E \in \mathcal{E}$, $\text{Supp}(P_E)$ is the set of most plausible states in E , that is,

$$\text{Supp}(P_E) = \{s \in E : s \succsim s' \text{ for every } s' \in E\},^a$$

2. there exists a full-support probability distribution P_0 on U such that, for every $E \in \mathcal{E}$, P_E is the probability distribution obtained by conditioning P_0 on $\text{Supp}(P_E)$.

^aThis condition says that if one defines the function $f : \mathcal{E} \rightarrow 2^U$ by $f(E) = \text{Supp}(P_E)$ then this function is an AGM belief revision function (see Definition 9.4.4).

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 9.8.4 at the end of this chapter.

9.5 Harsanyi consistency of beliefs or like-mindedness

[Note: the material of this section will not be needed until Chapter 14. It is presented here for completeness on the topic of beliefs.]

We can easily extend the analysis to the case of two or more individuals. We already know how to model interactive knowledge by means of information partitions; the addition of beliefs is a simple step: we merely add, for every individual and for every information set, a probability distribution over the elements of that information set.⁸ A two-person example is shown in Figure 9.4.




Figure 9.4: A two-person knowledge-belief structure.

⁸The probability distribution over an information set can be thought of as a probability distribution over the universal set U by assigning probability zero to every state which is not in the information set.

In this example, at every state, the two individuals hold different beliefs. For example, consider event $E = \{b, c\}$ and state a . Individual 1 attaches probability $\frac{1}{2}$ to E while Individual 2 attaches probability $\frac{1}{3}$ to E .

Can two “equally rational” individuals hold different beliefs? The answer is: of course! In the above example it is not surprising that the two individuals assign different probabilities to the same event E , because they have different information.

If the true state is a , then Individual 1’s information is that the true state is either a or b or c , while Individual 2 considers only a and b possible (Individual 2 knows more than Individual 1).

Is there a precise way of expressing the fact that two individuals assign different probabilities to an event *exclusively because they have different information*?

In the above example we could ask the hypothetical question: if Individual 1 had the same information as Individual 2, would he agree with Individual 2’s assessment that the probability of event $E = \{b, c\}$ is $\frac{1}{3}$?

This is, of course, a counterfactual question. The answer to this counterfactual question is affirmative: imagine giving Individual 1 the information that the true state is either a or b ; then – according to Definition 9.4.1 – he would update his beliefs from $\begin{pmatrix} a & b & c \\ \frac{1}{2} & \frac{1}{4} & \frac{1}{4} \end{pmatrix}$ to $\begin{pmatrix} a & b \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}$ and thus have the same beliefs as Individual 2.

We say that two individuals are *like-minded* if it is the case that they would have the same beliefs if they had the same information.

It is not straightforward how to turn this into a precise definition.

Consider, again, the example of Figure 9.4 and event $E = \{b, c\}$. Above we asked the question “what would Individual 1 believe if he knew as much as Individual 2?” This is a simple question because we can imagine giving more information to Individual 1 and have him update his beliefs based on that information.

However, we could also have asked the question “what would Individual 2 believe if he knew as little as Individual 1?” In this case we would have to imagine “taking away information” from Individual 2, by increasing his information set from $\{a, b\}$ to $\{a, b, c\}$. This is not the same as asking Individual 2 to update his beliefs based on information $\{a, b, c\}$, because updating on something you already know leaves your beliefs unchanged.

There is a sense in which the beliefs of the two individuals of Figure 9.4 are “in agreement”: for example, they both consider state a twice as likely as state b . One could try to use this condition to define like-mindedness: for every two states x and y , whenever two individuals consider both x and y as possible (given their information) then they agree on the relative likelihood of x versus y .

Unfortunately, this condition is too weak. To see this, consider the three-individual example of Figure 9.5.




Figure 9.5: A three-person knowledge-belief structure.

No matter what the true state is, we cannot find two individuals and two states x and y that both individuals consider possible. Thus any beliefs would make the three individuals like-minded. For example, consider the following beliefs:

- Individual 1 at his information set $\{b, c\}$ assigns equal probability to b and c (thus considering b to be as likely as c),
- Individual 2 at her information set $\{a, b\}$ assigns probability $\frac{3}{4}$ to a and $\frac{1}{4}$ to b (thus considering a to be three times more likely than b), and
- Individual 3 at his information set $\{a, c\}$ assigns probability $\frac{1}{4}$ to a and $\frac{3}{4}$ to c (thus considering c to be three times more likely than a).

Then, putting together the beliefs of Individuals 1 and 2, we would have that a is judged to be three times more likely than c (according to Individual 2, a is three times more likely than b and, according to Individual 1, b is just as likely as c), while Individual 3 has the opposite judgment that c is three times more likely than a .

In order to give a precise definition of like-mindedness we need to introduce some notation.

- Let there be n individuals ($n \geq 2$).
- Let U be a set of states and let \mathcal{I}_i be the information partition of individual i ($i \in \{1, \dots, n\}$).
- As usual, if s is a state, we denote by $I_i(s)$ the element (information set) of the partition \mathcal{I}_i that contains s .
- Let $P_{i,s}$ denote the beliefs of individual i at state s , that is, $P_{i,s}$ is a probability distribution over $I_i(s)$.
- Clearly, we can think of $P_{i,s}$ as a probability distribution over the entire set of states U satisfying the property that if $s' \notin I_i(s)$ then $P_{i,s}(s') = 0$.
- Note that, for every individual i , there is just one probability distribution over $I_i(s)$ and thus if $s' \in I_i(s)$ then $P_{i,s} = P_{i,s'}$.⁹

Definition 9.5.1 A probability distribution P over U is called a *common prior* if, for every individual i and for every state s ,

1. $P(I_i(s)) > 0$, and
2. updating P on $I_i(s)$ (see Definition 9.4.1) yields precisely $P_{i,s}$, that is, $P(s'|I_i(s)) = P_{i,s}(s')$, for every $s' \in I_i(s)$.

When a common prior exists we say that the individuals are *like-minded* or that the individuals' beliefs are *Harsanyi consistent*.^a

^aJohn Harsanyi, who in 1994 won the Nobel Memorial prize in Economics (together with Reinhardt Selten and John Nash), introduced the *theory of games of incomplete information* which will be the object of Part V. In that theory the notion of Harsanyi consistency plays a crucial role.

⁹Thus, every individual knows his own probabilistic beliefs.

For instance, in the example of Figure 9.6, which reproduces Figure 9.4, a common prior exists and thus the two individuals are like-minded. Indeed, a common prior is

$$\begin{pmatrix} a & b & c & d & e \\ \frac{2}{8} & \frac{1}{8} & \frac{1}{8} & \frac{2}{8} & \frac{2}{8} \end{pmatrix}$$

(the reader should convince himself/herself that, indeed, updating this probability distribution on each information set yields the probability distribution written inside that information set). How can we determine if a common prior exists? The issue of existence




Figure 9.6: Copy of Figure 9.4.

of a common prior can be reduced to the issue of whether a system of equations has a solution. To see this, let us go back to the example of Figure 9.6. A common prior would be a probability distribution

$$\begin{pmatrix} a & b & c & d & e \\ p_a & p_b & p_c & p_d & p_e \end{pmatrix}$$

that satisfies the following conditions:

1. Updating on information set $\{a, b, c\}$ of Individual 1 we need

$$\frac{p_b}{p_a + p_b + p_c} = \frac{1}{4} \quad \text{and} \quad \frac{p_c}{p_a + p_b + p_c} = \frac{1}{4}.$$

Note that from these two conditions the third condition follows, namely

$$\frac{p_a}{p_a + p_b + p_c} = \frac{1}{2}.$$

2. Updating on information set $\{d, e\}$ of Individual 1 we need $\frac{p_d}{p_d + p_e} = \frac{1}{2}$, from which it follows that $\frac{p_e}{p_d + p_e} = \frac{1}{2}$.

3. Updating on information set $\{a, b\}$ of Individual 2 we need $\frac{p_a}{p_a+p_b} = \frac{2}{3}$, from which it follows that $\frac{p_b}{p_a+p_b} = \frac{1}{3}$.
4. Updating on information set $\{c, d\}$ of Individual 2 we need $\frac{p_c}{p_c+p_d} = \frac{1}{3}$, from which it follows that $\frac{p_d}{p_c+p_d} = \frac{2}{3}$.

From the first condition we get $p_b = p_c$, from the second $p_d = p_e$, from the third $p_a = 2p_b$ and from the fourth $p_d = 2p_c$.

Adding to these three equalities the requirement that $p_a + p_b + p_c + p_d + p_e = 1$, we have a system of five equations in five unknowns, which admits a solution, namely

$$\begin{pmatrix} a & b & c & d & e \\ \frac{2}{8} & \frac{1}{8} & \frac{1}{8} & \frac{2}{8} & \frac{2}{8} \end{pmatrix}.$$

It is not always the case that a common prior exists. For instance, if we add to the example of Figure 9.5 the beliefs shown in Figure 9.7, then we get a situation where the individuals' beliefs are not Harsanyi consistent. In this case, from the updating conditions of Individuals 1 and 2 one would get that $p_a = p_b$ and $p_b = p_c$, from which it follows that $p_a = p_c$; however, from the updating condition for Individual 3 we get that $p_a = 3p_c$, yielding a contradiction.

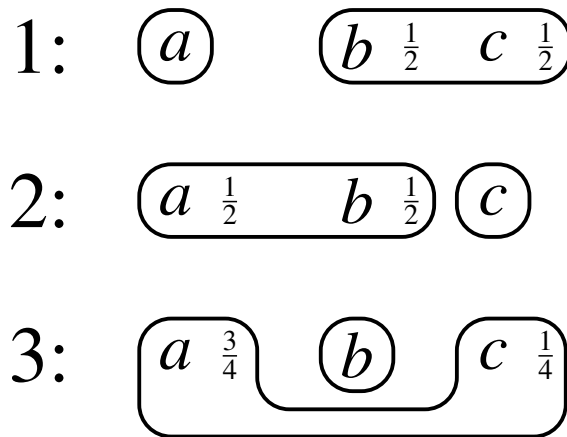


Figure 9.7: The structure of Figure 9.5 with the addition of beliefs.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 9.8.5 at the end of this chapter.

9.6 Agreeing to disagree

Can two rational and like-minded individuals agree to disagree? This question was raised in 1976 by Robert Aumann (who received the Nobel Memorial prize in Economics in 2005, together with Thomas Schelling). As remarked above, it is certainly quite possible for two rational individuals to have different beliefs about a particular event, because they might have different information. Let us go back to the example of Figure 9.4, reproduced in Figure 9.8 together with the common prior (showing that the individuals are like-minded).




Figure 9.8: The structure of Figure 9.4 with a common prior.

Suppose that the true state is *a* and consider what the two individuals believe about event $E = \{b, c\}$.

- ◇ Individual 1’s information set is $\{a, b, c\}$ and, given his beliefs at that information set, he attaches probability $\frac{1}{4} + \frac{1}{4} = \frac{1}{2}$ to E : let us denote this by “at state *a*, $P_1(E) = \frac{1}{2}$ ”.
- ◇ Individual 2’s information set is $\{a, b\}$ and, given her beliefs at that information set, she attaches probability $\frac{1}{3}$ to E : let us denote this by “at state *a*, $P_2(E) = \frac{1}{3}$ ”.
- ◇ Thus the two individuals disagree about the probability of event E . Furthermore, they know that they disagree.

To see this, let $\|P_1(E) = \frac{1}{2}\|$ be the event that Individual 1 attaches probability $\frac{1}{2}$ to E ; then

$$\|P_1(E) = \frac{1}{2}\| = \{a, b, c\}.$$

Similarly, let $\|P_2(E) = \frac{1}{3}\|$ be the event that Individual 2 attaches probability $\frac{1}{3}$ to E ; then

$$\|P_2(E) = \frac{1}{3}\| = \{a, b, c, d\}.$$

These are events and thus we can check at what states the two individuals know them.

Using Definition 8.1.3 (Chapter 8), we have that

$$K_1 \underbrace{\|P_2(E) = \frac{1}{3}\|}_{=\{a,b,c,d\}} = \{a, b, c\} \quad \text{and} \quad K_2 \underbrace{\|P_1(E) = \frac{1}{2}\|}_{=\{a,b,c\}} = \{a, b\}.$$

Hence

$$a \in \|P_1(E) = \frac{1}{2}\| \cap \|P_2(E) = \frac{1}{3}\| \cap K_1 \|P_2(E) = \frac{1}{3}\| \cap K_2 \|P_1(E) = \frac{1}{2}\|.$$

Thus at state a not only do the individuals disagree, but they know that they disagree. However, their disagreement is not common knowledge. Indeed, $K_1 K_2 \|P_1(E) = \frac{1}{2}\| = \emptyset$ and thus $a \notin K_1 K_2 \|P_1(E) = \frac{1}{2}\| = \emptyset$, that is, at state a it is not the case that Individual 1 knows that Individual 2 knows that Individual 1 assigns probability $\frac{1}{2}$ to event E . As the following theorem states, the opinions of two like-minded individuals about an event E cannot be in disagreement and, at the same time, commonly known.

The following theorem is proved in Section 9.7.

Theorem 9.6.1 — Agreement Theorem; Aumann, 1976. Let U be a set of states and consider a knowledge-belief structure with two individuals, 1 and 2. Let E be an event and let $p, q \in [0, 1]$. Suppose that at some state s it is common knowledge that Individual 1 assigns probability p to E and Individual 2 assigns probability q to E . Then, if the individuals' beliefs are Harsanyi consistent (Definition 9.5.1), $p = q$. In other words, two like-minded individuals *cannot agree to disagree* about the probability of an event. Formally, if there exists a common prior and $s \in CK(\|P_1(E) = p\| \cap \|P_2(E) = q\|)$ then $p = q$.

Another way to think about this result is to imagine that the two individuals communicate their opinions to each other. Hearing that the other person has a different opinion is in itself a valuable piece of information, which ought to be incorporated (by updating) into one's beliefs. Thus sequential communication leads to changing beliefs. If, at the end of this process, the beliefs become common knowledge, then they must be identical.¹⁰ We shall illustrate this with an example.

Imagine two scientists who agree that the laws of Nature are such that the true state of the world must be one of seven, call them a, b, c, d, e, f, g . They also agree on the relative likelihood of these possibilities, which they take to be as follows:

$$\begin{array}{ccccccc} a & b & c & d & e & f & g \\ \frac{4}{32} & \frac{2}{32} & \frac{8}{32} & \frac{5}{32} & \frac{7}{32} & \frac{2}{32} & \frac{4}{32} \end{array}$$

Experiments can be conducted to learn more. An experiment leads to a partition of the set of states. For example, if the true state is a and you perform an experiment then you might learn that the true state cannot be d or e or f or g but you still would not know which is the true state among the remaining ones. Suppose that the scientists agree that Scientist 1, from now on denoted by S_1 , will perform experiment 1 and Scientist 2, denoted by S_2 , will perform experiment 2. They also agree that each experiment would lead to a partition of the set of states as shown in Figure 9.9.




Figure 9.9: The partitions representing the two experiments.

¹⁰This line of reasoning was investigated by Geanakoplos and Polemarchakis (1982).

Suppose that the scientists are interested in establishing the truth of a proposition that is represented by the event $E = \{a, c, d, e\}$. Initially (given their shared probabilistic beliefs) they agree that the probability that E is true is 75%:

$$P(E) = P(a) + P(c) + P(d) + P(e) = \frac{4}{32} + \frac{8}{32} + \frac{5}{32} + \frac{7}{32} = \frac{24}{32} = \frac{3}{4} = 75\%.$$

Before they perform the experiments they also realize that, depending on what the true state is, after the experiment they will have an updated probability of event E conditional on what the experiment has revealed.

For example, they agree that if one performs Experiment 1 and the true state is b then the experiment will yield the information $F = \{a, b, c\}$ and $P(E|F)$ is given by

$$P(E|F) = \frac{P(E \cap F)}{P(F)} = \frac{P(a) + P(c)}{P(a) + P(b) + P(c)} = \frac{\frac{4}{32} + \frac{8}{32}}{\frac{4}{32} + \frac{2}{32} + \frac{8}{32}} = \frac{12}{14} = \frac{6}{7} = 86\%.$$

[Note the interesting fact that sometimes experiments, although they are informative – that is, they reduce one’s state of uncertainty – might actually induce one to become more confident of the truth of something that is false: in this case one would increase one’s subjective probability that E is true from 75% to 86%, although E is actually false if the true state is b , as we hypothesized.]

We can associate with every cell of each experiment (that is, with every possible state of information yielded by the experiment) a new updated probability of event E , as shown in Figure 9.10.




Figure 9.10: The probability of event E conditional on the result of each experiment.


Suppose now that each scientist goes to her laboratory and performs the respective experiment (Scientist 1 performs Experiment 1 and Scientist 2 performs Experiment 2). Assume also that

the true state is f

Suppose that the two scientists send each other an e-mail communicating their new subjective estimates of the truth of E . Scientist 1 writes that he now attaches probability $\frac{12}{14}$ to E , while Scientist 2 says that she attaches probability $\frac{15}{21}$ to E . So their estimates disagree (not surprisingly, since they have performed different experiments and have collected different information). Should they be happy with these estimates? The answer is negative.

Consider first S_1 (Scientist 1). He hears that S_2 has a revised probability of $\frac{15}{21}$ (recall our assumption that the true state is f). What does he learn from this? He learns that the true state cannot be d (if it had been, then he would have received an e-mail from S_2 saying “my new probability is $\frac{9}{11}$ ”). S_1 ’s new state of knowledge and corresponding probabilities after receiving S_2 ’s e-mail are then as shown in Figure 9.11.


Scientist 1 learns from Scientist 2’s email:




The new information partition of S_1 is explained as follows.
 If g is the true state, then experiment 1 reveals it.
 If either d, e or f is the true state, then experiment 1 reveals $\{d, e, f\}$.
 With e or f , S_2 ’s e-mail is “ $\frac{15}{21}$ ”, while with d it is “ $\frac{9}{11}$ ”.
 Thus S_2 ’s email splits $\{d, e, f\}$ into $\{d\}$ and $\{e, f\}$.
 If the true state is a, b or c , then Experiment 1 reveals $\{a, b, c\}$.
 With a or b , S_2 sends “ $\frac{9}{11}$ ” while with c she sends “ $\frac{15}{21}$ ”.
 Thus S_2 ’s email splits $\{a, b, c\}$ into $\{a, b\}$ and $\{c\}$.

Figure 9.11: S_1 ’s assessment of the probability of event E after receiving S_2 ’s first e-mail.

Consider now Scientist 2. From S_1 's e-mail she learns that S_1 has a new updated probability of $\frac{12}{14}$. What can she deduce from this? That the true state is *not* g (if it had been g then she would have received an e-mail from S_1 saying that her revised probability of E was zero). Thus she can revise her knowledge partition by eliminating g from her information set. A similar reasoning applies to the other states. S_2 's new state of knowledge and corresponding probabilities after receiving S_1 's e-mail are shown in Figure 9.12.



What Scientist 2 learns from Scientist 1's first e-mail:



The new information partition of S_2 is explained as follows.
 If either c, e, f or g is the true state, then Experiment 2 reveals $\{c, e, f, g\}$.
 With c or e , S_1 's e-mail is " $\frac{12}{14}$ ", while with g it is "0".
 Thus S_1 's email splits $\{c, e, f, g\}$ into $\{c, e, f\}$ and $\{g\}$.
 If the true state is a, b or d , then Experiment 2 reveals $\{a, b, d\}$ and, in all three cases, S_1 's email is " $\frac{12}{14}$ ". Thus S_2 learns nothing from S_1 's email.

Figure 9.12: S_2 's assessment of the probability of event E after receiving S_1 's first e-mail.

Thus the new situation after the first exchange of e-mails is as shown in Figure 9.13.




Figure 9.13: The new situation after the first e-mail exchange.

Now there is a second round of e-mails. S_1 communicates “in light of your e-mail, my new $P(E)$ is $\frac{7}{9}$ ”, while S_2 writes “after your e-mail I changed my $P(E)$ to $\frac{15}{17}$ ”. While S_1 learns nothing new from S_2 ’s second e-mail, S_2 learns that the true state cannot be c (the second e-mail would have been “ $P(E) = 1$ ” if c had been the true state; in the hypothetical case where S_2 ’s revised information was $\{a, b, d\}$ then after S_1 ’s second e-mail it would split into $\{a, b\}$ and $\{d\}$). Thus the new situation is as shown in Figure 9.14.

Now the two scientists have reached complete agreement: $P(E) = \frac{7}{9}$. Further exchanges do not convey any new information. Indeed it has become common knowledge (at state f) that both scientists estimate the probability of event E to be $\frac{7}{9} = 78\%$ (before the experiments the probability of E was judged to be $\frac{24}{32} = 75\%$; note, again, that with the experiments and the exchange of information they have gone farther from the truth than at the beginning!).




Figure 9.14: The new situation after the second e-mail exchange.

Notice that before the last step it was never common knowledge between the two what probability each scientist attached to E . When one scientist announced his subjective estimate, the other scientist found that announcement informative and revised her own estimate accordingly. After the exchange of two e-mails, the further announcement by one scientist of his/her estimate of the probability of E would not make the other scientist change his/her own estimate: the announcement would reveal nothing new.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 9.8.6 at the end of this chapter.

9.7 Proof of the Agreement Theorem

First we prove the following.

Lemma 9.1 Let U be a finite set of states, P a probability distribution on U and $E, F \subseteq U$ two events. Let $\{F_1, \dots, F_m\}$ be a partition of F (thus $F = F_1 \cup \dots \cup F_m$ and any two F_j and F_k with $j \neq k$ are non-empty and disjoint). Suppose that $P(E|F_j) = q$ for all $j = 1, \dots, m$. Then $P(E|F) = q$.

Proof. By definition of conditional probability, $P(E|F_j) = \frac{P(E \cap F_j)}{P(F_j)}$. Hence, since $P(E|F_j) = q$, we have that $P(E \cap F_j) = qP(F_j)$. Adding over j , the left-hand side becomes $P(E \cap F)$ [because $E \cap F = (E \cap F_1) \cap \dots \cap (E \cap F_m)$ and for any j and k with $j \neq k$, $(E \cap F_j) \cap (E \cap F_k) = \emptyset$, so that $P(E \cap F) = P(E \cap F_1) + \dots + P(E \cap F_m)$] and the right-hand side becomes $qP(F)$. Hence $P(E|F) = \frac{P(E \cap F)}{P(F)} = \frac{qP(F)}{P(F)} = q$. ■

Proof of Theorem 9.6.1. Suppose that $CK(\|P_1(E) = p\| \cap \|P_2(E) = q\|) \neq \emptyset$. Let P be a common prior. Select an arbitrary $s \in CK(\|P_1(E) = p\| \cap \|P_2(E) = q\|)$ and let $I_{CK}(s)$ be the cell of the common knowledge partition containing s . Consider Individual 1. $I_{CK}(s)$ is equal to the union of a collection of cells (information sets) of 1's information partition. On each such cell 1's conditional probability of E , using the common prior P , is p . Hence, by Lemma 9.1, $P(E|I_{CK}(s)) = p$. A similar reasoning for Individual 2 leads to $P(E|I_{CK}(s)) = q$. Hence $p = q$. ■

9.8 Exercises

9.8.1 Exercises for Section 9.1: Sets and probability

The answers to the following exercises are in Section 9.9 at the end of this chapter.

Exercise 9.1

Let U be the universal set (or sample space) and E and F two events. Let the complement of E be denoted by $\neg E$ and the complement of F by $\neg F$.

Suppose that $P(E) = \frac{3}{10}$, $P(F) = \frac{3}{5}$ and $P(\neg E \cup \neg F) = \frac{4}{5}$. What is the probability of $E \cup F$?

Exercise 9.2

Consider the following probability distribution: $\left(\begin{array}{cccccc} z_1 & z_2 & z_3 & z_4 & z_5 & z_6 & z_7 \\ \frac{3}{12} & \frac{1}{12} & 0 & \frac{3}{12} & \frac{2}{12} & 0 & \frac{3}{12} \end{array} \right)$.

What is the probability of the event $\{z_2, z_3, z_5, z_6, z_7\}$?

Exercise 9.3

Let the universal set be $U = \{z_1, z_2, z_3, z_4, z_5, z_6, z_7, z_8\}$. Let $A = \{z_2, z_4, z_5, z_7\}$, $B = \{z_3, z_6, z_8\}$, $C = \{z_2, z_6\}$, $D = \{z_3, z_4\}$ and $E = \{z_7, z_8\}$.

You are given the following data: $P(A \cup B) = \frac{21}{24}$, $P(A \cap C) = \frac{5}{24}$, $P(B \cap C) = \frac{3}{24}$, $P(A \cap D) = \frac{2}{24}$, $P(B \cap D) = \frac{3}{24}$, $P(B) = \frac{7}{24}$ and $P(E) = \frac{2}{24}$.

- Find the probability $P(z_i)$ for each $i = 1, \dots, 8$.
- Calculate $P((A \cup B) \cap (C \cup D))$.

Exercise 9.4

Let $U = \{a, b, c, d, e, f, g, h, i\}$ and consider the following probability distribution:

$$\left(\begin{array}{cccccc} a & b & c & d & e & f & g & h & i \\ \frac{11}{60} & 0 & \frac{7}{60} & \frac{9}{60} & \frac{16}{60} & \frac{5}{60} & \frac{4}{60} & \frac{8}{60} & 0 \end{array} \right)$$

- Let $E = \{a, f, g, h, i\}$. What is the probability of E ?
- List all the events that have probability 1.

Exercise 9.5

Let P be a probability measure on a finite set U and let A and B be two events (that is, subsets of U). Explain why $P(A \cup B) = P(A) + P(B) - P(A \cap B)$. ■

Exercise 9.6

You plan to toss a fair coin three times and record the sequence of Heads/Tails.

- (a) What is the set of possibilities (or universal set or sample space)?
- (b) Let E be the event that you will get at least one Heads. What is E ?
- (c) What is the probability of event E ?
- (d) Let F be the event that you will get Tails either in the first toss or in the third toss? [Note: this is *not* an exclusive 'or'.] What is event F ?
- (e) What is the probability of event F ? ■

9.8.2 Exercises for Section 9.2: Probabilistic beliefs

The answers to the following exercises are in Section 9.9 at the end of this chapter.

Exercise 9.7

Let the set of states be $W = \{a, b, c, d, e, f, g, h, i\}$.

Amy's initial beliefs are given by the following probability distribution:

<i>state</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>	<i>g</i>	<i>h</i>	<i>i</i>
<i>probability</i>	$\frac{11}{60}$	0	$\frac{7}{60}$	$\frac{9}{60}$	$\frac{16}{60}$	$\frac{5}{60}$	$\frac{4}{60}$	$\frac{8}{60}$	0

- (a) Let $E = \{a, f, g, h, i\}$. What is the probability of E ?
- (b) Find all the events that Amy is certain of. ■

9.8.3 Exercises for Section 9.3: Conditional probability and Bayes' rule

The answers to the following exercises are in Section 9.9 at the end of this chapter.

Exercise 9.8

Let A and B be two events such that $P(A) > 0$ and $P(B) > 0$. Prove that $P(A|B) = P(B|A)$ if and only if $P(A) = P(B)$. ■

Exercise 9.9

Two events A and B are *independent* if $P(A|B) = P(A)$.

Construct an example to show that $P(A|B) = P(B|A)$ and $P(A) = P(B)$ but A and B are **not** independent. ■

Exercise 9.10

There is an urn with 40 balls: 4 red, 16 white, 10 blue and 10 black. You close your eyes and pick a ball at random. Let E be the event “the selected ball is either red or white”.

- (a) What is the probability of E ?
- (b) Now somebody tells you: “the ball in your hand is not black”. How likely is it now that you picked either a red or a white ball? ■

Exercise 9.11

Suppose there are 3 individuals. It is known that one of them has a virus. A blood test can be performed to test for the virus. If an individual does have the virus, then the result of the test will be positive.

However, the test will be positive also for an individual who does not have the virus but has a particular defective gene. It is known that exactly one of the three individuals has this defective gene: it could be the same person who has the virus or somebody who does not have the virus.

A test result will come up positive if and only if either the patient has the virus or the defective gene (or both).

Suppose that Individual 1 takes the blood test and the result is positive. Assuming that all the states are equally likely, what is the probability that he has the virus? [Hint: think of the universal set (or sample space) U as a list of states and each state tells you which individual has the virus and which individual has the defective gene.] ■

Exercise 9.12

Let A and B be two events such that $P(A) = 0.2$, $P(B) = 0.5$ and $P(B|A) = 0.1$. Calculate $P(A|B)$. ■

Exercise 9.13

In a remote rural clinic with limited resources, a patient arrives complaining of low-abdomen pain. Based on all the information available, the doctor thinks that there are only four possible causes: a bacterial infection (b), a viral infection (v), cancer (c), internal bleeding (i).

Of the four, only the bacterial infection and internal bleeding are treatable at the clinic. In the past the doctor has seen 600 similar cases and they eventually turned out to be as follows:

b : bacterial infection	v : viral infection	c : cancer	i : internal bleeding
140	110	90	260

The doctor's probabilistic estimates are based on those past cases.

- (a) What is the probability that the patient has a treatable disease?

There are two possible ways of gathering more information: a blood test and an ultrasound. A positive blood test will reveal that there is an infection, however it could be either bacterial or viral; a negative blood test rules out an infection and thus leaves cancer and internal bleeding as the only possibilities. The ultrasound, on the other hand, will reveal if there is internal bleeding.

- (b) Suppose that the patient gets an ultrasound and it turns out that there is no internal bleeding. What is the probability that he does **not** have a treatable disease? What is the probability that he has cancer?
- (c) If instead of getting the ultrasound he had taken the blood test and it had been positive, what would the probability that he had a treatable disease have been?
- (d) Now let us go back to the hypothesis that the patient only gets the ultrasound and it turns out that there is no internal bleeding. He then asks the doctor: "if I were to take the blood test too (that is, in addition to the ultrasound), how likely is it that it would be positive?". What should the doctor's answer be?
- (e) Finally, suppose that the patient gets both the ultrasound and the blood test and the ultrasound reveals that there is no internal bleeding, while the blood test is positive. How likely is it that he has a treatable disease?

Exercise 9.14

A lab technician was asked to mark some specimens with two letters, the first from the set $\{A, B, C\}$ and the second from the set $\{E, F, G\}$. For example, a specimen could be labeled as AE or BG , etc. He had a total of 220 specimens. He has to file a report to his boss by filling in the table shown in Figure 9.15. Unfortunately, he does not remember all the figures. He had written some notes to himself, which are reproduced below. Fill in the table with the help of his notes and conditional probabilities.

Here are the technician's notes:

LABEL	Number
AE	
AF	
AG	
BE	
BF	
BG	
CE	
CF	
CG	

Figure 9.15: The specimen example of Exercise 9.14.

- (a) Of all the ones that he marked with an E , $\frac{1}{5}$ were also marked with an A and $\frac{1}{5}$ were marked with a B .
- (b) He marked 36 specimens with the label CE .
- (c) Of all the specimens that he marked with a C , the fraction $\frac{12}{23}$ were marked with a G .
- (d) Of all the specimens, the fraction $\frac{23}{55}$ were marked with a C .
- (e) The number of specimens marked BG was twice the number of specimens marked BE .
- (f) Of all the specimens marked with an A , the fraction $\frac{3}{20}$ were marked with an E .
- (g) Of all the specimens marked with an A , $\frac{1}{10}$ were marked with a G .

9.8.4 Exercises for Section 9.4: Changing beliefs in response to new information

The answers to the following exercises are in Section 9.9 at the end of this chapter.

Exercises for Section 9.4.1: Belief updating

Exercise 9.15

Let the set of states be $U = \{a, b, c, d, e, f, g\}$. Fran's initial beliefs are as follows:

<i>state</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>	<i>g</i>
<i>probability</i>	$\frac{3}{20}$	0	$\frac{7}{20}$	$\frac{1}{20}$	0	$\frac{4}{20}$	$\frac{5}{20}$

Consider the event $E = \{a, d, e, g\}$.

- (a) What probability does Fran attach to event E ?
- (b) Suppose that Fran is now informed that E is indeed true (that is, that the true state belongs to E). What are her updated beliefs?
- (c) Consider the event $D = \{a, b, c, f, g\}$. What probability does Fran assign to event D (1) initially and (2) after she is informed that E ?

Exercise 9.16

Consider again the example where there are only three students in a class: Ann, Bob and Carla and the professor tells them that in the last exam one of them got 95 points (out of 100), another 78 and the third 54. Ann's initial beliefs are as follows (where the triple (a, b, c) is interpreted as follows: a is Ann's score, b is Bob's score and c is Carla's score):

$$\begin{array}{cccccc} (95, 78, 54) & (95, 54, 78) & (78, 95, 54) & (54, 95, 78) & (78, 54, 95) & (54, 78, 95) \\ \frac{16}{32} & \frac{8}{32} & \frac{4}{32} & \frac{2}{32} & \frac{1}{32} & \frac{1}{32} \end{array}$$

- (a) Suppose that (before distributing the exams) the professor tells the students that Carla received a lower score than Bob. Let E be the event that represents this information. What is E ?
- (b) How should Ann update her beliefs in response to information E ?

Exercise 9.17

Let the set of states be $U = \{a, b, c, d, e, f, g\}$.

Bill's initial beliefs are as follows:

a	b	c	d	e	f	g
$\frac{3}{20}$	$\frac{2}{20}$	$\frac{5}{20}$	$\frac{1}{20}$	$\frac{1}{20}$	$\frac{3}{20}$	$\frac{5}{20}$

- (a) Suppose that Bill receives information $E = \{a, c, e, f, g\}$. What are his updated beliefs?
- (b) Suppose that, after receiving information E , he later learns a new piece of information, namely $F = \{b, d, e, f, g\}$. What are his final beliefs (that is, after updating first on E and then on F)?

Exercise 9.18

Inspector Gethem has been put in charge of a museum robbery that took place yesterday. Two precious items were stolen: a statuette and a gold tiara, which were displayed in the same room. Surveillance cameras show that only three people visited the room at the time the items disappeared: call them suspect A , suspect B and suspect C .

Let a state be a complete specification of who stole what (including the possibility that the same person stole both items).

- (a) List all the states.
- (b) Inspector Gethem recognizes the suspects and, based on what he knows about them, initially believes that the probability that suspect A stole both items is $\frac{1}{20}$, the probability that suspect B stole both items is $\frac{3}{20}$ and the probability that suspect C stole both items is $\frac{4}{20}$. Furthermore, he assigns equal probability to every other state.

What are his initial beliefs?

- (c) Suppose now that the inspector receives reliable information that suspect B did not steal the statuette and suspect C did not steal the tiara. What are his beliefs after he updates on this information? ■

Exercise 9.19

Let the set of states be $U = \{a, b, c, d, e, f, g\}$ and let $E = \{a, d, e, g\}$.

The individual's initial beliefs are given by the following probability distribution, call it P :

a	b	c	d	e	f	g
$\frac{3}{20}$	0	$\frac{7}{20}$	$\frac{1}{20}$	0	$\frac{4}{20}$	$\frac{5}{20}$

- (a) Calculate $P(E)$, $P(b|E)$ and $P(d|E)$.
 (b) Calculate the updated beliefs in response to information E . ■

Exercise 9.20

The instructor of a class has the following data on enrollment:

major	Economics	Mathematics	Philosophy	Psychology	Statistics
enrollment	35%	22%	18%	16%	9%

- (a) A student in her class, Jim, tells her that his major is neither Math nor Statistics. What are the instructor's beliefs about Jim's major upon learning this?
 (b) After a while Jim further informs the instructor that he is not an Economics major. What are the instructor's beliefs about Jim's major upon learning this second fact?
 (c) Finally, Jim tells the instructor that he is not a Philosophy major. What are the instructor's beliefs about Jim's major upon learning this third fact? ■

Exercises for Section 9.4.2: Belief revision

The answers to the following exercises are in Section 9.9 at the end of this chapter.

Exercise 9.21

Prove that an AGM belief revision function (Definition 9.4.4) satisfies Arrow's Axiom: if $E, F \in \mathcal{E}$, $E \subseteq F$ and $E \cap f(F) \neq \emptyset$ then $f(E) = E \cap f(F)$. ■

Exercise 9.22

Prove that the qualitative belief updating rule (9.9) (page 307) is implied by Arrow's Axiom. ■

Exercise 9.23

Let $U = \{a, b, c, d, e, g, h, k, m\}$ and let \succsim be the following plausibility order on U (as usual, we use the convention that if the row to which state s belongs is above the row to which state s' belongs then $s \succ s'$, and if s and s' belong to the same row then $s \sim s'$).

most plausible	b, g
	c, k, m
	d, h
	e
least plausible	a

Let $\mathcal{E} = \{\{a, e\}, \{d, e, k, m\}, \{b, d, e, k\}, U\}$.

Find the belief revision function $f: \mathcal{E} \rightarrow 2^U$ that is rationalized by \succsim . ■

Exercise 9.24

As in Exercise 9.23, let $U = \{a, b, c, d, e, g, h, k, m\}$ and

$$\mathcal{E} = \left\{ \underbrace{\{a, e\}}_E, \underbrace{\{d, e, k, m\}}_F, \underbrace{\{b, d, e, k\}}_G, U \right\}$$

Using the plausibility order of Exercise 9.23, namely

most plausible	b, g
	c, k, m
	d, h
	e
least plausible	a

find a collection of probability distributions $\{P_E, P_F, P_G, P_U\}$ that provides an AGM probabilistic belief revision policy (Definition 9.4.5). [There are many; find one.] ■

9.8.5 Exercises for Section 9.5: Harsanyi consistency of beliefs or like-mindedness

The answers to the following exercises are in Section 9.9 at the end of this chapter.

Exercise 9.25

Consider the knowledge-belief structure shown in Figure 9.16. Are the individuals' beliefs Harsanyi consistent? ■

Exercise 9.26

Consider the knowledge-belief structure shown in Figure 9.17.

- (a) Show that if $p = \frac{1}{13}$ then the beliefs are Harsanyi consistent.
- (b) Show that if $p \neq \frac{1}{13}$ then the beliefs are **not** Harsanyi consistent. ■

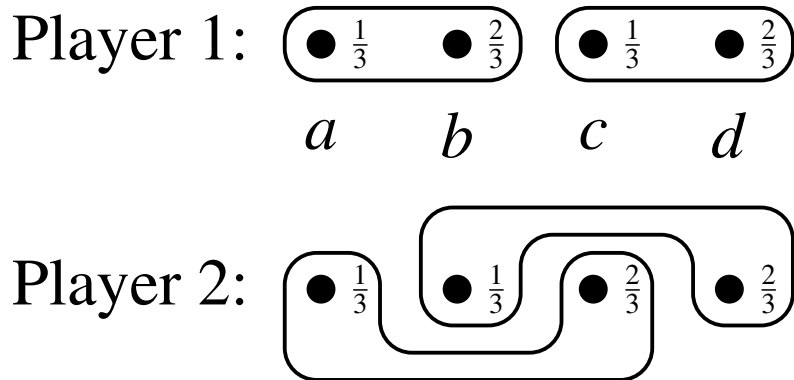


Figure 9.16: The knowledge-belief structure for Exercise 9.25.

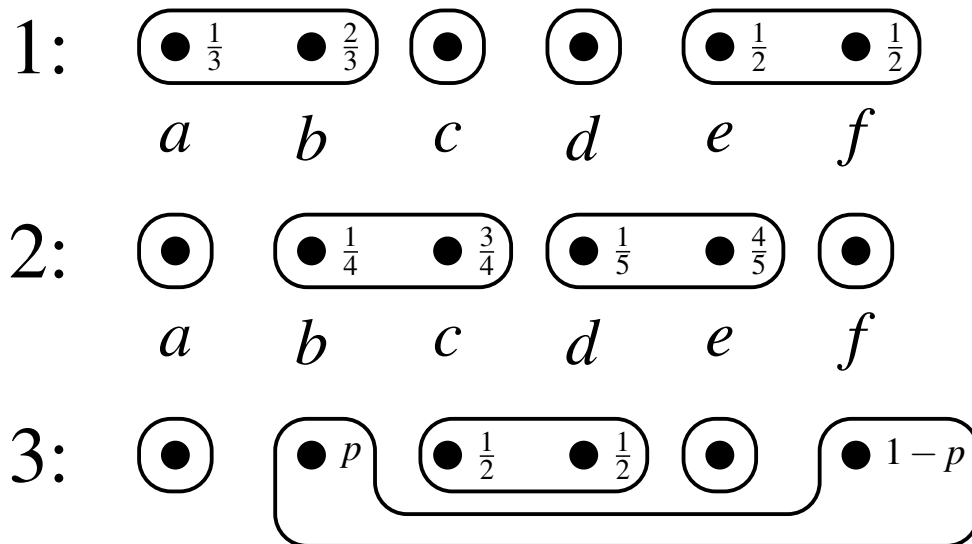


Figure 9.17: The knowledge-belief structure for Exercise 9.26.

9.8.6 Exercises for Section 9.6: Agreeing to disagree

The answers to the following exercises are in Section 9.9 at the end of this chapter.




Figure 9.18: The knowledge-belief structure for Exercise 9.27.

Exercise 9.27 Consider the knowledge-belief structure shown in Figure 9.18.

- (a) Find the common knowledge partition.
- (b) Find a common prior.
- (c) Let $E = \{b, d, f\}$.
 - (1) Is the probability that Individual 1 assigns to E common knowledge?
 - (2) Is the probability that Individual 2 assigns to E common knowledge?
- (d) Let $E = \{b, d, f\}$.
 - (1) At state b does Individual 1 know what probability Individual 2 assigns to E ?
 - (2) At state c does Individual 1 know what probability Individual 2 assigns to E ?
 - (3) At state f does Individual 1 know what probability Individual 2 assigns to E ?

Exercise 9.28 — ★★★ Challenging Question ★★★.

This is known as the *Monty Hall problem*.

You are a contestant in a show. You are shown three doors, numbered 1, 2 and 3. Behind one of them is a new car, which will be yours if you choose to open that door.

The door behind which the car was placed was chosen randomly with equal probability (a die was thrown, if it came up 1 or 2 then the car was placed behind Door 1, if it came up 3 or 4 then the car was placed behind Door 2 and if it came up 5 or 6 then the car was placed behind Door 3).

You have to choose one door. Suppose that you have chosen door number 1. Before the door is opened the host tells you that he knows where the car is and, to help you, he will open one of the other two doors, making sure that he opens a door behind which there is no car; if there are two such doors, then he will choose randomly with equal probability.

Afterwards he will give you a chance to change your mind and switch to the other closed door, but you will have to pay \$20 if you decide to switch.

Suppose that initially you chose Door 1 and the host opens Door 3 to show you that the car is not there. Assume that, if switching increases the probability of getting the car (relative to not switching), then you find it worthwhile to pay \$20 to switch.

Should you switch from Door 1 to Door 2?

Answer the question using two different approaches.

- (a) **Method 1.** Let D_n denote the event that the car is behind door n and let O_n denote the event that the host opens door n ($n \in \{1, 2, 3\}$).

The prior probabilities are $P(D_1) = P(D_2) = P(D_3) = \frac{1}{3}$. Compute $P(D_1|O_3)$ using Bayes' rule (if $P(D_1|O_3) \geq \frac{1}{2}$ then you should not switch, since there is a cost in switching).

- (b) **Method 2.** Draw (part of) an extensive form with imperfect information where Nature moves first and chooses where the car is, then you choose one door and then the host chooses which door to open (of course, the host's choice is made according to the rules specified above).

Reasoning about the information set you are in after you have pointed to door 1 and the host has opened door 3, determine if you should switch from door 1 to door 2.

■

9.9 Solutions to Exercises

Solution to Exercise 9.1. The general formula is $P(E \cup F) = P(E) + P(F) - P(E \cap F)$.

By The Morgan's Law, $\neg E \cup \neg F = \neg(E \cap F)$.

Thus, since $P(\neg(E \cap F)) = \frac{4}{5}$, we have that $P(E \cap F) = 1 - \frac{4}{5} = \frac{1}{5}$.

Hence, $P(E \cup F) = \frac{3}{10} + \frac{3}{5} - \frac{1}{5} = \frac{7}{10}$. □

Solution to Exercise 9.2.

$$P(\{z_2, z_3, z_5, z_6, z_7\}) = \sum_{i \in \{2, 3, 5, 6, 7\}} P(\{z_i\}) = \frac{1}{12} + 0 + \frac{2}{12} + 0 + \frac{3}{12} = \frac{1}{2}. \quad \square$$

Solution to Exercise 9.3.

(a) Since $\{z_1\}$ is the complement of $A \cup B$, $P(z_1) = 1 - \frac{21}{24} = \frac{3}{24}$.

Since $\{z_2\} = A \cap C$, $P(z_2) = \frac{5}{24}$.

Similarly, $P(z_6) = P(B \cap C) = \frac{3}{24}$, $P(z_3) = P(B \cap D) = \frac{3}{24}$ and

$P(z_4) = P(A \cap D) = \frac{2}{24}$.

Thus, $P(z_8) = P(B) - P(z_3) - P(z_6) = \frac{7}{24} - \frac{3}{24} - \frac{3}{24} = \frac{1}{24}$.

Hence, $P(z_7) = P(E) - P(z_8) = \frac{2}{24} - \frac{1}{24} = \frac{1}{24}$. Finally, $P(z_5) = 1 - \sum_{i \neq 5} P(z_i) = \frac{6}{24}$.

Thus, the probability distribution is:

$$\begin{pmatrix} z_1 & z_2 & z_3 & z_4 & z_5 & z_6 & z_7 & z_8 \\ \frac{3}{24} & \frac{5}{24} & \frac{3}{24} & \frac{2}{24} & \frac{6}{24} & \frac{3}{24} & \frac{1}{24} & \frac{1}{24} \end{pmatrix}$$

(b) $A \cup B = \{z_2, z_3, z_4, z_5, z_6, z_7, z_8\}$, $C \cup D = \{z_2, z_3, z_4, z_6\}$.

Hence, $(A \cup B) \cap (C \cup D) = C \cup D = \{z_2, z_3, z_4, z_6\}$ so that

$$P((A \cup B) \cap (C \cup D)) = P(z_2) + P(z_3) + P(z_4) + P(z_6) = \frac{5}{24} + \frac{3}{24} + \frac{2}{24} + \frac{3}{24} = \frac{13}{24}. \quad \square$$

Solution to Exercise 9.4. The given probability distribution is:

$$\begin{pmatrix} a & b & c & d & e & f & g & h & i \\ \frac{11}{60} & 0 & \frac{7}{60} & \frac{9}{60} & \frac{16}{60} & \frac{5}{60} & \frac{4}{60} & \frac{8}{60} & 0 \end{pmatrix}.$$

(a) Let $E = \{a, f, g, h, i\}$.

$$\text{Then } P(E) = P(a) + P(f) + P(g) + P(h) + P(i) = \frac{11}{60} + \frac{5}{60} + \frac{4}{60} + \frac{8}{60} + 0 = \frac{28}{60} = \frac{7}{15}.$$

(b) The events that have probability 1 are:

$$\{a, c, d, e, f, g, h\} = U \setminus \{b, i\}, \quad \{a, b, c, d, e, f, g, h\} = U \setminus \{i\},$$

$$\{a, c, d, e, f, g, h, i\} = U \setminus \{b\} \quad \text{and} \quad \{a, b, c, d, e, f, g, h, i\} = U. \quad \square$$

Solution to Exercise 9.5. Since $P(A) = \sum_{w \in A} P(w)$ and $P(B) = \sum_{w \in B} P(w)$, when adding

$P(A)$ to $P(B)$ the elements that belong to both A and B (that is, the elements of $A \cap B$) are

added twice and thus we need to subtract $\sum_{w \in A \cap B} P(w)$ from $P(A) + P(B)$ in order to get

$$\sum_{w \in A \cup B} P(w) = P(A \cup B). \quad \square$$

Solution to Exercise 9.6.

(a) There are 8 possibilities:

$$HHH \quad HHT \quad HTH \quad HTT \quad THH \quad THT \quad TTH \quad TTT$$

Since the coin is fair, each possibility has the same probability, namely $\frac{1}{8}$.

(b) $E = U \setminus \{TTT\}$, where U is the universal set (the set of 8 possibilities listed above).

(c) $P(E) = P(U) - P(TTT) = 1 - \frac{1}{8} = \frac{7}{8}$.

(d) $F = U \setminus \{HHH, HTH\}$

(e) $P(F) = P(U) - P(\{HHH, HTH\}) = 1 - \frac{1}{8} - \frac{1}{8} = \frac{6}{8} = \frac{3}{4}$. □

Solution to Exercise 9.7. The probability distribution is

<i>state</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>	<i>g</i>	<i>h</i>	<i>i</i>
<i>probability</i>	$\frac{11}{60}$	0	$\frac{7}{60}$	$\frac{9}{60}$	$\frac{16}{60}$	$\frac{5}{60}$	$\frac{4}{60}$	$\frac{8}{60}$	0

(a) Let $E = \{a, f, g, h, i\}$.

Then $P(E) = P(a) + P(f) + P(g) + P(h) + P(i) = \frac{11}{60} + \frac{5}{60} + \frac{4}{60} + \frac{8}{60} + 0 = \frac{28}{60} = \frac{7}{15}$.

(b) Amy is certain of all events that have probability 1, namely

$\{a, c, d, e, f, g, h\}$, $\{a, b, c, d, e, f, g, h\}$, $\{a, c, d, e, f, g, h, i\}$ and $\{a, b, c, d, e, f, g, h, i\}$. □

Solution to Exercise 9.8. Suppose that $P(A|B) = P(B|A)$.

Since $P(A|B) = \frac{P(A \cap B)}{P(B)}$ and $P(B|A) = \frac{P(A \cap B)}{P(A)}$ it follows that $P(A) = P(B)$.

Conversely, if $P(A) = P(B)$ then $P(A|B) = \frac{P(A \cap B)}{P(B)} = \frac{P(A \cap B)}{P(A)} = P(B|A)$. □

Solution to Exercise 9.9. Example 1. Let $P(A) = \frac{1}{2}$ and let $B = \neg A$.

Then $P(B) = 1 - P(A) = \frac{1}{2}$ and $A \cap B = \emptyset$ so that $P(A \cap B) = 0$

and thus $P(A|B) = \frac{P(A \cap B)}{P(B)} = \frac{0}{\frac{1}{2}} = 0 = P(B|A) = \frac{P(A \cap B)}{P(A)}$.

Thus $P(A|B) = P(B|A) = 0$ and $P(A) = P(B) = \frac{1}{2}$ and A and B are not independent since $P(A|B) \neq P(A)$.

Example 2. $U = \{a, b, c\}$, $P(a) = P(c) = \frac{2}{5}$ and $P(b) = \frac{1}{5}$. Let $A = \{a, b\}$ and $B = \{b, c\}$.

Then $P(A) = P(B) = \frac{3}{5}$ and $P(A|B) = \frac{P(A \cap B)}{P(B)} = \frac{P(b)}{P(b) + P(c)} = \frac{\frac{1}{5}}{\frac{3}{5}} = \frac{1}{3} = P(B|A)$.

Thus A and B are not independent since $P(A|B) \neq P(A)$. □

Solution to Exercise 9.10.

(a) $P(E) = \frac{4+16}{40} = \frac{1}{2}$.

(b) Let F be the event “the selected ball is not black”. Then, initially, $P(F) = \frac{30}{40} = \frac{3}{4}$.

Furthermore, $E \cap F = E$. Thus, $P(E|F) = \frac{P(E \cap F)}{P(F)} = \frac{P(E)}{P(F)} = \frac{\frac{1}{2}}{\frac{3}{4}} = \frac{2}{3}$. □

Solution to Exercise 9.11. First we list the possible states. A state is a complete description of the facts that are relevant: it tells you who has the virus and who has the gene.

Let us represent a state as a pair (x, y) interpreted as follows: individual x has the virus and individual y has the defective gene.

Then $U = \{a = (1, 1), b = (1, 2), c = (1, 3), d = (2, 1), e = (2, 2), f = (2, 3), g = (3, 1), h = (3, 2), i = (3, 3)\}$.

Let V_1 be the event “Individual 1 has the virus”. Then $V_1 = \{a, b, c\}$.

Let G_1 be the event “Individual 1 has the defective gene”. Then $G_1 = \{a, d, g\}$.

Since every state is assumed to have probability $\frac{1}{9}$, $P(V_1) = P(G_1) = \frac{1}{9} + \frac{1}{9} + \frac{1}{9} = \frac{1}{3}$.

Let 1_+ be the event that a blood test administered to Individual 1 comes up positive.

Then $1_+ = \{a, b, c, d, g\}$ and $P(1_+) = \frac{5}{9}$.

Now we can compute the requested conditional probability as follows (note that $V_1 \cap 1_+ = V_1$):

$$P(V_1|1_+) = \frac{P(V_1 \cap 1_+)}{P(1_+)} = \frac{P(V_1)}{P(1_+)} = \frac{\frac{1}{3}}{\frac{5}{9}} = \frac{3}{5} = 60\%.$$

□

Solution to Exercise 9.12.

Using Bayes’ rule, $P(A|B) = \frac{P(B|A)P(A)}{P(B)} = \frac{(0.1)(0.2)}{0.5} = 0.04 = 4\%$. □

Solution to Exercise 9.13. The probabilities are as follows:

$$\begin{array}{c|c|c|c} b & v & c & i \\ \hline \frac{140}{600} = \frac{14}{60} & \frac{110}{600} = \frac{11}{60} & \frac{90}{600} = \frac{9}{60} & \frac{260}{600} = \frac{26}{60} \end{array}$$

(a) The event that the patient has a treatable disease is $\{b, i\}$.

$$P(\{b, i\}) = P(b) + P(i) = \frac{14}{60} + \frac{26}{60} = \frac{2}{3}.$$

(b) A negative result of the ultrasound is represented by the event $\{b, v, c\}$.

A non-treatable disease is the event $\{v, c\}$. Thus,

$$P(\{v, c\}|\{b, v, c\}) = \frac{P(\{v, c\} \cap \{b, v, c\})}{P(\{b, v, c\})} = \frac{P(\{v, c\})}{P(\{b, v, c\})} = \frac{\frac{11}{60} + \frac{9}{60}}{\frac{14}{60} + \frac{11}{60} + \frac{9}{60}} = \frac{10}{17} = 58.82\%.$$

$$P(c|\{b, v, c\}) = \frac{P(c)}{P(\{b, v, c\})} = \frac{\frac{9}{60}}{\frac{14}{60} + \frac{11}{60} + \frac{9}{60}} = \frac{9}{34} = 26.47\%.$$

(c) A positive blood test is represented by the event $\{b, v\}$.

A treatable disease is the event $\{b, i\}$. Thus,

$$P(\{b, i\}|\{b, v\}) = \frac{P(\{b, i\} \cap \{b, v\})}{P(\{b, v\})} = \frac{P(b)}{P(\{b, v\})} = \frac{\frac{14}{60}}{\frac{14}{60} + \frac{11}{60}} = \frac{14}{25} = 56\%.$$

(d) Here we want

$$P(\{b, v\}|\{b, v, c\}) = \frac{P(\{b, v\})}{P(\{b, v, c\})} = \frac{\frac{14}{60} + \frac{11}{60}}{\frac{14}{60} + \frac{11}{60} + \frac{9}{60}} = \frac{25}{34} = 73.53\%.$$

(e) We are conditioning on $\{b, v\} \cap \{b, v, c\} = \{b, v\}$;

thus, we want $P(\{b, i\}|\{b, v\})$ which was calculated in Part (c) as $\frac{14}{25} = 56\%$. \square

Solution to Exercise 9.14.

Let $\#xy$ be the *number* of specimens that were marked xy (thus, $x \in \{A, B, C\}$ and $y \in \{D, E, F\}$) and $P(xy) = \frac{\#xy}{220}$ be the *fraction* of specimens that were marked xy ; let $\#z$ be the number of specimens whose label contains a $z \in \{A, B, C, D, E, F\}$ and let $P(z) = \frac{\#z}{220}$; finally, let $P(xy|z) = \frac{\#xy}{\#z}$.¹¹

With this notation we can re-write the information contained in the technician's notes as follows.

(a) $P(AE|E) = P(BE|E) = \frac{1}{5}$. It follows that the remaining three fifth were marked with a C, that is $P(CE|E) = \frac{3}{5}$

(b) $\#CE = 36$; thus, $P(CE) = \frac{36}{220}$. Since $P(CE|E) = \frac{P(CE)}{P(E)}$, using (a) we get $\frac{3}{5} = \frac{\frac{36}{220}}{P(E)}$, that is, $P(E) = \frac{36}{220} \left(\frac{5}{3}\right) = \frac{3}{11}$. Hence, the number of specimens marked with an E is $\frac{3}{11} 220 = 60$. Furthermore, since $P(AE|E) = \frac{P(AE)}{P(E)}$, using (a) we get $\frac{1}{5} = \frac{P(AE)}{\frac{3}{11}}$, that is, $P(AE) = \frac{3}{55}$. Thus, the number of specimens marked AE is $\frac{3}{55} 220 = 12$. The calculation for $P(BE|E)$ is identical; thus, the number of specimens marked BE is also 12. So far, we have:

¹¹This is a conditional probability, since $P(xy|z) = \frac{\#xy}{\#z} = \frac{\frac{\#xy}{220}}{\frac{\#z}{220}} = \frac{P(xy)}{P(z)}$.

LABEL	number
<i>AE</i>	12
<i>AF</i>	
<i>AG</i>	
<i>BE</i>	12
<i>BF</i>	
<i>BG</i>	
<i>CE</i>	36
<i>CF</i>	
<i>CG</i>	

(c) $P(CG|C) = \frac{12}{23}$. Since $P(CG|C) = \frac{P(CG)}{P(C)}$, it follows that $\frac{12}{23} = \frac{P(CG)}{P(C)}$.

(d) $P(C) = \frac{23}{55}$. Thus, using (c) we get $\frac{12}{23} = \frac{P(CG)}{\frac{23}{55}}$, that is, $P(CG) = \frac{12}{55}$.

Hence, the number of specimens marked *CG* is $\frac{12}{55} 220 = 48$.

Since $P(C) = \frac{23}{55}$, the total number of specimens marked with a *C* is $\frac{23}{55} 220 = 92$.

Since 36 were marked *CE* (see the above table) and 48 were marked *CG*, it follows that the number of specimens marked *CF* is $92 - 48 - 36 = 8$. Up to this point we have:

LABEL	number
<i>AE</i>	12
<i>AF</i>	
<i>AG</i>	
<i>BE</i>	12
<i>BF</i>	
<i>BG</i>	
<i>CE</i>	36
<i>CF</i>	8
<i>CG</i>	48

(e) The number of *BG*s is twice the number of *BE*s.

Since the latter is 12 (see the above table), the number of *BG*s is 24.

(f) $P(AE|A) = \frac{3}{20}$. Since $P(AE|A) = \frac{P(AE)}{P(A)}$ and, from (b), $P(AE) = \frac{3}{55}$, we have that

$$\frac{3}{20} = \frac{\frac{3}{55}}{P(A)}. \text{ Hence, } P(A) = \frac{3}{55} \left(\frac{20}{3}\right) = \frac{4}{11}.$$

Thus, the number of specimens marked with an *A* is $\frac{4}{11} 220 = 80$.

Since $P(A) = \frac{4}{11}$ and, from (d), $P(C) = \frac{23}{55}$, it follows that $P(B) = 1 - \frac{4}{11} - \frac{23}{55} = \frac{12}{55}$.

Thus, the number of specimens marked with a B is $\frac{12}{55} 220 = 48$. Of these, 12 were marked BE and 24 were marked BG . Thus, the number of specimens marked BF is $48 - 12 - 24 = 12$. So far, we have:

LABEL	number
AE	12
AF	
AG	
BE	12
BF	12
BG	24
CE	36
CF	8
CG	48

(g) $P(AG|A) = \frac{1}{10}$. Since $P(AG|A) = \frac{P(AG)}{P(A)}$, and from (f) we have that $P(A) = \frac{4}{11}$ it follows that $\frac{1}{10} = \frac{P(AG)}{\frac{4}{11}}$, that is, $P(AG) = \frac{1}{10} \left(\frac{4}{11} \right) = \frac{4}{110}$.

Thus, the number of specimens marked AG is $\frac{4}{110} 220 = 8$.

Since the number marked with an A is $\frac{4}{11} 220 = 80$ and the number of those marked AE is 12 and the number of those marked AG is 8, we get that the number of specimens marked AF is $80 - 12 - 8 = 60$.

Thus, we have completed the table:

LABEL	number
AE	12
AF	60
AG	8
BE	12
BF	12
BG	24
CE	36
CF	8
CG	48

□

Solution to Exercise 9.15. Fran's initial beliefs are:

<i>state</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>	<i>g</i>
<i>probability</i>	$\frac{3}{20}$	0	$\frac{7}{20}$	$\frac{1}{20}$	0	$\frac{4}{20}$	$\frac{5}{20}$

The events under consideration are: $E = \{a, d, e, g\}$ and $D = \{a, b, c, f, g\}$.

(a) $P(E) = P(a) + P(d) + P(e) + P(g) = \frac{3}{20} + \frac{1}{20} + 0 + \frac{5}{20} = \frac{9}{20}$.

(b) $P(b|E) = P(c|E) = P(f|E) = 0$ (since each of these states does not belong to E),

$P(a|E) = \frac{\frac{3}{20}}{\frac{9}{20}} = \frac{3}{9}$, $P(d|E) = \frac{\frac{1}{20}}{\frac{9}{20}} = \frac{1}{9}$, $P(e|E) = \frac{0}{9} = 0$ and $P(g|E) = \frac{\frac{5}{20}}{\frac{9}{20}} = \frac{5}{9}$.

Thus the operation of conditioning on event E yields the following updated probability distribution:

<i>state</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>	<i>g</i>
<i>probability</i>	$\frac{3}{9}$	0	0	$\frac{1}{9}$	0	0	$\frac{5}{9}$

(c) (1) Initially Fran attaches the following probability to event D :

$P(D) = P(a) + P(b) + P(c) + P(f) + P(g) = \frac{3}{20} + 0 + \frac{7}{20} + \frac{4}{20} + \frac{5}{20} = \frac{19}{20} = 95\%$.

(2) After updating on information E Fran attaches the following probability to event D :

$P(D|E) = \frac{P(D \cap E)}{P(E)} = \frac{P(\{a, g\})}{P(E)} = \frac{\frac{8}{20}}{\frac{9}{20}} = \frac{8}{9} = 88.89\%$. □

Solution to Exercise 9.16.

(a) $E = \{(95, 78, 54), (78, 95, 54), (54, 95, 78)\}$. Thus $P(E) = \frac{16}{32} + \frac{4}{32} + \frac{2}{32} = \frac{22}{32}$.

(b) Conditioning on E yields the following beliefs:

(95, 78, 54) (95, 54, 78) (78, 95, 54) (54, 95, 78) (78, 54, 95) (54, 78, 95)

$\frac{\frac{16}{32}}{\frac{22}{32}} = \frac{8}{11}$ 0 $\frac{\frac{4}{32}}{\frac{22}{32}} = \frac{2}{11}$ $\frac{\frac{2}{32}}{\frac{22}{32}} = \frac{1}{11}$ 0 0

□

Solution to Exercise 9.17.

(a) Updating on information $E = \{a, c, e, f, g\}$ yields the following beliefs:

<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>	<i>g</i>
$\frac{3}{17}$	0	$\frac{5}{17}$	0	$\frac{1}{17}$	$\frac{3}{17}$	$\frac{5}{17}$

(b) Updating the beliefs of Part (a) on information $F = \{b, d, e, f, g\}$ yields the following beliefs:

<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>	<i>g</i>
0	0	0	0	$\frac{1}{9}$	$\frac{3}{9}$	$\frac{5}{9}$

□

Solution to Exercise 9.18. Represent a state as a pair (x, y) where x is the suspect who stole the statuette and y is the suspect who stole the tiara.

- (a) The set of states is $U = \{(A, A), (A, B), (A, C), (B, A), (B, B), (B, C), (C, A), (C, B), (C, C)\}$.
 (b) The inspector's initial beliefs are:

$$\begin{array}{cccccccccc} (A, A) & (A, B) & (A, C) & (B, A) & (B, B) & (B, C) & (C, A) & (C, B) & (C, C) \\ \frac{1}{20} & \frac{2}{20} & \frac{2}{20} & \frac{2}{20} & \frac{3}{20} & \frac{2}{20} & \frac{2}{20} & \frac{2}{20} & \frac{4}{20} \end{array}$$

[Explanation: $P(A, A) + P(B, B) + P(C, C) = \frac{1}{20} + \frac{3}{20} + \frac{4}{20} = \frac{8}{20}$; thus $\frac{12}{20}$ remains to be distributed equally among the remaining six states, so that each receives $\frac{2}{20}$.]

- (c) The information is $F = \{(A, A), (A, B), (C, A), (C, B)\}$.

Updating on this information yields the following beliefs:

$$\begin{array}{cccccccccc} (A, A) & (A, B) & (A, C) & (B, A) & (B, B) & (B, C) & (C, A) & (C, B) & (C, C) \\ \frac{1}{7} & \frac{2}{7} & 0 & 0 & 0 & 0 & \frac{2}{7} & \frac{2}{7} & 0 \end{array} \quad \square$$

Solution to Exercise 9.19.

- (a) $P(E) = P(a) + P(d) + P(e) + P(g) = \frac{3}{20} + \frac{1}{20} + 0 + \frac{5}{20} = \frac{9}{20}$, $P(b|E) = 0$
 and $P(d|E) = \frac{\frac{1}{20}}{\frac{9}{20}} = \frac{1}{9}$.

- (b) The updated beliefs are as follows: $\begin{array}{cccccc} a & b & c & d & e & f & g \\ \frac{3}{9} & 0 & 0 & \frac{1}{9} & 0 & 0 & \frac{5}{9} \end{array} \quad \square$

Solution to Exercise 9.20. The initial beliefs are:

$$\begin{array}{ccccc} \text{Economics} & \text{Mathematics} & \text{Philosophy} & \text{Psychology} & \text{Statistics} \\ \frac{35}{100} & \frac{22}{100} & \frac{18}{100} & \frac{16}{100} & \frac{9}{100} \end{array}$$

- (a) Updating on {Economics, Philosophy, Psychology} yields the following beliefs:

$$\begin{array}{ccccc} \text{Economics} & \text{Mathematics} & \text{Philosophy} & \text{Psychology} & \text{Statistics} \\ \frac{35}{69} & 0 & \frac{18}{69} & \frac{16}{69} & 0 \end{array}$$

- (b) Updating the beliefs of Part (a) on {Philosophy, Psychology}¹² yields the following beliefs:

$$\begin{array}{ccccc} \text{Economics} & \text{Mathematics} & \text{Philosophy} & \text{Psychology} & \text{Statistics} \\ 0 & 0 & \frac{18}{34} & \frac{16}{34} & 0 \end{array}$$

- (c) Updating the beliefs of Part (b) on {Psychology} yields the following beliefs:

$$\begin{array}{ccccc} \text{Economics} & \text{Mathematics} & \text{Philosophy} & \text{Psychology} & \text{Statistics} \\ 0 & 0 & 0 & 1 & 0 \end{array}$$

that is, the instructor now knows that the student is a Psychology major. \square

¹²This is the intersection of the initial piece of information, namely {Economics, Philosophy, Psychology}, and the new piece of information, namely {Mathematics, Philosophy, Psychology, Statistics}. Updating the updated beliefs on {Mathematics, Philosophy, Psychology, Statistics} yields the same result as updating on {Philosophy, Psychology}. Indeed, one would obtain the same result by updating the *initial* beliefs on {Philosophy, Psychology}.

Solution to Exercise 9.21.

Let $f : \mathcal{E} \rightarrow 2^U$ be an AGM belief revision function.

Let $E, F \in \mathcal{E}$ be such that $E \subseteq F$ and $E \cap f(F) \neq \emptyset$. We need to show that $f(E) = E \cap f(F)$.

By definition of AGM belief revision function (Definition 9.4.2), there is a plausibility order \succsim on U such that

$$f(F) = \{s \in F : s \succsim s' \text{ for every } s' \in F\} \quad (9.10)$$

and

$$f(E) = \{s \in E : s \succsim s' \text{ for every } s' \in E\}. \quad (9.11)$$

Choose an arbitrary $s \in E \cap f(F)$.

Then, by (9.10) and the fact that $E \subseteq F$, $s \succsim s'$ for every $s' \in E$ and thus, by (9.11), $s \in f(E)$.

Hence, $E \cap f(F) \subseteq f(E)$.

Conversely, choose an arbitrary $s_1 \in f(E)$.

Then, since (by definition of belief revision function: Definition 9.4.2) $f(E) \subseteq E$, $s_1 \in E$.

We want to show that $s_1 \in f(F)$ [so that $s_1 \in E \cap f(F)$ and, therefore, $f(E) \subseteq E \cap f(F)$].

Suppose it is not true. Then, by (9.10), there exists an $s_2 \in F$ such that $s_2 \succ s_1$.

Select an $s_3 \in E \cap f(F)$ (recall that, by hypothesis, $E \cap f(F) \neq \emptyset$).

Then, by (9.10) (since $s_2, s_3 \in f(F)$), $s_3 \succsim s_2$, from which it follows (by transitivity of \succsim and the fact that $s_2 \succ s_1$) that $s_3 \succ s_1$.

But then, since $s_3 \in E$, it is not true that $s_1 \succsim s'$ for every $s' \in E$, contradicting – by (9.11) – the hypothesis that $s_1 \in f(E)$. \square

Solution to Exercise 9.22. For every event E (representing a possible item of information), let P_E be the probability distribution on E that represents the *revised* beliefs of the individual after receiving information E .

Let P be the probability distribution on U representing the individual's *initial* beliefs.

Define the following belief revision function $f : f(U) = \text{Supp}(P)$ and $f(E) = \text{Supp}(P_E)$.

Suppose that f satisfies Arrow's Axiom.

Then, for every event E , if $E \cap f(U) \neq \emptyset$ [that is, if $E \cap \text{Supp}(P) \neq \emptyset$ or $P(E) > 0$]

then $f(E) = E \cap f(U)$ [that is, $\text{Supp}(P_E) = E \cap \text{Supp}(P)$]. \square

Solution to Exercise 9.23. We have that $\mathcal{E} = \{\{a, e\}, \{d, e, k, m\}, \{b, d, e, k\}, U\}$ and \succsim is given by

most plausible b, g
 c, k, m
 d, h
 e
 least plausible a

Then the belief revision function rationalized by this plausibility order is given by: $f(\{a, e\}) = \{e\}$, $f(\{d, e, k, m\}) = \{k, m\}$, $f(\{b, d, e, k\}) = \{b\}$ and $f(U) = \{b, g\}$. \square

Solution to Exercise 9.24.

From Exercise 9.23 we get that $\{P_E, P_F, P_G, P_U\}$ must be such that

$Supp(P_E) = \{e\}$, $Supp(P_F) = \{k, m\}$, $Supp(P_G) = \{b\}$ and $Supp(P_U) = \{b, g\}$.

For every full-support probability distribution P_0 on U , there is a corresponding collection $\{P_E, P_F, P_G, P_U\}$. For example, if P_0 is the uniform distribution on U (that assigns probability $\frac{1}{9}$ to every state)

then the corresponding $\{P_E, P_F, P_G, P_U\}$ is given by:

state	a	b	c	d	e	g	h	k	m
P_E	0	0	0	0	1	0	0	0	0
P_F	0	0	0	0	0	0	0	$\frac{1}{2}$	$\frac{1}{2}$
P_G	0	1	0	0	0	0	0	0	0
P_U	0	$\frac{1}{2}$	0	0	0	$\frac{1}{2}$	0	0	0

As another example, if P_0 is the following probability distribution

state	a	b	c	d	e	g	h	k	m
P_0	$\frac{1}{50}$	$\frac{3}{50}$	$\frac{11}{50}$	$\frac{4}{50}$	$\frac{8}{50}$	$\frac{9}{50}$	$\frac{5}{50}$	$\frac{2}{50}$	$\frac{7}{50}$

then the corresponding $\{P_E, P_F, P_G, P_U\}$ is given by: P_E and P_G the same as above, and P_F and P_U as follows:

state	a	b	c	d	e	g	h	k	m
P_F	0	0	0	0	0	0	0	$\frac{2}{9}$	$\frac{7}{9}$ \square
P_U	0	$\frac{1}{4}$	0	0	0	$\frac{3}{4}$	0	0	0

Solution to Exercise 9.25. Yes. The following is a common prior: $\begin{pmatrix} a & b & c & d \\ \frac{1}{9} & \frac{2}{9} & \frac{2}{9} & \frac{4}{9} \end{pmatrix}$. \square

Solution to Exercise 9.26.

- (a) Assume that $p = \frac{1}{13}$ (so that $1 - p = \frac{12}{13}$).

From the updating conditions for Individual 1 we get $p_b = 2p_a$ and $p_e = p_f$;

from the updating conditions for Individual 2 we get $p_c = 3p_b$ and $p_e = 4p_d$;

from the updating conditions for Individual 3 we get $p_c = p_d$ and $p_f = 12p_b$.

All these equations have a solution, which constitutes a common prior, namely

$$\begin{pmatrix} a & b & c & d & e & f \\ \frac{1}{63} & \frac{2}{63} & \frac{6}{63} & \frac{6}{63} & \frac{24}{63} & \frac{24}{63} \end{pmatrix}.$$

- (b) From the equations of part (a) we have that $p_f = p_e$, $p_e = 4p_d$, $p_d = p_c$ and $p_c = 3p_b$, from which it follows that $p_f = 12p_b$.

Thus, updating on the information set $\{b, f\}$ of Individual 3 we get that we must have $1 - p = \frac{p_f}{p_b + p_f} = \frac{12p_b}{p_b + 12p_b} = \frac{12}{13}$ and thus $p = \frac{1}{13}$. Hence, if $p \neq \frac{1}{13}$ there is no common prior. \square

Solution to Exercise 9.27.

- (a) The common knowledge partition is the trivial partition shown in Figure 9.19.
- (b) The common prior is also shown shown in Figure 9.19.
- (c) (1) Individual 1 assigns probability $\frac{1}{3}$ to $E = \{b, d, f\}$ at every state. Thus his assessment of the probability of E is common knowledge at every state.
- (2) At state a Individual 2 assigns probability $\frac{1}{2}$ to E , while at state b she assigns probability $\frac{2}{3}$ to E . Thus there is no state where her assessment of the probability of E is common knowledge.
- (d) (1) At state b Individual 1's information set is $\{a, b\}$; since at a Individual 2 assigns probability $\frac{1}{2}$ to E , while at b she assigns probability $\frac{2}{3}$ to E , it follows that at b it is not the case that Individual 1 knows Individual 2's assessment of the probability of E .
- (2) At state c Individual 1's information set is $\{c, d, e\}$; since at c Individual 2 assigns probability $\frac{2}{3}$ to E , while at d she assigns probability $\frac{1}{2}$ to E , it follows that at c it is not the case that Individual 1 knows Individual 2's assessment of the probability of E .
- (3) At state f Individual 1's information set is $\{f, g\}$; since at f Individual 2 assigns probability $\frac{1}{2}$ to E , while at g she assigns probability 0 to E , it follows that at f it is not the case that Individual 1 knows Individual 2's assessment of the probability of E . \square

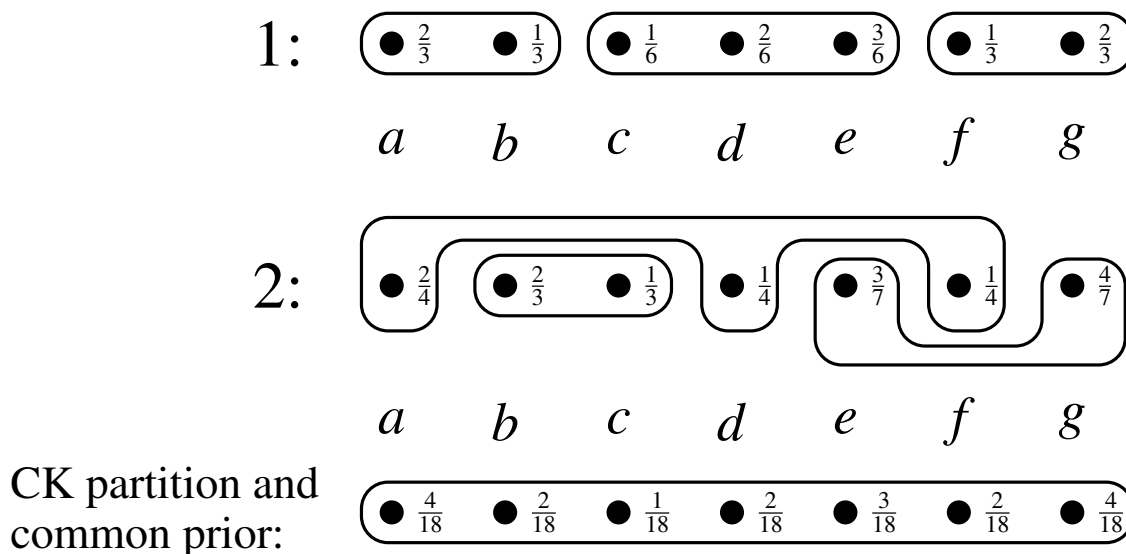


Figure 9.19: The information structure for Exercise 9.27.

Solution to Exercise 9.28.

(a) **Method 1.** We solve this problem using Bayes' formula.

For every $n \in \{1, 2, 3\}$, let D_n denote the event that the car is behind door n and let O_n denote the event that the host opens Door n .

The initial probabilities are $P(D_1) = P(D_2) = P(D_3) = \frac{1}{3}$.

We want to compute $P(D_1|O_3)$; if $P(D_1|O_3) \geq \frac{1}{2}$ then you should **not** switch, since there is a cost in switching (recall that Door 1 is your initial choice).

By Bayes' rule, $P(D_1|O_3) = \frac{P(O_3|D_1)P(D_1)}{P(O_3)}$.

We know that $P(D_1) = \frac{1}{3}$ and $P(O_3|D_1) = \frac{1}{2}$ (when the car is behind Door 1 then the host has a choice between opening Door 2 and opening Door 3 and he chooses with equal probability).

Thus,

$$P(D_1|O_3) = \frac{\frac{1}{2} \times \frac{1}{3}}{P(O_3)} = \frac{\frac{1}{6}}{P(O_3)}. \quad (9.12)$$

We need to compute $P(O_3)$:

$$\begin{aligned} P(O_3) &= P(O_3|D_1)P(D_1) + P(O_3|D_2)P(D_2) + P(O_3|D_3)P(D_3) \\ &= P(O_3|D_1)\frac{1}{3} + P(O_3|D_2)\frac{1}{3} + P(O_3|D_3)\frac{1}{3} \\ &= \frac{1}{2}\left(\frac{1}{3}\right) + 1\left(\frac{1}{3}\right) + 0\left(\frac{1}{3}\right) = \frac{1}{6} + \frac{1}{3} = \frac{1}{2} \end{aligned}$$

because $P(O_3|D_1) = \frac{1}{2}$, $P(O_3|D_2) = 1$ (if the car is behind Door 2 then the host has to open Door 3, since he cannot open the door that you chose, namely Door 1) and $P(O_3|D_3) = 0$ (if the car is behind Door 3 then the host cannot open that door).

Substituting $\frac{1}{2}$ for $P(O_3)$ in (9.12) we get that $P(D_1|O_3) = \frac{1}{3}$.

Hence, the updated probability that the car is behind the other door (Door 2) is $\frac{2}{3}$ and therefore you should switch.

(b) **Method 2.** The extensive form is shown in Figure 9.20

(‘cbn’ means ‘the car is behind door n ’, ‘chn’ means ‘you choose door n ’).

The hypothesized sequence of events leads to either node x or node y (you first choose Door 1 and then the host opens Door 3).

The prior probability of getting to x is $\frac{1}{3} \times \frac{1}{2} = \frac{1}{6}$

while the prior probability of getting to y is $\frac{1}{3} \times 1 = \frac{1}{3}$.

The information you have is $\{x, y\}$ and, using the conditional probability rule,

$$P(\{x\}|\{x, y\}) = \frac{P(\{x\} \cap \{x, y\})}{P(\{x, y\})} = \frac{P(\{x\})}{P(\{x, y\})} = \frac{\frac{1}{6}}{\frac{1}{6} + \frac{1}{3}} = \frac{1}{3}.$$

Thus you should switch. □

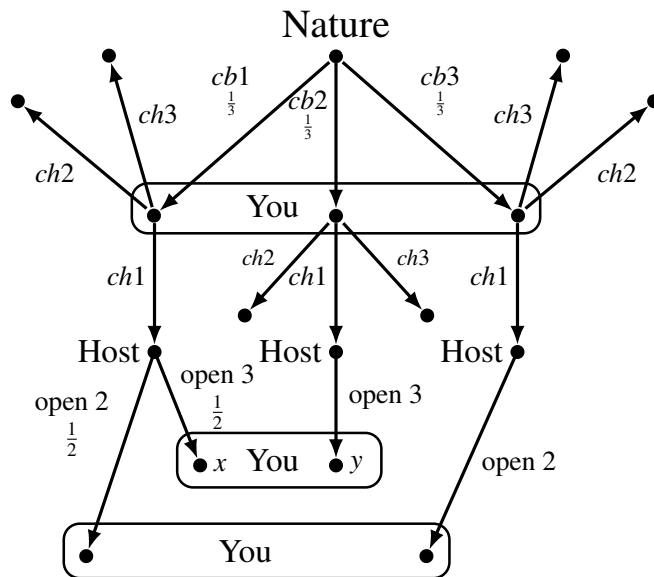


Figure 9.20: The Monty Hall problem represented using an imperfect-information frame.

10. Rationality

10.1 Models of strategic-form games

In Chapter 2 (Section 2.5) and Chapter 6 (Section 6.4) we discussed the procedure of iterative elimination of strictly dominated strategies and claimed that it captures the notion of common knowledge of rationality in strategic-form games. We now have the tools to state this precisely.

The *epistemic foundation program* in game theory aims to identify, for every game, the strategies that might be chosen by rational and intelligent players who know the structure of the game and the preferences of their opponents and who recognize each other's rationality. The two central questions are thus:

- (1) Under what circumstances is a player rational?
- (2) What does 'mutual recognition of rationality' mean?

A natural interpretation of the latter notion is in terms of 'common knowledge of rationality'.¹ We already know how to model the notion of common knowledge (Chapter 8); thus, we only need to define what it means for a player to be rational. Intuitively, a player is rational if she chooses an action which is "best" given what she believes. In order to make this more precise we need to introduce the notion of a model of a game. We shall focus on strategic-form games with cardinal payoffs (Definition 6.1.2, Chapter 6).

¹As pointed out in Chapter 8, a defining characteristic of knowledge is truth: if, at a state, a player knows event E then, at that state, E must be true. A more general notion is that of belief, which allows for the possibility of mistakes: believing E is compatible with E being false. Thus a more appealing notion is that of *common belief* of rationality; however, in order to simplify the exposition, we shall restrict attention to the notions of knowledge and common knowledge developed in Chapter 8. For the analysis of common belief and an overview of the epistemic foundation program in game theory the reader is referred to Battigalli and Bonanno (1999).

The definition of a strategic-form game specifies the choices available to the players and what motivates those choices (their preferences over the possible outcomes); however, it leaves out an important factor in the determination of players' choices, namely what they believe about the other players. Adding a specification of the players' knowledge and beliefs determines the context in which a particular game is played; this can be done with the help of an interactive knowledge-belief structure.

Recall that an interactive knowledge-belief structure consists of

- ◇ a set of states W ,
- ◇ n partitions $\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_n$ of W ,
- ◇ a collection of probability distributions on W , one for each information set of each partition, whose support is a subset of that information set.

The probability distributions encode the beliefs of the players in each possible state of knowledge.

Recall also that if w is a state we denote by $I_i(w)$ the element of the partition \mathcal{I}_i that contains w , that is, the information set of Player i at state w and by $P_{i,w}$ the probability distribution on $I_i(w)$, representing the beliefs of Player i at state w .

R In Chapter 6 we used the symbol σ_i to denote a mixed strategy of Player i . In this chapter the symbol σ_i will have a different meaning (see the following definition). Since mixed strategies play no role in this chapter, there is no possibility of confusion. Only pure strategies are considered throughout this chapter.

Definition 10.1.1 Let G be a n -player strategic-form game with cardinal payoffs. A *model* of G is an interactive knowledge-belief structure together with a function

$$\sigma_i : W \rightarrow S_i$$

(for every Player i) that associates with every state a pure strategy of Player i (recall that S_i denotes the set of pure strategies of Player i). The interpretation of $s_i = \sigma_i(w)$ is that, at state w , Player i plays (or has chosen) strategy s_i . We impose the restriction that a player always knows what strategy she has chosen, that is, the function σ_i is constant on each information set of Player i :

$$\text{if } w' \in I_i(w) \text{ then } \sigma_i(w') = \sigma_i(w).$$

The addition of the functions σ_i to an interactive knowledge-belief structure yields an interpretation of events in terms of propositions about what strategies the players choose, thereby giving content to players' knowledge and beliefs.

Figure 10.1 shows a strategic-form game with cardinal payoffs and Figure 10.2 a model of it.

		Player 2					
		<i>L</i>		<i>C</i>		<i>R</i>	
Player 1	<i>T</i>	4	6	3	2	8	0
	<i>M</i>	0	9	0	0	4	12
	<i>B</i>	8	3	2	4	0	0

Figure 10.1: A strategic-form game with cardinal payoffs.

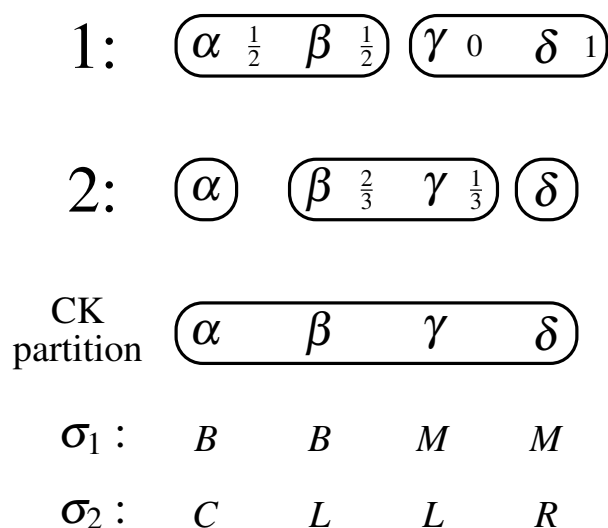


Figure 10.2: A model of the game of Figure 10.1.

State β in Figure 10.2 describes the following situation:

- Player 1 plays B ($\sigma_1(\beta) = B$) and Player 2 plays L ($\sigma_2(\beta) = L$).
- Player 1 (whose information set is $\{\alpha, \beta\}$) is uncertain as to whether Player 2 has chosen to play C ($\sigma_2(\alpha) = C$) or L ($\sigma_2(\beta) = L$); furthermore, he attaches probability $\frac{1}{2}$ to each of these two possibilities.
- Player 2 (whose information set is $\{\beta, \gamma\}$) is uncertain as to whether Player 1 has chosen to play B ($\sigma_1(\beta) = B$) or M ($\sigma_1(\gamma) = M$); furthermore, she attaches probability $\frac{2}{3}$ to Player 1 playing B and $\frac{1}{3}$ to Player 1 playing M .

Are the two players' choices rational at state β ? Definition 10.1.2 provides an answer to this question.

To summarize: one can view a strategic-form game as only a partial description of an interactive situation: it specifies who the players are, what actions they can take and how they rank the possible outcomes; a model of the game completes this description by also specifying what each player actually does and what she believes about what the other players are going to do. Once we know what a player does and what she believes, then we are in a position to judge her choice to be either rational or irrational.

The following definition says that at a state (of a model) a player is rational if her choice at that state is optimal, given her beliefs, that is, if there is no other pure strategy that would give her a higher expected payoff, given what she believes about the choices of her opponents.

Given a state w and a Player i , we denote by $\sigma_{-i}(w)$ the profile of strategies chosen at w by the players other than i : $\sigma_{-i}(w) = (\sigma_1(w), \dots, \sigma_{i-1}(w), \sigma_{i+1}(w), \dots, \sigma_n(w))$. Furthermore, as in Chapter 8, we denote by $P_{i,w}$ the beliefs of Player i at state w , that is, $P_{i,w}$ is a probability distribution over $I_i(w)$ (although we can think of $P_{i,w}$ as a probability distribution over the entire set of states W satisfying the property that if $w' \notin I_i(w)$ then $P_{i,w}(w') = 0$). Note that, for every Player i , there is just one probability distribution over $I_i(w)$ that is, if $w' \in I_i(w)$ then $P_{i,w'} = P_{i,w}$.

Definition 10.1.2 Player i is *rational* at state w if, given her beliefs $P_{i,w}$, there is no pure strategy $s'_i \in S_i$ of hers that yields a higher expected payoff than $\sigma_i(w)$ (recall that $\sigma_i(w)$ is the strategy chosen by Player i at state w): for all $s'_i \in S_i$,

$$\sum_{w' \in I_i(w)} P_{i,w}(w') \times \pi_i(\sigma_i(w), \sigma_{-i}(w')) \geq \sum_{w' \in I_i(w)} P_{i,w}(w') \times \pi_i(s'_i, \sigma_{-i}(w')).$$

To illustrate Definition 10.1.2 let us re-consider the example of Figures 10.1 and 10.2, which are reproduced on the following page.

Consider first Player 1. Given his beliefs and his choice of B , his expected payoff is: $\frac{1}{2} \pi_1(B, C) + \frac{1}{2} \pi_1(B, L) = \frac{1}{2}(2) + \frac{1}{2}(8) = 5$. Given his beliefs about Player 2, could he obtain a higher expected payoff with a different choice?

With M his expected payoff would be $\frac{1}{2} \pi_1(M, C) + \frac{1}{2} \pi_1(M, L) = \frac{1}{2}(0) + \frac{1}{2}(0) = 0$, while with T his expected payoff would be $\frac{1}{2} \pi_1(T, C) + \frac{1}{2} \pi_1(T, L) = \frac{1}{2}(3) + \frac{1}{2}(4) = 3.5$. Thus, given his beliefs, Player 1's choice of B is optimal and we can say that Player 1 is rational at state β .

Consider now Player 2. Given her beliefs and her choice of L , her expected payoff is: $\frac{2}{3} \pi_2(B, L) + \frac{1}{3} \pi_2(M, L) = \frac{2}{3}(3) + \frac{1}{3}(9) = 5$. Given her beliefs about Player 1, could she obtain a higher expected payoff with a different choice?

With C her expected payoff would be $\frac{2}{3} \pi_2(B, C) + \frac{1}{3} \pi_2(M, C) = \frac{2}{3}(4) + \frac{1}{3}(0) = \frac{8}{3}$, while with R her expected payoff would be $\frac{2}{3} \pi_2(B, R) + \frac{1}{3} \pi_2(M, R) = \frac{2}{3}(0) + \frac{1}{3}(12) = 4$.

Thus, given her beliefs, Player 2's choice of L is optimal and we can say that also Player 2 is rational at state β .

		Player 2					
		L		C		R	
Player 1	T	4	6	3	2	8	0
	M	0	9	0	0	4	12
	B	8	3	2	4	0	0

$$1: \quad \left(\alpha \ \frac{1}{2} \ \beta \ \frac{1}{2} \right) \quad \left(\gamma \ 0 \ \delta \ 1 \right)$$

$$2: \quad \left(\alpha \right) \quad \left(\beta \ \frac{2}{3} \ \gamma \ \frac{1}{3} \right) \quad \left(\delta \right)$$

$$\text{CK partition} \quad \left(\alpha \ \beta \ \gamma \ \delta \right)$$

$$\sigma_1 : \quad B \quad B \quad M \quad M$$

$$\sigma_2 : \quad C \quad L \quad L \quad R$$

Given a model of a game, using Definition 10.2 we can determine, for every player, the set of states where that player is rational. Let \mathbf{R}_i be the event that (that is, the set of states at which) Player i is rational and let $\mathbf{R} = \mathbf{R}_1 \cap \dots \cap \mathbf{R}_n$ be the event that all players are rational.

In the example of Figures 10.1 and 10.2 we have that $\mathbf{R}_1 = \{\alpha, \beta\}$.

Indeed, it was shown above that $\beta \in \mathbf{R}_1$; since Player 1's choice and beliefs at α are the same as at β , it follows that also $\alpha \in \mathbf{R}_1$.

On the other hand, Player 1 is not rational at state γ because he is certain (that is, attaches probability 1 to the fact) that Player 2 plays R and his choice of M is not optimal against R (the unique best response to R is T). The same is true of state δ .

For Player 2 we have that $\mathbf{R}_2 = \{\alpha, \beta, \gamma, \delta\}$, that is, Player 2 is rational at every state (the reader should convince himself/herself of this).

Thus $\mathbf{R} = \mathbf{R}_1 \cap \mathbf{R}_2 = \{\alpha, \beta\} \cap \{\alpha, \beta, \gamma, \delta\} = \{\alpha, \beta\}$.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 10.5.1 at the end of this chapter.

10.2 Common knowledge of rationality in strategic-form games

Since we have defined the rationality of a player and the rationality of all the players as events, we can apply the knowledge operator and the common knowledge operator to these events; thus we can determine, for example, if there are any states where the rationality of all the players is common knowledge and see what strategy profiles are compatible with common knowledge of rationality.

R Note that, for every player i , $K_i \mathbf{R}_i = \mathbf{R}_i$, that is, every player is rational if and only if she knows it (this is a consequence of the assumption that a player always knows what strategy she has chosen and knows her own beliefs:
if $w' \in I_i(w)$ then $\sigma_i(w') = \sigma_i(w)$ and $P_{i,w'} = P_{i,w}$).

Note also that $CK\mathbf{R} \subseteq \mathbf{R}$, that is, if it is common knowledge that all the players are rational, then they are indeed rational (while, in general, it is **not** true that \mathbf{R} is a subset of $CK\mathbf{R}$, that is, it is possible for all the players to be rational without it being common knowledge).

In Chapters 2 (for ordinal strategic-form games) and 6 (for cardinal strategic-form games) we claimed that the iterated deletion of strictly dominated strategies corresponds to the notion of common knowledge of rationality. We are now able to state this precisely. The following two theorems establish the correspondence between the two notions for the case of strategic-form games with cardinal payoffs. A corresponding characterization holds for strategic-form games with ordinal payoffs.² The proofs of Theorems 10.2.1 and 10.2.2 are given in Section 10.4.

Theorem 10.2.1 Given a finite strategic-form game with cardinal payoffs G , the following is true: for any model of G and any state w in that model, if $w \in CK\mathbf{R}$ (that is, at w there is common knowledge of rationality) then the pure-strategy profile associated with w must be one that survives the Cardinal Iterated Deletion of Strictly Dominated Strategies (Definition 6.4.1, Chapter 6).

Theorem 10.2.2 Given a finite strategic-form game with cardinal payoffs G , the following is true: if $s \in S$ is a pure-strategy profile that survives the Cardinal Iterated Deletion of Strictly Dominated Strategies, then there is a model of G and a state w in that model, such that $w \in CK\mathbf{R}$ and the strategy profile associated with w is s .

²See Bonanno (2015).

As an application of Theorem 10.2.1 consider the game of Figure 10.3 (which reproduces Figure 10.2).

		Player 2					
		<i>L</i>		<i>C</i>		<i>R</i>	
Player 1	<i>T</i>	4	6	3	2	8	0
	<i>M</i>	0	9	0	0	4	12
	<i>B</i>	8	3	2	4	0	0

Figure 10.3: Copy of Figure 10.2.

In this game the iterated deletion of strictly dominated strategies leads to the following strategy profiles: (T, L) , (T, C) , (B, L) and (B, C) .³ Thus these are the only strategy profiles that are compatible with common knowledge of rationality: by Theorem 10.1, at a state in a model of this game where there is common knowledge of rationality the players can play only one of these strategy profiles. Furthermore, by Theorem 10.2.2, any of these four strategy profiles can in fact be played in a situation where there is common knowledge of rationality.

It is worth stressing that common knowledge of rationality does *not* imply that the players play a Nash equilibrium: indeed, none of the above four strategy profiles is a Nash equilibrium. However, since a pure-strategy Nash equilibrium always survives the iterated deletion of strictly dominated strategies, the pure-strategy Nash equilibria (if any) are always compatible with common knowledge of rationality.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 10.5.2 at the end of this chapter.

³For Player 1, M is strictly dominated by any mixed strategy $\begin{pmatrix} T & B \\ p & 1-p \end{pmatrix}$ with $p > \frac{1}{2}$. After deletion of M , for Player 2 R becomes strictly dominated by either of the other two strategies.

10.3 Common knowledge of rationality in extensive-form games

So far we have only discussed the implications of common knowledge of rationality in strategic-form games, where the players make their choices simultaneously (or in ignorance of the other players' choices). We now turn to a brief discussion of the issues that arise when one attempts to determine the implications of common knowledge of rationality in dynamic (or extensive-form) games.

How should a model of a dynamic game be constructed? One approach in the literature has been to consider models of the corresponding strategic-form. However, there are several conceptual issues that arise in this context. In the models considered in Section 10.1 the interpretation of $s_i = \sigma_i(w)$ is that at state w Player i "plays" or "chooses" strategy s_i .

Consider the perfect-information game shown in Figure 10.4 and any model of the strategic-form game associated with it.

Let w be a state in that model where

$$\sigma_1(w) = (d_1, a_3).$$

What does it mean to say that Player 1 "chooses" strategy (d_1, a_3) ? The first part of the strategy, namely d_1 , can be interpreted as a description of Player 1's actual behavior (what he actually does: he plays d_1), but the second part of the strategy, namely a_3 , has no such interpretation: if Player 1 in fact plays d_1 then he knows that he will not have to make any further choices and thus it is not clear what it means for him to "choose" to play a_3 in a situation that is made impossible by his decision to play d_1 .




Figure 10.4: A perfect-information game.

Thus it does not seem to make sense to interpret $\sigma_1(w) = (d_1, a_3)$ as 'at state w Player 1 chooses (d_1, a_3) '. Perhaps the correct interpretation is in terms of a more complex sentence such as 'Player 1 chooses to play d_1 and if – contrary to this plan – he were to play a_1 and Player 2 were to follow with a_2 , then Player 1 would play a_3 .'

While in a simultaneous game the association of a strategy of Player i to a state can be interpreted as a description of Player i 's actual behavior at that state, in the case of dynamic games this interpretation is no longer valid, since one would end up describing not only the actual behavior of Player i at that state but also his counterfactual behavior.

Methodologically, this is not satisfactory: if it is considered to be necessary to specify what a player would do in situations that do not occur at the state under consideration, then one should model the counterfactual explicitly.

But why should it be necessary to specify at state w (where Player 1 is playing d_1) what he would do at the counterfactual node following actions a_1 and a_2 ? Perhaps what matters is not so much what Player 1 would actually do in that situation but what Player 2 believes that Player 1 would do: after all, Player 2 might not know that Player 1 has decided to play d_1 and she needs to consider what to do in the eventuality that Player 1 actually ends up playing a_1 .

So, perhaps, the strategy of Player 1 is to be interpreted as having two components: (1) a description of Player 1's behavior and (2) a conjecture in the mind of Player 2 about what Player 1 would do. If this is the correct interpretation, then one could object – again from a methodological point of view – that it would be preferable to disentangle the two components and model them explicitly.⁴

An alternative – although less common – approach in the literature dispenses with strategies and considers models of games where

1. states are described in terms of players' *actual behavior* and
2. players' conjectures concerning the actions of their opponents in various hypothetical situations are modeled by means of a generalization of the knowledge-belief structures considered in Section 10.1. The generalization is obtained by encoding not only the initial beliefs of the players (at each state) but also their *dispositions to revise those beliefs* under various hypotheses.⁵

A third approach has been to move away from static models (like the models considered in Section 10.1) and consider dynamic models where time is introduced explicitly into the analysis. These are behavioral models where strategies play no role and the only beliefs that are specified are the actual beliefs of the players at the time of choice. Thus players' beliefs are modeled as temporal, rather than conditional, beliefs and rationality is defined in terms of actual choices, rather than hypothetical plans.⁶

A discussion of the implications of common knowledge (or common belief) of rationality in dynamic games would require the introduction of several new concepts and definitions. For the sake of brevity, we shall not pursue this topic in this book.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 10.5.3 at the end of this chapter.

⁴For an extensive discussion of these issues see Bonanno (2015).

⁵The interested reader is referred to Perea (2012).

⁶The interested reader is referred to Bonanno (2014).

10.4 Proofs of Theorems

In order to prove Theorem 10.2.1 we will need the following extension of Theorem 6.4.1 (Chapter 6) to the case of an arbitrary number of players. A proof of Theorem 10.4.1 can be found in Osborne and Rubinstein (1994, Lemma 60.1, p. 60).

Theorem 10.4.1 Consider a finite n -player game in strategic form with cardinal payoffs $\langle I, (S_i)_{i \in I}, (\pi_i)_{i \in I} \rangle$.

Select an arbitrary Player i and a pure strategy s_i of Player i .

For every Player $j \neq i$, let $S'_j \subseteq S_j$ be a subset of j 's set of pure strategies

and let $S'_{-i} = S'_1 \times \cdots \times S'_{i-1} \times S'_{i+1} \times \cdots \times S'_n$ be the Cartesian product of the sets S'_j ($j \neq i$).

Then the following are equivalent:

(1) There is a belief of Player i on S'_{-i} (that is, a probability distribution $P : S'_{-i} \rightarrow [0, 1]$) that makes s_i a best response, that is, for every $x \in S_i$,

$$\sum_{s'_{-i} \in S_{-i}} \pi_i(s_i, s'_{-i}) P(s'_{-i}) \geq \sum_{s'_{-i} \in S'_{-i}} \pi_i(x, s'_{-i}) P(s'_{-i}),$$

(2) s_i is not strictly dominated by a mixed strategy of Player i in the restriction of the game to the sets of pure strategies $S'_1, \dots, S'_{i-1}, S_i, S'_{i+1}, \dots, S'_n$.

Given a finite strategic-form game with cardinal payoffs $G = \langle I, (S_i)_{i \in I}, (\pi_i)_{i \in I} \rangle$ we shall denote by S_i^∞ the set of pure strategies of Player i that survive the Cardinal Iterated Deletion of Strictly Dominated Strategies (Definition 6.4.1, Chapter 6) and by S^∞ the corresponding set of strategy profiles (that is, $S^\infty = S_1^\infty \times \cdots \times S_n^\infty$).

Proof of Theorem 10.2.1. Consider a finite strategic-form game G with cardinal payoffs and an epistemic model of G .

Let w be a state in that model such that $w \in CKR$.

We want to show that $\sigma(w) = (\sigma_1(w), \dots, \sigma_n(w)) \in S^\infty$.

We shall prove the stronger claim that, for every Player i and for every $w' \in I_{CK}(w)$ (recall that $I_{CK}(w)$ is the cell of the common knowledge partition that contains w) and for every $m \geq 0$, $\sigma_i(w') \notin D_i^m$, where $D_i^m \subseteq S_i$ is the set of pure strategies of Player i that are strictly dominated in game G^m (G^m is the subgame of G obtained at step m of the iterated deletion procedure; we define G^0 to be G itself). Note that this is equivalent to stating that, for every $w' \in I_{CK}(w)$, $\sigma(w') \in S^\infty$. Since $w \in I_{CK}(w)$, it follows from this that $\sigma(w) \in S^\infty$. The proof is by induction.

1. **BASE STEP** ($m = 0$). Select an arbitrary $w' \in I_{CK}(w)$ and an arbitrary Player i .
 Since $w \in CK\mathbf{R}$, $I_{CK}(w) \subseteq \mathbf{R}$ and thus $w' \in \mathbf{R}$;
 furthermore, since $\mathbf{R} \subseteq \mathbf{R}_i$, $w' \in \mathbf{R}_i$.
 Thus $\sigma_i(w')$ is a best response to the beliefs of Player i at state w' .
 Hence, by Theorem 10.4.1, $\sigma_i(w') \notin D_i^0$.

2. **INDUCTIVE STEP**: assuming that the statement is true for all $k \leq m$ (for some $m \geq 0$), we prove that it is true for $k + 1$.
 The hypothesis is that, for every Player i and for every $w' \in I_{CK}(w)$, $\sigma_i(w') \notin D_i^k$;
 that is, for every $w' \in I_{CK}(w)$, $\sigma_i(w') \in S^{k+1}$.
 Select an arbitrary $w' \in I_{CK}(w)$ and an arbitrary Player i .
 Since $w \in CK\mathbf{R}$, $w' \in \mathbf{R}_i$
 and thus $\sigma_i(w')$ is a best reply to Player i 's beliefs at w' which, by hypothesis, attach
 positive probability only to strategy profiles of the other players that belong to
 S_{-i}^{k+1} (note that, since $w' \in I_{CK}(w)$, $I_{CK}(w') = I_{CK}(w)$ and, by definition of common
 knowledge partition, $I_i(w') \subseteq I_{CK}(w')$).
 By Theorem 10.4.1, it follows that $\sigma_i(w') \notin D_i^{k+1}$.

Proof of Theorem 10.2.1. Given a strategic-form game $G = \langle I, (S_i)_{i \in I}, (\pi_i)_{i \in I} \rangle$ with cardinal payoffs, construct the following epistemic model of G :

- $W = S^\infty$,
- for every Player i and every $s \in S^\infty$, $I_i(s) = \{s' \in S^\infty : s'_i = s_i\}$ (that is, the information set of Player i that contains s consists of all the strategy profiles in S^∞ that have the same strategy of Player i as in s);
- finally, let $\sigma_i(s) = s_i$ (s_i is the i^{th} component of s).

Select an arbitrary $s \in S^\infty$ and an arbitrary Player i .

By definition of S^∞ , it is not the case that s_i is strictly dominated in the restriction of G to the sets of pure strategies $S_1^\infty, \dots, S_{i-1}^\infty, S_i, S_{i+1}^\infty, \dots, S_n^\infty$.

Thus, by Theorem 10.4.1, there is a probability distribution over S_{-i}^∞ that makes $s_i = \sigma_i(s)$ a best reply.

Choose one such probability distribution and let that probability distribution give Player i 's beliefs at s .

Then $s \in \mathbf{R}_i$.

Since i was chosen arbitrarily, $s \in \mathbf{R}$;

hence, since $s \in S^\infty$ was chosen arbitrarily, $\mathbf{R} = S^\infty$.

It follows that $s \in CK\mathbf{R}$ for every $s \in S^\infty$.

10.5 Exercises

10.5.1 Exercises for Section 10.1: Model of strategic-form games

The answers to the following exercises are in Section 10.6 at the end of this chapter.

Exercise 10.1

Consider the game of Figure 10.5 (where the payoffs are von Neumann-Morgenstern payoffs) and the model of it shown in Figure 10.6.

- Find the event \mathbf{R}_1 (that is, the set of states where Player 1 is rational).
- Find the event \mathbf{R}_2 (the set of states where Player 2 is rational).
- Find the event \mathbf{R} (the set of states where both players are rational).

		Player 2					
		<i>L</i>		<i>M</i>		<i>R</i>	
Player 1	<i>A</i>	3	5	2	0	2	2
	<i>B</i>	5	2	1	2	2	1
	<i>C</i>	9	0	1	5	2	2

Figure 10.5: A strategic-form game with cardinal payoffs.




Figure 10.6: A model of the game of Figure 10.5.

Exercise 10.2

Consider the three-player game and model shown in Figures 10.7 and 10.8.

- (a) Find the event \mathbf{R}_1 (that is, the set of states where Player 1 is rational).
- (b) Find the event \mathbf{R}_2 (that is, the set of states where Player 2 is rational).
- (c) Find the event \mathbf{R}_3 (that is, the set of states where Player 3 is rational).
- (d) Find the event \mathbf{R} (that is, the set of states where all players are rational).

		Player 2					
		C			D		
Player 1	A	2	3	2	1	0	4
	B	3	6	4	0	8	0
		Player 3: E					

		Player 2				
		C			D	
A	0	0	3	2	9	7
	B	2	1	3	0	3
		Player 3: F				

Figure 10.7: A three-player strategic-form game with cardinal payoffs.

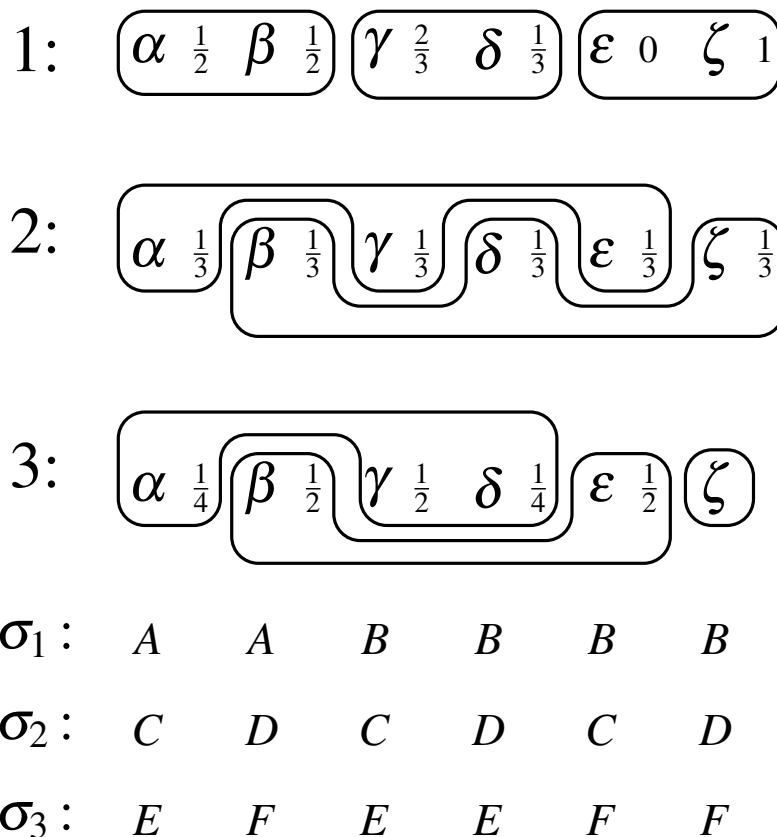


Figure 10.8: A model of the game of Figure 10.7.

10.5.2 Exercises for Section 10.2: Common knowledge of rationality in strategic-form games

The answers to the following exercises are in Section 10.6 at the end of this chapter.

Exercise 10.3

- (a) For the game and model of Exercise 10.1 find the following events:
- (i) $K_1\mathbf{R}_2$,
 - (ii) $K_2\mathbf{R}_1$,
 - (iii) $K_2K_1\mathbf{R}_2$,
 - (iv) $K_1K_2\mathbf{R}_1$,
 - (v) $CK\mathbf{R}_1$,
 - (vi) $CK\mathbf{R}_2$,
 - (vii) $CK\mathbf{R}$.
- (b) Suppose that you found a model of the game of Exercise 10.1 (Figure 10.5) and a state w in that model such that $w \in CK\mathbf{R}$.
What strategy profile could you find at w ?

Exercise 10.4

- (a) For the game and model of Exercise Exercise 10.2 find the following events:
- (i) $K_1\mathbf{R}$,
 - (ii) $K_2\mathbf{R}$,
 - (iii) $K_3\mathbf{R}$,
 - (iv) $CK\mathbf{R}_3$,
 - (v) $CK\mathbf{R}$.
- (b) Suppose that you found a model of the game of Exercise 10.2 (Figure 10.7) and a state w in that model such that $w \in CK\mathbf{R}$.
What strategy profile could you find at w ?

Exercise 10.5

For the game of Exercise 10.1 construct a model where there is a state at which there is common knowledge of rationality and the strategy profile played there is (C, L) .
[Hints: (1) four states are sufficient, (2) it is easiest to postulate degenerated beliefs where a player assigns probability 1 to a particular state in his information set.]

Exercise 10.6

For the game of Exercise 10.2 construct a model where there is a state at which there is common knowledge of rationality. [Hint: think carefully, you don't need many states!]

10.5.3 Exercises for Section 10.3: Common knowledge of rationality in extensive-form games

Exercise 10.7 — *Challenging Question***.**

Consider the extensive-form game shown in Figure 10.9 (where o_1, o_2, o_3 and o_4 are the possible outcomes).

- (a) What is the backward-induction outcome?
- (b) Write the strategic-form game associated with the extensive-form game and find all the outcomes that are supported by a pure-strategy Nash equilibrium.
- (c) Considering models of the strategic form, what strategy profiles are compatible with common knowledge of rationality?
- (d) Choose a strategy profile which is not a Nash equilibrium and construct a model of the strategic form where at a state there is common knowledge of rationality and the associated strategy profile is the one you selected.

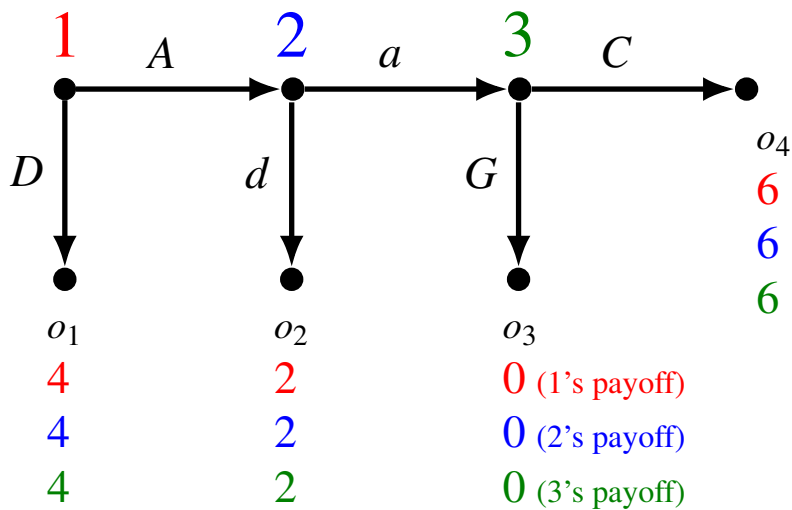


Figure 10.9: A perfect-information game with cardinal payoffs.

10.6 Solutions to Exercises

Solutions to Exercise 10.1

- (a) $\mathbf{R}_1 = \{a, b\}$. At states a and b Player 1 believes that Player 2 is playing either L or M with equal probability.

Thus his expected payoff from playing A is $\frac{1}{2}(3) + \frac{1}{2}(2) = 2.5$,

from playing B is $\frac{1}{2}(5) + \frac{1}{2}(1) = 3$

and from playing C is $\frac{1}{2}(9) + \frac{1}{2}(1) = 5$.

Hence the best choice is C and this is indeed his choice at those two states.

Thus $a, b \in \mathbf{R}_1$.

At states c and d Player 1 plays A with an expected payoff of $\frac{1}{2}(3) + \frac{1}{2}(2) = 2.5$

but he could get a higher expected payoff with C (namely $\frac{1}{2}(9) + \frac{1}{2}(3) = 6$);

thus $c, d \notin \mathbf{R}_1$. Hence $\mathbf{R}_1 = \{a, b\}$.

- (b) $\mathbf{R}_2 = \{a, b, c\}$.

- (c) $\mathbf{R} = \{a, b\}$. □

Solutions to Exercise 10.2

- (a) For each state let us calculate Player 1's expected payoff (denoted by $\mathbb{E}\pi_1$) from playing A and from playing B .

At states α and β ,

$$\mathbb{E}\pi_1(A) = \frac{1}{2}\pi_1(A, C, E) + \frac{1}{2}\pi_1(A, D, F) = \frac{1}{2}(2) + \frac{1}{2}(2) = 2,$$

$$\mathbb{E}\pi_1(B) = \frac{1}{2}\pi_1(B, C, E) + \frac{1}{2}\pi_1(B, D, F) = \frac{1}{2}(3) + \frac{1}{2}(0) = 1.5.$$

Thus A is optimal and hence, since $\sigma_1(\alpha) = \sigma_1(\beta) = A$, $\alpha, \beta \in \mathbf{R}_1$.

At states γ and δ ,

$$\mathbb{E}\pi_1(A) = \frac{2}{3}\pi_1(A, C, E) + \frac{1}{3}\pi_1(A, D, E) = \frac{2}{3}(2) + \frac{1}{3}(1) = \frac{5}{3} \text{ and}$$

$$\mathbb{E}\pi_1(B) = \frac{2}{3}\pi_1(B, C, E) + \frac{1}{3}\pi_1(B, D, E) = \frac{2}{3}(3) + \frac{1}{3}(0) = \frac{6}{3}.$$

Thus B is optimal and hence, since $\sigma_1(\gamma) = \sigma_1(\delta) = B$, $\gamma, \delta \in \mathbf{R}_1$.

At states ε and ζ ,

Player 1 assigns probability 1 to (D, F) against which A is the unique best reply and yet $\sigma_1(\varepsilon) = \sigma_1(\zeta) = B$;

hence $\varepsilon, \zeta \notin \mathbf{R}_1$. Thus $\mathbf{R}_1 = \{\alpha, \beta, \gamma, \delta\}$.

- (b) For each state let us calculate Player 2's expected payoff (denoted by $\mathbb{E}\pi_2$) from playing C and from playing D .

At states α, γ and ε ,

$$\mathbb{E}\pi_2(C) = \frac{1}{3}\pi_2(A, C, E) + \frac{1}{3}\pi_2(B, C, E) + \frac{1}{3}\pi_2(B, C, F) = \frac{1}{3}(3) + \frac{1}{3}(6) + \frac{1}{3}(1) = \frac{10}{3},$$

$$\mathbb{E}\pi_2(D) = \frac{1}{3}\pi_2(A, D, E) + \frac{1}{3}\pi_2(B, D, E) + \frac{1}{3}\pi_2(B, D, F) = \frac{1}{3}(0) + \frac{1}{3}(8) + \frac{1}{3}(3) = \frac{11}{3}.$$

Thus D is optimal and hence, since $\sigma_2(\alpha) = \sigma_2(\gamma) = \sigma_2(\varepsilon) = C$, $\alpha, \gamma, \varepsilon \notin \mathbf{R}_2$.

At states β, δ and ζ ,

$$\mathbb{E}\pi_2(C) = \frac{1}{3}\pi_2(A, C, F) + \frac{1}{3}\pi_2(B, C, E) + \frac{1}{3}\pi_2(B, C, F) = \frac{1}{3}(0) + \frac{1}{3}(6) + \frac{1}{3}(1) = \frac{7}{3},$$

$$\mathbb{E}\pi_2(D) = \frac{1}{3}\pi_2(A, D, F) + \frac{1}{3}\pi_2(B, D, E) + \frac{1}{3}\pi_2(B, D, F) = \frac{1}{3}(9) + \frac{1}{3}(8) + \frac{1}{3}(3) = \frac{20}{3}.$$

Thus D is optimal and hence, since $\sigma_2(\beta) = \sigma_2(\delta) = \sigma_2(\zeta) = D$, $\beta, \delta, \zeta \in \mathbf{R}_2$.

Thus $\mathbf{R}_2 = \{\beta, \delta, \zeta\}$.

- (c) For each state let us calculate Player 3's expected payoff (denoted by $\mathbb{E}\pi_3$) from playing E and from playing F .

At states α, γ and δ

$$\begin{aligned}\mathbb{E}\pi_3(E) &= \frac{1}{4}\pi_3(A, C, E) + \frac{1}{2}\pi_3(B, C, E) + \frac{1}{4}\pi_3(B, D, E) = \frac{1}{4}(2) + \frac{1}{2}(4) + \frac{1}{4}(0) = 2.5, \\ \mathbb{E}\pi_3(F) &= \frac{1}{4}\pi_3(A, C, F) + \frac{1}{2}\pi_3(B, C, F) + \frac{1}{4}\pi_3(B, D, F) = \frac{1}{4}(3) + \frac{1}{2}(3) + \frac{1}{4}(1) = 2.5.\end{aligned}$$

Thus both E and F are optimal and hence $\alpha, \gamma, \delta \in \mathbf{R}_3$.

At states β , and ε

$$\begin{aligned}\mathbb{E}\pi_3(E) &= \frac{1}{2}\pi_3(A, D, E) + \frac{1}{2}\pi_3(B, C, E) = \frac{1}{2}(4) + \frac{1}{2}(4) = 4, \\ \mathbb{E}\pi_3(F) &= \frac{1}{2}\pi_3(A, D, F) + \frac{1}{2}\pi_3(B, C, F) = \frac{1}{2}(7) + \frac{1}{2}(3) = 5.\end{aligned}$$

Thus F is optimal and hence, since $\sigma_3(\beta) = \sigma_3(\varepsilon) = F$, $\beta, \varepsilon \in \mathbf{R}_3$.

At state ζ Player 3 knows that Players 1 and 2 play (B, D) and she is best replying with F . Thus $\zeta \in \mathbf{R}_3$.

Hence $\mathbf{R}_3 = \{\alpha, \beta, \gamma, \delta, \varepsilon, \zeta\}$.

- (d) $\mathbf{R} = \mathbf{R}_1 \cap \mathbf{R}_2 \cap \mathbf{R}_3 = \{\beta, \delta\}$. □

Solutions to Exercise 10.3

In Exercise 10.1 we determined that $\mathbf{R}_1 = \{a, b\}$, $\mathbf{R}_2 = \{a, b, c\}$ and $\mathbf{R} = \{a, b\}$. Thus

- (a) (i) $K_1\mathbf{R}_2 = \{a, b\}$, (ii) $K_2\mathbf{R}_1 = \{a\}$, (iii) $K_2K_1\mathbf{R}_2 = \{a\}$, (iv) $K_1K_2\mathbf{R}_1 = \emptyset$.

The common knowledge partition consists of a single information set containing all the states. Thus (v) $CK\mathbf{R}_1 = \emptyset$, (vi) $CK\mathbf{R}_2 = \emptyset$, (vii) $CK\mathbf{R} = \emptyset$.

- (b) By Theorem 10.2.1, at a state at which there is common knowledge of rationality one can only find a strategy profile that survives the cardinal iterated deletion of strictly dominated strategies.

In this game, for Player 1 strategy B is strictly dominated by the mixed strategy

$$\begin{pmatrix} A & C \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix};$$

after deleting B , for Player 2 strategy R is strictly dominated by the mixed strategy

$$\begin{pmatrix} L & M \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}.$$

Thus the iterated deletion of strictly dominated strategies yields the set of strategy profiles $\{(A, L), (A, M), (C, L), (C, M)\}$. Hence at a state where there is common knowledge of rationality one could only find one of these four strategy profiles. □

Solutions to Exercise 10.4

In Exercise 10.2 we determined that $\mathbf{R} = \{\beta, \delta\}$.

- (a) (i) $K_1\mathbf{R} = \emptyset$, (ii) $K_2\mathbf{R} = \emptyset$, (iii) $K_3\mathbf{R} = \emptyset$. The common knowledge partition consists of a single information set containing all the states.

Thus (since $\mathbf{R}_3 = \{\alpha, \beta, \gamma, \delta, \varepsilon, \zeta\}$)

- (iv) $CK\mathbf{R}_3 = \{\alpha, \beta, \gamma, \delta, \varepsilon, \zeta\}$, (v) $CK\mathbf{R} = \emptyset$.

- (b) By Theorem 10.2.1, at a state at which there is common knowledge of rationality one can only find a strategy profile that survives the cardinal iterated deletion of strictly dominated strategies. Since in this game there are no strictly dominated strategies, at such a state one could find *any* strategy profile. \square

Solutions to Exercise 10.5 In the model shown in Figure 10.10 there is common knowledge of rationality at every state. At state α the strategy profile played is (C, L) . \square




Figure 10.10: A model of the game of Figure 10.5 where there is common knowledge of rationality at every state.

Solutions to Exercise 10.6 Whenever a game has a Nash equilibrium in pure strategies, a one-state model where that Nash equilibrium is played is such that there is common knowledge of rationality at that state. Thus the model shown in Figure 10.11 provides a possible answer. \square

1: α
 2: α
 3: α
 A
 D
 F

Figure 10.11: A one-state model of the game of Figure 10.5 where there is common knowledge of rationality and the strategy profile is (A, D, F) .

Solutions to Exercise 10.7

- (a) The backward induction outcome is o_4 , with associated strategy profile (A, a, C) and payoffs $(6, 6, 6)$.
- (b) The strategic form is shown in Figure 10.12 (the Nash equilibria have been highlighted):
 The Nash equilibria are: (D, d, G) , (D, a, G) , (D, d, C) and (A, a, C) .
 Thus the only two outcomes sustained by a Nash equilibrium are o_1 , with payoffs $(4, 4, 4)$, and o_4 , with payoffs $(6, 6, 6)$.

		Player 2										
		<i>d</i>			<i>a</i>							
Player 1	<i>D</i>	4	4	4	4	4	4		Player 2	<i>d</i>	<i>a</i>	
	<i>A</i>	2	2	2	0	0	0				<i>D</i>	4
		Player 3: <i>G</i>								Player 3: <i>C</i>		

Figure 10.12: The strategic-form associated with the extensive-form game of Figure 10.9.

- (c) By Theorem 10.2.1, the outcomes that are compatible with common knowledge of rationality are those associated with strategy profiles that survive the iterated deletion of strictly dominated strategies. Since no player has any strictly dominated strategies, all the outcomes are compatible with common knowledge of rationality. (Note that, for Player 3, *C* weakly dominates *G*, but not strictly.)

- (d) In the model shown in Figure 10.13 at state α the players choose (A, d, G) , which is not a Nash equilibrium; furthermore, there is no Nash equilibrium whose associated outcome is o_2 (with payoffs $(2,2,2)$).

In this model $\alpha \in CKR$ (in particular, Player 1's choice of A is rational, given his belief that Player 2 plays d and a with equal probability). \square

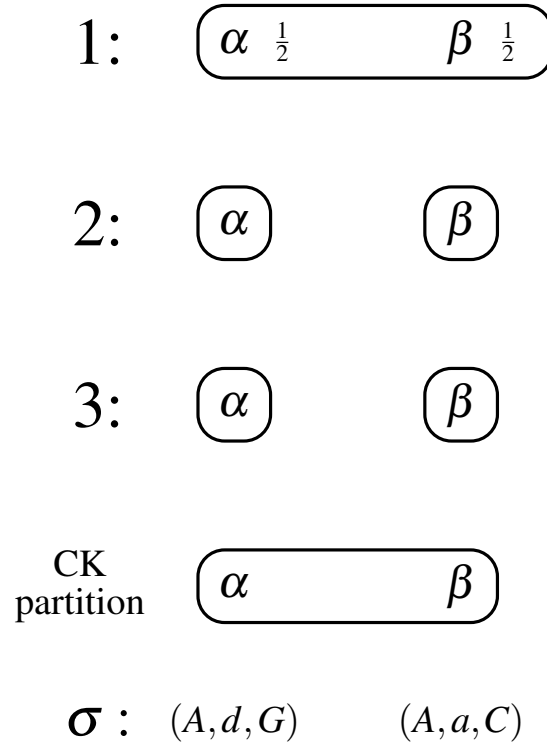


Figure 10.13: A model of the strategic-form of the game of Figure 10.9 where the strategy profile (A, d, G) , which is not a Nash equilibrium, is played under common knowledge of rationality.

IV

Refinements of Subgame Perfection

11 Weak Sequential Equilibrium	369
11.1 Assessments and sequential rationality	
11.2 Bayesian updating at reached information sets	
11.3 A first attempt: Weak sequential equilibrium	
11.4 Exercises	
11.5 Solutions to Exercises	
12 Sequential Equilibrium	403
12.1 Consistent assessments	
12.2 Sequential equilibrium	
12.3 Is 'consistency' a good notion?	
12.4 Exercises	
12.5 Solutions to Exercises	
13 Perfect Bayesian Equilibrium	429
13.1 Belief revision and AGM consistency	
13.2 Bayesian consistency	
13.3 Perfect Bayesian equilibrium	
13.4 Adding independence	
13.5 Characterization of SE in terms of PBE	
13.6 History-based definition of extensive-form game	
13.7 Proofs	
13.8 Exercises	
13.9 Solutions to Exercises	

11. Weak Sequential Equilibrium

11.1 Assessments and sequential rationality

At the end of Chapter 6 (Section 6.4) we showed that, although the notion of subgame-perfect equilibrium is a refinement of Nash equilibrium, it is not strong enough to eliminate all “unreasonable” Nash equilibria. One reason for this is that a subgame-perfect equilibrium σ allows a player’s strategy to include a strictly dominated choice at an information set that is not reached by the play induced by σ . In order to eliminate such possibilities we need to define the notion of equilibrium for dynamic games in terms of a more complex object than merely a strategy profile. We need to add a description of what the players believe when it is their turn to move.

Definition 11.1.1 Given an extensive-form game G , an *assessment* for G is a pair (σ, μ) , where σ is a profile of behavioral strategies and μ is a list of probability distributions, one for every information set, over the nodes in that information set. We call μ a *system of beliefs*.

The system of beliefs μ specifies, for every information set, the beliefs – about past moves – that the player who moves at that information set would have if told that her information set had been reached.

Consider, for example, the extensive-form game of Figure 11.1.




Figure 11.1: An extensive-form game with cardinal payoffs.

A possible assessment for this game is (σ, μ) with

$$\sigma = \left(\begin{array}{ccc|cc|cc|cc} a & b & c & f & g & d & e & h & k \\ \frac{1}{8} & \frac{3}{8} & \frac{4}{8} & 1 & 0 & \frac{3}{4} & \frac{1}{4} & \frac{1}{5} & \frac{4}{5} \end{array} \right) \quad \text{and} \quad \mu = \left(\begin{array}{cc|cc} x & y & w & z \\ \frac{2}{3} & \frac{1}{3} & \frac{1}{2} & \frac{1}{2} \end{array} \right).$$

Note that, typically, we will not bother to include in μ the trivial probability distributions over singleton information sets.¹

The interpretation of this assessment is that

- ◇ Player 1 plans to play a with probability $\frac{1}{8}$, b with probability $\frac{3}{8}$ and c with probability $\frac{4}{8}$;
- ◇ Player 2 plans to play f if her information set $\{x, y\}$ is reached and to mix between d and e with probabilities $\frac{3}{4}$ and $\frac{1}{4}$, respectively, if her decision node t is reached;
- ◇ Player 3 plans to mix between h and k with probabilities $\frac{1}{5}$ and $\frac{4}{5}$, respectively, if his information set $\{w, z\}$ is reached;
- ◇ Player 2 – if informed that her information set $\{x, y\}$ had been reached – would attach probability $\frac{2}{3}$ to Player 1 having played a (node x) and probability $\frac{1}{3}$ to Player 1 having played b (node y);

¹A complete specification of μ would be $\left(\begin{array}{c|c|cc|cc} s & t & x & y & w & z \\ 1 & 1 & \frac{2}{3} & \frac{1}{3} & \frac{1}{2} & \frac{1}{2} \end{array} \right)$ where s is the decision node of Player 1 (the root) and t is the decision node of Player 2 following choice c .

- ◇ Player 3 – if informed that his information set $\{w, z\}$ had been reached – would attach probability $\frac{1}{2}$ to node w (that is, to the sequence of moves cd) and probability $\frac{1}{2}$ to node z (that is, to the sequence of moves ce).

In order for an assessment (σ, μ) to be considered “reasonable” we will impose two requirements:

1. The choices specified by σ should be optimal given the beliefs specified by μ . We call this requirement *sequential rationality*.
2. The beliefs specified by μ should be consistent with the strategy profile σ . We call this requirement *Bayesian updating*.

Before we give a precise definition of these concepts, we shall illustrate them with reference to the assessment considered in the game of Figure 11.1, namely

$$\sigma = \left(\begin{array}{ccc|cc} a & b & c & f & g \\ \frac{1}{8} & \frac{3}{8} & \frac{4}{8} & 1 & 0 \end{array} \middle| \begin{array}{cc|cc} d & e & h & k \\ \frac{3}{4} & \frac{1}{4} & \frac{1}{5} & \frac{4}{5} \end{array} \right) \quad \text{and} \quad \mu = \left(\begin{array}{cc|cc} x & y & w & z \\ \frac{2}{3} & \frac{1}{3} & \frac{1}{2} & \frac{1}{2} \end{array} \right).$$

This assessment fails to satisfy sequential rationality, because, for example, at Player 3’s information set $\{w, z\}$ the planned mixed strategy $\left(\begin{array}{cc} h & k \\ \frac{1}{5} & \frac{4}{5} \end{array} \right)$ yields Player 3 – given her beliefs $\left(\begin{array}{cc} w & z \\ \frac{1}{2} & \frac{1}{2} \end{array} \right)$ – an expected payoff of $\frac{1}{2} \left[\frac{1}{5}(3) + \frac{4}{5}(1) \right] + \frac{1}{2} \left[\frac{1}{5}(0) + \frac{4}{5}(1) \right] = \frac{11}{10}$, while she could get a higher expected payoff, namely $\frac{1}{2}(3) + \frac{1}{2}(0) = \frac{3}{2} = \frac{15}{10}$, by playing h with probability 1.

This assessment also fails the rule for belief updating (Definition 9.4.1, Chapter 9); for example, given σ the prior probability that node x is reached is $P(x) = \frac{1}{8}$ (it is the probability with which Player 1 plays a) and the prior probability that node y is reached is $P(y) = \frac{3}{8}$ (the probability with which Player 1 plays b), so that updating on information $\{x, y\}$ one gets

$$P(x|\{x, y\}) = \frac{P(x)}{P(\{x, y\})} = \frac{P(x)}{P(x) + P(y)} = \frac{\frac{1}{8}}{\frac{1}{8} + \frac{3}{8}} = \frac{1}{4}$$

and $P(y|\{x, y\}) = \frac{3}{4}$. Thus, in order to be consistent with the rule for belief updating, Player 2’s beliefs should be

$$\left(\begin{array}{cc} x & y \\ \frac{1}{4} & \frac{3}{4} \end{array} \right).$$

Now we can turn to the formal definitions. First we need to introduce some notation. If σ is a profile of behavior strategies and a is a choice of Player i (at some information set of Player i), we denote by $\sigma(a)$ the probability that σ_i (the strategy of Player i that is part of σ) assigns to a .

For example, for the game of Figure 11.1, if

$$\sigma = \left(\begin{array}{ccc|cc|cc|cc} a & b & c & f & g & d & e & h & k \\ \frac{1}{8} & \frac{3}{8} & \frac{4}{8} & 1 & 0 & \frac{3}{4} & \frac{1}{4} & \frac{1}{5} & \frac{4}{5} \end{array} \right)$$

then $\sigma(b) = \frac{3}{8}$, $\sigma(g) = 0$, $\sigma(d) = \frac{3}{4}$, etc.

Similarly, if μ is a system of beliefs and x is a decision node, then we denote by $\mu(x)$ the probability that the relevant part of μ assigns to x . For example, if

$$\mu = \left(\begin{array}{c|c|cc|cc} s & t & x & y & w & z \\ 1 & 1 & \frac{2}{3} & \frac{1}{3} & \frac{1}{2} & \frac{1}{2} \end{array} \right)$$

or, written more succinctly,

$$\mu = \left(\begin{array}{cc|cc} x & y & w & z \\ \frac{2}{3} & \frac{1}{3} & \frac{1}{2} & \frac{1}{2} \end{array} \right)$$

then $\mu(s) = 1$, $\mu(y) = \frac{1}{3}$, $\mu(w) = \frac{1}{2}$, etc.

Recall that Z denotes the set of terminal nodes and, for every Player i , $\pi_i : Z \rightarrow \mathbb{R}$ is the payoff function of Player i .

- Given a decision node x , let $Z(x) \subseteq Z$ be the set of terminal nodes that can be reached starting from x . For example, in the game of Figure 11.1, $Z(t) = \{z_5, z_6, z_7, z_8\}$.
- Given a behavior strategy profile σ and a decision node x , let $\mathbb{P}_{x,\sigma}$ be the probability distribution over $Z(x)$ induced by σ , that is, if $z \in Z(x)$ and $\langle a_1, \dots, a_m \rangle$ is the sequence of choices that leads from x to z then $\mathbb{P}_{x,\sigma}(z)$ is the product of the probabilities of those choices: $\mathbb{P}_{x,\sigma}(z) = \sigma(a_1) \times \sigma(a_2) \times \dots \times \sigma(a_m)$.

For example, in the game of Figure 11.1, if σ is the strategy profile

$$\sigma = \left(\begin{array}{ccc|cc|cc|cc} a & b & c & f & g & d & e & h & k \\ \frac{1}{8} & \frac{3}{8} & \frac{4}{8} & 1 & 0 & \frac{3}{4} & \frac{1}{4} & \frac{1}{5} & \frac{4}{5} \end{array} \right)$$

and t is Player 2's decision node after choice c of Player 1, then

$$\mathbb{P}_{t,\sigma}(z_5) = \sigma(d) \sigma(h) = \frac{3}{4} \left(\frac{1}{5} \right) = \frac{3}{20}.$$

If H is an information set of Player i we denote by $\pi_i(H|\sigma, \mu)$ the expected payoff of Player i starting from information set H , given the beliefs specified by μ at H and given the choices prescribed by σ at H and at the information sets that come after H , that is,

$$\pi_i(H|\sigma, \mu) = \sum_{x \in H} \left[\mu(x) \left(\sum_{z \in Z(x)} \mathbb{P}_{x, \sigma}(z) \pi_i(z) \right) \right].$$

For example, in the game of Figure 11.1, if

$$\sigma = \left(\begin{array}{ccc|cc} a & b & c & f & g \\ \frac{1}{8} & \frac{3}{8} & \frac{4}{8} & 1 & 0 \end{array} \middle| \begin{array}{cc|cc} d & e & h & k \\ \frac{3}{4} & \frac{1}{4} & \frac{1}{5} & \frac{4}{5} \end{array} \right) \quad \text{and} \quad \mu = \left(\begin{array}{cc|cc} x & y & w & z \\ \frac{2}{3} & \frac{1}{3} & \frac{1}{2} & \frac{1}{2} \end{array} \right),$$

then (as we computed above)

$$\begin{aligned} \pi_3(\{w, z\}|\sigma, \mu) &= \\ &= \mu(w) \left(\mathbb{P}_{w, \sigma}(z_5) \pi_3(z_5) + \mathbb{P}_{w, \sigma}(z_6) \pi_3(z_6) \right) + \mu(z) \left(\mathbb{P}_{z, \sigma}(z_7) \pi_3(z_7) + \mathbb{P}_{z, \sigma}(z_8) \pi_3(z_8) \right) \\ &= \frac{1}{2} \left[\frac{1}{5}(3) + \frac{4}{5}(1) \right] + \frac{1}{2} \left[\frac{1}{5}(0) + \frac{4}{5}(1) \right] = \frac{11}{10}. \end{aligned}$$

Recall that if σ is a strategy profile and i is a player, then σ_{-i} denotes the profile of strategies of the players other than i and we can use (σ_i, σ_{-i}) as an alternative way of denoting σ ; furthermore, if τ_i is a strategy of Player i , we denote by (τ_i, σ_{-i}) the strategy profile obtained from σ by replacing σ_i with τ_i (and leaving everything else unchanged).

Definition 11.1.2 Fix an extensive-form game and an *assessment* (σ, μ) . We say that Player i 's behavior strategy σ_i is *sequentially rational* if, for every information set H of Player i , $\pi_i(H|(\sigma_i, \sigma_{-i}), \mu) \geq \pi_i(H|(\tau_i, \sigma_{-i}), \mu)$, for every behavior strategy τ_i of Player i . We say that σ is *sequentially rational* if, for every Player i , σ_i is sequentially rational.

Note that for Player i 's strategy σ_i to be sequentially rational it is not sufficient (although it is necessary) that at every information set H of Player i the choice(s) at H prescribed by σ_i be optimal (given the choices of the other players specified by σ_{-i}): we need to check if Player i could improve her payoff by changing her choice(s) not only at H but also at information sets of hers that follow H . To see this, consider the game as shown in Figure 11.2.




Figure 11.2: A strategic-form game with cardinal payoffs.

Let (σ, μ) be the assessment where σ is the pure-strategy profile $(a, (c, e))$ and

$$\mu = \left(\begin{array}{cc|cc} s & t & x & y \\ 1 & 0 & 0 & 1 \end{array} \right).$$

For Player 2, e is rational at information set $\{x, y\}$ because – given her belief that she is making her choice at node y – e gives her a payoff of 1 while f would give her a payoff of 0;

furthermore, at information set $\{s, t\}$ – given her belief that she is making her choice at node s and given her future choice of e at $\{x, y\}$ – choice c is better than choice d because the former gives her a payoff of 1 while the latter gives her a payoff of 0.

However, the strategy (c, e) (while sequentially rational at $\{x, y\}$) is not sequentially rational at $\{s, t\}$ because – given her belief that she is making her choice at node s – with (c, e) she gets a payoff of 1 but if she switched to (d, f) she would get a payoff of 2; in other words, Player 2 can improve her payoff by changing her choice at $\{s, t\}$ from c to d and also her future planned choice at $\{x, y\}$ from e to f .²

Note that the pure-strategy profile $(a, (c, e))$ is not a Nash equilibrium: for Player 2 the unique best reply to a is (d, f) .

²It is possible to impose restrictions on the system of beliefs μ such that sequential rationality as defined in Definition 11.1.2 is equivalent to the weaker condition that, at every information set H , the corresponding player cannot increase her payoff by changing her choice(s) at H only. We will not discuss this condition here. The interested reader is referred to Hendon et al (1996) and Perea (2002).

An example of a sequentially rational assessment for the game shown in Figure 11.3 is

$$\sigma = \left(\begin{array}{ccc|cc} L & M & R & A & B \\ 0 & 0 & 1 & \frac{1}{2} & \frac{1}{2} \\ \hline & & & c & d \end{array} \right) \quad \text{and} \quad \mu = \left(\begin{array}{cc|ccc} x & y & u & v & w \\ \frac{1}{4} & \frac{3}{4} & \frac{1}{5} & \frac{3}{5} & \frac{1}{5} \\ \hline & & & & \end{array} \right).$$




Figure 11.3: The assessment $\sigma = (R, (\frac{1}{2}A, \frac{1}{2}B), c)$, $\mu = ((\frac{1}{4}x, \frac{3}{4}y), (\frac{1}{5}u, \frac{3}{5}v, \frac{1}{5}w))$ is sequentially rational.

Let us first verify sequential rationality of Player 3's strategy.

At her information set $\{u, v, w\}$, given her beliefs,

- c gives a payoff of $\frac{1}{5}(5) + \frac{3}{5}(1) + \frac{1}{5}(2) = 2$,
- while d gives a payoff of $\frac{1}{5}(2) + \frac{3}{5}(2) + \frac{1}{5}(2) = 2$.
- Thus c is optimal (as would be d and any randomization over c and d).

Now consider Player 2:

At his information set $\{x, y\}$, given his beliefs and given the strategy of Player 3,

- A gives a payoff of $\frac{1}{4}(5) + \frac{3}{4}(3) = 3.5$
- and B gives a payoff of $\frac{1}{4}(2) + \frac{3}{4}(4) = 3.5$;
- thus any mixture of A and B is optimal, in particular, the mixture $\begin{pmatrix} A & B \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$ is optimal.

Finally, at the root, R gives Player 1 a payoff of 3, L a payoff of $\frac{1}{2}(4) + \frac{1}{2}(1) = 2.5$ and M a payoff of $\frac{1}{2}(0) + \frac{1}{2}(6) = 3$; thus R is optimal (as would be any mixture over M and R).

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 11.4.1 at the end of this chapter.

11.2 Bayesian updating at reached information sets

The second requirement for an assessment to be “reasonable” is that the beliefs encoded in μ should be consistent with the behavior postulated by σ in the sense that the rule for belief updating (Definition 9.4.1, Chapter 9), should be used to form those beliefs, whenever it is applicable. We shall call this rule *Bayesian updating*.

Let x be a node that belongs to information set H and let $\mathbb{P}_{root,\sigma}(x)$ be the probability that node x is reached (from the root of the tree) if σ is implemented.³

Let $\mathbb{P}_{root,\sigma}(H) = \sum_{y \in H} \mathbb{P}_{root,\sigma}(y)$ be the probability that information set H is reached (that is, the probability that some node in H is reached).

Then Bayesian updating requires that the probability that is assigned to x – given the information that H has been reached – be given by the conditional probability $\mathbb{P}_{root,\sigma}(x|H) = \frac{\mathbb{P}_{root,\sigma}(x)}{\mathbb{P}_{root,\sigma}(H)}$,

of course, this conditional probability is well defined if and only if $\mathbb{P}_{root,\sigma}(H) > 0$.

Definition 11.2.1 Consider an extensive-form game and a behavioral strategy profile σ . We say that an information set H is *reachable by σ* if $\mathbb{P}_{root,\sigma}(H) > 0$. That is, H is *reachable by σ* if at least one node in H is reached with positive probability when the game is played according to σ .

For example, in the game-frame of Figure 11.4 (which partially reproduces of Figure 11.1 by omitting the payoffs), if

$$\sigma = \left(\begin{array}{ccc|cc} a & b & c & f & g \\ \frac{1}{3} & \frac{2}{3} & 0 & 1 & 0 \end{array} \middle| \begin{array}{cc|cc} d & e & h & k \\ \frac{3}{4} & \frac{1}{4} & 1 & 0 \end{array} \right)$$

then $\{x, y\}$ is reachable (with probability 1) by σ , while $\{w, z\}$ is not.⁴

Definition 11.2.2 Given an extensive-form game and an assessment (σ, μ) , we say that (σ, μ) satisfies *Bayesian updating at reachable information sets* if, for every information set H and every node $x \in H$, if H is reachable by σ (that is, if $\mathbb{P}_{root,\sigma}(H) > 0$) then $\mu(x) = \frac{\mathbb{P}_{root,\sigma}(x)}{\mathbb{P}_{root,\sigma}(H)}$.

For example, in the game-frame of Figure 11.4, the assessment

$$\sigma = \left(\begin{array}{ccc|cc} a & b & c & f & g \\ \frac{1}{9} & \frac{5}{9} & \frac{3}{9} & 0 & 1 \end{array} \middle| \begin{array}{cc|cc} d & e & h & k \\ \frac{3}{4} & \frac{1}{4} & 1 & 0 \end{array} \right) \quad \text{and} \quad \mu = \left(\begin{array}{cc|cc} x & y & w & z \\ \frac{1}{6} & \frac{5}{6} & \frac{3}{4} & \frac{1}{4} \end{array} \right)$$

satisfies Bayesian updating at reachable information sets.

³That is, if $\langle a_1, \dots, a_m \rangle$ is the sequence of choices that leads from the root to x then $\mathbb{P}_{root,\sigma}(x) = \sigma(a_1) \times \dots \times \sigma(a_m)$.

⁴ $\mathbb{P}_{root,\sigma}(x) = \frac{1}{3}$, $\mathbb{P}_{root,\sigma}(y) = \frac{2}{3}$ and thus $\mathbb{P}_{root,\sigma}(\{x, y\}) = \frac{1}{3} + \frac{2}{3} = 1$; on the other hand, $\mathbb{P}_{root,\sigma}(w) = \sigma(c) \sigma(d) = 0 \left(\frac{3}{4}\right) = 0$ and $\mathbb{P}_{root,\sigma}(z) = \sigma(c) \sigma(e) = 0 \left(\frac{1}{4}\right) = 0$, so that $\mathbb{P}_{root,\sigma}(\{w, z\}) = 0$.




Figure 11.4: A game-frame in extensive form.

In fact, we have that

$$\mathbb{P}_{root,\sigma}(x) = \frac{1}{9}, \quad \mathbb{P}_{root,\sigma}(y) = \frac{5}{9}, \quad \mathbb{P}_{root,\sigma}(\{x,y\}) = \frac{1}{9} + \frac{5}{9} = \frac{6}{9},$$

$$\mathbb{P}_{root,\sigma}(w) = \frac{3}{9} \left(\frac{3}{4}\right) = \frac{9}{36}, \quad \mathbb{P}_{root,\sigma}(z) = \frac{3}{9} \left(\frac{1}{4}\right) = \frac{3}{36}$$

and $\mathbb{P}_{root,\sigma}(\{w,x\}) = \frac{9}{36} + \frac{3}{36} = \frac{12}{36}$.

Thus

$$\frac{\mathbb{P}_{root,\sigma}(x)}{\mathbb{P}_{root,\sigma}(\{x,y\})} = \frac{\frac{1}{9}}{\frac{6}{9}} = \frac{1}{6} = \mu(x), \quad \frac{\mathbb{P}_{root,\sigma}(y)}{\mathbb{P}_{root,\sigma}(\{x,y\})} = \frac{\frac{5}{9}}{\frac{6}{9}} = \frac{5}{6} = \mu(y),$$

$$\frac{\mathbb{P}_{root,\sigma}(w)}{\mathbb{P}_{root,\sigma}(\{w,z\})} = \frac{\frac{9}{36}}{\frac{12}{36}} = \frac{3}{4} = \mu(w) \quad \text{and} \quad \frac{\mathbb{P}_{root,\sigma}(z)}{\mathbb{P}_{root,\sigma}(\{w,z\})} = \frac{\frac{3}{36}}{\frac{12}{36}} = \frac{1}{4} = \mu(z).$$

R Note that the condition “Bayesian updating at reachable information sets” is trivially satisfied at every information set H that is *not* reachable by σ that is, at every information set H such that $\mathbb{P}_{root,\sigma}(H) = 0$.⁵

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 11.4.2 at the end of this chapter.

⁵In logic a proposition of the form “if A then B ” (A is called the *antecedent* and B is called the *consequent*) is false only when A is true and B is false. Thus, in particular, the proposition “if A then B ” is true whenever A is false (whatever the truth value of B). In our case the antecedent is “ $\mathbb{P}_{root,\sigma}(H) > 0$ ”.

11.3 A first attempt: Weak sequential equilibrium

Our objective is to find a refinement of subgame-perfect equilibrium that rules out “unreasonable” equilibria. Moving from strategy profiles to assessments allowed us to judge the rationality of a choice at an information set independently of whether that information set is reached by the strategy profile under consideration: the requirement of sequential rationality rules out choices that are strictly dominated at an information set. As a first attempt we define a notion of equilibrium based on the two requirements of sequential rationality and Bayesian updating at reached information sets.

Definition 11.3.1 A *weak sequential equilibrium* is an assessment (σ, μ) which satisfies two requirements: (1) sequential rationality and (2) Bayesian updating at reached information sets.

Before we discuss the properties of weak sequential equilibria we illustrate the type of reasoning that one needs to go through in order to find weak sequential equilibria. Consider first the game of Figure 11.5.




Figure 11.5: An extensive-form game with cardinal payoffs.

Let us see if there is a weak sequential equilibrium (σ, μ) where Player 1’s strategy in σ is a pure strategy. The set of pure strategies of Player 1 is $S_1 = \{(A, C), (B, D), (B, C), (A, D)\}$. The strategy of Player 1 determines the beliefs of Player 2 at her information set $\{x, y\}$. Let us consider the four possibilities.

- If Player 1’s strategy is (A, C) , then Player 2’s information set $\{x, y\}$ is reached with positive probability and the only beliefs that are consistent with Bayesian updating are $\begin{pmatrix} x & y \\ 0 & 1 \end{pmatrix}$, so that – by sequential rationality – Player 2 must choose E .

However, if Player 2’s strategy is E then at node s it is not sequentially rational for Player 1 to choose A . Thus there is no weak sequential equilibrium where Player 1’s strategy is (A, C) .

- If Player 1's strategy is (B, D) , then Player 2's information set $\{x, y\}$ is reached with positive probability and the only beliefs that are consistent with Bayesian updating are $\begin{pmatrix} x & y \\ 1 & 0 \end{pmatrix}$, so that – by sequential rationality – Player 2 must choose F .

However, if Player 2's strategy is F then at node s it is not sequentially rational for Player 1 to choose B . Thus there is no weak sequential equilibrium where Player 1's strategy is (B, D) .

- If Player 1's strategy is (B, C) , then Player 2's information set $\{x, y\}$ is reached (with probability 1) and the only beliefs that are consistent with Bayesian updating are $\begin{pmatrix} x & y \\ \frac{1}{4} & \frac{3}{4} \end{pmatrix}$.

- Given these beliefs, Player 2's payoff from playing E is $\frac{1}{4}(0) + \frac{3}{4}(1) = \frac{3}{4}$ and her payoff from playing F is $\frac{1}{4}(3) + \frac{3}{4}(0) = \frac{3}{4}$.

- Thus any mixed strategy is sequentially rational for Player 2, that is, for any $p \in [0, 1]$, $\begin{pmatrix} E & F \\ p & 1-p \end{pmatrix}$ is sequentially rational.

- At node s choice B is sequentially rational for Player 1 if and only if the expected payoff from playing B is at least 1 (which is the payoff from playing A): $2p + 0(1 - p) \geq 1$, that is, $p \geq \frac{1}{2}$.

- At node t choice C is sequentially rational for Player 1 if and only if the expected payoff from playing C is at least 2 (which is the payoff from playing D): $3p + 0(1 - p) \geq 2$, that is, $p \geq \frac{2}{3}$.

- Hence, if $p \geq \frac{2}{3}$ then both B and C are sequentially rational. Thus we have an infinite number of weak sequential equilibria at which Player 1's strategy is (B, C) : for every $p \in [\frac{2}{3}, 1]$, (σ, μ) is a weak sequential equilibrium, where

$$\sigma = \left(\begin{array}{cc|cc|cc} A & B & C & D & E & F \\ 0 & 1 & 1 & 0 & p & 1-p \end{array} \right) \quad \text{and} \quad \mu = \begin{pmatrix} x & y \\ \frac{1}{4} & \frac{3}{4} \end{pmatrix}$$

- If Player 1's strategy is (A, D) , then Player 2's information set $\{x, y\}$ is not reached and thus, according to the notion of weak sequential equilibrium, any beliefs are allowed there. Let $\begin{pmatrix} E & F \\ p & 1-p \end{pmatrix}$ be Player 2's strategy. From previous calculations we have that if $p \leq \frac{1}{2}$ then A is sequentially rational at node s and D is sequentially rational at node t .

One possibility is to set $p = 0$; this means that Player 2 chooses the pure strategy F and this is sequentially rational if and only if her beliefs are $\begin{pmatrix} x & y \\ q & 1-q \end{pmatrix}$ with

$$\underbrace{3q + 0(1-q)}_{\text{payoff from } F} \geq \underbrace{0q + 1(1-q)}_{\text{payoff from } E}, \quad \text{that is,} \quad q \geq \frac{1}{4}.$$

Thus we have an infinite number of weak sequential equilibria: for every $q \in [\frac{1}{4}, 1]$, (σ, μ) is a weak sequential equilibrium, where

$$\sigma = \left(\begin{array}{cc|cc|cc} A & B & C & D & E & F \\ 1 & 0 & 0 & 1 & 0 & 1 \end{array} \right) \quad \text{and} \quad \mu = \begin{pmatrix} x & y \\ q & 1-q \end{pmatrix}.$$

Next we consider the more complex game shown in Figure 11.6.

For simplicity, let us limit the search to pure-strategy weak sequential equilibria. How should we proceed? The first step is to see if the game itself can be simplified by checking if there are any information sets where there is a strictly dominant choice: if there is such a choice then, no matter what beliefs the relevant player has at that information set, sequential rationality requires that choice to be selected.

In the game of Figure 11.6 there is indeed such an information set, namely information set $\{w, z\}$ of Player 2: here L is strictly better than R for Player 2 at both nodes, that is, R is strictly dominated by L .

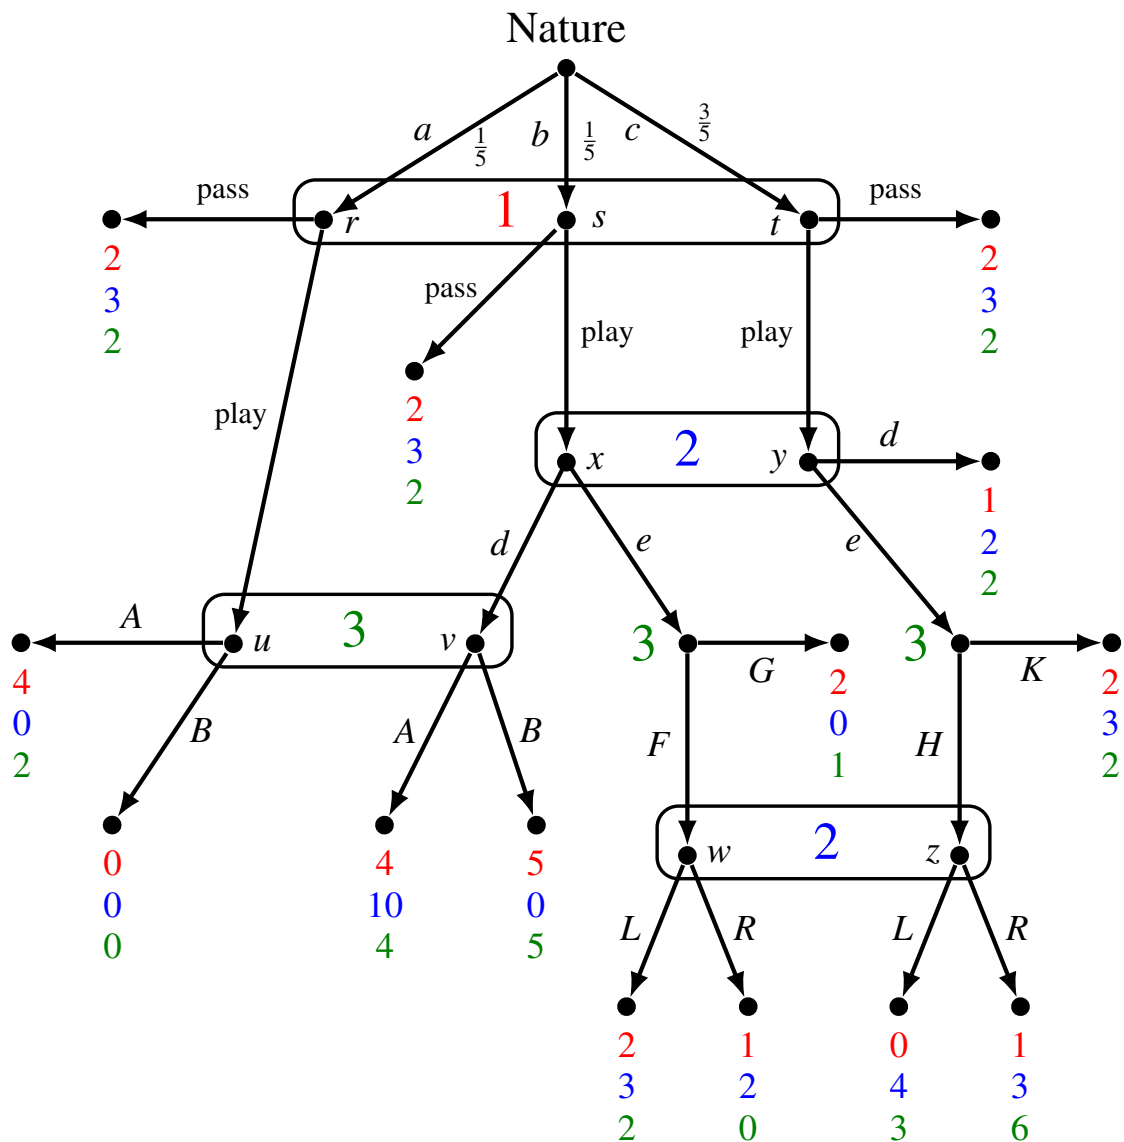


Figure 11.6: An extensive-form game with cardinal payoffs.

Thus we can simplify the game by

1. removing the information set $\{w, z\}$,
2. converting nodes w and z into terminal nodes, and
3. assigning to these newly created terminal nodes the payoffs associated with choice L .

The simplified game is shown in Figure 11.7.




Figure 11.7: The game of Figure 11.6 simplified.

Applying the same reasoning to the simplified game of Figure 11.7, we can delete the two singleton decision nodes of Player 3 and replace the one on the left with the payoff associated with choice F and the one on the right with the payoff associated with choice H , thus obtaining the further simplification shown in Figure 11.8.

In this game there are no more information sets with strictly dominant choices. Thus we have to proceed by trial and error.

Note first that, at any weak sequential equilibrium, by Bayesian updating Player 1's beliefs must be $\begin{pmatrix} r & s & t \\ \frac{1}{5} & \frac{1}{5} & \frac{3}{5} \end{pmatrix}$.

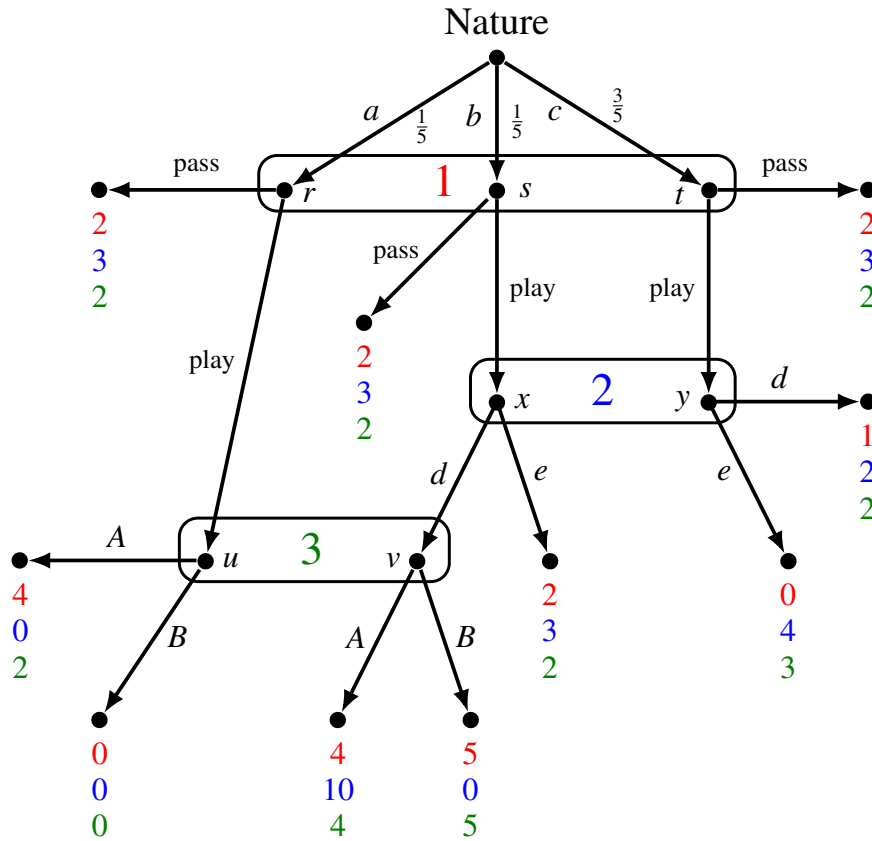


Figure 11.8: The game of Figure 11.7 simplified.

Let us see if in the simplified game of Figure 11.8 there is a pure-strategy weak sequential equilibrium where Player 1's strategy is "play".

When Player 1 chooses "play" then, by Bayesian updating, Player 2's beliefs must be

$$\begin{pmatrix} x & y \\ \frac{1}{4} & \frac{3}{4} \end{pmatrix}.$$

- Let us first try the hypothesis that there is a weak sequential equilibrium where (Player 1's strategy is "play" and) Player 2's strategy is e . Then, by Bayesian updating, Player 3's beliefs must be $\begin{pmatrix} u & v \\ 1 & 0 \end{pmatrix}$, making A the only sequentially rational choice at information set $\{u, v\}$.

However, if Player 3's strategy is A then Player 2, at her information set $\{x, y\}$, gets a payoff of $\frac{1}{4}(10) + \frac{3}{4}(2) = \frac{16}{4}$ from playing d and a payoff of $\frac{1}{4}(3) + \frac{3}{4}(4) = \frac{15}{4}$ from playing e .

Thus e is not sequentially rational and we have reached a contradiction.

- Let us now try the hypothesis that there is a weak sequential equilibrium where (Player 1's strategy is "play" and) Player 2's strategy is d .

Then, by Bayesian updating, Player 3's beliefs must be $\begin{pmatrix} u & v \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$;

hence Player 3's expected payoff from playing A is $\frac{1}{2}(2) + \frac{1}{2}(4) = 3$

and his expected payoff from playing B is $\frac{1}{2}(0) + \frac{1}{2}(5) = 2.5$.

Thus A is the only sequentially rational choice at information set $\{u, v\}$.

Hence, by the previous calculations for Player 2, d is indeed sequentially rational for Player 2.

It only remains to check if Player 1 indeed wants to choose "play".

Given the strategies d and A of Players 2 and 3, respectively, Player 1 gets a payoff of 2 from "pass" and a payoff of $\frac{1}{5}(4) + \frac{1}{5}(4) + \frac{3}{5}(1) = \frac{11}{5} > 2$ from "play".

Thus "play" is indeed sequentially rational.

Thus we have found a pure-strategy weak sequential equilibrium of the game of Figure 11.8, namely

$$\sigma = \left(\begin{array}{cc|cc|cc} \text{pass} & \text{play} & d & e & A & B \\ 0 & 1 & 1 & 0 & 1 & 0 \end{array} \right) \quad \mu = \left(\begin{array}{ccc|cc|cc} r & s & t & x & y & u & v \\ \frac{1}{5} & \frac{1}{5} & \frac{3}{5} & \frac{1}{4} & \frac{3}{4} & \frac{1}{2} & \frac{1}{2} \end{array} \right)$$

This equilibrium can be extended to a weak sequential equilibrium of the original game of Figure 11.6 by adding the choices that led to the simplified game of Figure 11.8 and arbitrary beliefs at information set $\{w, z\}$: for any $p \in [0, 1]$,

$$\sigma = \left(\begin{array}{cc|cc|cc|cc|cc} \text{pass} & \text{play} & d & e & A & B & F & G & H & K & L & R \\ 0 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \end{array} \right)$$

$$\text{and } \mu = \left(\begin{array}{ccc|cc|cc|cc} r & s & t & x & y & u & v & w & z \\ \frac{1}{5} & \frac{1}{5} & \frac{3}{5} & \frac{1}{4} & \frac{3}{4} & \frac{1}{2} & \frac{1}{2} & p & 1-p \end{array} \right),$$

Are there any other pure-strategy weak sequential equilibria? This question is addressed in Exercise 11.7.

We now explore the relationship between the notion of weak sequential equilibrium and other equilibrium concepts. The proof of the following theorem is omitted.

Theorem 11.3.1 Given an extensive-form game with cardinal payoffs G , if (σ, μ) is a weak sequential equilibrium of G then σ is a Nash equilibrium of G .

In general, not every Nash equilibrium can be part of a weak sequential equilibrium. To see this, consider the extensive-form game of Figure 11.9.




Figure 11.9: A game with a Nash equilibrium which is not part of any weak sequential equilibria.

The pure-strategy profile (b, f, d) is a Nash equilibrium, but it cannot be part of any weak sequential equilibrium, because – no matter what beliefs Player 3 has at her information set – choice d is not sequentially rational (it is strictly dominated by c at both nodes in that information set).

Thus weak sequential equilibrium is a strict refinement of Nash equilibrium. Does it also refine the notion of subgame-perfect equilibrium? Unfortunately, the answer is negative: it is possible for (σ, μ) to be a weak sequential equilibrium without σ being a subgame-perfect equilibrium.

To see this, consider the game of Figure 11.10 (which reproduces Figure 11.1).




Figure 11.10: A game with a weak sequential equilibrium whose strategy profile is not a subgame-perfect equilibrium.

The assessment

$$\sigma = (b, (f, e), h) \quad \text{and} \quad \mu = \left(\begin{array}{cc|cc} x & y & w & z \\ 0 & 1 & 1 & 0 \end{array} \right)$$

is a weak sequential equilibrium (the reader should verify this; in particular, note that information set $\{w, z\}$ is not reached and thus Bayesian updating allows for arbitrary beliefs at that information set).

However, $(b, (f, e), h)$ is not a subgame-perfect equilibrium, because the restriction of this strategy profile to the proper subgame that starts at node t of Player 2, namely (e, h) , is not a Nash equilibrium of that subgame: h is not a best reply to e .

The relationship between the three notions of Nash equilibrium, subgame-perfect equilibrium and weak sequential equilibrium is illustrated in the Venn diagram of Figure 11.11.

- Every subgame-perfect equilibrium is a Nash equilibrium. However, there are games in which there is a subgame-perfect equilibrium that is not part of any weak sequential equilibrium. For example, in the game of Figure 11.9, (b, f, d) is a Nash equilibrium which is also subgame-perfect (because there are no proper subgames); however, choice d is strictly dominated and thus is not sequentially rational for Player 3, no matter what beliefs he has at his information set. Thus (b, f, d) cannot be part of a weak sequential equilibrium.
- Every weak sequential equilibrium is a Nash equilibrium. However, as shown in the example of Figure 11.10, there are games in which there is a weak sequential equilibrium whose strategy is not a subgame-perfect equilibrium.

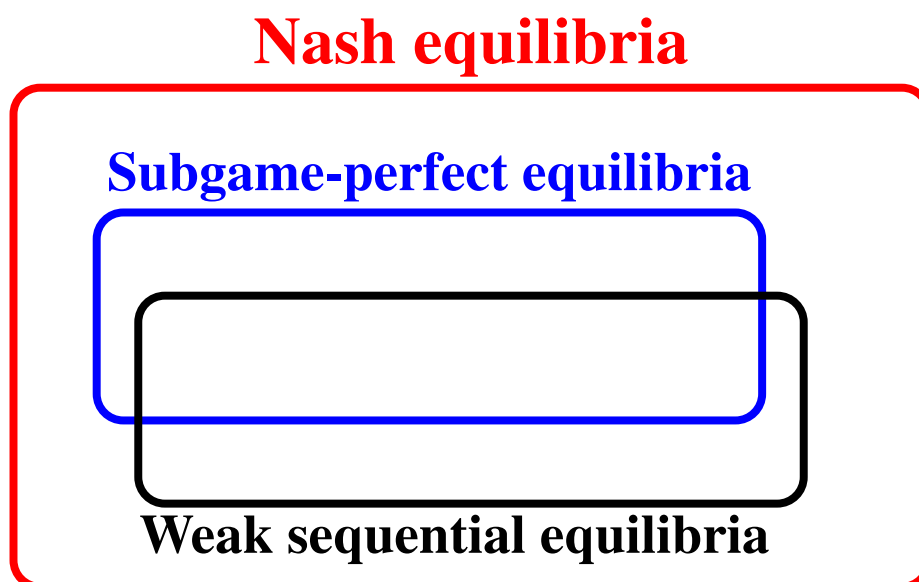


Figure 11.11: The relationship between Nash, subgame-perfect and weak sequential equilibrium.

In the next chapter we define the notion of sequential equilibrium and state two results: (1) a sequential equilibrium is a weak sequential equilibrium and (2) every finite extensive-form game with cardinal payoffs has at least one sequential equilibrium. As a corollary we obtain the following result.

Theorem 11.3.2 Every finite extensive-form game with cardinal payoffs has at least one weak sequential equilibrium (possibly in mixed strategies).

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 11.4.3 at the end of this chapter.

11.4 Exercises

11.4.1 Exercises for Section 11.1: Assessments and sequential rationality

The answers to the following exercises are in Section 11.5 at the end of this chapter.

Exercise 11.1

For the game of Figure 11.12 (which reproduces Figure 11.1), check whether the following assessment is sequentially rational.

$$\sigma = \left(\begin{array}{ccc|cc|cc|cc} a & b & c & f & g & d & e & h & k \\ \frac{1}{8} & \frac{3}{8} & \frac{4}{8} & 1 & 0 & 0 & 1 & 1 & 0 \end{array} \right),$$

$$\mu = \left(\begin{array}{cc|cc} x & y & w & z \\ \frac{1}{3} & \frac{2}{3} & \frac{1}{2} & \frac{1}{2} \end{array} \right)$$




Figure 11.12: Copy of Figure 11.1.

Exercise 11.2

Consider the game of Figure 11.13, obtained from the game of Figure 11.12 by replacing Player 3 with Player 1 at information set $\{w, z\}$

(note that in the game of Figure 11.12 the payoffs of Players 1 and 3 are identical).

Is the following assessment sequentially rational?

$$\sigma = \left(\begin{array}{ccc|cc|cc|cc} a & b & c & f & g & d & e & h & k \\ \frac{1}{8} & \frac{3}{8} & \frac{4}{8} & 1 & 0 & 0 & 1 & 1 & 0 \end{array} \right),$$

$$\mu = \left(\begin{array}{cc|cc} x & y & w & z \\ \frac{1}{3} & \frac{2}{3} & \frac{1}{2} & \frac{1}{2} \end{array} \right)$$




Figure 11.13: A two-player version of the game of Figure 11.12.

Exercise 11.3

Consider the game shown in Figure 11.14. Is the assessment

$$\sigma = \left(\begin{array}{ccc|cc} a & b & r & c & d \\ 0 & 0 & 1 & \frac{2}{5} & \frac{3}{5} \\ e & f & & \frac{1}{3} & \frac{2}{3} \end{array} \right) \quad \text{and} \quad \mu = \left(\begin{array}{cc|cc} s & t & x & y \\ \frac{1}{2} & \frac{1}{2} & \frac{3}{4} & \frac{1}{4} \end{array} \right)$$

sequentially rational? ■




Figure 11.14: The game for Exercise 11.3.

11.4.2 Exercises for Section 11.2:

Bayesian updating at reached information sets

The answers to the following exercises are in Section 11.5 at the end of this chapter.

Exercise 11.4

For the game of Figure 11.12 find a system of beliefs μ such that (σ, μ) satisfies Bayesian updating at reached information sets (Definition 11.2.2), where

$$\sigma = \left(\begin{array}{ccc|cc|cc|cc} a & b & c & f & g & d & e & h & k \\ \frac{1}{8} & \frac{3}{8} & \frac{4}{8} & 1 & 0 & \frac{3}{4} & \frac{1}{4} & \frac{1}{5} & \frac{4}{5} \end{array} \right).$$

■

Exercise 11.5

In the game shown in Figure 11.15, let

$$\sigma = \left(\begin{array}{ccc|cc|cc} a & b & r & c & d & e & f \\ \frac{2}{10} & \frac{1}{10} & \frac{7}{10} & 1 & 0 & \frac{1}{3} & \frac{2}{3} \end{array} \right).$$

Find all the systems of beliefs which, combined with σ yield assessments that satisfy Bayesian updating at reached information sets.




Figure 11.15: The game for Exercise 11.5.

11.4.3 Exercises for Section 11.3: Weak sequential equilibrium

The answers to the following exercises are in Section 11.5 at the end of this chapter.

Exercise 11.6

Find all the pure-strategy weak sequential equilibria of the game shown in Figure 11.16.




Figure 11.16: The game for Exercise 11.6.

Exercise 11.7

In the game of Figure 11.17 (which reproduces Figure 11.8), is there a pure-strategy weak sequential equilibrium where Player 1 chooses “pass”?




Figure 11.17: The game for Exercise 11.7.

Exercise 11.8 — ***Challenging Question***.

Player 1 can take action C or L and Player 2 can take action c or f . The von Neumann-Morgenstern payoffs are as shown in Figure 11.18.

The game, however, is more complex than the strategic form shown in Figure 11.18.

- Player 1 moves first and chooses between C and L .
- He then sends an e-mail to Player 2 telling her truthfully what choice he made.
- However, it is commonly known between the two players that a hacker likes to intercept e-mails and change the text. The hacker is a computer program that, with probability $(1 - \varepsilon)$, leaves the text unchanged and, with probability ε , changes the sentence “I chose C ” into the sentence “I chose L ” and the sentence “I chose L ” into the sentence “I chose C ”. This is commonly known.
- The value of ε is also commonly known. Assume that $\varepsilon \in (0, \frac{1}{4})$.

- (a) Draw the extensive-form game.
- (b) Find all the pure-strategy weak sequential equilibria.
- (c) Are there any weak sequential equilibria (pure or mixed) in which Player 2, when he receives a message from Player 1 saying “I chose L ”, plays f with probability 1?

		Player 2			
		c		f	
Player 1	C	4	4	6	3
	L	3	1	5	2

Figure 11.18: The payoffs for Exercise 11.8.

11.5 Solutions to Exercises

Solutions to Exercise 11.1 The assessment

$$\sigma = \left(\begin{array}{ccc|cc} a & b & c & f & g \\ \frac{1}{8} & \frac{3}{8} & \frac{4}{8} & 1 & 0 \end{array} \middle| \begin{array}{cc|cc} d & e & h & k \\ 0 & 1 & 1 & 0 \end{array} \right), \quad \text{and} \quad \mu = \left(\begin{array}{cc|cc} x & y & w & z \\ \frac{1}{3} & \frac{2}{3} & \frac{1}{2} & \frac{1}{2} \end{array} \right)$$

is sequentially rational. In fact,

- at the root, a gives Player 1 a payoff of 0 and so do b and c (given the strategies of Players 2 and 3). Thus any mixture of a , b and c is sequentially rational; in particular, $\left(\begin{array}{ccc} a & b & c \\ \frac{1}{8} & \frac{3}{8} & \frac{4}{8} \end{array} \right)$;
- at Player 2's node t , given Player 3's strategy, d gives Player 2 a payoff of 1 while e gives a payoff of 2; thus e is sequentially rational;
- given μ at information set $\{x, y\}$, f gives Player 2 a payoff of $\frac{1}{3}(3) + \frac{2}{3}(3) = 3$ while g gives $\frac{1}{3}(6) + \frac{2}{3}(0) = 2$; thus f is sequentially rational;
- given μ at information set $\{w, z\}$, h gives Player 3 a payoff of $\frac{1}{2}(3) + \frac{1}{2}(0) = 1.5$, while k gives 1; thus h is sequentially rational. \square

Solutions to Exercise 11.2 It might seem that the answer is the same as in the previous exercise, because the calculations at the various information sets remain the same; however, in this game checking sequential rationality at the root involves checking whether, given the strategy of Player 2 (namely (f, e)), Player 1 can increase his payoff by changing his entire strategy, that is by changing his choices at both the root and at information set $\{w, z\}$. Indeed, if Player 1 changes his strategy from

$$\left(\begin{array}{ccc|cc} a & b & c & h & k \\ \frac{1}{8} & \frac{3}{8} & \frac{4}{8} & 1 & 0 \end{array} \right) \quad \text{to} \quad \left(\begin{array}{ccc|cc} a & b & c & h & k \\ 0 & 0 & 1 & 0 & 1 \end{array} \right)$$

his payoff increases from 0 to 1.

Hence $\left(\begin{array}{ccc|cc} a & b & c & h & k \\ \frac{1}{8} & \frac{3}{8} & \frac{4}{8} & 1 & 0 \end{array} \right)$ is not sequentially rational at the root. \square

Solutions to Exercise 11.3 The game under consideration is shown in Figure 11.19.




Figure 11.19: The game for Exercise 11.3.

The assessment

$$\sigma = \left(\begin{array}{ccc|cc|cc} a & b & r & c & d & e & f \\ 0 & 0 & 1 & \frac{2}{5} & \frac{3}{5} & \frac{1}{3} & \frac{2}{3} \end{array} \right) \quad \text{and} \quad \mu = \left(\begin{array}{cc|cc} s & t & x & y \\ \frac{1}{2} & \frac{1}{2} & \frac{3}{4} & \frac{1}{4} \end{array} \right)$$

is sequentially rational, as the following calculations show.

- At Player 3's information set $\{x, y\}$, e gives Player 3 a payoff of $\frac{3}{4}(0) + \frac{1}{4}(12) = 3$ and f a payoff of $\frac{3}{4}(4) + \frac{1}{4}(0) = 3$; thus both e and f are optimal and so is any mixture of e and f ; in particular, the mixture $\begin{pmatrix} e & f \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$.
- At Player 2's information set $\{s, t\}$, c gives Player 2 a payoff of $\frac{1}{2}(2) + \frac{1}{2}(0) = 1$ and (given the strategy of Player 3) d a payoff of $\frac{1}{2} \left[\frac{1}{3}(1) + \frac{2}{3}(0) \right] + \frac{1}{2} \left[\frac{1}{3}(0) + \frac{2}{3} \left(\frac{5}{2} \right) \right] = 1$; thus both c and d are optimal and so is any mixture of c and d ; in particular the mixture $\begin{pmatrix} c & d \\ \frac{2}{5} & \frac{3}{5} \end{pmatrix}$.
- At the root, r gives Player 1 a payoff of 2 and (given the strategies of Players 2 and 3) a gives a payoff of $\frac{2}{5}(1) + \frac{3}{5} \left[\frac{1}{3}(2) + \frac{2}{3}(0) \right] = \frac{4}{5}$ and b a payoff of $\frac{2}{5}(0) + \frac{3}{5} \left[\frac{1}{3}(0) + \frac{2}{3}(2) \right] = \frac{4}{5}$. Thus r is sequentially rational. \square

Solutions to Exercise 11.4 The game under consideration is shown in Figure 11.20.




Figure 11.20: The game for Exercise 11.4.

The system of beliefs is $\mu = \left(\begin{array}{cc|cc} x & y & w & z \\ \frac{1}{4} & \frac{3}{4} & \frac{3}{4} & \frac{1}{4} \end{array} \right)$.

In fact, we have that

$$\mathbb{P}_{root,\sigma}(x) = \frac{1}{8}, \quad \mathbb{P}_{root,\sigma}(y) = \frac{3}{8}, \quad \mathbb{P}_{root,\sigma}(\{x,y\}) = \frac{1}{8} + \frac{3}{8} = \frac{4}{8},$$

$$\mathbb{P}_{root,\sigma}(w) = \frac{4}{8} \left(\frac{3}{4} \right) = \frac{3}{8}, \quad \mathbb{P}_{root,\sigma}(z) = \frac{4}{8} \left(\frac{1}{4} \right) = \frac{1}{8} \quad \text{and} \quad \mathbb{P}_{root,\sigma}(\{w,z\}) = \frac{3}{8} + \frac{1}{8} = \frac{4}{8}.$$

Thus

$$\mu(x) = \frac{\mathbb{P}_{root,\sigma}(x)}{\mathbb{P}_{root,\sigma}(\{x,y\})} = \frac{\frac{1}{8}}{\frac{4}{8}} = \frac{1}{4}, \quad \mu(y) = \frac{\mathbb{P}_{root,\sigma}(y)}{\mathbb{P}_{root,\sigma}(\{x,y\})} = \frac{\frac{3}{8}}{\frac{4}{8}} = \frac{3}{4},$$

$$\mu(w) = \frac{\mathbb{P}_{root,\sigma}(w)}{\mathbb{P}_{root,\sigma}(\{w,z\})} = \frac{\frac{3}{8}}{\frac{4}{8}} = \frac{3}{4}, \quad \mu(z) = \frac{\mathbb{P}_{root,\sigma}(z)}{\mathbb{P}_{root,\sigma}(\{w,z\})} = \frac{\frac{1}{8}}{\frac{4}{8}} = \frac{1}{4}$$

□

Solutions to Exercise 11.5 The game under consideration is shown in Figure 11.21.




Figure 11.21: The game for Exercise 11.5.

Let

$$\sigma = \left(\begin{array}{ccc|cc|cc} a & b & r & c & d & e & f \\ \frac{2}{10} & \frac{1}{10} & \frac{7}{10} & 1 & 0 & \frac{1}{3} & \frac{2}{3} \end{array} \right).$$

Since only information set $\{s, t\}$ is reached by σ , no restrictions are imposed on the beliefs at information set $\{x, y\}$.

Thus, for every p such that $0 \leq p \leq 1$, the system of beliefs

$$\mu = \left(\begin{array}{cc|cc} s & t & x & y \\ \frac{2}{3} & \frac{1}{3} & p & 1-p \end{array} \right)$$

combined with σ yields an assessment that satisfies Bayesian updating at reached information sets. \square

Solutions to Exercise 11.6 The game under consideration is shown in Figure 11.22.




Figure 11.22: The game for Exercise 11.6.

We restrict attention to pure strategies. First of all, note that – for Player 1 – A is sequentially rational no matter what strategy Player 2 chooses and, similarly, C is sequentially rational no matter what strategy Player 2 chooses. The strategy of Player 1 determines the beliefs of Player 2 at her information set $\{x, y\}$.

Let us consider the four possibilities (recall that $S_1 = \{(A, C), (B, D), (B, C), (A, D)\}$).

- If Player 1's strategy is (A, C) , then Player 2's information set $\{x, y\}$ is reached with positive probability and the only beliefs that are consistent with Bayesian updating are $\begin{pmatrix} x & y \\ 0 & 1 \end{pmatrix}$, so that both E and F are sequentially rational for Player 2.

By our preliminary observation it follows that $((A, C), E)$ with beliefs $\begin{pmatrix} x & y \\ 0 & 1 \end{pmatrix}$ is a weak sequential equilibrium and so is $((A, C), F)$ with beliefs $\begin{pmatrix} x & y \\ 0 & 1 \end{pmatrix}$.

- If Player 1's strategy is (B, D) , then Player 2's information set $\{x, y\}$ is reached with positive probability and the only beliefs that are consistent with Bayesian updating are $\begin{pmatrix} x & y \\ 1 & 0 \end{pmatrix}$,

so that – by sequential rationality – Player 2 must choose E .

However, if Player 2's strategy is E then at node t it is not sequentially rational for Player 1 to choose D .

Thus there is no pure-strategy weak sequential equilibrium where Player 1's strategy is (B, D) .

- If Player 1's strategy is (B, C) , then Player 2's information set $\{x, y\}$ is reached (with probability 1) and the only beliefs that are consistent with Bayesian updating are $\begin{pmatrix} x & y \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$.

Given these beliefs, E is the only sequentially rational choice for Player 2 (her payoff from playing E is $\frac{1}{2}(2) + \frac{1}{2}(0) = 1$, while her payoff from playing F is 0).

Thus $((B, C), E)$ with beliefs $\begin{pmatrix} x & y \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$ is a weak sequential equilibrium.

- If Player 1's strategy is (A, D) , then Player 2's information set $\{x, y\}$ is not reached and thus, according to the notion of weak sequential equilibrium, any beliefs are allowed there.

In order for D to be sequentially rational for Player 1 it must be that Player 2's pure strategy is F .

In order for F to be sequentially rational for Player 2, her beliefs must be $\begin{pmatrix} x & y \\ 0 & 1 \end{pmatrix}$.

Thus $((A, D), F)$ with beliefs $\begin{pmatrix} x & y \\ 0 & 1 \end{pmatrix}$ is a weak sequential equilibrium.

Summarizing, there are four pure-strategy weak sequential equilibria:

1. $((A, C), E)$ with beliefs $\begin{pmatrix} x & y \\ 0 & 1 \end{pmatrix}$,
2. $((A, C), F)$ with beliefs $\begin{pmatrix} x & y \\ 0 & 1 \end{pmatrix}$,
3. $((B, C), E)$ with beliefs $\begin{pmatrix} x & y \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$,
4. $((A, D), F)$ with beliefs $\begin{pmatrix} x & y \\ 0 & 1 \end{pmatrix}$.

□

Solutions to Exercise 11.7 The game under consideration is shown in Figure 11.23.

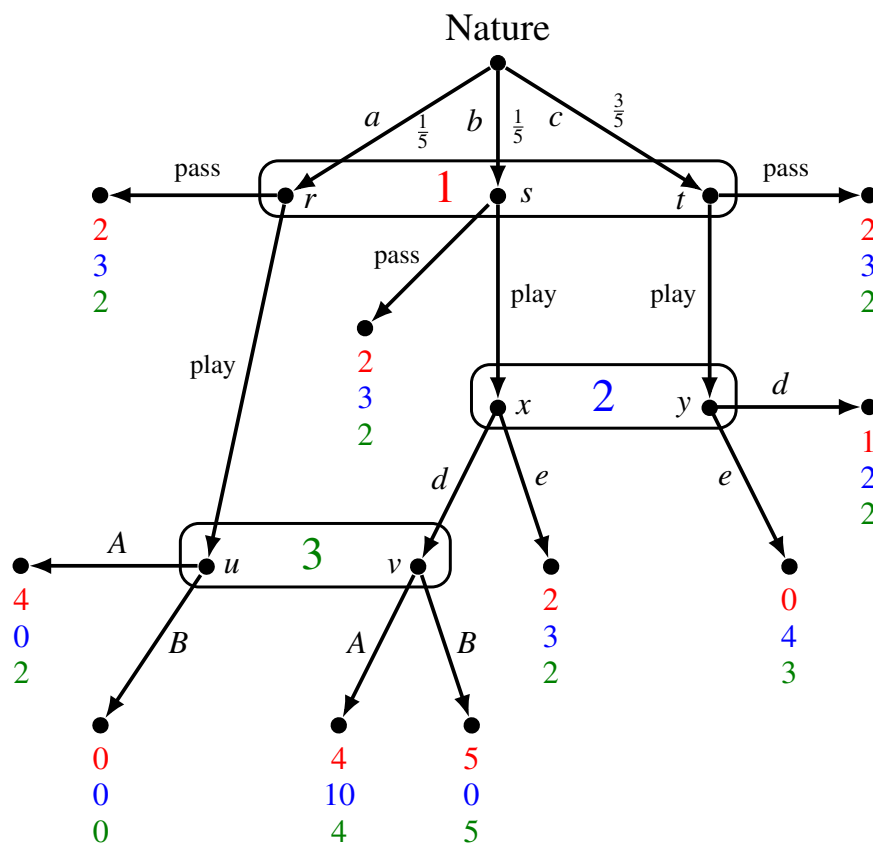


Figure 11.23: The game for Exercise 11.7.

When Player 1’s strategy is “pass” then it is much easier to construct a weak sequential equilibrium, because there are no restrictions on the beliefs at the information sets of Players 2 and 3.

For example, we can choose beliefs $\begin{pmatrix} x & y \\ 0 & 1 \end{pmatrix}$ for Player 2, which make e the only sequentially rational choice, and beliefs $\begin{pmatrix} u & v \\ 0 & 1 \end{pmatrix}$ for Player 3, which make B the only sequentially rational choice.

It only remains to check sequential rationality of “pass”: if Player 1 chooses “pass” he gets a payoff of 2, while if he chooses “play” he gets a payoff of $\frac{1}{5}(0) + \frac{1}{5}(2) + \frac{3}{5}(0) = \frac{2}{5} < 2$, so that “pass” is indeed the better choice.

Thus we have found the following weak sequential equilibrium:

$$\sigma = \left(\begin{array}{cc|cc|cc} \textit{pass} & \textit{play} & \textit{d} & \textit{e} & \textit{A} & \textit{B} \\ 1 & 0 & 0 & 1 & 0 & 1 \end{array} \right) \quad \text{and} \quad \mu = \left(\begin{array}{ccc|cc|cc} r & s & t & x & y & u & v \\ \frac{1}{5} & \frac{1}{5} & \frac{3}{5} & 0 & 1 & 0 & 1 \end{array} \right).$$

(Note, however, that this is just one of several weak sequential equilibria). □

Solutions to Exercise 11.8

(a) The extensive form is shown in Figure 11.24.




Figure 11.24: The game for Exercise 11.8.

(b) Let us consider the pure strategies of Player 1. If Player 1 plays L then Player 2 assigns probability 1 to the bottom node of each information set and responds with f with probability 1, but this makes Player 1 want to deviate to C to get a payoff of 6.⁶ Thus there is no pure-strategy sequential equilibrium where Player 1 plays L .

On the other hand, if Player 1 plays C , then Player 2 assigns probability 1 to the top node of each information set and thus sequential rationality requires her to respond with c at each information set, which makes playing C optimal for Player 1. Thus $(C, (c, c))$, with beliefs that assign probability one to the top node of each information set, is the only pure-strategy weak sequential equilibrium.

⁶In other words, if Player 1 plays L , then Bayesian updating requires Player 2 to assign probability 1 to the bottom node of each information set and then sequential rationality requires Player 2 to play f at each information set, so that $(L, (f, f))$ is the only candidate for a weak sequential equilibrium where Player 1 plays L . But L is not a best reply to (f, f) and thus $(L, (f, f))$ is not a Nash equilibrium and hence cannot be part of a weak sequential equilibrium.

(c) Suppose that Player 2 plays f after reading the message “I chose L ”, that is, at her information set on the right. We know from the argument in part (b) that there are no equilibria of this kind in which Player 1 chooses a pure strategy, so Player 1 must be playing a mixed strategy. For him to be willing to do so, he must receive the same payoff from playing C or L .

If we let p be the probability with which Player 2 plays c if she receives the message “I

$$\underbrace{(1 - \varepsilon)[4p + 6(1 - p)] + \varepsilon 6}_{=\pi_1(C)} = \underbrace{\varepsilon[3p + 5(1 - p)] + (1 - \varepsilon)5}_{=\pi_1(L)}$$

that is, when

$$p = \frac{1}{2 - 4\varepsilon}.$$

Note that, since $\varepsilon \in (0, \frac{1}{4})$, it follows that

$$p \in (\frac{1}{2}, 1).$$

Thus Player 2 randomizes after reading “I chose C ”. For Player 2 to be willing to do this, she must be indifferent between c and f in this event.

- Let $q \in (0, 1)$ be the probability with which Player 1 plays C (so that $1 - q$ is the probability of L);

- then Bayesian updating requires that Player 2 assign probability $\frac{(1-\varepsilon)q}{(1-\varepsilon)q+\varepsilon(1-q)}$ to the top node of her information set on the **left** and probability $\frac{\varepsilon(1-q)}{(1-\varepsilon)q+\varepsilon(1-q)}$ to the bottom node.

- Then, for Player 2, the expected payoff from playing c at the information set on the left is $\pi_2(c) = \frac{(1-\varepsilon)q}{(1-\varepsilon)q+\varepsilon(1-q)}4 + \frac{\varepsilon(1-q)}{(1-\varepsilon)q+\varepsilon(1-q)}1$

and the expected payoff from playing f is $\pi_2(f) = \frac{(1-\varepsilon)q}{(1-\varepsilon)q+\varepsilon(1-q)}3 + \frac{\varepsilon(1-q)}{(1-\varepsilon)q+\varepsilon(1-q)}2$.

- Player 2 is indifferent if these two are equal, that is if

$$4q(1 - \varepsilon) + (1 - q)\varepsilon = 3q(1 - \varepsilon) + 2(1 - q)\varepsilon,$$

which is true if and only if $q = \varepsilon$.

We have now specified behavior at all information sets. To ensure that the specified behavior constitutes an equilibrium, we need to check that f is optimal for Player 2 if she receives the message “I chose L ”.

This will be true if

$$\pi_2(c|I \text{ chose } L) \leq \pi_2(f|I \text{ chose } L)$$

if and only if
since $\varepsilon=q$

$$4q\varepsilon + 1(1 - q)(1 - \varepsilon) \leq 3q\varepsilon + 2(1 - q)(1 - \varepsilon)$$

if and only if

$$4\varepsilon^2 + (1 - \varepsilon)^2 \leq 3\varepsilon^2 + 2(1 - \varepsilon)^2$$

if and only if

$$\varepsilon \leq \frac{1}{2}.$$

Since we have assumed that $\varepsilon < \frac{1}{4}$, Player 2 strictly prefers to play f after receiving the message “I chose L ”. Thus we have constructed the following weak sequential equilibrium:

- Behavior strategy of Player 1:

$$\begin{pmatrix} C & L \\ \varepsilon & 1 - \varepsilon \end{pmatrix}.$$

- Behavior strategy of Player 2: at the information set on the left (where she receives the message “I chose C ”):

$$\begin{pmatrix} c & f \\ \frac{1}{2-4\varepsilon} & \frac{1-4\varepsilon}{2-4\varepsilon} \end{pmatrix}.$$

at the information set on the right (where she receives the message “I chose L ”):

$$\begin{pmatrix} c & f \\ 0 & 1 \end{pmatrix}.$$

- Player 2’s beliefs *at the information set on the left* assign probability

$$\frac{\varepsilon(1 - \varepsilon)}{\varepsilon(1 - \varepsilon) + \varepsilon(1 - \varepsilon)} = \frac{1}{2} \quad \text{to the top node}$$

and probability $\frac{1}{2}$ to the bottom node

and her beliefs *at the information set on the right* assign probability

$$\frac{\varepsilon^2}{\varepsilon^2 + (1 - \varepsilon)^2} \quad \text{to the top node}$$

and probability

$$\frac{(1 - \varepsilon)^2}{\varepsilon^2 + (1 - \varepsilon)^2} \quad \text{to the bottom node.} \quad \square$$

12. Sequential Equilibrium

12.1 Consistent assessments

The stated goal at the beginning of Chapter 11 was to seek a refinement of the notion of subgame-perfect equilibrium that would rule out strictly dominated choices at unreached information sets.

The notion of weak sequential equilibrium achieved the goal of ruling out strictly dominated choices, by means of the requirement of sequential rationality. According to this requirement, a choice at an information set of Player i must be optimal given Player i 's beliefs at that information set; for a strictly dominated choice there can be no beliefs that make it optimal.

However, the notion of weak sequential equilibrium turned out *not* to be a refinement of subgame-perfect equilibrium: as shown in Section 11.3 (Chapter 11), it is possible for the strategy profile in a weak sequential equilibrium not to be a subgame-perfect equilibrium. The reason for this is that the only restriction on beliefs that is incorporated in the notion of weak sequential equilibrium is Bayesian updating at reached information sets. At an information set that is *not* reached by the strategy profile under consideration any beliefs whatsoever are allowed, even if those beliefs are at odds with the strategy profile.

To see this, consider the game of Figure 12.1 and the assessment consisting of the pure-strategy profile $\sigma = (c, d, f)$ (highlighted as double edges) and the system of beliefs that assigns probability 1 to node u : $\begin{pmatrix} s & t & u \\ 0 & 0 & 1 \end{pmatrix}$.




Figure 12.1: A system of beliefs μ that attaches probability 1 to node u is at odds with the strategy profile $\sigma = (c, d, f)$.

Since Player 3's information set $\{s, t, u\}$ is not reached by σ , Bayesian updating imposes no restrictions on beliefs at that information set. However, attaching probability 1 to node u is at odds with σ because in order for node u to be reached the play must have gone through Player 2's node and there, according to σ , Player 2 should have played d with probability 1, making it impossible for node u to be reached.

Thus we need to impose some restrictions on beliefs to ensure that they are consistent with the strategy profile with which they are paired (in the assessment under consideration). At reached information sets this is achieved by requiring Bayesian updating, but so far we have imposed no restriction on beliefs at unreached information sets. We want these restrictions to be “just like Bayesian updating”. Kreps and Wilson (1982) proposed a restriction on beliefs that they called *consistency*, which is stated formally in Definition 12.1.1. To understand the rationale behind this notion, note that if σ is a completely mixed strategy profile (in the sense that $\sigma(a) > 0$, for every choice a) then the issue disappears, because every information set is reached with positive probability and Bayesian updating yields unique beliefs at every information set.

For example, in the game of Figure 12.1, if Player 1 uses the completely mixed strategy

$$\begin{pmatrix} a & b & c \\ p_a & p_b & 1 - p_a - p_b \end{pmatrix} \text{ with } p_a, p_b \in (0, 1) \text{ and } p_a + p_b < 1$$

and Player 2 uses the completely mixed strategy $\begin{pmatrix} d & e \\ p_d & 1 - p_d \end{pmatrix}$ with $p_d \in (0, 1)$

then, by Bayesian updating, Player 3's beliefs must be

$$\mu(s) = \frac{p_a}{p_a + p_b p_d + p_b(1 - p_d)}$$

$$\mu(t) = \frac{p_b p_d}{p_a + p_b p_d + p_b(1 - p_d)}$$

$$\mu(u) = \frac{p_b(1 - p_d)}{p_a + p_b p_d + p_b(1 - p_d)}.$$

In the case of a completely mixed strategy profile σ , it is clear what it means for a system of beliefs μ to be consistent with the strategy profile σ : μ must be the unique system of beliefs obtained from σ by applying Bayesian updating.

What about assessments (σ, μ) where σ is such that some information sets are not reached? How can we decide, in such cases, whether μ is consistent with σ ? Kreps and Wilson (1982) proposed the following criterion: the assessment (σ, μ) is consistent if there is a completely mixed strategy profile σ' which is arbitrarily close to σ and whose associated unique system of beliefs μ' (obtained by applying Bayesian updating) is arbitrarily close to μ . In mathematics "arbitrary closeness" is captured by the notion of limit.

Definition 12.1.1 Given an extensive game, an assessment (σ, μ) is *consistent* if there is a sequence of completely mixed strategy profiles $\langle \sigma_1, \sigma_2, \dots, \sigma_n, \dots \rangle$ such that:

1. the sequence converges to σ as n tends to infinity, that is, $\lim_{n \rightarrow \infty} \sigma_n = \sigma$, and
2. letting μ_n be the unique system of beliefs obtained from σ_n by using Bayesian updating, the sequence $\langle \mu_1, \mu_2, \dots, \mu_n, \dots \rangle$ converges to μ as n tends to infinity, that is, $\lim_{n \rightarrow \infty} \mu_n = \mu$.

For example, consider the extensive form shown in Figure 12.2.




Figure 12.2: The assessment $\sigma = (c, d, f)$ and $\mu(s) = \frac{3}{8}$, $\mu(t) = \frac{5}{8}$ is consistent.

The assessment $\sigma = (c, d, f)$, $\mu = \begin{pmatrix} s & t \\ \frac{3}{8} & \frac{5}{8} \end{pmatrix}$ is consistent. To see this, let

$$\sigma_n = \left(\begin{array}{ccc|cc|cc} a & b & c & d & e & f & g \\ \frac{3}{n} & \frac{5}{n} & 1 - \frac{8}{n} & 1 - \frac{1}{n} & \frac{1}{n} & 1 - \frac{1}{n} & \frac{1}{n} \end{array} \right).$$

Then, as n tends to infinity, all of $\frac{3}{n}$, $\frac{5}{n}$, $\frac{1}{n}$ tend to 0 and both $1 - \frac{8}{n}$ and $1 - \frac{1}{n}$ tend to 1:

$$\lim_{n \rightarrow \infty} \sigma_n = \left(\begin{array}{ccc|cc|cc} a & b & c & d & e & f & g \\ 0 & 0 & 1 & 1 & 0 & 1 & 0 \end{array} \right) = \sigma.$$

Furthermore, $\mu_n(s) = \frac{\frac{3}{n}}{\frac{3}{n} + \frac{5}{n}(1 - \frac{1}{n})} = \frac{3}{8 - \frac{5}{n}}$, which tends to $\frac{3}{8}$ as n tends to infinity and

$\mu_n(t) = \frac{\frac{5}{n}(1 - \frac{1}{n})}{\frac{3}{n} + \frac{5}{n}(1 - \frac{1}{n})} = \frac{5 - \frac{5}{n}}{8 - \frac{5}{n}}$, which tends to $\frac{5}{8}$ as n tends to infinity, so that

$$\lim_{n \rightarrow \infty} \mu_n = \begin{pmatrix} s & t \\ \frac{3}{8} & \frac{5}{8} \end{pmatrix} = \mu.$$

The notion of consistent assessment (σ, μ) was meant to capture an extension of the requirement of Bayesian updating that would apply also to information sets that have zero probability of being reached (when the play of the game is according to the strategy profile σ). However, Definition 12.1.1 is rather technical and not easy to apply. Showing that an assessment is consistent requires displaying an appropriate sequence and showing that the sequence converges to the given assessment. This is relatively easy as compared to the considerably more difficult task of proving that an assessment is *not* consistent: this requires showing that *every possible sequence* that one could construct would not converge.

To see the kind of reasoning involved, consider the extensive form of Figure 12.3.




Figure 12.3: The assessment $\sigma = (a, T, f, L)$, $\mu(r) = \mu(s) = \frac{1}{2}$, $\mu(t) = \frac{1}{5}$, $\mu(u) = \mu(v) = \frac{2}{5}$ is not consistent.

We want to show that the following assessment is *not* consistent:

$$\sigma = \left(\begin{array}{cccc|cc|cc|cc} a & b & c & d & T & B & e & f & L & R \\ 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 1 & 0 \end{array} \right) \quad \text{and} \quad \mu = \left(\begin{array}{cc|cc|cc} r & s & t & u & v \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{5} & \frac{2}{5} & \frac{2}{5} \end{array} \right)$$

Suppose that (σ, μ) is consistent. Then there must be a sequence $\langle \sigma_n \rangle_{n=1,2,\dots}$ of completely mixed strategies that converges to σ , whose corresponding sequence of systems of beliefs $\langle \mu_n \rangle_{n=1,2,\dots}$ (obtained by applying Bayesian updating) converges to μ . Let the n^{th} term of this sequence of completely mixed strategies be:

$$\sigma_n = \left(\begin{array}{cccc|cc|cc|cc} a & b & c & d & T & B & e & f & L & R \\ p_n^a & p_n^b & p_n^c & p_n^d & p_n^T & p_n^B & p_n^e & p_n^f & p_n^L & p_n^R \end{array} \right).$$

Thus $p_n^a + p_n^b + p_n^c + p_n^d = p_n^T + p_n^B = p_n^e + p_n^f = p_n^L + p_n^R = 1$,

$p_n^x \in (0, 1)$ for all $x \in \{a, b, c, d, B, T, e, f, L, R\}$,

$\lim_{n \rightarrow \infty} p_n^x = 0$ for $x \in \{b, c, d, B, e, R\}$

and $\lim_{n \rightarrow \infty} p_n^x = 1$ for $x \in \{a, T, f, L\}$.

From $\sigma_n = \left(\begin{array}{cccc|cc|cc|cc} a & b & c & d & T & B & e & f & L & R \\ p_n^a & p_n^b & p_n^c & p_n^d & p_n^T & p_n^B & p_n^e & p_n^f & p_n^L & p_n^R \end{array} \right)$ we obtain, by Bayesian

updating, the following system of beliefs μ_n :

$$\left(\begin{array}{cc} r & s \\ \frac{p_n^b}{p_n^b+p_n^c} & \frac{p_n^c}{p_n^b+p_n^c} \end{array} \right) \text{ and } \left(\begin{array}{ccc} t & u & v \\ \frac{p_n^b p_n^B}{p_n^b p_n^B + p_n^c p_n^B p_n^f + p_n^d} & \frac{p_n^c p_n^B p_n^f}{p_n^b p_n^B + p_n^c p_n^B p_n^f + p_n^d} & \frac{p_n^d}{p_n^b p_n^B + p_n^c p_n^B p_n^f + p_n^d} \end{array} \right)$$

Note that $\frac{\mu_n(s)}{\mu_n(r)} = \frac{p_n^c}{p_n^b}$.

By hypothesis, $\lim_{n \rightarrow \infty} \mu_n(s) = \mu(s) = \frac{1}{2}$ and $\lim_{n \rightarrow \infty} \mu_n(r) = \mu(r) = \frac{1}{2}$ and thus

$$\lim_{n \rightarrow \infty} \frac{p_n^c}{p_n^b} = \lim_{n \rightarrow \infty} \frac{\mu_n(s)}{\mu_n(r)} = \frac{\lim_{n \rightarrow \infty} \mu_n(s)}{\lim_{n \rightarrow \infty} \mu_n(r)} = \frac{\mu(s)}{\mu(r)} = \frac{\frac{1}{2}}{\frac{1}{2}} = 1.$$

On the other hand, $\frac{\mu_n(u)}{\mu_n(t)} = \frac{p_n^c}{p_n^b} p_n^f$, so that $\lim_{n \rightarrow \infty} \frac{\mu_n(u)}{\mu_n(t)} = \lim_{n \rightarrow \infty} \left(\frac{p_n^c}{p_n^b} p_n^f \right)$.

By hypothesis, $\lim_{n \rightarrow \infty} \mu_n(u) = \mu(u) = \frac{2}{5}$ and $\lim_{n \rightarrow \infty} \mu_n(t) = \mu(t) = \frac{1}{5}$, so that

$$\lim_{n \rightarrow \infty} \frac{\mu_n(u)}{\mu_n(t)} = \frac{\lim_{n \rightarrow \infty} \mu_n(u)}{\lim_{n \rightarrow \infty} \mu_n(t)} = \frac{\mu(u)}{\mu(t)} = 2.$$

However, $\lim_{n \rightarrow \infty} \left(\frac{p_n^c}{p_n^b} p_n^f \right) = \underbrace{\left(\lim_{n \rightarrow \infty} \frac{p_n^c}{p_n^b} \right)}_{=1} \underbrace{\left(\lim_{n \rightarrow \infty} p_n^f \right)}_{=0} = (1)(0) = 0$, yielding a contradiction.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 12.4.1 at the end of this chapter.

12.2 Sequential equilibrium

The notion of sequential equilibrium was introduced by Kreps and Wilson (1982).

Definition 12.2.1 Given an extensive game, an assessment (σ, μ) is a *sequential equilibrium* if it is consistent (Definition 12.1.1) and sequentially rational (Definition 11.1.2, Chapter 11).

For an example of a sequential equilibrium consider the extensive game of Figure 12.4.




Figure 12.4: An extensive-form game with cardinal payoffs.

Let us show that the following assessment is a sequential equilibrium:

$$\sigma = \left(\begin{array}{ccc|cc} L & M & R & A & B \\ 0 & 0 & 1 & \frac{1}{2} & \frac{1}{2} \end{array} \middle| \begin{array}{cc|cc} c & d & c & d \\ 1 & 0 & 3 & 0 \end{array} \right) \quad \text{and} \quad \mu = \left(\begin{array}{cc|ccc} x & y & u & v & w \\ \frac{1}{4} & \frac{3}{4} & \frac{1}{7} & \frac{3}{7} & \frac{3}{7} \end{array} \right).$$

Let us first verify sequential rationality of Player 3's strategy. At her information set – given her beliefs – c gives a payoff of $\frac{1}{7}(5) + \frac{3}{7}(1) + \frac{3}{7}(2) = 2$ and d gives a payoff of $\frac{1}{7}(2) + \frac{3}{7}(2) + \frac{3}{7}(2) = 2$. Thus c is optimal (as would be d and any randomization over c and d).

For Player 2, at his information set $\{x, y\}$ – given his beliefs and given the strategy of Player 3 – A gives a payoff of $\frac{1}{4}(5) + \frac{3}{4}(3) = 3.5$ and B gives a payoff of $\frac{1}{4}(2) + \frac{3}{4}(4) = 3.5$; thus any mixture of A and B is optimal, in particular, the mixture $\left(\begin{array}{cc} A & B \\ \frac{1}{2} & \frac{1}{2} \end{array} \right)$ is optimal.

Finally, at the root, R gives Player 1 a payoff of 3, L a payoff of $\frac{1}{2}(4) + \frac{1}{2}(1) = 2.5$ and M a payoff of $\frac{1}{2}(0) + \frac{1}{2}(6) = 3$; thus R is optimal (as would be any mixture of M and R).

Next we show consistency. Consider the sequence of completely mixed strategies $\langle \sigma_n \rangle_{n=1,2,\dots}$ where

$$\sigma_n = \left(\begin{array}{ccc|cc|cc} L & M & R & A & B & c & d \\ \frac{1}{n} & \frac{3}{n} & 1 - \frac{4}{n} & \frac{1}{2} & \frac{1}{2} & 1 - \frac{1}{n} & \frac{1}{n} \end{array} \right).$$

Clearly $\lim_{n \rightarrow \infty} \sigma_n = \sigma$. The corresponding sequence of systems of beliefs $\langle \mu_n \rangle_{n=1,2,\dots}$ is given by the following constant sequence, which obviously converges to μ :

$$\mu_n = \left(\begin{array}{cc|c} \text{x} & \text{y} & \text{u} \\ \frac{\frac{1}{n}}{\frac{1}{n} + \frac{3}{n}} = \frac{1}{4} & \frac{\frac{3}{n}}{\frac{1}{n} + \frac{3}{n}} = \frac{3}{4} & \frac{\frac{1}{n}(\frac{1}{2})}{\frac{1}{n}(\frac{1}{2}) + \frac{3}{n}(\frac{1}{2}) + \frac{3}{n}(\frac{1}{2})} = \frac{1}{7} \\ & & \frac{\frac{3}{n}(\frac{1}{2})}{\frac{1}{n}(\frac{1}{2}) + \frac{3}{n}(\frac{1}{2}) + \frac{3}{n}(\frac{1}{2})} = \frac{3}{7} \quad \frac{3}{7} \\ & & \text{v} & \text{w} \end{array} \right).$$

- R** Since consistency implies Bayesian updating at reached information sets, every sequential equilibrium is a weak sequential equilibrium.

We now turn to the properties of sequential equilibria.

Theorem 12.2.1 — Kreps and Wilson, 1982. Given an extensive game, if (σ, μ) is a sequential equilibrium then σ is a subgame-perfect equilibrium.

Theorem 12.2.2 — Kreps and Wilson, 1982. Every finite extensive-form game with cardinal payoffs has at least one sequential equilibrium.

By Theorem 12.1, the notion of sequential equilibrium achieves the objective of refining the notion of subgame-perfect equilibrium. The relationship between the various solution concepts considered so far is shown in the Venn diagram of Figure 12.5.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 12.4.2 at the end of this chapter.

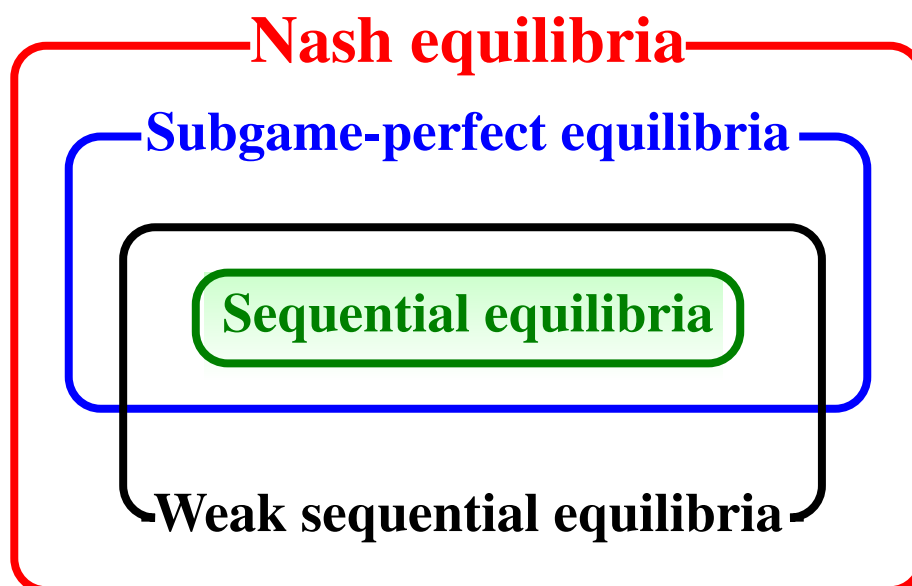


Figure 12.5: The relationship between Nash equilibrium, subgame-perfect equilibrium, weak sequential equilibrium and sequential equilibrium.

12.3 Is 'consistency' a good notion?

The notion of consistency (Definition 12.1) is unsatisfactory in two respects.

- ◇ From a practical point of view, consistency is computationally hard to prove, since one has to construct a sequence of completely mixed strategies, calculate the corresponding Bayesian beliefs and take the limit of the two sequences. The larger and more complex the game, the harder it is to establish consistency.
- ◇ From a conceptual point of view, it is not clear how one should interpret, or justify, the requirement of taking the limit of sequences of strategies and beliefs.

Concerning the latter point, Kreps and Wilson themselves express dissatisfaction with their definition of sequential equilibrium:

“We shall proceed here to develop the properties of sequential equilibrium as defined above; however, we do so with some doubts of our own concerning what 'ought' to be the definition of a consistent assessment that, with sequential rationality, will give the 'proper' definition of a sequential equilibrium.” (Kreps and Wilson, 1982, p. 876.)

In a similar vein, Osborne and Rubinstein (1994, p. 225) write

“We do not find the consistency requirement to be natural, since it is stated in terms of limits; it appears to be a rather opaque technical assumption.”

In the next chapter we will introduce a simpler refinement of subgame-perfect equilibrium which has a clear interpretation in terms of the AGM theory of belief revision: we call it *Perfect Bayesian Equilibrium*. We will also show that one can use this notion to provide a characterization of sequential equilibrium that does not require the use of limits of sequences of completely mixed strategies.

We conclude this chapter by observing that while the notion of sequential equilibrium eliminates strictly dominated choices at information sets (even if they are reached with zero probability), it is not strong enough to eliminate *weakly* dominated choices. To see this, consider the “simultaneous” game shown in Figure 12.6.




Figure 12.6: A game where a weakly dominated strategy can be part of a sequential equilibrium.

In this game there are two Nash (and subgame-perfect) equilibria: (A, D) and (B, C) . Note that C is a weakly dominated strategy for Player 2. The only beliefs of Player 2 that rationalize choosing C is that Player 1 chose B with probability 1 (if Player 2 attaches any positive probability, no matter how small, to Player 1 choosing A , then D is the only sequentially rational choice). Nevertheless, both Nash equilibria are sequential equilibria.

For example, it is straightforward to check that the “unreasonable” Nash equilibrium

$\sigma = (B, C)$, when paired with beliefs $\mu = \begin{pmatrix} x & y \\ 0 & 1 \end{pmatrix}$, constitutes a sequential equilibrium.

Consistency of this assessment is easily verified by considering the sequence $\sigma_n = \left(\begin{array}{cc|cc} A & B & C & D \\ \frac{1}{n} & 1 - \frac{1}{n} & 1 - \frac{1}{n} & \frac{1}{n} \end{array} \right)$ whose associated beliefs are $\mu_n = \begin{pmatrix} x & y \\ \frac{1}{n} & 1 - \frac{1}{n} \end{pmatrix}$ and sequential rationality is clearly satisfied.

Many game theorists feel that it is “irrational” to choose a weakly dominated strategy; thus further refinements beyond sequential equilibrium have been proposed. A stronger notion of equilibrium, which is a strict refinement of sequential equilibrium, is the notion of *trembling-hand perfect equilibrium*. This notion, due to Reinhardt Selten (who also introduced the notion of subgame-perfect equilibrium) precedes chronologically the notion of sequential equilibrium (Selten, 1975). *Trembling-hand perfect equilibrium* does in fact eliminate weakly dominated strategies. This topic is outside the scope of this book.¹

12.4 Exercises

12.4.1 Exercises for Section 12.1: Consistent assessments

The answers to the following exercises are in Section 12.5 at the end of this chapter.

Exercise 12.1

Consider the extensive form shown in Figure 12.7.

Consider the following (partial) behavior strategy profile $\sigma = \left(\begin{array}{ccc|cc} a & b & c & d & e \\ \frac{1}{5} & \frac{3}{5} & \frac{1}{5} & \frac{1}{4} & \frac{3}{4} \end{array} \right)$.

Find the corresponding system of beliefs obtained by Bayesian updating. ■




Figure 12.7: The extensive form for Exercise 12.1.

¹The interested reader is referred to van Damme (2002).

Exercise 12.2

Consider the extensive game shown in Figure 12.8.

- Write the corresponding strategic form.
- Find all the pure-strategy Nash equilibria.
- Find all the pure-strategy subgame-perfect equilibria.
- Which of the pure-strategy subgame-perfect equilibria can be part of a consistent assessment? Give a proof for each of your claims.




Figure 12.8: The extensive form for Exercise 12.2.

12.4.2 Exercises for Section 12.2: Sequential equilibrium

Exercise 12.3

Consider the extensive-form game shown in Figure 12.9. For each pure-strategy Nash equilibrium determine whether it is part of an assessment which is a sequential equilibrium.

Exercise 12.4

Consider the game shown in Figure 12.10.

- Find three subgame-perfect equilibria. [Use pure strategies wherever possible.]
- For each of the equilibria you found in Part (a), explain if it can be written as part of a weak sequential equilibrium.
- Find a sequential equilibrium. [Use pure strategies wherever possible.]




Figure 12.9: The extensive-form game for Exercise 12.3.




Figure 12.10: The extensive-form game for Exercise 12.4

Exercise 12.5

An electric circuit connects two switches and a light. One switch is in Room 1, the second switch is in Room 2 and the light is in Room 3.

Player 1 is in Room 1, Player 2 in Room 2 and Player 3 in Room 3.

The two switches are now in the Off position. The light in Room 3 comes on if and only if **both** switches are in the On position. Players 1 and 2 act simultaneously and independently: each is allowed only two choices, namely whether to leave her switch in the Off position or turn it to the On position.

If the light comes on in Room 3 then the game ends and Players 1 and 2 get \$100 each while Player 3 gets \$300.

If the light in Room 3 stays off, then Player 3 (not knowing what the other players did) has to make a guess as to what Players 1 and 2 did (thus, for example, one possible guess is “both players left their respective switches in the Off position”).

Then the payoffs are as follows:

- if Player 3’s guess turns out to be correct then each player gets \$200,
- if Player 3 makes one correct guess but the other wrong (e.g. he guesses that both Player 1 and Player 2 chose “Off” and, as a matter of fact, Player 1 chose “Off” while Player 2 chose “On”), then Player 3 gets \$50, the player whose action was guessed correctly gets \$100 and the remaining player gets nothing (in the previous example, Player 1 gets \$100, Player 2 gets nothing and Player 3 gets \$50) and
- if Player 3’s guess is entirely wrong then all the players get nothing.

All the players are selfish and greedy (that is, each player only cares about how much money he/she gets and prefers more money to less) and risk neutral.

- (a) Represent this situation as an extensive-form game.
- (b) Write the corresponding strategic form, assigning the rows to Player 1, the columns to Player 2, etc.
- (c) Find all the pure-strategy Nash equilibria.
- (d) For at least one pure-strategy Nash equilibrium prove that it cannot be part of a sequential equilibrium.



Exercise 12.6 — ***Challenging Question***.

A buyer and a seller are bargaining over an object owned by the seller. The value of the object to the buyer is known to her but not to the seller. The buyer is drawn randomly from a population with the following characteristics: the fraction λ value the object at $\$H$ while the remaining fraction value the object at $\$L$, with $H > L > 0$.

The bargaining takes place over two periods. In the first period the seller makes a take-it-or-leave-it offer (i.e. names the price) and the buyer accepts or rejects. If the buyer accepts, the transaction takes place and the game ends. If the buyer rejects, then the seller makes a new take-it-or-leave-it offer and the buyer accepts or rejects. In either case the game ends.

Payoffs are as follows:

- (1) if the seller's offer is accepted (whether it was made in the first period or in the second period), the seller's payoff is equal to the price agreed upon and the buyer's payoff is equal to the difference between the value of the object to the buyer and the price paid;
- (2) if the second offer is rejected both players get a payoff of 0.

Assume that both players discount period 2 payoffs with a discount factor $\delta \in (0, 1)$, that is, from the point of view of period 1, getting $\$x$ in period 2 is considered to be the same as getting δx in period 1. For example, if the seller offers price p in the first period and the offer is accepted, then the seller's payoff is p , whereas if the same price p is offered and accepted in the second period, then the seller's payoff, viewed from the standpoint of period 1, is δp .

Assume that these payoffs are von Neumann-Morgenstern payoffs; assume further that $H = 20$, $L = 10$, $\delta = \frac{3}{4}$ and $\lambda = \frac{2}{3}$.

- (a) Draw the extensive form of this game for the case where, in both periods, the seller can only offer one of two prices: $\$10$ or $\$12$. Nature moves first and selects the value for the buyer; the buyer is informed, while the seller is not. It is common knowledge between buyer and seller that Nature will pick H with probability λ and L with probability $(1-\lambda)$.
- (b) For the game of part (a) find a pure-strategy sequential equilibrium. Prove that what you suggest is indeed a sequential equilibrium.

12.5 Solutions to Exercises

Solutions to Exercise 12.1 The extensive form under consideration is shown in Figure 12.11.




Figure 12.11: The extensive form for Exercise 12.1.

The system of beliefs obtained, by Bayesian updating, from the (partial) behavior strategy profile $\sigma = \left(\begin{array}{ccc|cc} a & b & c & d & e \\ \frac{1}{5} & \frac{3}{5} & \frac{1}{5} & \frac{1}{4} & \frac{3}{4} \end{array} \right)$ is as follows:

$$\begin{aligned} \mu(r) &= \frac{\frac{1}{5}}{\frac{1}{5} + \frac{3}{5}} = \frac{1}{4}, & \mu(s) &= \frac{\frac{3}{5}}{\frac{1}{5} + \frac{3}{5}} = \frac{3}{4}, & \mu(x) &= \frac{\frac{3}{5}(\frac{3}{4})}{\frac{3}{5}(\frac{3}{4}) + \frac{1}{5}} = \frac{9}{13}, & \mu(y) &= \frac{\frac{1}{5}}{\frac{3}{5}(\frac{3}{4}) + \frac{1}{5}} = \frac{4}{13} \\ \mu(t) &= \frac{\frac{1}{5}(\frac{1}{4})}{\frac{1}{5}(\frac{1}{4}) + \frac{1}{5}(\frac{3}{4}) + \frac{3}{5}(\frac{1}{4})} = \frac{1}{7} & \mu(u) &= \frac{\frac{1}{5}(\frac{3}{4})}{\frac{1}{5}(\frac{1}{4}) + \frac{1}{5}(\frac{3}{4}) + \frac{3}{5}(\frac{1}{4})} = \frac{3}{7} \\ \mu(v) &= \frac{\frac{3}{5}(\frac{1}{4})}{\frac{1}{5}(\frac{1}{4}) + \frac{1}{5}(\frac{3}{4}) + \frac{3}{5}(\frac{1}{4})} = \frac{3}{7}. \end{aligned} \quad \square$$

Solutions to Exercise 12.2

- (a) The game under consideration is shown in Figure 12.12. The corresponding strategic form is shown in Figure 12.13.
- (b) The pure-strategy Nash equilibria are: $(L, bl), (R, tl), (R, tr), (E, tr), (E, br)$.
- (c) There is only one proper subgame, namely the one that starts at node x_1 ; in that subgame the unique Nash equilibrium is t .
Thus only the following are subgame-perfect equilibria: $(R, tl), (R, tr), (E, tr)$.
- (d) Each of the above subgame-perfect equilibria can be part of a consistent assessment.

Let $\sigma = (R, tl), \mu = \begin{pmatrix} x_2 & x_3 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$; this assessment is consistent as the following sequences of completely mixed strategies and corresponding system of beliefs show:

$$\begin{aligned} \sigma_n &= \left(\begin{array}{ccc|cc} L & W & E & R & b & t & l & r \\ \frac{1}{n} & \frac{1}{n} & \frac{1}{n} & 1 - \frac{3}{n} & \frac{1}{n} & 1 - \frac{1}{n} & 1 - \frac{1}{n} & \frac{1}{n} \end{array} \right), \\ \mu_n &= \left(\begin{array}{cc} x_2 & x_3 \\ \frac{\frac{1}{n}}{\frac{1}{n} + \frac{1}{n}} = \frac{1}{2} & \frac{\frac{1}{n}}{\frac{1}{n} + \frac{1}{n}} = \frac{1}{2} \end{array} \right). \end{aligned} \quad \text{Clearly } \lim_{n \rightarrow \infty} \sigma_n = \sigma \text{ and } \lim_{n \rightarrow \infty} \mu_n = \mu.$$

The proof for (R, tr) is similar.

Let $\sigma = (E, tr)$, $\mu = \begin{pmatrix} x_2 & x_3 \\ 0 & 1 \end{pmatrix}$; this assessment is consistent as the following sequences of completely mixed strategies and corresponding beliefs show:

$$\sigma_n = \left(\begin{array}{cccc|cc|cc} L & W & E & R & b & t & l & r \\ \frac{1}{n} & \frac{1}{n} & 1 - \frac{3}{n} & \frac{1}{n} & \frac{1}{n} & 1 - \frac{1}{n} & \frac{1}{n} & 1 - \frac{1}{n} \end{array} \right), \mu_n = \begin{pmatrix} x_2 & x_3 \\ \frac{1}{n} & 1 - \frac{3}{n} \\ \frac{1}{n+1-\frac{3}{n}} & \frac{1}{n+1-\frac{3}{n}} \end{pmatrix}.$$

Clearly $\lim_{n \rightarrow \infty} \sigma_n = \sigma$ and $\lim_{n \rightarrow \infty} \mu_n = \mu$. □




Figure 12.12: The extensive form for Exercise 12.2.

		Player 2			
		<i>tl</i>	<i>tr</i>	<i>bl</i>	<i>br</i>
Player 1	<i>L</i>	0 10	0 10	0 10	0 10
	<i>R</i>	1 1	1 1	-1 -1	-1 -1
	<i>W</i>	-1 0	0 1	-1 0	0 1
	<i>E</i>	-1 1	1 2	-1 1	1 2

Figure 12.13: The strategic-form of the game of Figure 12.12.

Solutions to Exercise 12.3 The game under consideration is shown in Figure 12.14.




Figure 12.14: The extensive form for Exercise 12.3.

(a) The corresponding strategic form is shown in Figure 12.15.

		Player 2					
		<i>l</i>			<i>r</i>		
Player 1	<i>R</i>	1	0	0	1	0	0
	<i>M</i>	0	1	1	2	2	1
	<i>L</i>	0	1	0	0	0	0

Player 3: *a*

		Player 2					
		<i>l</i>			<i>r</i>		
Player 1	<i>R</i>	1	0	0	1	0	0
	<i>M</i>	0	1	0	0	0	0
	<i>L</i>	0	1	0	0	0	0

Player 3: *b*

Figure 12.15: The strategic-form of the game of Figure 12.14.

(b) The pure-strategy Nash equilibria are: (R, l, a) , (M, r, a) , (R, l, b) and (R, r, b) . They are all subgame perfect because there are no proper subgames. (R, l, b) and (R, r, b) cannot be part of a sequential equilibrium because b is a strictly dominated choice at Player 3's information set and, therefore, (R, l, b) and (R, r, b) would violate sequential rationality (with any system of beliefs). On the other hand, both (R, l, a) and (M, r, a) can be part of an assessment which is a sequential equilibrium.

First we show that $\sigma = (R, l, a)$ together with the system of beliefs

$\mu = \left(\begin{array}{cc|cc} x_1 & x_2 & x_3 & x_4 \\ \frac{2}{3} & \frac{1}{3} & 1 & 0 \end{array} \right)$ is a sequential equilibrium.

Consider the sequence of completely mixed strategy profiles whose n^{th} term is

$\sigma_n = \left(\begin{array}{ccc|cc} L & M & R & l & r \\ \frac{2}{n} & \frac{1}{n} & 1 - \frac{3}{n} & 1 - \frac{1}{n} & \frac{1}{n} \\ & & & 1 - \frac{1}{n} & \frac{1}{n} \end{array} \right)$. Clearly $\lim_{n \rightarrow \infty} \sigma_n = \sigma$.

The corresponding Bayesian system of beliefs has n^{th} term

$$\mu_n = \left(\begin{array}{cc|cc} x_1 & x_2 & x_3 & x_4 \\ \frac{\frac{2}{n}}{\frac{2}{n} + \frac{1}{n}} = \frac{2}{3} & \frac{\frac{1}{n}}{\frac{2}{n} + \frac{1}{n}} = \frac{1}{3} & \frac{\frac{1}{n}(1 - \frac{1}{n})}{\frac{1}{n}(1 - \frac{1}{n}) + \frac{1}{n}(\frac{1}{n})} = 1 - \frac{1}{n} & \frac{\frac{1}{n}(\frac{1}{n})}{\frac{1}{n}(1 - \frac{1}{n}) + \frac{1}{n}(\frac{1}{n})} = \frac{1}{n} \end{array} \right)$$

Clearly $\lim_{n \rightarrow \infty} \mu_n = \mu$. Thus the assessment is consistent (Definition 12.1.1).

Sequential rationality is easily checked: given the strategy profile and the system of beliefs,

- (1) for Player 3, a yields 1, while b yields 0,
- (2) for Player 2, l yields 1, while r yields $\frac{2}{3}(0) + \frac{1}{3}(2) = \frac{2}{3}$,
- (3) for Player 1, R yields 1, while L and M yield 0.

Next we show that $\sigma = (M, r, a)$ together with $\mu = \left(\begin{array}{cc|cc} x_1 & x_2 & x_3 & x_4 \\ 0 & 1 & 0 & 1 \end{array} \right)$ is a sequential equilibrium.

Consider the sequence of completely mixed strategy profiles whose n^{th} term is

$$\sigma_n = \left(\begin{array}{ccc|cc} L & M & R & l & r \\ \frac{1}{2n} & 1 - \frac{1}{n} & \frac{1}{2n} & \frac{1}{n} & 1 - \frac{1}{n} \\ & & & 1 - \frac{1}{n} & \frac{1}{n} \end{array} \right)$$

Clearly $\lim_{n \rightarrow \infty} \sigma_n = \sigma$. The corresponding Bayesian system of beliefs has n^{th} term

$$\mu_n = \left(\begin{array}{cc|cc} x_1 & x_2 & x_3 & x_4 \\ \frac{\frac{1}{2n}}{\frac{1}{2n} + (1 - \frac{1}{n})} = \frac{1}{2n-1} & 1 - \frac{1}{2n-1} & \frac{\frac{1}{n}(1 - \frac{1}{n})}{\frac{1}{n}(1 - \frac{1}{n}) + (1 - \frac{1}{n})^2} = \frac{1}{n} & 1 - \frac{1}{n} \end{array} \right)$$

Clearly $\lim_{n \rightarrow \infty} \mu_n = \mu$. Thus the assessment is consistent (Definition 12.1.1). Sequential rationality is easily checked: given the strategy profile and the system of beliefs,

- (1) for Player 3, a yields 1, while b yields 0,
- (2) for Player 2, l yields 1, while r yields 2 and
- (3) for Player 1, M yields 2, while R yields 1 and L yields 0.

□

Solutions to Exercise 12.4 The game under consideration is shown in Figure 12.16.




Figure 12.16: The extensive-form game for Exercise 12.4.

- (a) First we solve the subgame on the left, whose strategic form is shown in Figure 12.17.

		Player 2	
		C	D
Player 1	A	1 0	0 1
	B	0 2	2 0

Figure 12.17: The strategic-form of the subgame that starts after choice L of Player 1.

There is no pure-strategy Nash equilibrium. Let p be the probability of A and q the probability of C.

Then at a Nash equilibrium it must be that $q = 2(1 - q)$ and $2(1 - p) = p$.

Thus there is a unique Nash equilibrium given by $\left(\begin{array}{cc|cc} A & B & C & D \\ \frac{2}{3} & \frac{1}{3} & \frac{2}{3} & \frac{1}{3} \end{array} \right)$ with an expected payoff of $\frac{2}{3}$ for both players.

Next consider the subgame on the right. In this subgame the following are pure-strategy Nash equilibria: (F, H, M) (where Player 1's payoff is 1), (E, L, H) (where Player 1's payoff is 2), and (E, L, G) (where Player 1's payoff is 2).

Thus the following are subgame-perfect equilibria:

$$\left(\begin{array}{cc|cc|cc|cc|cc|cc} L & R & A & B & C & D & E & F & G & H & L & M \\ 0 & 1 & \frac{2}{3} & \frac{1}{3} & \frac{2}{3} & \frac{1}{3} & 0 & 1 & 0 & 1 & 0 & 1 \end{array} \right),$$

$$\left(\begin{array}{cc|cc|cc|cc|cc|cc} L & R & A & B & C & D & E & F & G & H & L & M \\ 0 & 1 & \frac{2}{3} & \frac{1}{3} & \frac{2}{3} & \frac{1}{3} & 1 & 0 & 0 & 1 & 1 & 0 \end{array} \right),$$

$$\left(\begin{array}{cc|cc|cc|cc|cc|cc} L & R & A & B & C & D & E & F & G & H & L & M \\ 0 & 1 & \frac{2}{3} & \frac{1}{3} & \frac{2}{3} & \frac{1}{3} & 1 & 0 & 1 & 0 & 1 & 0 \end{array} \right).$$

- (b) $\left(\begin{array}{cc|cc|cc|cc|cc|cc} L & R & A & B & C & D & E & F & G & H & L & M \\ 0 & 1 & \frac{2}{3} & \frac{1}{3} & \frac{2}{3} & \frac{1}{3} & 0 & 1 & 0 & 1 & 0 & 1 \end{array} \right)$ cannot be part of a weak sequential equilibrium, because choice M is strictly dominated by L and thus there are no beliefs at Player 3's information set that justify choosing M .

$\left(\begin{array}{cc|cc|cc|cc|cc|cc} L & R & A & B & C & D & E & F & G & H & L & M \\ 0 & 1 & \frac{2}{3} & \frac{1}{3} & \frac{2}{3} & \frac{1}{3} & 1 & 0 & 0 & 1 & 1 & 0 \end{array} \right)$ cannot be part of a weak sequential equilibrium, because (given that Player 3 chooses L) H is not a sequentially rational choice for Player 2 at his singleton node.

$\left(\begin{array}{cc|cc|cc|cc|cc|cc} L & R & A & B & C & D & E & F & G & H & L & M \\ 0 & 1 & \frac{2}{3} & \frac{1}{3} & \frac{2}{3} & \frac{1}{3} & 1 & 0 & 1 & 0 & 1 & 0 \end{array} \right)$ is a weak sequential equilibrium with the system of beliefs $\left(\begin{array}{cc|cc} x & y & w & z \\ \frac{2}{3} & \frac{1}{3} & 1 & 0 \end{array} \right)$.

- (c) $\left(\begin{array}{cc|cc|cc|cc|cc|cc} L & R & A & B & C & D & E & F & G & H & L & M \\ 0 & 1 & \frac{2}{3} & \frac{1}{3} & \frac{2}{3} & \frac{1}{3} & 1 & 0 & 1 & 0 & 1 & 0 \end{array} \right)$ together with beliefs $\mu = \left(\begin{array}{cc|cc} x & y & w & z \\ \frac{2}{3} & \frac{2}{3} & 1 & 0 \end{array} \right)$ is a sequential equilibrium.

Consistency can be verified with the sequence of completely mixed strategies

$$\sigma_n = \left(\begin{array}{cc|cc|cc|cc|cc|cc} L & R & A & B & C & D & E & F & G & H & L & M \\ \frac{1}{n} & 1 - \frac{1}{n} & \frac{2}{3} & \frac{1}{3} & \frac{2}{3} & \frac{1}{3} & 1 - \frac{1}{n} & \frac{1}{n} & 1 - \frac{1}{n} & \frac{1}{n} & 1 - \frac{1}{n} & \frac{1}{n} \end{array} \right).$$

For example,

$$\mu_n(w) = \frac{(1 - \frac{1}{n})(1 - \frac{1}{n})}{(1 - \frac{1}{n})(1 - \frac{1}{n}) + (1 - \frac{1}{n})\frac{1}{n}(1 - \frac{1}{n})} = \frac{1}{1 + \frac{1}{n}} \rightarrow 1 \text{ as } n \rightarrow \infty.$$

Sequential rationality is easily verified. \square

Solutions to Exercise 12.5

(a) The extensive form is shown in Figure 12.18, where FF means 1Off-2Off, FN means 1Off-2On, etc.




Figure 12.18: The extensive form for Exercise 12.5.

(b) The corresponding strategic form is shown in Figure 12.19.

		Player 2					
		On			Off		
Player 1	On	100	100	300	200	200	200
	Off	0	0	0	0	100	50

Player 3: 1-On/2-Off

		Player 2					
		On			Off		
Player 1	On	100	100	300	0	0	0
	Off	200	200	200	100	0	50

Player 3: 1-Off/2-On

		Player 2					
		On			Off		
Player 1	On	100	100	300	0	100	100
	Off	100	0	50	200	200	200

Player 3: both Off

Figure 12.19: The strategic form of the game of Figure 12.18.

- (c) The Nash equilibria are highlighted in the strategic form: (On, Off, 1On-2Off), (Off, On, 1Off-2On), (On, On, both-Off) and (Off, Off, both-Off).
- (d) (On, On, both-Off) cannot be part of a sequential equilibrium. First of all, for Player 3 'both-Off' is a sequentially rational choice only if Player 3 attaches (sufficiently high) positive probability to node x .

However, consistency does not allow beliefs with $\mu(x) > 0$. To see this, consider a sequence of completely mixed strategies $\{p_n, q_n\}$ for Players 1 and 2, where p_n is the probability with which Player 1 chooses Off and q_n is the probability with which Player 2 chooses Off and $\lim_{n \rightarrow \infty} p_n = \lim_{n \rightarrow \infty} q_n = 0$.

Then, by Bayesian updating,

$$P_n(x | \{x, y, z\}) = \frac{p_n q_n}{p_n q_n + p_n(1 - q_n) + q_n(1 - p_n)}. \quad (\diamond)$$

If q_n goes to 0 as fast as, or faster than, p_n (that is, if

$$\lim_{n \rightarrow \infty} \frac{q_n}{p_n}$$

is finite), then divide numerator and denominator of (\diamond) by p_n to get

$$P_n(x | \{x, y, z\}) = \frac{q_n}{q_n + (1 - q_n) + \frac{q_n}{p_n}(1 - p_n)}$$

Taking the limit as $n \rightarrow \infty$ we get

$$\frac{0}{0 + 1 + \left(\lim_{n \rightarrow \infty} \frac{q_n}{p_n}\right)(1)} = 0.$$

[If p_n goes to 0 as fast as or faster than q_n then repeat the above argument by dividing by q_n .] Thus a consistent assessment must assign zero probability to node x . □

Solutions to Exercise 12.6

(a) The extensive-form game is shown in Figure 12.20.

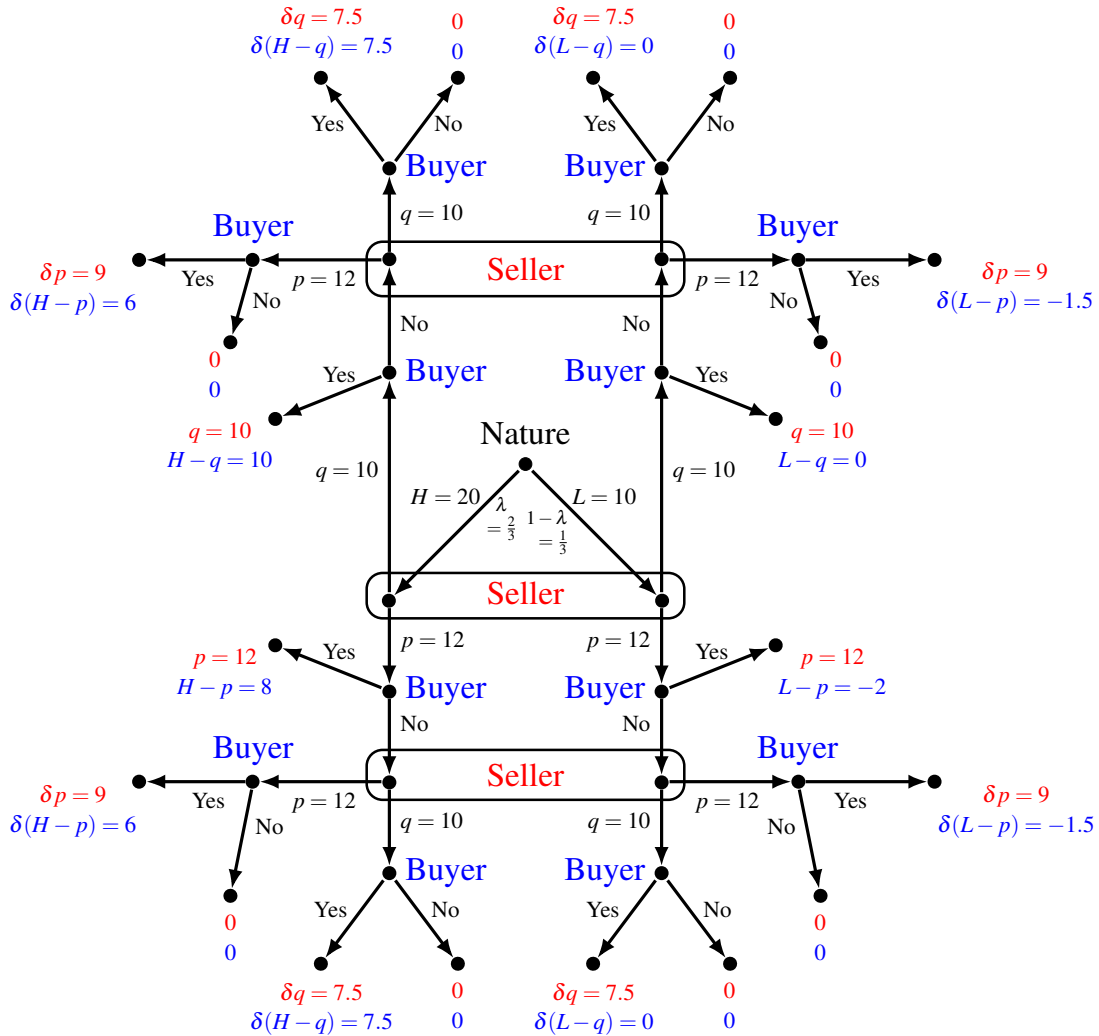


Figure 12.20: The buyer-seller game of Exercise 12.6.

(b) The following assessment (σ, μ) is a sequential equilibrium. The pure-strategy profile σ is as follows:

The seller offers $p = 12$ in period 1 and, if his offer is rejected, adjusts his offer to $q = 10$ in period 2; furthermore, if the first-period offer had been $q = 10$ and it had been rejected then he would have offered $q = 10$ again in the second period.

The H buyer (that is, the buyer at information sets that follow Nature's choice of H) always says Yes to any offer of the seller.

The L buyer (that is, the buyer at information sets that follow Nature's choice of L) always says No to an offer of p and always says Yes to an offer of q . The system

of beliefs is as follows (where TL means the left node of the top information set of the seller, TR the right node of that information set, ML means the left node of the middle information set of the seller, MR the right node of that information set, BL means the left node of the bottom information set of the seller, BR the right node of that information set):

$$\mu = \left(\begin{array}{cc|cc|cc} TL & TR & ML & MR & BL & BR \\ 0 & 1 & \frac{2}{3} & \frac{1}{3} & 0 & 1 \end{array} \right).$$

Let us first check sequential rationality. The seller's payoff is (recall that $\delta = \frac{3}{4}$)

$$\left(\frac{2}{3}\right) 12 + \left(\frac{1}{3}\right) \left(\frac{3}{4}\right) 10 = \frac{63}{6}$$

which is greater than the payoff he would get if he offered $q = 10$, namely a payoff of 10.

The H -type's payoff is $20 - 12 = 8$, while if she said No to $p = 12$ and then Yes to $q = 10$ in period 2 her payoff would be $\frac{3}{4}(20 - 10) = \frac{15}{2} = 7.5$.

Furthermore, for the H type, at every node of hers, saying Yes is always strictly better than saying No.

The L -type's payoff is 0, while if she said Yes to $p = 12$ then her payoff would be -2 .

At every node of the L type after having been offered $p = 12$ saying No is strictly better than saying Yes and at every node after having been offered $q = 10$ saying No gives the same payoff as saying Yes, namely 0.

To check consistency, construct the following completely mixed strategy profile $\langle \sigma_n \rangle_{n=1,2,\dots}$:

- (1) for the seller and for the L -buyer, any choice that has zero probability in σ is assigned probability $\frac{1}{n}$ in σ_n and any choice that has probability 1 in σ is assigned probability $1 - \frac{1}{n}$ in σ_n ,
- (2) for the H -buyer, any choice that has zero probability in σ is assigned probability $\frac{1}{n^2}$ in σ_n and any choice that has probability 1 in σ is assigned probability $1 - \frac{1}{n^2}$ in σ_n .

Let us compute the corresponding beliefs μ_n at the top and at the bottom information sets of the seller:

$$\mu_n(TL) = \frac{\frac{1}{n} \left(\frac{1}{n^2} \right)}{\frac{1}{n} \left(\frac{1}{n^2} \right) + \frac{1}{n} \left(\frac{1}{n} \right)} = \frac{1}{1+n} \quad \text{and} \quad \mu_n(TR) = \frac{\frac{1}{n} \left(\frac{1}{n} \right)}{\frac{1}{n} \left(\frac{1}{n^2} \right) + \frac{1}{n} \left(\frac{1}{n} \right)} = \frac{1}{\frac{1}{n} + 1}.$$

Thus $\lim_{n \rightarrow \infty} \mu_n(TL) = 0 = \mu(TL)$ and $\lim_{n \rightarrow \infty} \mu_n(TR) = 1 = \mu(TR)$.

$$\mu_n(BL) = \frac{\left(1 - \frac{1}{n}\right) \left(\frac{1}{n}\right)}{\left(1 - \frac{1}{n}\right) \left(\frac{1}{n}\right) + \left(1 - \frac{1}{n}\right) \left(1 - \frac{1}{n}\right)} = \frac{1}{n}$$

$$\text{and } \mu_n(BR) = \frac{\left(1 - \frac{1}{n}\right) \left(1 - \frac{1}{n}\right)}{\left(1 - \frac{1}{n}\right) \left(\frac{1}{n}\right) + \left(1 - \frac{1}{n}\right) \left(1 - \frac{1}{n}\right)} = 1 - \frac{1}{n};$$

thus $\lim_{n \rightarrow \infty} \mu_n(BL) = 0 = \mu(BL)$ and $\lim_{n \rightarrow \infty} \mu_n(BR) = 1 = \mu(BR)$.

□

13. Perfect Bayesian Equilibrium

13.1 Belief revision and AGM consistency

Any attempt to refine the notion of subgame-perfect equilibrium in extensive-form (or dynamic) games must deal with the issue of *belief revision*: how should a player revise her beliefs when informed that she has to make a choice at an information set of hers to which she initially assigned zero probability? As we saw in the previous chapter, Kreps and Wilson (1982) suggested the notion of a consistent assessment (Definition 12.1.1, Chapter 12) to deal with this issue. From now on, we shall refer to the notion of consistency proposed by Kreps and Wilson, as *KW-consistency* (KW stands for ‘Kreps-Wilson’), in order to distinguish it from a different notion of consistency, called *AGM-consistency*, that will be introduced in this section. We shall make use of concepts developed in Section 9.4 (Chapter 9): the reader might want to review that material before continuing.

In this chapter it will be more convenient to use the so called “history-based” definition of extensive-form game, which is spelled out in Section 13.6. Essentially it boils down to identifying a node in the tree with the sequence of actions leading from the root to it. We call a sequence of actions starting from the root of the tree a *history*.

For example, consider the extensive form of Figure 13.1.




Figure 13.1: A game-frame in extensive form.

Node x_0 (the root of the tree) is identified with the null or empty history \emptyset , decision node x_6 with the history (or sequence of actions) adf , terminal node z_8 with history $adfn$, etc. Thus it will no longer be necessary to label the nodes of the tree, because we can refer to them by naming the corresponding histories.

If h is a decision history we denote by $A(h)$ the set of actions (or choices) available at h . For example, in the game of Figure 13.1, $A(a) = \{c, d\}$, $A(ac) = \{e, f\}$, $A(adf) = \{m, n\}$, etc. If h is a history and a is an action available at h (that is $a \in A(h)$), then we denote by ha the history obtained by appending a to h .

Definition 13.1.1 Given a set H , a total pre-order on H is a binary relation $\succsim \subseteq H \times H$ which is complete ($\forall h, h' \in H$, either $h \succsim h'$ or $h' \succsim h$, or both) and transitive ($\forall h, h', h'' \in H$, if $h \succsim h'$ and $h' \succsim h''$ then $h \succsim h''$).

We write

$h \sim h'$ as a short-hand for “ $h \succsim h'$ and $h' \succsim h$ ”, and

$h \succ h'$ as a short-hand for “ $h \succsim h'$ and $h' \not\succsim h$ ”.

We saw in Chapter 9 (Section 9.4.2) that the AGM theory of belief revision (introduced by Alchourrón *et al.*, 1985) is intimately linked to the notion of a plausibility order. This is reflected in the following definition.

Definition 13.1.2 Given an extensive form, a *plausibility order* is a total pre-order on the set of histories H that satisfies the following properties (D denotes the set of decision histories, $A(h)$ the set of actions available at decision history h and $I(h)$ the information set that contains decision history h): for all $h \in D$,

PL1. $h \succsim ha$, for all $a \in A(h)$.

PL2. (i) There exists an $a \in A(h)$ such that $ha \sim h$,
(ii) for all $a \in A(h)$, if $ha \sim h$ then $h'a \sim h'$, for all $h' \in I(h)$.

PL3. If history h is assigned to chance (or Nature), then $ha \sim h$, for all $a \in A(h)$.

The interpretation of $h \succsim h'$ is that history h is *at least as plausible as* history h' (thus the interpretation of $h \sim h'$ is that h is *just as plausible as* h' and the interpretation of $h \succ h'$ is that h is *more plausible than* h').

- Property *PL1* of Definition 13.1.2 says that adding an action to a decision history h cannot yield a more plausible history than h itself.
- Property *PL2* says that at every decision history h there is at least one action a which is “plausibility preserving” in the sense that adding a to h yields a history which is as plausible as h ; furthermore, any such action a performs the same role with any other history that belongs to the same information set as h .
- Property *PL3* says that all the actions at a history assigned to chance are plausibility preserving.

Definition 13.1.3 Given an extensive-form, an assessment (σ, μ) (see Definition 11.1.1, Chapter 11) is *AGM-consistent* if there exists a plausibility order \succsim on H such that:

(i) the actions that are assigned positive probability by σ are precisely the plausibility-preserving actions: for all $h \in D$ and for all $a \in A(h)$,

$$\sigma(a) > 0 \text{ if and only if } h \sim ha, \quad (P1)$$

(ii) the histories that are assigned positive probability by μ are precisely those that are most plausible within the corresponding information set: for all $h \in D$,

$$\mu(h) > 0 \text{ if and only if } h \succsim h', \text{ for all } h' \in I(h). \quad (P2)$$

If \succsim satisfies properties (P1) and (P2) with respect to (σ, μ) , we say that \succsim *rationalizes* (σ, μ) .

In conjunction with sequential rationality, the notion of AGM-consistency is sufficient to rule out some subgame-perfect equilibria. Consider, for example, the extensive game of Figure 13.2 and the pure-strategy profile $\sigma = (c, d, f)$ (highlighted by double edges), which constitutes a Nash equilibrium of the game (and also a subgame-perfect equilibrium since there are no proper subgames).




Figure 13.2: The strategy profile (c, d, f) cannot be part of an AGM-consistent assessment

Can $\sigma = (c, d, f)$ be part of a sequentially rational AGM-consistent assessment (σ, μ) ? Since, for Player 3, choice f is rational only if the player assigns (sufficiently high) positive probability to history be (at histories a and bd , g yields a higher payoff than f for Player 3), sequential rationality requires that $\mu(be) > 0$; however, any such assessment is *not* AGM-consistent.

In fact, if there were a plausibility order \succsim that satisfied Definition 13.1.3, then, by $P1$, $b \sim bd$ (since $\sigma(d) = 1 > 0$) and $b \succ be$ (since $\sigma(e) = 0$) and, by $P2$,

$$be \succsim bd \tag{13.1}$$

(since – by hypothesis – μ assigns positive probability to be). By transitivity of \succsim , from $b \succ be$ and $b \sim bd$ it follows that $bd \succ be$, contradicting (13.1).

On the other hand, the Nash equilibrium

$$\sigma = \left(\begin{array}{ccc|cc|cc} a & b & c & d & e & f & g \\ 0 & 0 & 1 & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \end{array} \right)$$

together with the system of beliefs

$$\mu = \left(\begin{array}{ccc} a & bd & be \\ 0 & \frac{1}{2} & \frac{1}{2} \end{array} \right)$$

forms a sequentially rational, AGM-consistent assessment: it can be rationalized by the following plausibility order (we use the convention that if the row to which history h belongs is above the row to which history h' belongs, then $h \succ h'$, that is, h is more plausible than h' , and if h and h' belong to the same row then $h \sim h'$, that is, h and h' are equally plausible; as usual, \emptyset denotes the null history, that is, the root of the tree):

$$\left(\begin{array}{ll} \text{most plausible} & \emptyset, c \\ & b, bd, be, bdf, bdg, bef, beg \\ \text{least plausible} & a, af, ag \end{array} \right)$$

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 13.8.1 at the end of this chapter.

13.2 Bayesian consistency

The definition of AGM-consistency pertains to the supports of a given assessment, that is, with the actions that are assigned positive probability by the strategy profile σ and the histories that are assigned positive probability by the system of beliefs μ . In this sense it is a *qualitative* (as opposed to quantitative) property: *how the probabilities are distributed on those supports is irrelevant for AGM-consistency*. However, AGM-consistency is not sufficient: we also need to impose quantitative restrictions concerning the actual probabilities. The reason for this is that we want the given assessment to satisfy “Bayesian updating as long as possible”. By this we mean the following:

1. when information causes no surprises, because the play of the game is consistent with the most plausible play(s) (that is, when information sets are reached that have positive prior probability), then beliefs should be formed using Bayesian updating (Definition 9.4.1, Chapter 9), and
2. when information is surprising (that is, when an information set is reached that had zero prior probability) then new beliefs can be formed in an arbitrary way, but from then on Bayesian updating should be used to update those new beliefs, whenever further information is received that is consistent with those new beliefs.

The next definition formalizes the above requirements.

Definition 13.2.1 Given a finite extensive form, let \succsim be a plausibility order that rationalizes the assessment (σ, μ) . We say that (σ, μ) is *Bayesian* (or *Bayes*) *consistent* relative to \succsim if for every equivalence class E of \succsim that contains some decision history h with $\mu(h) > 0$ (that is, $E \cap D_\mu^+ \neq \emptyset$, where $D_\mu^+ = \{h \in D : \mu(h) > 0\}$) there exists a probability distribution $v_E : H \rightarrow [0, 1]$ (recall that H is a finite set) such that:

- B1. $v_E(h) > 0$ if and only if $h \in E \cap D_\mu^+$.
- B2. If $h, h' \in E \cap D_\mu^+$ and $h' = ha_1 \dots a_m$ (that is, h is a prefix of h') then $v_E(h') = v_E(h) \times \sigma(a_1) \times \dots \times \sigma(a_m)$.
- B3. If $h \in E \cap D_\mu^+$ then, for every $h' \in I(h)$, $\mu(h') = v_E(h'|I(h)) \stackrel{def}{=} \frac{v_E(h')}{\sum_{h'' \in I(h)} v_E(h')}$.

Property B1 requires that $v_E(h) > 0$ if and only if $h \in E$ and $\mu(h) > 0$.¹

Property B2 requires v_E to be consistent with the strategy profile σ in the sense that if $h, h' \in E$, $v_E(h) > 0$, $v_E(h') > 0$ and $h' = ha_1 \dots a_m$ then the probability of h' (according to v_E) is equal to the probability of h multiplied by the probabilities (according to σ) of the actions that lead from h to h' .²

Property B3 requires the system of beliefs μ to satisfy Bayes' rule in the sense that if $h \in E$ and $\mu(h) > 0$ (so that E is the equivalence class of the most plausible elements of the information set $I(h)$) then, for every history $h' \in I(h)$, $\mu(h')$ (the probability assigned to h' by μ) coincides with the probability of h' conditional on $I(h)$, using the probability measure v_E .

For an example of an AGM-consistent and Bayesian-consistent assessment, consider the extensive form of Figure 13.3. The assessment

$$\sigma = \left(\begin{array}{ccc|cc} a & b & c & d & e \\ 0 & 0 & 1 & 1 & 0 \end{array} \middle| \begin{array}{cc} f & g \\ \frac{1}{3} & \frac{2}{3} \end{array} \right) \quad \mu = \left(\begin{array}{ccc|ccc} ad & ae & b & a & bf & bg \\ \frac{1}{4} & 0 & \frac{3}{4} & \frac{1}{4} & \frac{1}{4} & \frac{2}{4} \end{array} \right)$$

is AGM-consistent because it is rationalized by the following plausibility order:

$$\left(\begin{array}{l} \text{most plausible} \\ \text{least plausible} \end{array} \begin{array}{l} \emptyset, c \\ a, b, ad, bf, bg, adf, adg, bfd, bgd \\ ae, aef, aeg, bfe, bge \end{array} \right)$$

Furthermore, it is Bayesian relative to this plausibility order. First of all, note that $D_\mu^+ = \{\emptyset, a, ad, b, bf, bg\}$.³ Let $E = \{\emptyset, c\}$ be the top equivalence class, $F = \{a, b, ad, bf, bg, adf, adg, bfd, bgd\}$ the middle one and $G = \{ae, aef, aeg, bfe, bge\}$ the bottom one. Then only E and F have a non-empty intersection with D_μ^+ and thus, by Definition 13.2.1, we only need to specify two probability distributions: v_E and v_F . The first one is trivial: since $D_\mu^+ \cap E = \{\emptyset\}$, it must be $v_E(\emptyset) = 1$ (and $v_E(c) = 0$). Since $D_\mu^+ \cap F = \{a, ad, b, bf, bg\}$, by B1 of Definition 13.2.1 it must be that $v_F(h) > 0$ if and only if $h \in \{a, ad, b, bf, bg\}$.

¹Thus $v_E(h) = 0$ if and only if either $h \in H \setminus E$ or $\mu(h) = 0$

²Note that if $h, h' \in E$ and $h' = ha_1 \dots a_m$, then $\sigma(a_j) > 0$, for all $j = 1, \dots, m$. In fact, since $h' \sim h$, every action a_j is plausibility preserving and therefore, by Property (P1) of Definition 13.1.3, $\sigma(a_j) > 0$.

³Recall that $D_\mu^+ \stackrel{def}{=} \{h \in D : \mu(h) > 0\}$.




Figure 13.3: The assessment $(c, d, (\frac{1}{3}f, \frac{2}{3}g))$, $\mu(ad) = \frac{1}{4}$, $\mu(b) = \frac{3}{4}$, $\mu(a) = \mu(bf) = \frac{1}{4}$, $\mu(bg) = \frac{2}{4}$ is both AGM-consistent and Bayes-consistent.

Consider the following probability distribution:

$$v_F = \left(\begin{array}{ccccc} a & ad & b & bf & bg \\ \frac{1}{8} & \frac{1}{8} & \frac{3}{8} & \frac{1}{8} & \frac{2}{8} \end{array} \right) \quad (\text{and } v_F(h) = 0 \text{ for every other history } h).$$

Property B2 of Definition 13.2.1 is satisfied, because

$$v_F(ad) = \underbrace{\frac{1}{8}}_{=v_F(a)} \times \underbrace{1}_{\sigma(d)}, \quad v_F(bf) = \underbrace{\frac{3}{8}}_{=v_F(b)} \times \underbrace{\frac{1}{3}}_{\sigma(f)} \quad \text{and} \quad v_F(bg) = \underbrace{\frac{3}{8}}_{=v_F(b)} \times \underbrace{\frac{2}{3}}_{\sigma(g)}.$$

To check that Property B3 of Definition 13.2.1 is satisfied, let $I_2 = \{a, bf, bg\}$ be the information set of Player 2 and $I_3 = \{b, ad, ae\}$ be the information set of Player 3. Then

$$v_F(I_2) = v_F(a) + v_F(bf) + v_F(bg) = \frac{4}{8} \quad \text{and}$$

$$v_F(I_3) = v_F(b) + v_F(ad) + v_F(ae) = \frac{3}{8} + \frac{1}{8} + 0 = \frac{4}{8}.$$

Thus

$$\frac{v_F(a)}{v_F(I_2)} = \frac{\frac{1}{8}}{\frac{4}{8}} = \frac{1}{4} = \mu(a), \quad \frac{v_F(bf)}{v_F(I_2)} = \frac{\frac{1}{8}}{\frac{4}{8}} = \frac{1}{4} = \mu(bf), \quad \frac{v_F(bg)}{v_F(I_2)} = \frac{\frac{2}{8}}{\frac{4}{8}} = \frac{2}{4} = \mu(bg)$$

$$\frac{v_F(b)}{v_F(I_3)} = \frac{\frac{3}{8}}{\frac{4}{8}} = \frac{3}{4} = \mu(b), \quad \frac{v_F(ad)}{v_F(I_3)} = \frac{\frac{1}{8}}{\frac{4}{8}} = \frac{1}{4} = \mu(ad), \quad \frac{v_F(ae)}{v_F(I_3)} = \frac{0}{\frac{4}{8}} = 0 = \mu(ae).$$

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 13.8.2 at the end of this chapter.

13.3 Perfect Bayesian equilibrium

By adding sequential rationality (Definition 11.1.2, Chapter 11) to AGM-consistency (Definition 13.1.3) and Bayesian consistency (Definition 13.2.1) we obtain a new notion of equilibrium for extensive-form games.

Definition 13.3.1 Given a finite extensive-form game, an assessment (σ, μ) is a *perfect Bayesian equilibrium* (PBE) if it is sequentially rational, it is rationalized by a plausibility order on the set of histories and is Bayesian relative to that plausibility order.

For an example of a perfect Bayesian equilibrium, consider the game of Figure 13.4 (which is based on the game-frame of Figure 13.3) and the assessment considered in the previous section, namely

$$\sigma = \left(\begin{array}{ccc|cc|cc} a & b & c & d & e & f & g \\ 0 & 0 & 1 & 1 & 0 & \frac{1}{3} & \frac{2}{3} \end{array} \right) \quad \text{and} \quad \mu = \left(\begin{array}{ccc|ccc} ad & ae & b & a & bf & bg \\ \frac{1}{4} & 0 & \frac{3}{4} & \frac{1}{4} & \frac{1}{4} & \frac{2}{4} \end{array} \right).$$




Figure 13.4: A game based on the game-frame of Figure 13.3.

We showed in the previous section that the assessment under consideration is AGM-consistent and Bayes-consistent. Thus we only need to verify sequential rationality.

For Player 3 the expected payoff from playing f is $(\frac{1}{4})0 + (0)1 + (\frac{3}{4})1 = \frac{3}{4}$

and the expected payoff from playing g is $(\frac{1}{4})3 + (0)2 + (\frac{3}{4})0 = \frac{3}{4}$;

thus any mixture of f and g is sequentially rational, in particular the mixture $\begin{pmatrix} f & g \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$.

For Player 2 the expected payoff from playing d is $\frac{1}{4}[(\frac{1}{3})2 + (\frac{2}{3})2] + (\frac{1}{4})2 + (\frac{2}{4})2 = 2$,

while the expected payoff from playing e is $\frac{1}{4}[(\frac{1}{3})3 + (\frac{2}{3})3] + (\frac{1}{4})0 + (\frac{2}{4})1 = \frac{5}{4}$;

thus d is sequentially rational.

For Player 1, c gives a payoff of 2 while a gives a payoff of $\frac{1}{3}$ and b a gives a payoff of $\frac{5}{3}$;

thus c is sequentially rational.

We now turn to the properties of Perfect Bayesian equilibria.

Theorem 13.3.1 — Bonanno, 2013. Consider a finite extensive-form game and an assessment (σ, μ) . If (σ, μ) is a perfect Bayesian equilibrium then

- (1) σ is a subgame-perfect equilibrium and
- (2) (σ, μ) is a weak sequential equilibrium.

The example of Figure 13.2 showed that not every subgame-perfect equilibrium can be part of a perfect Bayesian equilibrium. Thus, by Theorem 13.3.1, the notion of perfect Bayesian equilibrium is a strict refinement of the notion of subgame-perfect equilibrium.

Theorem 13.3.2 — Bonanno, 2013. Consider a finite extensive-form game. If the assessment (σ, μ) is sequential equilibrium then it is a perfect Bayesian equilibrium.

The next example shows that the notion of sequential equilibrium is a strict refinement of perfect Bayesian equilibrium.

Not every perfect Bayesian equilibrium is a sequential equilibrium. To see this, consider the game of Figure 13.5.




Figure 13.5: A game with a perfect Bayesian equilibrium which is not a sequential equilibrium.

A perfect Bayesian equilibrium of this game is given by the pure-strategy profile $\sigma = (c, d, g)$ (highlighted by double edges), together with the degenerate beliefs $\mu(a) = \mu(be) = 1$.

In fact, (σ, μ) is sequentially rational and, furthermore, it is rationalized by the plausibility order (\blacktriangle) below and is Bayesian relative to it (the probability distributions on the equivalence classes that contain histories h with $\mu(h) > 0$ are written next to the order):

$$\left(\begin{array}{l} \text{most plausible} \\ \emptyset, c \\ a, ad \\ b, bd \\ be, beg \\ ae, aeg \\ bef \\ \text{least plausible} \\ aef \end{array} \right) \quad \left(\begin{array}{l} v_{\{\emptyset, c\}}(\emptyset) = 1 \\ v_{\{a, ad\}}(a) = 1 \\ - \\ v_{\{be, beg\}}(be) = 1 \\ - \\ - \\ - \end{array} \right) \quad (\blacktriangle)$$

The belief revision policy encoded in a perfect Bayesian equilibrium can be interpreted either as the point of view of an external observer or as a belief revision policy which is shared by all the players. For example, the perfect Bayesian equilibrium under consideration (for the game of Figure 13.5), namely $\sigma = (c, d, g)$ and $\mu(a) = \mu(be) = 1$, reflects the following belief revision policy:

- the initial beliefs are that Player 1 will play c ;
- conditional on learning that Player 1 did not play c , the observer would become convinced that Player 1 played a (that is, she would judge a to be strictly more plausible than b) and would expect Player 2 to play d ;
- upon learning that Player 1 did not play c and Player 2 did not play d , the observer would become convinced that Player 1 played b and Player 2 played e , hence judging history be to be strictly more plausible than history ae , thereby reversing her earlier belief that a was strictly more plausible than b .

Such a belief revision policy is consistent with the AGM rationality axioms (Alchourrón *et al.*, 1985) but is incompatible with the notion of sequential equilibrium. In fact, (σ, μ) is *not* KW-consistent (Definition 12.1.1, Chapter 12). To see this, consider an arbitrary sequence $\langle \sigma_n \rangle_{n=1,2,\dots}$ that converges to σ :

$$\sigma_n = \left(\begin{array}{ccc|cc|cc} a & b & c & d & e & f & g \\ p_n & q_n & 1-p_n-q_n & 1-r_n & r_n & t_n & 1-t_n \end{array} \right)$$

with

$$\lim_{n \rightarrow \infty} p_n = \lim_{n \rightarrow \infty} q_n = \lim_{n \rightarrow \infty} r_n = \lim_{n \rightarrow \infty} t_n = 0.$$

Then the corresponding sequence $\langle \mu_n \rangle_{n=1,2,\dots}$ of beliefs obtained by Bayesian updating is given by

$$\mu_n = \left(\begin{array}{cc|cc} a & b & ae & be \\ \frac{p_n}{p_n+q_n} & \frac{q_n}{p_n+q_n} & \frac{p_n r_n}{p_n r_n + q_n r_n} = \frac{p_n}{p_n+q_n} & \frac{q_n r_n}{p_n r_n + q_n r_n} = \frac{q_n}{p_n+q_n} \end{array} \right).$$

Since $\mu_n(a) = \mu_n(ae)$, if $\lim_{n \rightarrow \infty} \mu_n(a) = \mu(a) = 1$ then $\lim_{n \rightarrow \infty} \mu_n(ae) = 1 \neq \mu(ae) = 0$.

By Theorem 12.2 (Chapter 12), every finite extensive-form game with cardinal payoffs has at least one sequential equilibrium and, by Theorem 13.3.2, every sequential equilibrium is a perfect Bayesian equilibrium. Thus the following theorem follows as a corollary of these two results.

Theorem 13.3.3 Every finite extensive-form game with cardinal payoffs has at least one perfect Bayesian equilibrium.

The relationship among the different notions of equilibrium introduced so far (Nash equilibrium, subgame-perfect equilibrium, weak sequential equilibrium, perfect Bayesian equilibrium and sequential equilibrium) is shown in the Venn diagram of Figure 13.6.

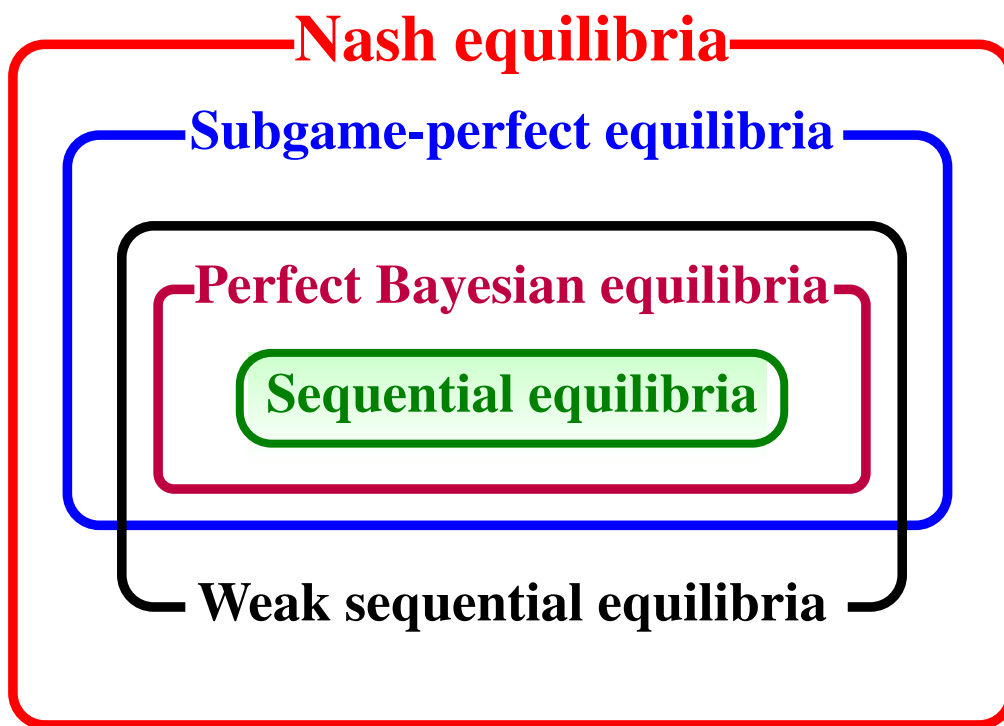


Figure 13.6: The relationship among Nash equilibria, subgame-perfect equilibria, weak sequential equilibria, perfect Bayesian equilibria and sequential equilibria.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 13.8.3 at the end of this chapter.

13.4 Adding independence

13.4.1 Weak independence

The notion of perfect Bayesian equilibrium imposes relatively mild restrictions on beliefs at information sets that are reached with zero probability. Those restrictions require consistency between the strategy profile σ and the system of belief μ , as well as the requirement of Bayesian updating “as long as possible” (that is, also after beliefs have been revised at unreached information sets). The example of Figure 13.5 showed that perfect Bayesian equilibrium is compatible with a belief revision policy that allows a reversal of judgment concerning the behavior of one player after observing an unexpected move by a different player. One might want to rule out such forms of belief revision. In this section we introduce and discuss further restrictions on belief revision that incorporate some form of independence.

“Questionable” belief revision policies like the one illustrated in the previous section⁴ are ruled out by imposing the following restriction on the plausibility order:

if h and h' belong to the same information set (that is, $h' \in I(h)$)
and a is an action available at h ($a \in A(h)$), then

$$h \succsim h' \text{ if and only if } ha \succsim h'a. \quad (IND_1)$$

(IND_1) says that if h is deemed to be at least as plausible as h' then the addition of any available action a must preserve this judgment, in the sense that ha must be deemed to be at least as plausible as $h'a$, and *vice versa*; it can also be viewed as an “independence” condition, in the sense that observation of a new action cannot lead to a change in the relative plausibility of previous histories.⁵

Any plausibility order that rationalizes the assessment $\sigma = (c, d, g)$ and $\mu(a) = \mu(be) = 1$ for the game of Figure 13.5 violates (IND_1), since it must be such that $a \succ b$ (because $\mu(a) > 0$ while $\mu(b) = 0$: see Definition 13.1.3) and also that $be \succ ae$ (because $\mu(be) > 0$ while $\mu(ae) = 0$: see Definition 13.1.3).

Property (IND_1) is a qualitative property (that is, a property that pertains to the plausibility order). We can add to it a quantitative condition on the probabilities to obtain a refinement of the notion of perfect Bayesian equilibrium. This quantitative property is given in Definition 13.4.1 and is a strengthening of the notion of Bayesian consistency introduced in Definition 13.2.1. First we need to define a “full-support common prior”.

Let (σ, μ) be an assessment which is rationalized by a plausibility order \succsim . As before, let D_μ^+ be the set of decision histories to which μ assigns positive probability: $D_\mu^+ = \{h \in D : \mu(h) > 0\}$. Let \mathcal{E}_μ^+ be the set of equivalence classes of \succsim that have a non-empty intersection with D_μ^+ . Clearly \mathcal{E}_μ^+ is a non-empty, finite set. Suppose that (σ, μ) is Bayesian relative to \succsim and let $\{v_E\}_{E \in \mathcal{E}_\mu^+}$ be a collection of probability density functions that satisfy the properties of Definition 13.2.1.

We call a probability density function $v : D \rightarrow (0, 1]$ a *full-support common prior* of $\{v_E\}_{E \in \mathcal{E}_\mu^+}$ if, for every $E \in \mathcal{E}_\mu^+$, $v_E(\cdot) = v(\cdot | E \cap D_\mu^+)$, that is,

$$\text{for all } h \in E \cap D_\mu^+, \quad v_E(h) = \frac{v(h)}{\sum_{h' \in E \cap D_\mu^+} v(h')}.$$

Note that a full support common prior assigns positive probability to *all* decision histories, not only to those in D_μ^+ .

⁴Reversal of relative likelihood judgments implied by the belief revision policy encoded in the assessment $\sigma = (c, d, g)$ and $\mu(a) = \mu(be) = 1$ for the game of Figure 13.5.

⁵Note, however, that (IND_1) is compatible with the following: $a \succ b$ (with $b \in I(a)$) and $bc \succ ad$ (with $bc \in I(ad)$, $c, d \in A(a)$, $c \neq d$).

Definition 13.4.1 Consider an extensive form. Let (σ, μ) be an assessment which is rationalized by the plausibility order \succsim and is Bayesian relative to it and let $\{v_E\}_{E \in \mathcal{E}_\mu^+}$ be a collection of probability density functions that satisfy the properties of Definition 13.2.1. We say that (σ, μ) is *uniformly Bayesian relative to* \succsim if there exists a full-support common prior $v : D \rightarrow (0, 1]$ of $\{v_E\}_{E \in \mathcal{E}_\mu^+}$ that satisfies the following properties.

UB1. If $a \in A(h)$ and $ha \in D$, then

(i) $v(ha) \leq v(h)$ and, (ii) if $\sigma(a) > 0$ then $v(ha) = v(h) \times \sigma(a)$.

UB2. If $a \in A(h)$, h and h' belong to the same information set and $ha, h'a \in D$

then $\frac{v(h)}{v(h')} = \frac{v(ha)}{v(h'a)}$.

We call such a function v a *uniform full-support common prior* of $\{v_E\}_{E \in \mathcal{E}_\mu^+}$.

Property *UB1* requires that the common prior v be consistent with the strategy profile σ , in the sense that if $\sigma(a) > 0$ then $v(ha) = v(h) \times \sigma(a)$ (thus extending Property *B2* of Definition 13.2.1 from D_μ^+ to D).

Property *UB2* requires that the relative probability, according to the common prior v , of any two histories that belong to the same information set remain unchanged by the addition of the same action.

We can obtain a strengthening of the notion of perfect Bayesian equilibrium (Definition 13.3.1) by (1) adding property (*IND*₁) and (2) strengthening Bayes consistency (Definition 13.2.1) to uniform Bayesian consistency (Definition 13.4.1).

Definition 13.4.2 Given an extensive-form game, an assessment (σ, μ) is a *weakly independent perfect Bayesian equilibrium* if it is sequentially rational, it is rationalized by a plausibility order that satisfies (*IND*₁) and is uniformly Bayesian relative to that plausibility order.

As an example of a weakly independent PBE consider the game of Figure 13.7 and the assessment (σ, μ) where $\sigma = (c, d, g, \ell)$ (highlighted by double edges) and $\mu(b) = \mu(ae) = \mu(bf) = 1$ (thus $\mu(a) = \mu(af) = \mu(be) = 0$; the decision histories x such that $\mu(x) > 0$ are shown as black nodes and the decision histories x such that $\mu(x) = 0$ are shown as gray nodes).

This assessment is sequentially rational and is rationalized by the plausibility order (13.2).

It is straightforward to check that plausibility order (13.2) satisfies (*IND*₁).⁶

To see that (σ, μ) is uniformly Bayesian relative to plausibility order (13.2), note that $D_\mu^+ = \{\emptyset, b, ae, bf\}$ and thus the only equivalence classes that have a non-empty intersection with D_μ^+ are $E_1 = \{\emptyset, c\}$, $E_2 = \{b, bd\}$, $E_3 = \{ae, aeg\}$ and $E_4 = \{bf, bfl\}$.

Letting $v_{E_1}(\emptyset) = 1$, $v_{E_2}(b) = 1$, $v_{E_3}(ae) = 1$ and $v_{E_4}(bf) = 1$, this collection of probability distributions satisfies the Properties of Definition 13.2.1. Let v be the uniform distribution over the set of decision histories $D = \{\emptyset, a, b, ae, af, be, bf\}$ (thus $v(h) = \frac{1}{7}$ for every $h \in D$). Then v is a full support common prior of the collection $\{v_{E_i}\}_{i \in \{1, 2, 3, 4\}}$ and satisfies Properties *UB1* and *UB2* of Definition 13.4.1.

⁶ We have that (1) $b \succ a$, $bd \succ ad$, $be \succ ae$ and $bf \succ af$, (2) $ae \succ af$, $aeg \succ afg$ and $aek \succ afk$, (3) $bf \succ be$, $bfl \succ bel$ and $bfm \succ bem$.

$$\left(\begin{array}{l} \text{most plausible} \\ \emptyset, c \\ b, bd \\ a, ad \\ bf, bfl \\ be, bel \\ ae, aeg \\ af, afg \\ bfm \\ bem \\ aek \\ \text{least plausible} \\ afk \end{array} \right) \tag{13.2}$$




Figure 13.7: The assessment $\sigma = (c, d, g, \ell), \mu(b) = \mu(ae) = \mu(bf) = 1$ is a weakly independent PBE.

Note, however, that (σ, μ) is not a sequential equilibrium. The reader is asked to prove this in Exercise 13.9.

13.4.2 Strong independence

A second independence condition (besides (IND_1)) is Property (IND_2) below, which says that if action a is implicitly judged to be at least as plausible as action b , conditional on history h (that is, $ha \succsim hb$), then the same judgment must be made conditional on any other history that belongs to the same information set as h : if $h' \in I(h)$ and $a, b \in A(h)$, then

$$ha \succsim hb \text{ if and only if } h'a \succsim h'b. \quad (IND_2)$$

- R** The two properties (IND_1) and (IND_2) are independent of each other, in the sense that there are plausibility orders that satisfy one of the two properties but not the other (see Exercises 13.7 and 13.8).

Adding Property (IND_2) to the properties given in Definition 13.4.2 we obtain a refinement of the notion of weakly independent perfect Bayesian equilibrium.

Definition 13.4.3 Given an extensive-form game, an assessment (σ, μ) is a *strongly independent perfect Bayesian equilibrium* if it is sequentially rational, it is rationalized by a plausibility order that satisfies Properties (IND_1) and (IND_2) , and is uniformly Bayesian relative to that plausibility order.

The notion of strongly independent PBE is a refinement of the notion of weakly independent PBE. To see this, consider again the game of Figure 13.7 and the assessment (σ, μ) where $\sigma = (c, d, g, \ell)$ and $\mu(b) = \mu(ae) = \mu(bf) = 1$ (thus $\mu(a) = \mu(af) = \mu(be) = 0$). It was shown in the previous section that (σ, μ) is a weakly independent PBE; however, it is not a strongly independent PBE because any plausibility order \succsim that rationalizes (σ, μ) must violate (IND_2) . In fact, since $\mu(ae) > 0$ and $\mu(af) = 0$ it follows from Property P2 of Definition 13.1.3 that

$$ae \succ af. \quad (13.3)$$

Similarly, since $\mu(bf) > 0$ and $\mu(be) = 0$ it follows from Property P2 of Definition 13.1.3 that

$$bf \succ be. \quad (13.4)$$

Since a and b belong to the same information set, (13.3) and (13.4) constitute a violation of (IND_2) .

The following result states that the notions of weakly/strongly independent PBE identify two (nested) solution concepts that lie strictly in the gap between PBE and sequential equilibrium.

Theorem 13.4.1 Consider a finite extensive-form game and an assessment (σ, μ) .

- (A) If (σ, μ) is a sequential equilibrium then it is a strongly independent perfect Bayesian equilibrium.
- (B) There are games where there is a strongly independent perfect Bayesian equilibrium which is not a sequential equilibrium.

Part **(A)** of Theorem 13.4.1 is proved in Section 13.7.

For an example of a strongly independent PBE which is not a sequential equilibrium, consider the game of Figure 13.8 and the assessment (σ, μ) with

$$\sigma = (M, \ell, a, c, e) \quad (\text{highlighted by double edges}) \quad \text{and}$$

$$\mu = \left(\begin{array}{ccc|cc|cc|cc} L & M & R & Ll & Mr & Lm & Rr & Mm & Rl \\ \hline 0 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 \end{array} \right)$$

(the decision histories h such that $\mu(h) = 1$ are denoted by large black dots and the decision histories h such that $\mu(h) = 0$ are denoted by small grey dots). In Exercise 13.13 the reader is asked to prove that (σ, μ) is a strongly independent perfect Bayesian equilibrium but not a sequential equilibrium.




Figure 13.8: A game with a strongly independent PBE which is not a sequential equilibrium.

The relationship between the notions of (weakly/strongly independent) perfect Bayesian equilibrium and sequential equilibrium is illustrated in the Venn diagram of Figure 13.9.

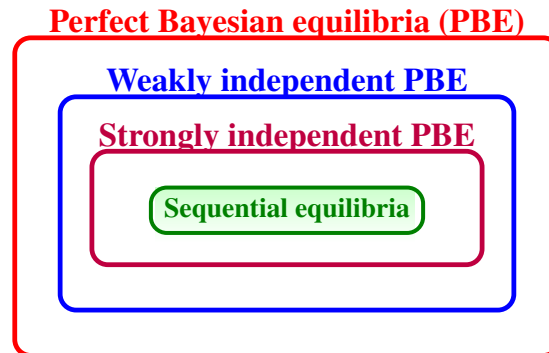


Figure 13.9: The relationship between (weakly/strongly independent) PBE and sequential equilibrium.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 13.8.4 at the end of this chapter.

13.5 Characterization of SE in terms of PBE

Besides sequential rationality, the notion of perfect Bayesian equilibrium (Definition 13.3.1) is based on two elements:

1. the qualitative notions of plausibility order and AGM-consistency and
2. the notion of Bayesian consistency relative to the plausibility order.

By adding two further qualitative properties to the plausibility order – properties (IND_1) and (IND_2) – and by strengthening Bayesian consistency to uniform Bayesian consistency (Definition 13.4.1) we then obtained two nested refinements of the notion of PBE: weakly independent and strongly independent PBE. However, we also noted that the notion of sequential equilibrium is stronger than the notion of (weakly/strictly) independent PBE.

In this section we show that by introducing a further property of the plausibility order – which is a strengthening of both (IND_1) and (IND_2) – we can obtain a characterization of sequential equilibrium. The new requirement is that the plausibility order that rationalizes the given assessment have a “cardinal” numerical representation that can be interpreted as measuring the plausibility distance between histories in a way that is preserved by the addition of a common action.

Definition 13.5.1 Given a plausibility order \succsim on a finite set of histories H , a function $F : H \rightarrow \mathbb{N}$ (where \mathbb{N} denotes the set of non-negative integers) is said to be an *ordinal integer-valued representation* of \succsim if, for every $h, h' \in H$,

$$F(h) \leq F(h') \text{ if and only if } h \succsim h'. \quad (13.5)$$



- An ordinal integer-valued representation of a plausibility order is analogous to an ordinal utility function for a preference relation: it is just a numerical representation of the order. Note that, in the case of a plausibility order, we find it more convenient to assign *lower* values to *more plausible* histories (while a utility function assigns higher values to more preferred outcomes).
- If $F : H \rightarrow \mathbb{N}$ is an integer-valued representation of a plausibility order \succsim on the set of histories H , without loss of generality we can assume that $F(\emptyset) = 0$ (recall that \emptyset denotes the null history, which is always one of the most plausible histories).⁷
- Since H is a finite set, an integer-valued representation of a plausibility order \succsim on H always exists. A natural integer-valued representation is the following, which we shall call the *canonical integer-valued representation*.
 - Define $H_0 = \{h \in H : h \succsim x, \text{ for all } x \in H\}$ (thus H_0 is the set of most plausible histories in H), and $H_1 = \{h \in H \setminus H_0 : h \succsim x, \text{ for all } x \in H \setminus H_0\}$ (thus H_1 is the set of most plausible histories among the ones that remain after removing the set H_0 from H).
 - In general, for every integer $k \geq 1$, define $H_k = \{h \in H \setminus (H_0 \cup \dots \cup H_{k-1}) : h \succsim x, \text{ for all } x \in H \setminus (H_0 \cup \dots \cup H_{k-1})\}$. Since H is finite, there is an $m \in \mathbb{N}$ such that $\{H_0, \dots, H_m\}$ is a partition of H and, for every $j, k \in \mathbb{N}$, with $j < k \leq m$, and for every $h, h' \in H$, if $h \in H_j$ and $h' \in H_k$ then $h \succ h'$.
 - Define $\hat{F} : H \rightarrow \mathbb{N}$ as follows: $\hat{F}(h) = k$ if and only if $h \in H_k$; then the function \hat{F} so defined is an integer-valued representation of \succsim and $\hat{F}(\emptyset) = 0$.

Instead of an ordinal representation of the plausibility order \succsim one could seek a *cardinal* representation which, besides (13.5), satisfies the following property: if h and h' belong to the same information set and $a \in A(h)$, then

$$F(h') - F(h) = F(h'a) - F(ha). \quad (CM)$$

If we think of F as measuring the “plausibility distance” between histories, then we can interpret (CM) as a distance-preserving condition: the plausibility distance between two histories in the same information set is preserved by the addition of the same action.

Definition 13.5.2 A plausibility order \succsim on the set of histories H is *choice measurable* if it has at least one integer-valued representation that satisfies property (CM).

⁷Let $\bar{F} : H \rightarrow \mathbb{N}$ be an integer-valued representation of a plausibility order \succsim and define $F : H \rightarrow \mathbb{N}$ as follows: $F(h) = \bar{F}(h) - \bar{F}(\emptyset)$. Then F is also an integer-valued representation of \succsim and $F(\emptyset) = 0$.

For example, consider the extensive form of Figure 13.10 and the assessment consisting of the pure-strategy profile $\sigma = (a, e)$ (highlighted by double edges) and the system of beliefs $\mu(b) = 1, \mu(c) = 0$ (the grey node in Player 2’s information set represents the history which is assigned zero probability).

This assessment is rationalized by the plausibility order shown below together with two integer-valued representations: \hat{F} is the canonical representation explained in the above remark (page 447), while F is an alternative representation:

$$\left(\begin{array}{cc|cc} & \succsim & \hat{F} & F \\ \hline \text{most plausible} & \emptyset, a & 0 & 0 \\ & b, be & 1 & 1 \\ & bf & 2 & 3 \\ & c, ce & 3 & 4 \\ & d & 4 & 5 \\ \text{least plausible} & cf & 5 & 6 \end{array} \right)$$

While \hat{F} does not satisfy Property (CM) (since $\hat{F}(c) - \hat{F}(b) = 3 - 1 = 2 \neq \hat{F}(cf) - \hat{F}(bf) = 5 - 2 = 3$), F does satisfy Property (CM). Since there is at least one integer-valued representation that satisfies Property (CM), by Definition 13.5.2 the above plausibility order is choice measurable.




Figure 13.10: The assessment $\sigma = (a, e), \mu(b) = 1$ is choice measurable.

We can now state the main result of this section, namely that choice measurability and uniform Bayesian consistency are necessary and sufficient for a perfect Bayesian equilibrium to be a sequential equilibrium.

Theorem 13.5.1 — Bonanno, 2016. Consider a finite extensive-form game and an assessment (σ, μ) . The following are equivalent:

- (A) (σ, μ) is a sequential equilibrium,
- (B) (σ, μ) is a perfect Bayesian equilibrium that is rationalized by a choice-measurable plausibility order and is uniformly Bayesian relative to it.

Theorem 13.5.1 provides a characterization (or understanding) of sequential equilibrium which is free of the questionable requirement of taking a limit of sequences of completely mixed strategies and associated systems of beliefs (see the remarks in Chapter 12, Section 12.3).

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 13.8.5 at the end of this chapter.

13.6 History-based definition of extensive-form game

If A is a set, we denote by A^* the set of finite sequences in A . If $h = \langle a_1, \dots, a_k \rangle \in A^*$ and $1 \leq j \leq k$, the sequence $h' = \langle a_1, \dots, a_j \rangle$ is called a *prefix* of h (a *proper prefix* of h if $j < k$). If $h = \langle a_1, \dots, a_k \rangle \in A^*$ and $a \in A$, we denote the sequence $\langle a_1, \dots, a_k, a \rangle \in A^*$ by ha .

A *finite extensive form* is a tuple $\langle A, H, I, \iota, \{\approx_i\}_{i \in I} \rangle$ whose elements are:

- A finite set of actions A .
- A finite set of histories $H \subseteq A^*$ which is closed under prefixes (that is, if $h \in H$ and $h' \in A^*$ is a prefix of h , then $h' \in H$). The null (or empty) history $\langle \rangle$, denoted by \emptyset , is an element of H and is a prefix of every history (the null history \emptyset represents the root of the tree).

A history $h \in H$ such that, for every $a \in A$, $ha \notin H$, is called a *terminal history*. The set of terminal histories is denoted by Z .

$D = H \setminus Z$ denotes the set of non-terminal or *decision* histories. For every decision history $h \in D$, we denote by $A(h)$ the set of actions available at h , that is, $A(h) = \{a \in A : ha \in H\}$.

- A finite set $I = \{1, \dots, n\}$ of players. In some cases there is also an additional, fictitious, player called *chance*.
- A function $\iota : D \rightarrow I \cup \{\text{chance}\}$ that assigns a player to each decision history. Thus $\iota(h)$ is the player who moves at history h .

A game is said to be *without chance moves* if $\iota(h) \in I$ for every $h \in D$.

For every $i \in I \cup \{\text{chance}\}$, let $D_i = \iota^{-1}(i)$ be the histories assigned to Player i . Thus $\{D_{\text{chance}}, D_1, \dots, D_n\}$ is a partition of D . If history h is assigned to chance, then a probability distribution over $A(h)$ is given that assigns positive probability to every $a \in A(h)$.

- For every player $i \in I$, \approx_i is an equivalence relation on D_i . The interpretation of $h \approx_i h'$ is that, when choosing an action at history $h \in D_i$, Player i does not know whether she is moving at h or at h' .

The equivalence class of $h \in D_i$ is denoted by $I_i(h)$ and is called an *information set of Player i* ; thus $I_i(h) = \{h' \in D_i : h \approx_i h'\}$. The following restriction applies: if $h' \in I_i(h)$ then $A(h') = A(h)$, that is, the set of actions available to a player is the same at any two histories that belong to the same information set of that player.

- The following property, known as *perfect recall*, is assumed: for every player $i \in I$, if $h_1, h_2 \in D_i$, $a \in A(h_1)$ and h_1a is a prefix of h_2 then for every $h' \in I_i(h_2)$ there exists an $h \in I_i(h_1)$ such that ha is a prefix of h' . Intuitively, perfect recall requires a player to remember what she knew in the past and what actions she took previously.

Given an extensive form, one obtains an *extensive game* by adding, for every Player $i \in I$, a *utility* (or *payoff*) function $U_i : Z \rightarrow \mathbb{R}$ (where \mathbb{R} denotes the set of real numbers; recall that Z is the set of terminal histories).

Figure 13.11 shows an extensive form without chance moves where⁸

$$I = \{1, 2, 3, 4\}, \quad A = \{a, b, s, c, d, e, f, g, h, m, n\},$$

$$Z = \{s, ace, acfg, acfh, adeg, adeh, adfm, adfn, bm, bn\},$$

$$D = \{\emptyset, a, b, ac, ad, acf, ade, adf\}, \quad H = D \cup Z,$$

$$A(\emptyset) = \{a, b, s\}, \quad A(a) = \{c, d\}, \quad A(ac) = A(ad) = \{e, f\},$$

$$A(acf) = A(ade) = \{g, h\}, \quad A(adf) = A(b) = \{m, n\},$$

$$I(\emptyset) = 1, \quad I(a) = 2, \quad I(ac) = I(ad) = 3, \quad I(acf) = I(ade) = I(adf) = I(b) = 4,$$

$$\approx_1 = \{(\emptyset, \emptyset)\}, \quad \approx_2 = \{(a, a)\}, \quad \approx_3 = \{(ac, ac), (ac, ad), (ad, ac), (ad, ad)\},$$

$$\approx_4 = \{(acf, acf), (acf, ade), (ade, acf), (ade, ade), \\ (adf, adf), (adf, b), (b, adf), (b, b)\}.$$

The information sets containing more than one history (for example, $I_4(b) = \{adf, b\}$) are shown as rounded rectangles. The root of the tree represents the null history \emptyset .




Figure 13.11: Extensive form without chance moves.

⁸In order to simplify the notation we write a instead of $\langle \emptyset, a \rangle$, ac instead of $\langle \emptyset, a, c \rangle$, etc.

If h and h' are decision histories not assigned to chance, we write $h' \in I(h)$ as a shorthand for $h' \in I_{i(h)}(h)$. Thus $h' \in I(h)$ means that h and h' belong to the same information set (of the player who moves at h). If h is a history assigned to chance, we use the convention that $I(h) = \{h\}$.

Given an extensive form, a *pure strategy* of player $i \in I$ is a function that associates with every information set of Player i an action at that information set, that is, a function $s_i : D_i \rightarrow A$ such that **(1)** $s_i(h) \in A(h)$ and **(2)** if $h' \in I_i(h)$ then $s_i(h') = s_i(h)$.

For example, one of the pure strategies of Player 4 in the extensive form illustrated in Figure 13.11 is $s_4(acf) = s_4(ade) = g$ and $s_4(adf) = s_4(b) = m$.

A *behavior strategy* of player i is a collection of probability distributions, one for each information set, over the actions available at that information set; that is, a function $\sigma_i : D_i \rightarrow \Delta(A)$ (where $\Delta(A)$ denotes the set of probability distributions over A) such that **(1)** $\sigma_i(h)$ is a probability distribution over $A(h)$ and **(2)** if $h' \in I_i(h)$ then $\sigma_i(h') = \sigma_i(h)$.

- If the game does not have chance moves, we define a behavior strategy *profile* as an n -tuple $\sigma = (\sigma_1, \dots, \sigma_n)$ where, for every $i \in I$, σ_i is a behavior strategy of Player i .
- If the game has chance moves then we use the convention that a behavior strategy profile is an $(n+1)$ -tuple $\sigma = (\sigma_1, \dots, \sigma_n, \sigma_{chance})$ where, if h is a history assigned to chance and $a \in A(h)$ then $\sigma_{chance}(h)(a)$ is the probability associated with a .
- When there is no risk of ambiguity (e.g. because no action is assigned to more than one information set) we shall denote by $\sigma(a)$ the probability assigned to action a by the relevant component of the strategy profile σ .⁹

Note that a pure strategy is a special case of a behavior strategy where each probability distribution is degenerate.

A behavior strategy is *completely mixed at history* $h \in D$ if, for every $a \in A(h)$, $\sigma(a) > 0$.

For example, in the extensive form of Figure 13.11:

a possible behavior strategy for Player 1 is $\begin{pmatrix} a & b & s \\ 0 & 0 & 1 \end{pmatrix}$, which we will more simply denote by s (it coincides with a pure strategy of Player 1)

and a possible behavior strategy of Player 2 is $\begin{pmatrix} c & d \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$, which is a completely mixed strategy.

⁹If $h \in D_i$ and σ_i is the i^{th} component of σ , then $\sigma_i(h)$ is a probability distribution over $A(h)$ and if $a \in A(h)$ then $\sigma_i(h)(a)$ is the probability assigned to action a by $\sigma_i(h)$. Thus we denote $\sigma_i(h)(a)$ more simply by $\sigma(a)$.

13.7 Proofs

In this section we prove Theorem 13.4.1. To simplify the notation, we will assume that no action is available at more than one information set, that is, that if $a \in A(h) \cap A(h')$ then $h' \in I(h)$ (this is without loss of generality, because we can always rename some of the actions). We start with a preliminary result.

Lemma 13.1 Let \succsim be a plausibility order over the set H of histories of an extensive form and let $F : H \rightarrow \mathbb{N}$ be an integer-valued representation of \succsim (that is, for all $h, h' \in H$, $F(h) \leq F(h')$ if and only if $h \succsim h'$). Then the following are equivalent:

- (A) F satisfies Property (CM).
- (B) F satisfies the following property: for all $h, h' \in H$ and $a, b \in A(h)$, if $h' \in I(h)$ then

$$F(hb) - F(ha) = F(h'b) - F(h'a). \quad (CM')$$

Proof. Let \succsim be a plausibility order on the set of histories H and let $F : H \rightarrow \mathbb{N}$ be an integer-valued representation of \succsim that satisfies Property (CM).

Without loss of generality (see Remark on page 447), we can assume that $F(\emptyset) = 0$.

For every decision history h and action $a \in A(h)$, define $\lambda(a) = F(ha) - F(h)$. The function $\lambda : A \rightarrow \mathbb{N}$ is well defined, since, by assumption, no action is available at more than one information set and, by (CM), if $h' \in I(h)$ then $F(h'a) - F(h') = F(ha) - F(h)$.

Then, for every history $h = \langle a_1, a_2, \dots, a_m \rangle$, $F(h) = \sum_{i=1}^m \lambda(a_i)$. In fact,

$$\begin{aligned} & \lambda(a_1) + \lambda(a_2) + \dots + \lambda(a_m) \\ &= (F(a_1) - F(\emptyset)) + (F(a_1a_2) - F(a_1)) + \dots + (F(a_1a_2 \dots a_m) - F(a_1a_2 \dots a_{m-1})) \\ &= F(a_1a_2 \dots a_m) = F(h) \end{aligned}$$

(recall that $F(\emptyset) = 0$). Thus, for every $h \in D$ and $a \in A(h)$, $F(ha) = F(h) + \lambda(a)$.

Hence, $F(hb) - F(ha) = F(h) + \lambda(b) - (F(h) + \lambda(a)) = \lambda(b) - \lambda(a)$

and $F(h'b) - F(h'a) = F(h') + \lambda(b) - (F(h') + \lambda(a)) = \lambda(b) - \lambda(a)$

so that $F(hb) - F(ha) = F(h'b) - F(h'a)$.

Thus we have shown that (CM) implies (CM').

Now we show the converse, namely that (CM') implies (CM).

Let \succsim be a plausibility order on the set of histories H and let $F : H \rightarrow \mathbb{N}$ be an integer-valued representation of \succsim that satisfies (CM').

Select arbitrary $h' \in I(h)$ and $a \in A(h)$.

Let $b \in A(h)$ be a plausibility-preserving action at h (there must be at least one such action: see Definition 13.1.2); then, $h \sim hb$ and $h' \sim h'b$.

Hence, since F is a representation of \succsim , $F(hb) = F(h)$ and $F(h'b) = F(h')$ so that $F(h') - F(h) = F(h'b) - F(hb)$.

By (CM'), $F(h'b) - F(hb) = F(h'a) - F(ha)$.

From the last two equalities it follows that $F(h') - F(h) = F(h'a) - F(ha)$,

that is, (CM) holds. ■

Proof of Theorem 13.4.1 Consider a finite extensive-form game and an assessment (σ, μ) . Suppose that (σ, μ) is a sequential equilibrium. We want to show that (σ, μ) is a strongly independent PBE, that is, that (σ, μ) is sequentially rational, is rationalized by a plausibility order that satisfies properties (IND_1) and (IND_2) and is uniformly Bayesian relative to that plausibility order.

By Theorem 13.5.1 (σ, μ) is sequentially rational, is rationalized by a plausibility order that satisfies property (CM) and is uniformly Bayesian relative to that plausibility order.

Thus it is sufficient to prove that property (CM) implies properties (IND_1) and (IND_2) .

- Proof that CM implies IND_1 .

Let \succsim be a plausibility order and F an integer valued representation of \succsim that satisfies property (CM) .

Let histories h and h' belong to the same information set ($h' \in I(h)$) and let a be an action available at h ($a \in A(h)$).

We need to show that

$$h \succsim h' \quad \text{if and only if} \quad ha \succsim h'a.$$

Suppose that $h \succsim h'$; then (by Definition 13.5.1) $F(h) \leq F(h')$, that is,

$$F(h) - F(h') \leq 0. \tag{13.6}$$

By (CM) it follows from (13.6) that $F(ha) - F(h'a) \leq 0$, that is, $F(ha) \leq F(h'a)$ and thus, by Definition 13.5.1, $ha \succsim h'a$.

Conversely, if $ha \succsim h'a$ then $F(ha) \leq F(h'a)$ and thus, by (CM) , $F(h) \leq F(h')$, which implies that $h \succsim h'$.

- Proof that CM implies IND_2 .

Let $h' \in I(h)$ and $a, b \in A(h)$. We need to show that

$$ha \succsim hb \quad \text{if and only if} \quad h'a \succsim h'b.$$

By Lemma 13.1, $F(hb) - F(ha) = F(h'b) - F(h'a)$.

If $F(hb) \leq F(ha)$ then $F(h'b) \leq F(h'a)$ and thus (by Definition 13.5.1) $hb \succsim ha$ and $h'b \succsim h'a$;

if $F(ha) \leq F(hb)$ then $F(h'a) \leq F(h'b)$ and thus (by Definition 13.5.1) $ha \succsim hb$ and $h'a \succsim h'b$. ■

13.8 Exercises

13.8.1 Exercises for Section 13.1: Belief revision and AGM consistency

The answers to the following exercises are given in Section 13.9

Exercise 13.1

Consider the game shown in Figure 13.12. Determine if there is a plausibility order that rationalizes the following assessment (Definition 13.1.3):

$$\sigma = \left(\begin{array}{ccc|cc|cc|cc} a & b & s & c & d & e & f & g & h \\ \frac{1}{3} & 0 & \frac{2}{3} & 0 & 1 & \frac{1}{2} & \frac{1}{2} & 1 & 0 \end{array} \right) \mu = \left(\begin{array}{cc|ccc} ac & ad & acf & ade & adf & b \\ 0 & 1 & \frac{1}{5} & \frac{1}{5} & \frac{3}{5} & 0 \end{array} \right)$$




Figure 13.12: The game for Exercise 13.1.

Exercise 13.2

Consider again the game of Exercise 13.1 (Figure 13.12). Find all the assessments that are rationalized by the following plausibility order:

$$\left(\begin{array}{l} \text{most plausible } \emptyset, s \\ a, ac, ad, ace, ade, adeg, b, bg \\ acf, adf, acfg, adfg \\ adeh, bh \\ \text{least plausible } acfh, adfh \end{array} \right)$$

13.8.2 Exercises for Section 13.2: Bayesian consistency

The answers to the following exercises are given in Section 13.9.

Exercise 13.3

Consider the game of Figure 13.13 (which reproduces Figure 13.4) and the assessment

$$\sigma = \left(\begin{array}{ccc|cc|cc} a & b & c & d & e & f & g \\ 0 & 0 & 1 & 1 & 0 & \frac{1}{3} & \frac{2}{3} \end{array} \right), \quad \mu = \left(\begin{array}{ccc|ccc} ad & ae & b & a & bf & bg \\ \frac{1}{4} & 0 & \frac{3}{4} & \frac{1}{4} & \frac{1}{4} & \frac{2}{4} \end{array} \right)$$

which is rationalized by the plausibility order

$$\left(\begin{array}{l} \text{most plausible} \\ \text{least plausible} \end{array} \begin{array}{l} \emptyset, c \\ a, ad, adf, adg, b, bf, bg, bfd, bgd \\ ae, aef, aeg, bfe, bge \end{array} \right).$$

Let $F = \{a, ad, adf, adg, b, bf, bg, bfd, bgd\}$ be the middle equivalence class of the plausibility order. It was shown in Section 13.2 that the probability distribution

$$\nu_F = \left(\begin{array}{ccccc} a & ad & b & bf & bg \\ \frac{1}{8} & \frac{1}{8} & \frac{3}{8} & \frac{1}{8} & \frac{2}{8} \end{array} \right)$$

satisfies the properties of Definition 13.2.1.

Show that there is no other probability distribution on F that satisfies those properties. ■




Figure 13.13: The game for Exercise 13.3.

Exercise 13.4

Consider the extensive form shown in Figure 13.14 and the assessment

$$\sigma = \left(f, A, \left(\begin{array}{cc|cc} L & R & \ell & r \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{5} & \frac{4}{5} \end{array} \right), D \right), \quad \mu = \left(\begin{array}{ccc|cc} a & b & c & d & e \\ 0 & \frac{1}{3} & \frac{2}{3} & \frac{3}{4} & \frac{1}{4} \end{array} \right).$$

(a) Verify that the given assessment is rationalized by the following plausibility order:

$$\left(\begin{array}{l} \text{most plausible} \\ \text{least plausible} \end{array} \begin{array}{l} \emptyset, f, fA \\ b, c, d, e, bL, bR, cL, cR, d\ell, dr, e\ell, er, erD \\ a, aL, aR, fB, erC \end{array} \right)$$

(b) Let $E = \{\emptyset, f, fA\}$ be the top equivalence class of the plausibility order. Show that there is a unique probability distribution v_E on E that satisfies the properties of Definition 13.2.1.

(c) Let $F = \{b, c, d, e, bL, bR, cL, cR, d\ell, dr, e\ell, er, erD\}$ be the middle equivalence class of the plausibility order. Show that both of the following probability distributions satisfy the properties of Definition 13.2.1:

$$v_F = \left(\begin{array}{ccccc} b & c & d & e & er \\ \frac{20}{132} & \frac{40}{132} & \frac{45}{132} & \frac{15}{132} & \frac{12}{132} \end{array} \right) \quad \hat{v}_F = \left(\begin{array}{ccccc} b & c & d & e & er \\ \frac{5}{87} & \frac{10}{87} & \frac{45}{87} & \frac{15}{87} & \frac{12}{87} \end{array} \right)$$




Figure 13.14: The game for Exercise 13.4.

13.8.3 Exercises for Section 13.3: Perfect Bayesian equilibrium

The answers to the following exercises are given in Section 13.9

Exercise 13.5

Consider the extensive form shown in Figure 13.15.

Find two perfect Bayesian equilibria (σ, μ) and (σ', μ') where both σ and σ' are pure-strategy profiles and $\sigma \neq \sigma'$.




Figure 13.15: The game for Exercise 13.5.

Exercise 13.6

Consider the extensive form shown in Figure 13.16.

Prove that the following assessment is a perfect Bayesian equilibrium:

$$\sigma = \left(s, \left| \begin{array}{cc|cc|cc|cc} c & d & e & f & g & h \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{3} & \frac{2}{3} & \frac{1}{2} & \frac{1}{2} \end{array} \right. \right), \quad \mu = \left(\begin{array}{cc|cc|cc|cc} ac & ad & acf & ade & adf & b \\ \frac{1}{2} & \frac{1}{2} & \frac{2}{11} & \frac{1}{11} & \frac{2}{11} & \frac{6}{11} \end{array} \right).$$




Figure 13.16: The game for Exercise 13.6.

13.8.4 Exercises for Section 13.4: Adding independence

The answers to the following exercises are given in Section 13.9

Exercise 13.7

Draw an extensive form where Player 1 moves first and Player 2 moves second without being informed of Player 1's choice.

Player 1 chooses between a and b , while Player 2's choices are c , d and e .

Find an assessment (σ, μ) which is rationalized by a plausibility order that satisfies Property (IND_1) but fails Property (IND_2) .

Exercise 13.8

Find an extensive form and an assessment (σ, μ) which is rationalized by a plausibility order that violates Property (IND_1) but satisfies Property (IND_2) as well as the following property:

$$\text{if } h' \in I(h), \quad a \in A(h), \quad h'a \in I(ha) \quad \text{and} \quad \{h, h', ha, h'a\} \subseteq D_\mu^+$$


$$\text{then } \frac{\mu(h)}{\mu(h')} = \frac{\mu(ha)}{\mu(h'a)}.$$

13.8.5 Exercises for Section 13.5: Characterization of sequential equilibrium in terms of PBE

The answers to the following exercises are given in Section 13.9

Exercise 13.9

Consider the game of Figure 13.7, reproduced below, and the assessment (σ, μ) where $\sigma = (c, d, g, \ell)$ (highlighted by double edges) and $\mu(b) = \mu(ae) = \mu(bf) = 1$ (thus $\mu(a) = \mu(af) = \mu(be) = 0$; the decision histories x such that $\mu(x) > 0$ are shown as black nodes and the decision histories x such that $\mu(x) = 0$ are shown as gray nodes). It was shown in Section 13.4.1 that (σ, μ) is a weakly independent perfect Bayesian equilibrium. Using Theorem 13.5.1 prove that (σ, μ) is not a sequential equilibrium. ■



Exercise 13.10

Consider the (partial) extensive form shown in Figure 13.17.

Using Theorem 13.5.1, prove that there is no sequential equilibrium (σ, μ) where $\sigma = (a, g, r, \dots)$ (that is, σ assigns probability 1 to a, g and r), $\mu(c) = \mu(e) = \mu(ch) = 0$ and $\mu(h) > 0$ for every other decision history h .

[Hint: consider all the possible plausibility orders that rationalize (σ, μ) .] ■




Figure 13.17: The partial extensive form for Exercise 13.10.

Exercise 13.11

Consider the game shown in Figure 13.18.

Let (σ, μ) be an assessment with $\sigma = (a, T, f, L)$ (highlighted by double edges; note that σ is a subgame-perfect equilibrium), $\mu(b) > 0$ and $\mu(c) > 0$.

- (a) Prove that (σ, μ) can be rationalized by a choice-measurable plausibility order only if μ satisfies the following condition:

$$\mu(bB) > 0 \quad \text{if and only if} \quad \mu(cBf) > 0.$$

- (b) Prove that if, besides from being rationalized by a choice-measurable plausibility order \succsim , (σ, μ) is also uniformly Bayesian relative to \succsim (Definition 13.4.1), then μ satisfies the following condition:

$$\text{if } \mu(bB) > 0 \quad \text{then} \quad \frac{\mu(cBf)}{\mu(bB)} = \frac{\mu(c)}{\mu(b)}.$$




Figure 13.18: The extensive form for Exercise 13.11.

Exercise 13.12

Consider again the game of Exercise 13.11 (shown in Figure 13.18). Let (σ, μ) be an assessment with

$$\sigma = (a, T, f, L) \quad \mu = \left(\begin{array}{cc|ccc} b & c & bB & cBf & d \\ \frac{7}{10} & \frac{3}{10} & \frac{7}{18} & \frac{3}{18} & \frac{8}{18} \end{array} \right).$$

Prove that (σ, μ) is a sequential equilibrium by using Theorem 13.5.1.

Exercise 13.13 — ***Challenging Question***.

Consider the game shown in Figure 13.19. Let (σ, μ) be the following assessment:

$$\sigma = (M, \ell, a, c, e) \quad (\text{highlighted by double edges})$$

$$\mu = \left(\begin{array}{ccc|cc|cc|cc} L & M & R & Ll & Mr & Lm & Rr & Mm & Rl \\ \hline 0 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 \end{array} \right)$$

(the decision histories h such that $\mu(h) = 1$ are denoted by large black dots and the decision histories h such that $\mu(h) = 0$ are denoted by small grey dots).

- (a) Prove that (σ, μ) is a strongly independent perfect Bayesian equilibrium (Definition 13.4.3).
- (b) Prove that (σ, μ) is not a sequential equilibrium.
 [Hint: prove that (σ, μ) cannot be rationalized by a choice-measurable plausibility order and then invoke Theorem 13.5.1.]




Figure 13.19: The extensive form for Exercise 13.13.

13.9 Solutions to Exercises

Solutions to Exercise 13.1

From σ we get that the plausibility preserving actions are a, s, d, e, f and g .

Thus ade is as plausible as adf (by transitivity, since each of them is as plausible as ad) and each of them is more plausible than:

- acf [since (1) ad is more plausible than ac , (2) ade and adf are as plausible as ad and (3) acf is as plausible as ac]
- and b [since (1) a is more plausible than b and (2) ade and adf are as plausible as ad , which in turn is as plausible as a].

Hence the most plausible histories in Player 4's information set are ade and adf .

It follows that it must be that $\mu(ade) > 0, \mu(adf) > 0, \mu(acf) = 0$ and $\mu(b) = 0$.

Hence the given assessment cannot be rationalized by a plausibility order. □

Solutions to Exercise 13.2

The assessment under consideration is:

$$\left(\begin{array}{l} \text{most plausible} \\ \emptyset, s \\ a, ac, ad, ace, ade, adeg, b, bg \\ acf, adf, acfg, adfg \\ adeh, bh \\ \text{least plausible} \\ acfh, adfh \end{array} \right)$$

The plausibility-preserving actions are s, c, d, e and g .

Thus the strategy profile must be of the form

$$\sigma = \left(\begin{array}{ccc|cc|cc|cc} a & b & s & c & d & e & f & g & h \\ 0 & 0 & 1 & p & 1-p & 1 & 0 & 1 & 0 \end{array} \right)$$

with $0 < p < 1$.

At Player 3's information set ac and ad are equally plausible and

at Player 4's information set ade and b are equally plausible and each is more plausible than acf and adf .

Thus the system of beliefs must be of the form

$$\mu = \left(\begin{array}{cc|cccc} ac & ad & acf & ade & adf & b \\ q & 1-q & 0 & r & 0 & 1-r \end{array} \right)$$

with $0 < q < 1$ and $0 < r < 1$. □

Solutions to Exercise 13.3

Since $D_\mu^+ \cap F = \{a, ad, b, bf, bg\}$, Property B1 of Definition 13.2.1 requires that $v_F(h) > 0$ if and only if $h \in \{a, ad, b, bf, bg\}$.

Let $v_F(a) = p \in (0, 1)$ and $v_F(b) = q \in (0, 1)$.

Then, by Property B2 of Definition 13.2.1,

$$v_F(ad) = v_F(a) \times \sigma(d) = p \times 1 = p, \quad v_F(bf) = v_F(b) \times \sigma(f) = q \times \frac{1}{3}$$

and $v_F(bg) = v_F(b) \times \sigma(g) = q \times \frac{2}{3}$.

$$\text{Thus } v_F = \begin{pmatrix} a & ad & b & bf & bg \\ p & p & q & \frac{1}{3}q & \frac{2}{3}q \end{pmatrix}$$

and the sum of these probabilities must be 1: $2p + 2q = 1$, that is $p + q = \frac{1}{2}$.

Let $I_2 = \{a, bf, bg\}$ be the information set of Player 2

and $I_3 = \{b, ad, ae\}$ the information set of Player 3.

$$\text{Then } v_F(I_2) = v_F(a) + v_F(bf) + v_F(bg) = p + \frac{1}{3}q + \frac{2}{3}q = p + q = \frac{1}{2}$$

$$\text{and } v_F(I_3) = v_F(b) + v_F(ad) + v_F(ae) = q + p + 0 = \frac{1}{2}.$$

Thus $\frac{v_F(a)}{v_F(I_2)} = \frac{p}{\frac{1}{2}} = 2p$ and, by Property B3, we need this to be equal to $\mu(a) = \frac{1}{4}$;

solving $2p = \frac{1}{4}$ we get $p = \frac{1}{8}$.

Similarly, $\frac{v_F(b)}{v_F(I_3)} = \frac{q}{\frac{1}{2}} = 2q$ and, by Property B3, we need this to be equal to $\mu(b) = \frac{3}{4}$;

solving $2q = \frac{3}{4}$ we get $q = \frac{3}{8}$

(alternatively, we could have derived $q = \frac{3}{8}$ from $p = \frac{1}{8}$ and $p + q = \frac{1}{2}$). □

Solutions to Exercise 13.4

- (a) The plausibility-preserving actions are f, A, L, R, ℓ, r and D and these are precisely the actions that are assigned positive probability by σ .

Furthermore, the most plausible histories in information set $\{a, b, c\}$ are b and c and the two histories in information set $\{d, e\}$ are equally plausible.

This is consistent with the fact that $D_\mu^+ = \{\emptyset, f, b, c, d, e, er\}$.

Thus the properties of Definition 13.1.3 are satisfied.

- (b) By Property B1 of Definition 13.2.1 the support of ν_E must be $D_\mu^+ \cap E = \{\emptyset, f\}$ and by Property B2 it must be that $\nu_E(f) = \nu_E(\emptyset) \times \sigma(f) = \nu_E(\emptyset) \times 1 = \nu_E(\emptyset)$. Thus the only solution is $\nu_E(\emptyset) = \nu_E(f) = \frac{1}{2}$.

- (c) Since $D_\mu^+ \cap F = \{b, c, d, e, er\}$:

- By Property B1 of Definition 13.2.1 the support of the probability distribution must coincide with the set $\{b, c, d, e, er\}$, which is indeed true for both ν_F and $\hat{\nu}_F$.
- By Property B2, the probability of er must be equal to the probability of e times $\sigma(r) = \frac{4}{5}$ and this is indeed true for both ν_F and $\hat{\nu}_F$.

- By Property B3, the conditional probability $\frac{\nu_F(b)}{\nu_F(a) + \nu_F(b) + \nu_F(c)} = \frac{\frac{20}{132}}{0 + \frac{20}{132} + \frac{40}{132}} = \frac{1}{3}$

must be equal to $\mu(b)$ and this is indeed true;

similarly, the conditional probability $\frac{\nu_F(d)}{\nu_F(d) + \nu_F(e)} = \frac{\frac{45}{132}}{\frac{45}{132} + \frac{15}{132}} = \frac{3}{4}$ must be equal to $\mu(d)$ and this is also true.

Similar computations show that $\frac{\hat{\nu}_F(b)}{\hat{\nu}_F(a) + \hat{\nu}_F(b) + \hat{\nu}_F(c)} = \mu(b)$

and $\frac{\hat{\nu}_F(d)}{\hat{\nu}_F(d) + \hat{\nu}_F(e)} = \mu(d)$. □

Solutions to Exercise 13.5 The game under consideration is reproduced in Figure 13.20.




Figure 13.20: The extensive game for Exercise 13.5.

One perfect Bayesian equilibrium is

$$\sigma = (b, c, e, g), \quad \mu = \left(\begin{array}{cc|cc} ac & ad & acf & ade & adf \\ 1 & 0 & \frac{1}{2} & \frac{1}{2} & 0 \end{array} \right).$$

Let us first verify sequential rationality.

For Player 1, at the root, b gives a payoff of 2, while a gives a payoff of 1; thus b is sequentially rational.

For Player 2, c gives a payoff of 1 and d a payoff of 0; thus c is sequentially rational.

For Player 3, e gives a payoff of 3, while f gives a payoff of 2; thus e is sequentially rational.

For Player 1 at his bottom information set, g gives a payoff of $(\frac{1}{2})3 + (\frac{1}{2})3 = 3$ and h gives a payoff of $(\frac{1}{2})1 + (\frac{1}{2})3 = 2$; thus g is sequentially rational.

The following plausibility order rationalizes the above assessment:

$$\left(\begin{array}{l} \text{most plausible} \\ \emptyset, b \\ a, ac, ace \\ acf, acfg, ad, ade, adeg \\ adf, adfg, acfh, adeh \\ \text{least plausible} \\ adfh \end{array} \right)$$

Name the equivalence classes of this order E_1, E_2, \dots, E_5 (with E_1 being the top one and E_5 the bottom one).

Then the following probability distributions on the sets $E_i \cap D_\mu^+$ ($i = 1, 2, 3$) satisfy the properties of Definition 13.2.1 (note that $E_i \cap D_\mu^+ \neq \emptyset$ if and only if $i \in \{1, 2, 3\}$, since $D_\mu^+ = \{\emptyset, a, ac, acf, ade\}$):

$$v_{E_1}(\emptyset) = 1, \quad v_{E_2} = \left(\begin{array}{cc} a & ac \\ \frac{1}{2} & \frac{1}{2} \end{array} \right), \quad v_{E_3} = \left(\begin{array}{ccc} acf & ad & ade \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{array} \right).$$

Another perfect Bayesian equilibrium is:

$$\sigma = (a, d, e, h), \quad \mu = \left(\begin{array}{cc|ccc} ac & ad & acf & ade & adf \\ 0 & 1 & 0 & 1 & 0 \end{array} \right).$$

Sequential rationality is easily verified. The assessment is rationalized by the following plausibility order:

$$\left(\begin{array}{l} \text{most plausible} \\ \emptyset, a, ad, ade, adeh \\ b, ac, ace, adf, adfh \\ acf, acfh, adeg \\ \text{least plausible} \\ acfg, adfg \end{array} \right)$$

Only the top equivalence class $E_1 = \{\emptyset, a, ad, ade, adeh\}$ has a non-empty intersection with $D_\mu^+ = \{\emptyset, a, ad, ade\}$.

The following probability distribution satisfies the properties of Definition 13.2.1:

$$v_{E_1} = \left(\begin{array}{cccc} \emptyset & a & ad & ade \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \end{array} \right).$$

□

Solutions to Exercise 13.6 The game under consideration is reproduced in Figure 13.21.




Figure 13.21: The extensive game for Exercise 13.6.

We need to show that the following is a perfect Bayesian equilibrium:

$$\sigma = \left(s, \begin{array}{c|c} c & d \\ \hline \frac{1}{2} & \frac{1}{2} \end{array} \mid \begin{array}{c|c} e & f \\ \hline \frac{1}{3} & \frac{2}{3} \end{array} \mid \begin{array}{c|c} g & h \\ \hline \frac{1}{2} & \frac{1}{2} \end{array} \right), \quad \mu = \left(\begin{array}{c|c} ac & ad \\ \hline \frac{1}{2} & \frac{1}{2} \end{array} \mid \begin{array}{c|c|c} acf & ade & adf \\ \hline \frac{2}{11} & \frac{1}{11} & \frac{2}{11} \end{array} \mid \begin{array}{c} b \\ \hline \frac{6}{11} \end{array} \right).$$

First we verify sequential rationality.

For Player 1 the possible payoffs are:

from s : 3

from b : $\frac{1}{2}(3) + \frac{1}{2}(2) = 2.5$

from a : $\frac{1}{2} \left[\frac{1}{3}(1) + \frac{2}{3} \left(\frac{1}{2}(2) + \frac{1}{2}(1) \right) \right] + \frac{1}{2} \left[\frac{1}{3}(3) + \frac{2}{3} \left(\frac{1}{2}(1) + \frac{1}{2}(3) \right) \right] = \frac{11}{6}$

Thus s is sequentially rational.

For Player 2 the possible payoffs are:

from c : $\frac{1}{3}(2) + \frac{2}{3}(2) = 2$

from d : $\frac{1}{3} \left[\frac{1}{2}(0) + \frac{1}{2}(6) \right] + \frac{2}{3} \left[\frac{1}{2}(2) + \frac{1}{2}(1) \right] = 2$.

Thus both c and d are sequentially rational and so is any mixture of the two, in particular

the mixture $\begin{pmatrix} c & d \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$.

For Player 3 the possible payoffs are:

$$\begin{aligned} \text{from } e: & \quad \frac{1}{2}(2) + \frac{1}{2} \left[\frac{1}{2}(0) + \frac{1}{2}(4) \right] = 2 \\ \text{from } f: & \quad \frac{1}{2} \left[\frac{1}{2}(3) + \frac{1}{2}(1) \right] + \frac{1}{2} \left[\frac{1}{2}(1) + \frac{1}{2}(3) \right] = 2. \end{aligned}$$

Thus both e and f are sequentially rational and so is any mixture of the two, in particular the mixture $\begin{pmatrix} e & f \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$.

For Player 4 the possible payoffs are:

$$\begin{aligned} \text{from } g: & \quad \frac{2}{11}(2) + \frac{1}{11}(0) + \frac{2}{11}(2) + \frac{6}{11}(3) = \frac{26}{11}, \\ \text{from } h: & \quad \frac{2}{11}(3) + \frac{1}{11}(2) + \frac{2}{11}(3) + \frac{6}{11}(2) = \frac{26}{11}. \end{aligned}$$

Thus both g and h are rational and so is any mixture of the two, in particular the mixture

$$\begin{pmatrix} g & h \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}.$$

For AGM consistency, note that all of the actions, except a and b , are plausibility preserving and, furthermore, all decision histories are assigned positive probability by μ . Thus there is only one plausibility order that rationalizes the given assessment, namely the one that has only two equivalence classes: the top one being $\{\emptyset, s\}$ and the other one consisting of all the remaining histories.

For Bayesian consistency, the probability distribution for the top equivalence class is the trivial one that assigns probability 1 to the null history \emptyset .

Let $E = H \setminus \{\emptyset, s\}$ be the other equivalence class and note that $E \cap D_m^+ = \{a, ac, ad, acf, ade, adf, b\}$.

In order for a probability distribution ν_E to satisfy the properties of Definition 13.2.1, the support of ν_E must be the set $\{a, ac, ad, acf, ade, adf, b\}$ (Property B1).

Let $\nu_E(a) = p$; then by Property B2 it must be that

$$\begin{aligned} \nu_E(ac) &= \nu_E(ad) = \frac{p}{2} \quad (\text{since } \sigma(c) = \sigma(d) = \frac{1}{2}), \\ \nu_E(acf) &= \nu_E(adf) = \frac{p}{3} \quad (\text{since } \sigma(c) \times \sigma(f) = \sigma(d) \times \sigma(f) = \frac{1}{2} \times \frac{2}{3} = \frac{1}{3}) \text{ and} \\ \nu_E(ade) &= \frac{p}{6} \quad (\text{since } \sigma(d) \times \sigma(e) = \frac{1}{2} \times \frac{1}{3} = \frac{1}{6}). \end{aligned}$$

Thus ν_E must be of the form

$$\begin{pmatrix} a & ac & ad & acf & ade & adf & b \\ p & \frac{p}{2} & \frac{p}{2} & \frac{p}{3} & \frac{p}{6} & \frac{p}{3} & q \end{pmatrix}$$

with the sum equal to 1, that is, $\frac{17}{6}p + q = 1$. Furthermore, by Property B3, it must be that

$$\frac{\nu_E(b)}{\nu_E(acf) + \nu_E(ade) + \nu_E(adf) + \nu_E(b)} = \frac{q}{\frac{p}{3} + \frac{p}{6} + \frac{p}{3} + q} = \underbrace{\frac{6}{11}}_{=\mu(b)}.$$

The solution to the pair of equations $\frac{17}{6}p + q = 1$ and $\frac{q}{\frac{p}{3} + \frac{p}{6} + \frac{p}{3} + q} = \frac{6}{11}$ is $p = q = \frac{6}{23}$, yielding

$$v_E = \begin{pmatrix} a & ac & ad & acf & ade & adf & b \\ \frac{6}{23} & \frac{3}{23} & \frac{3}{23} & \frac{2}{23} & \frac{1}{23} & \frac{2}{23} & \frac{6}{23} \end{pmatrix}.$$

Thus we have shown that the given assessment is AGM-consistent, Bayes-consistent and sequentially rational, hence a perfect Bayesian equilibrium. \square

Solutions to Exercise 13.7 The extensive form is shown in Figure 13.22.




Figure 13.22: The extensive form for Exercise 13.7.

Let

$$\sigma = \left(\begin{array}{cc|ccc} a & b & c & d & e \\ 1 & 0 & 1 & 0 & 0 \end{array} \right) \quad \mu = \begin{pmatrix} a & b \\ 1 & 0 \end{pmatrix}.$$

This assessment is rationalized by the following plausibility order:

$$\left(\begin{array}{l} \text{most plausible} \\ \emptyset, a, ac \\ b, bc \\ ad \\ ae \\ be \\ \text{least plausible} \\ bd \end{array} \right)$$

which satisfies Property (IND_1) (since $a \succ b$ and $ac \succ bc$, $ad \succ bd$, $ae \succ be$) but fails Property (IND_2) since $b \in I(a)$ and $ad \succ ae$ (implying that – conditional on a – d is more plausible than e) but $be \succ bd$ (implying that – conditional on b – e is more plausible than d). \square

Solutions to Exercise 13.8 The game of Figure 13.5, reproduced in Figure 13.23, provides such an example.




Figure 13.23: Copy of Figure 13.5.

Consider the assessment $\sigma = (c, d, g)$ (highlighted by double edges), together with the system of beliefs $\mu = \left(\begin{array}{cc|cc} a & b & ae & be \\ 1 & 0 & 0 & 1 \end{array} \right)$.

This assessment is rationalized by the plausibility order

$$\left(\begin{array}{l} \text{most plausible} \\ \emptyset, c \\ a, ad \\ b, bd \\ be, beg \\ ae, aeg \\ bef \\ \text{least plausible} \\ aef \end{array} \right)$$
 which violates Property (IND_1) since $a \succ b$ and yet $be \succ ae$.

On the other hand, the above plausibility order satisfies Property (IND_2) since **(1)** $ad \succ ae$ and $bd \succ be$ and **(2)** $aeg \succ aef$ and $beg \succ bef$.

Furthermore, the above assessment trivially satisfies the additional property, since $a, be \in D_\mu^+$ but $ae, b \notin D_\mu^+$ (we have that $b \in I(a)$, $e \in A(a)$, $be \in I(ae)$ but it is not true that $\{a, b, ae, be\} \subseteq D_\mu^+ = \{\emptyset, a, be\}$). □

Solutions to Exercise 13.9 By Appealing to Theorem 13.5.1 it is sufficient to show that any plausibility order that rationalizes (σ, μ) cannot be choice measurable.

Let \succsim be a plausibility order that rationalizes (σ, μ) . Then

- since $\mu(ae) > 0$ while $\mu(af) = 0$ and ae and af belong to the same information set, by P2 of Definition 13.1.3

$$ae \succ af, \tag{13.7}$$

- since $\mu(bf) > 0$ while $\mu(be) = 0$ and bf and be belong to the same information set, by P2 of Definition 13.1.3

$$bf \succ be. \tag{13.8}$$

Let F be an integer-valued representation of \succsim that satisfies (CM) .

Then, since a and b belong to the same information set, it must be that

$$F(b) - F(a) = F(be) - F(ae) \text{ and } F(b) - F(a) = F(bf) - F(af),$$

so that $F(be) - F(ae) = F(bf) - F(af)$ and thus

$$F(af) - F(ae) = F(bf) - F(be). \tag{13.9}$$

However, by definition of integer-valued representation (Definition 13.5.1), from (13.7) we get that $F(af) - F(ae) > 0$ and from (13.8) we get that $F(bf) - F(be) < 0$, contradicting (13.9). \square

Solutions to Exercise 13.10 The extensive form under consideration is shown in Figure 13.24.

Consider an arbitrary assessment of the form (σ, μ) where

$\sigma = (a, g, r, \dots)$, $\mu(c) = \mu(e) = \mu(ch) = 0$ and $\mu(h) > 0$ for every other decision history h .

There are plausibility orders that rationalize (σ, μ) ; for example, either of the following:

$$\left(\begin{array}{l} \text{most plausible } \emptyset, a \\ b, bg, d, dr \\ bh, ds \\ c, cg, e, er \\ es \\ \text{least plausible } ch \end{array} \right) \quad \text{or} \quad \left(\begin{array}{l} \text{most plausible } \emptyset, a \\ b, bg, d, dr \\ c, cg, e, er \\ bh, ds \\ es \\ \text{least plausible } ch \end{array} \right).$$

First we show that any plausibility order \succsim that rationalizes (σ, μ) must satisfy the following properties:

$c \sim e$ because, by P2 of Definition 13.1.3, $cg \sim er$ (since $cg \in I(er)$, $\mu(cg) > 0$ and $\mu(er) > 0$) and, by P1 of Definition 13.1.3, $c \sim cg$ and $e \sim er$ (since both g and r are plausibility preserving) and thus, by transitivity, $c \sim e$,




Figure 13.24: The extensive form for Exercise 13.10.

$es \succ ch$ this follows from P2 of Definition 13.1.3, since es and ch belong to the same information set and $\mu(es) > 0$, while $\mu(ch) = 0$,

$b \sim d$ because, by P2 of Definition 13.1.3, $bg \sim dr$ (since $bg \in I(dr)$, $\mu(bg) > 0$ and $\mu(dr) > 0$) and, by P1 of Definition 13.1.3, $b \sim bg$ and $d \sim dr$ (since both g and r are plausibility preserving) and thus, by transitivity, $b \sim d$,

$bh \sim ds$ by P2 of Definition 13.1.3, because $bh \in I(ds)$, $\mu(bh) > 0$ and $\mu(ds) > 0$.

Next we show that no plausibility order that rationalizes (σ, μ) is choice measurable. Select an arbitrary plausibility order that rationalizes (σ, μ) and let F be an integer valued representation of it.

Then the following must be true:

1. $F(e) - F(es) > F(c) - F(ch)$ (because $c \sim e$, implying that $F(c) = F(e)$, and $es \succ ch$, implying that $F(es) < F(ch)$),
2. $F(b) - F(bh) = F(d) - F(ds)$ (because $bh \sim ds$, implying that $F(bh) = F(ds)$, and $b \sim d$, implying that $F(b) = F(d)$).

Thus if, as required by choice measurability, $F(c) - F(ch) = F(b) - F(bh)$ then, by Points 1 and 2 above, $F(e) - F(es) > F(d) - F(ds)$, which violates choice measurability. It follows from Theorem 13.5.1 that, since any plausibility ordering that rationalizes (σ, μ) is not choice measurable, (σ, μ) cannot be a sequential equilibrium. \square

Solutions to Exercise 13.11 The game under consideration is shown in Figure 13.25.




Figure 13.25: The game for Exercise 13.11.

Let (σ, μ) be an assessment with $\sigma = (a, T, f, L)$, $\mu(b) > 0$ and $\mu(c) > 0$.

- (a) We have to prove that (σ, μ) can be rationalized by a choice-measurable plausibility order only if μ satisfies the following condition:

$$\mu(bB) > 0 \quad \text{if and only if} \quad \mu(cBf) > 0.$$

Let \succsim be a choice measurable plausibility order that rationalizes (σ, μ) and let F be an integer-valued representation of \succsim that satisfies choice measurability. Since $\mu(b) > 0$ and $\mu(c) > 0$, by P2 of Definition 13.1.3, $b \sim c$ and thus $F(b) = F(c)$; by choice measurability, $F(b) - F(c) = F(bB) - F(cB)$ and thus $F(bB) = F(cB)$, so that $bB \sim cB$.

Since $\sigma(f) > 0$, by P1 of Definition 13.1.3, $cB \sim cBf$ and therefore, by transitivity of \succsim , $bB \sim cBf$.

Hence if $\mu(bB) > 0$ then, by P2 of Definition 13.1.3, $bB \succsim h$ for every $h \in \{bB, cBf, d\}$ and thus (since $cBf \sim bB$) $cBf \succsim h$ for every $h \in \{bB, cBf, d\}$ so that, by P2 of Definition 13.1.3, $\mu(cBf) > 0$.

The proof that if $\mu(cBf) > 0$ then $\mu(bB) > 0$ is analogous.

- (b) We have to prove that if, besides from being rationalized by a choice-measurable plausibility order \succsim , (σ, μ) is also uniformly Bayesian relative to \succsim (Definition 13.4.1), then μ satisfies the following condition:

$$\text{if } \mu(bB) > 0 \text{ then } \frac{\mu(cBf)}{\mu(bB)} = \frac{\mu(c)}{\mu(b)}.$$

Suppose that $\mu(b) > 0$, $\mu(c) > 0$ (so that $b \sim c$) and $\mu(bB) > 0$.

Let ν be a full-support common prior that satisfies the properties of Definition 13.4.1.

Then, by UB2, $\frac{\nu(c)}{\nu(b)} = \frac{\nu(cB)}{\nu(bB)}$ and, by UB1, since $\sigma(f) = 1$, $\nu(cBf) = \nu(cB) \times \sigma(f) = \nu(cB)$.

Let E be the equivalence class that contains b .

Then $E \cap D_\mu^+ = \{b, c\}$.

Since $\nu_E(\cdot) = \nu(\cdot \mid E \cap D_\mu^+)$, by B3 of Definition 13.2.1, $\mu(b) = \frac{\nu(b)}{\nu(b)+\nu(c)}$ and $\mu(c) = \frac{\nu(c)}{\nu(b)+\nu(c)}$, so that $\frac{\mu(c)}{\mu(b)} = \frac{\nu(c)}{\nu(b)}$.

Let G be the equivalence class that contains bB .

Then, since – by hypothesis – $\mu(bB) > 0$, it follows from the condition proved in Part (a) that either $G \cap D_\mu^+ = \{bB, cBf\}$ or $G \cap D_\mu^+ = \{bB, cBf, d\}$.

Since $\nu_G(\cdot) = \nu(\cdot \mid G \cap D_\mu^+)$, by B3 of Definition 13.2.1,

in the former case $\mu(bB) = \frac{\nu(bB)}{\nu(bB)+\nu(cBf)}$ and $\mu(cBf) = \frac{\nu(cBf)}{\nu(bB)+\nu(cBf)}$ and in the latter case $\mu(bB) = \frac{\nu(bB)}{\nu(bB)+\nu(cBf)+\nu(d)}$ and $\mu(cBf) = \frac{\nu(cBf)}{\nu(bB)+\nu(cBf)+\nu(d)}$; thus in both cases $\frac{\mu(cBf)}{\mu(bB)} = \frac{\nu(cBf)}{\nu(bB)}$.

Hence, since $\nu(cBf) = \nu(cB)$, $\frac{\mu(cBf)}{\mu(bB)} = \frac{\nu(cB)}{\nu(bB)}$ and, therefore, since (as shown above) $\frac{\nu(cB)}{\nu(bB)} = \frac{\nu(c)}{\nu(b)}$ and $\frac{\nu(c)}{\nu(b)} = \frac{\mu(c)}{\mu(b)}$, we have that $\frac{\mu(cBf)}{\mu(bB)} = \frac{\mu(c)}{\mu(b)}$. \square

Solutions to Exercise 13.12 The game under consideration is reproduced in Figure 13.26.




Figure 13.26: The game for Exercise 13.12.

We have to show that the assessment (σ, μ) with

$$\sigma = (a, T, f, L) \quad \mu = \left(\begin{array}{cc|ccc} b & c & bB & cBf & d \\ \frac{7}{10} & \frac{3}{10} & \frac{7}{18} & \frac{3}{18} & \frac{8}{18} \end{array} \right)$$

is a sequential equilibrium. By Theorem 13.5.1 it is sufficient to show that (σ, μ) is sequentially rational, is rationalized by a choice measurable plausibility order and is uniformly Bayesian relative to it.

First we verify sequential rationality.

- For Player 1 at the root the possible payoffs are: 1 with a , 0 with b , 0 with c (more precisely, with either ce or cf) and 0 with d ; thus a is sequentially rational.
- For Player 1 at history cB the possible payoffs are: 1 with e and 2 with f ; thus f is sequentially rational.

- For Player 2 the possible payoffs are:
 $\frac{7}{10}(3) + \frac{3}{10}(0) = \frac{21}{10}$ with T
 and $\frac{7}{10}(0) + \frac{3}{10}(3) = \frac{9}{10}$ with B ;
 thus T is sequentially rational.
- For Player 3 the possible payoffs are:
 $\frac{7}{18}(1) + \frac{3}{18}(0) + \frac{8}{18}(1) = \frac{15}{18}$ with L
 and $\frac{7}{18}(0) + \frac{3}{18}(1) + \frac{8}{18}(0) = \frac{3}{18}$ with R ;
 thus L is sequentially rational.

Next we show that (σ, μ) is rationalized by a choice-measurable plausibility order. It is straightforward to verify that (σ, μ) is rationalized by the plausibility order shown below together with a choice-measurable integer-valued representation F :

\approx :		F :	
most plausible	\emptyset, a		0
	b, c, bT, cT		1
	$d, bB, cB, cBf, dL, bBL, cBfL$		2
least plausible	$dR, bBR, cBfR, cBe$		3

To see that (σ, μ) is uniformly Bayesian relative to this plausibility order, let E_1, E_2 and E_3 be the top three equivalence classes of the order and consider the following probability distributions, which satisfy the properties of Definition 13.2.1 (so that (σ, μ) is Bayesian relative to this plausibility order):

$$v_{E_1}(\emptyset) = 1, \quad v_{E_2} = \begin{pmatrix} b & c \\ \frac{7}{10} & \frac{3}{10} \end{pmatrix} \quad \text{and} \quad v_{E_3} = \begin{pmatrix} bB & cB & cBf & d \\ \frac{7}{21} & \frac{3}{21} & \frac{3}{21} & \frac{8}{21} \end{pmatrix}.$$

A full support common prior that satisfies the properties of Definition 13.4.1 is the following:

$$v = \begin{pmatrix} \emptyset & b & bB & c & cB & cBf & d \\ \frac{9}{40} & \frac{7}{40} & \frac{7}{40} & \frac{3}{40} & \frac{3}{40} & \frac{3}{40} & \frac{8}{40} \end{pmatrix}.$$

Solutions to Exercise 13.13 The extensive form under consideration is shown in Figure 13.27.

Let $\sigma = (M, \ell, a, c, e)$, $\mu = \begin{pmatrix} L & M & R & | & Ll & Mr & | & Lm & Rr & | & Mm & Rl \\ 0 & 1 & 0 & | & 0 & 1 & | & 1 & 0 & | & 0 & 1 \end{pmatrix}$.




Figure 13.27: The extensive form for Exercise 13.13.

- (a) We have to show that (σ, μ) is a strongly independent perfect Bayesian equilibrium. Sequential rationality is straightforward to verify. It is also straightforward to verify that the following plausibility order rationalizes (σ, μ) :

$$\left(\begin{array}{l} \text{most plausible} \\ \emptyset, M, M\ell \\ R, R\ell, R\ell e \\ Mm, Mme \\ Mr, Mra \\ L, L\ell, L\ell a \\ Rm \\ Lm, Lmc \\ Rr, Rrc \\ Lr \\ R\ell f \\ Mmf \\ Lmd \\ Rrd \\ Mrb \\ L\ell b \\ \text{least plausible} \end{array} \right)$$

Let us check that the above plausibility order satisfies properties (IND_1) and (IND_2) . For (IND_1) first note that there are no two decision histories h and h' that belong to the same information set and are such that $h \sim h'$; thus we only need to check that if h and h' belong to the same information set and $h \succ h'$ then $ha \succ h'a$ for every $a \in A(h)$. This is indeed true:

1. $M \succ R$ and $M\ell \succ R\ell$, $Mm \succ Rm$ and $Mr \succ Rr$,
2. $M \succ L$ and $M\ell \succ L\ell$, $Mm \succ Lm$ and $Mr \succ Lr$,
3. $R \succ L$ and $R\ell \succ L\ell$, $Rm \succ Lm$ and $Rr \succ Lr$,
4. $Mr \succ L\ell$ and $Mra \succ L\ell a$ and $Mrb \succ L\ell b$,
5. $Lm \succ Rr$ and $Lmc \succ Rrc$ and $Lmd \succ Rrd$,
6. $R\ell \succ Mm$ and $R\ell e \succ Mme$ and $R\ell f \succ Mmf$.

For (IND_2) first note that there is no decision history h that $ha \sim hb$ for $a, b \in A(h)$ with $a \neq b$; thus we only need to show that if $ha \succ hb$ and $h' \in I(h)$ then $h'a \succ h'b$. This is indeed true:

1. $M\ell \succ Mm$ and $L\ell \succ Lm$ and $R\ell \succ Rm$,
2. $Mm \succ Mr$ and $Lm \succ Lr$ and $Rm \succ Rr$,
3. $M\ell \succ Mr$ $L\ell \succ Lr$ and $R\ell \succ Rr$,

and the rest is trivial, since at the other information sets there are only two actions, one of which is plausibility preserving and the other is not.

It remains to show that (σ, μ) is uniformly Bayesian relative to the plausibility order given above. First of all, note that, for every equivalence class E of the order, $E \cap D_\mu^+$ is either empty or a singleton. Thus as a full support common prior one can take, for example, the uniform distribution over the set of decision histories: $v(h) = \frac{1}{10}$ for every $h \in \{\emptyset, L, M, R, Ll, Mr, Lm, Rr, Mm, Rl\}$.

(b) To prove that (σ, μ) is not a sequential equilibrium it would not be sufficient to prove that the plausibility order given above is not choice measurable (although it is indeed true that it is not choice measurable), because in principle there could be another plausibility order which is choice measurable and rationalizes (σ, μ) . Thus we need to show that *any* plausibility order that rationalizes (σ, μ) is not choice measurable. Let \succsim be a plausibility order that rationalizes (σ, μ) ; then it must satisfy the following properties:

- $Lm \succ Rr$ (because they belong to the same information set and $\mu(Lm) > 0$ while $\mu(Rr) = 0$). Thus if F is any integer-valued representation of \succsim it must be that

$$F(Lm) < F(Rr). \quad (13.10)$$

- $Mr \succ Ll \sim L$ ($Mr \succ Ll$ because Mr and Ll belong to the same information set and $\mu(Mr) > 0$ while $\mu(Ll) = 0$; $Ll \sim L$ because l is a plausibility-preserving action since $\sigma(l) > 0$). Thus if F is any integer-valued representation of \succsim it must be that

$$F(Mr) < F(L). \quad (13.11)$$

- $R \sim Rl \succ Mm$ ($R \sim Rl$ because l is a plausibility-preserving action; $Rl \succ Mm$ because Rl and Mm belong to the same information set and $\mu(Rl) > 0$ while $\mu(Mm) = 0$). Thus if F is any integer-valued representation of \succsim it must be that

$$F(R) < F(Mm). \quad (13.12)$$

Now suppose that \succsim is choice measurable and let F be an integer-valued representation of it that satisfies choice measurability. From (13.10) and (13.11) we get that

$$F(Lm) - F(L) < F(Rr) - F(Mr) \quad (13.13)$$

and by choice measurability

$$F(Rr) - F(Mr) = F(R) - F(M). \quad (13.14)$$

It follows from (13.13) and (13.14) that

$$F(Lm) - F(L) < F(R) - F(M). \quad (13.15)$$

Subtracting $F(M)$ from both sides of (13.12) we obtain

$$F(R) - F(M) < F(Mm) - F(M). \quad (13.16)$$

It follows from (13.15) and (13.16) that $F(Lm) - F(L) < F(Mm) - F(M)$, which can be written as $F(M) - F(L) < F(Mm) - F(Lm)$, yielding a contradiction, because choice measurability requires that $F(M) - F(L) = F(Mm) - F(Lm)$. \square



Incomplete Information

14 Static Games	483
14.1 Interactive situations with incomplete information	
14.2 One-sided complete information	
14.3 Two-sided incomplete information	
14.4 Multi-sided incomplete information	
14.5 Exercises	
14.6 Solutions to Exercises	
15 Dynamic Games	521
15.1 One-sided incomplete information	
15.2 Multi-sided incomplete information	
15.3 Exercises	
15.4 Solutions to Exercises	
16 The Type-Space Approach	567
16.1 Types of players	
16.2 Types that know their payoffs	
16.3 The general case	
16.4 Exercises	
16.5 Solutions to Exercises	
Index	589

14. Static Games

14.1 Interactive situations with incomplete information

An implicit assumption in game theory is that the game being played is common knowledge among the players. The expression “incomplete information” refers to situations where some of the elements of the game (e.g. the actions available to the players, the possible outcomes, the players’ preferences, etc.) are *not* common knowledge. In such situations the knowledge and beliefs of the players about the game need to be made an integral part of the description of the situation. Pioneering work in this direction was done by John Harsanyi (1967, 1968), who was the recipient of the 1994 Nobel Memorial prize in economics (together with John Nash and Reinhard Selten). Harsanyi suggested a method for converting a situation of incomplete information into an extensive game with imperfect information (this is the so-called *Harsanyi transformation*). The theory of games of incomplete information has been developed for the case of von Neumann-Morgenstern payoffs and the solution concept proposed by Harsanyi is *Bayesian Nash equilibrium*, which is merely Nash equilibrium of the imperfect-information game that is obtained by applying the Harsanyi transformation.

Although the approach put forward by Harsanyi was in terms of “types” of players and of probability distributions over types, we shall develop the theory using the so-called “state-space” approach, which makes use of the interactive knowledge-belief structures developed in Chapters 8 and 9. In Chapter 16 we will explain the “type-space” approach and show how to convert one type of structure into the other.

The distinction between imperfect and incomplete information is not at all the same as that between perfect and imperfect information. To say that a game (in extensive form) has *imperfect* information is to say that there is at least one player who may have to make a choice in a situation where she does not know what choices were made previously by other players. To say that there is *incomplete* information is to say that there is at least one player who does not quite know what game she is playing.

Consider, for example, the players' preferences. In some situations it is not overly restrictive to assume that the players' preferences are common knowledge: for instance, in the game of chess, it seems quite plausible to assume that it is common knowledge that both players prefer winning to either drawing or losing (and prefer drawing to losing).

But think of a contractual dispute. Somebody claims that you owe him a sum of money and threatens to sue you if you don't pay. We can view this situation as a two-player game: you have two strategies, "pay" and "not pay", and he has two strategies, "sue (if no payment)" and "not sue (if no payment)". In order to determine your best choice, you need to know how he will respond to your refusal to pay. A lawsuit is costly for him as well as for you; if he is the "aggressive" type he will sue you; if he is the "not aggressive" type he will drop the dispute. If you don't know what type he is, you are in a situation of incomplete information.

As explained above, we have a situation of incomplete information whenever at least one of the players does not know some aspects of the game: it could be the preferences of her opponents, the choices available to her opponents or the set of possible outcomes. We shall focus almost exclusively on the case where the uncertainty concerns the preferences of the players, while everything else will be assumed to be common knowledge (for an exception see Exercise 14.2). Harsanyi argued that every situation can be reduced to this case. For example, he argued that if Player 1 does not know whether Player 2 has available only choices a and b or also choice c , we can model this as a situation where there are two possible "states of the world" in both of which Player 2 has three choices available, a , b and c , but in one of the two states choice c gives an extremely low payoff to Player 2, so that she would definitely not choose c .

The interactive knowledge-belief structures developed in Chapters 8 and 9 are sufficiently rich to model situations of incomplete information. The states in these structures can be used to express any type of uncertainty. In Chapter 10 we interpreted the states in terms of the actual choices of the players, thus representing uncertainty in the mind of a player about the behavior of another player. In that case the game was assumed to be common knowledge among the players (that is, it was the same game at every state) and what varied from one state to another was the choice of at least one player. If we want to represent a situation where one player is not sure what game she is playing, all we have to do is interpret the states in terms of games, that is, assign different games to different states.

We shall begin with games in strategic form with cardinal payoffs where only one player is uncertain about what game is being played. We use the expression "one-sided incomplete information" to refer to these situations.

14.2 One-sided complete information

Let us model the following situation: the game being played is the one shown in Figure 14.1; call this the “true” game.

		Player 2			
		L	R		
Player 1	T	6	3	0	9
	B	3	3	3	0

Figure 14.1: The “true” game.

Player 1 knows that she is playing this game, while Player 2 is uncertain as to whether she is playing this game or a different game, shown in Figure 14.2, where the payoffs of Player 1 are different.

		Player 2			
		L	R		
Player 1	T	0	3	3	9
	B	3	3	0	0

Figure 14.2: The alternative game in Player 2’s mind.

For convenience, let us refer to the “true game” of Figure 14.1 as the game where Player 1 is of type b and the game of Figure 14.2 as the game where Player 1 is of type a . Suppose that Player 2 assigns probability $\frac{2}{3}$ to Player 1 being of type a and probability $\frac{1}{3}$ to Player 1 being of type b .

The description is not complete yet because we need to specify whether Player 1 has any uncertainty concerning the beliefs of Player 2; let us assume that the beliefs of Player 2 are common knowledge. We also need to specify whether Player 2 is uncertain about the state of mind of Player 1; let us assume that it is common knowledge that Player 2 knows that Player 1 knows what game is being played. This is a long verbal description! A picture is worth a thousand words and indeed the simple knowledge-belief structure shown in Figure 14.3 captures all of the above.

There are two states, α and β : α is interpreted as a (counterfactual) state where the game of Figure 14.2 is played, while β is interpreted as the state where the “true” game of Figure 14.1 is played. We capture the fact that latter is the “true” game, that is, the game which is actually played, by designating state β as “the true state”.¹

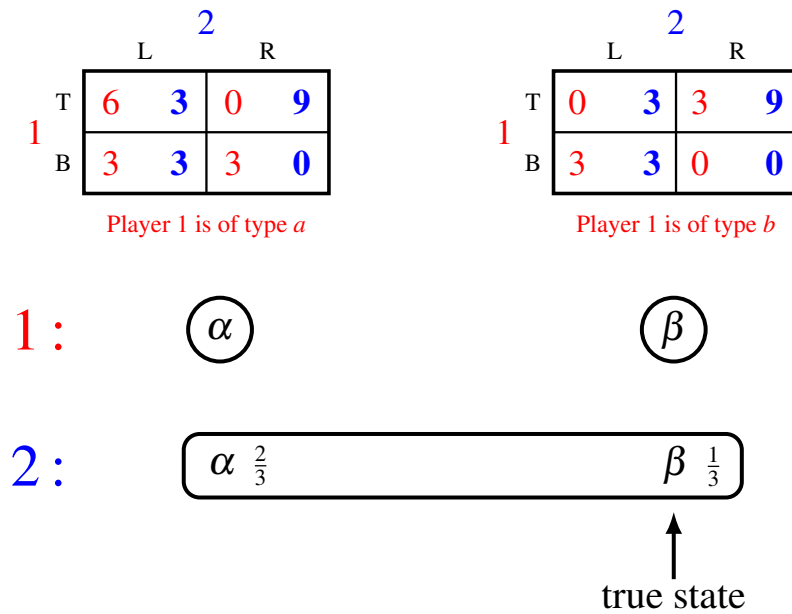


Figure 14.3: A situation of one-sided incomplete information

It is easy to check that, in the knowledge-belief structure of Figure 14.3, at state β all of the elements of the verbal description given above are true. Anything that is constant across states is common knowledge: thus (i) the payoffs of Player 2, (ii) the beliefs of Player 2, namely $\left(\begin{matrix} \alpha & \beta \\ \frac{2}{3} & \frac{1}{3} \end{matrix} \right)$, and (iii) the fact that Player 1 knows what game is being played, that is, that

- if the game being played is the one of Figure 14.2 (associated with state α) then Player 1 knows that they are playing that game and
- if the game being played is the one of Figure 14.1 (associated with state β) then Player 1 knows that they are playing that game.

Note that, at the true state β , Player 1 knows more than Player 2, namely which of the two games is actually being played.

R In the situation illustrated in Figure 14.3, at every state each player knows his/her own payoffs. It may seem natural to make this a requirement of rationality: shouldn't a rational player know her own payoffs? The answer is: Yes for preferences and No for payoffs. A rational player should know how she ranks the possible outcomes (her preferences over outcomes); however, it is perfectly rational to be uncertain about

¹This is something that is almost never done in the literature, but it is an important element of a description of a situation of incomplete information: what is the actual state of affairs?

what outcome will follow a particular action and thus about one's own payoff. For example, you know that you prefer accepting a wrapped box containing a thoughtful gift (payoff of 1) to accepting a wrapped box containing an insulting gift (payoff of -1): you know your preferences over these two outcomes. However, you may be uncertain about the intentions of the gift-giver and thus you are uncertain about what outcome will follow if you accept the gift: you don't know if your payoff will be 1 or -1 . Thus, although you know your payoff *function*, you may be uncertain about what your *payoff* will be if you accept the gift. This example is developed in Exercise 14.2.

Figure 14.3 illustrates a situation that has to do with games, but is not a game. Harsanyi's insight was that we can transform that situation into an extensive-form game with imperfect information, as follows. We start with a chance move, where Nature chooses the state; then Player 1 is informed of Nature's choice and makes her decision between T and B ; Player 2 then makes his choice between L and R , without being informed of Nature's choice (to capture her uncertainty about the game) and without being informed of Player 1's choice (to capture the simultaneity of the game). The game is shown in Figure 14.4.




Figure 14.4: The extensive-form game obtained from the incomplete-information situation of Figure 14.3 by applying the Harsanyi transformation.

The reduction of the situation illustrated in Figure 14.3 to the extensive-form game of Figure 14.4 is called the *Harsanyi transformation*. Once a game has been obtained by means of the Harsanyi transformation, it can be solved using, for example, the notion of Nash equilibrium. The Nash equilibria of a game corresponding to a situation of incomplete information are called *Bayesian Nash equilibria*.

Note that these are nothing more than Nash equilibria: the extra term 'Bayesian' is merely a hint that the game being solved is meant to represent a situation of incomplete information. Note also that the Harsanyi transformation involves some loss of information: in particular, in the resulting game one can no longer tell what the true state is, that is, what the actual game being played is.

To find the Bayesian Nash equilibria of the game of Figure 14.4 we can construct the corresponding strategic-form game, which is shown in Figure 14.5.²

		Player 2	
		L	R
Player 1	T if type <i>a</i> T if type <i>b</i>	2 3	2 9
	T if type <i>a</i> B if type <i>b</i>	1 3	3 6
	B if type <i>a</i> T if type <i>b</i>	4 3	0 3
	B if type <i>a</i> B if type <i>b</i>	3 3	1 0

Figure 14.5: The strategic-form game corresponding to the game of Figure 14.4.

There are two pure-strategy Bayesian Nash equilibria: $((T, B), R)$ and $((B, T), L)$ (where (T, B) means “*T* if type *a* and *B* if type *b*” and (B, T) means “*B* if type *a* and *T* if type *b*”). How should we interpret them?

Let us consider, for example, the Bayesian Nash equilibrium $((T, B), R)$. This is a Nash equilibrium of the game of Figure 14.4; however, we assumed that the true state was β , where Player 1 is of type *b* and thus, in the actual game (the one associated with state β), the actual play is (B, R) , which is not a Nash equilibrium of that game (because, while *B* is a best reply to *R*, *R* is not a best reply to *B*).

This is not surprising, since Player 1 knows that she is playing that game, while Player 2 attaches only probability $\frac{1}{3}$ to that game being the actual game. Thus the first observation is that a Bayesian Nash equilibrium of a “game of incomplete information” does not imply that the players play a Nash equilibrium in the actual (or “true”) game.

The second observation has to do with Player 1’s strategy. By definition, Player 1’s strategy in the game of Figure 14.4 consists of a pair of choices, one for the case where she is informed that her type is *a* and the other for the case where she is informed that her type is *b*: in the Bayesian Nash equilibrium under consideration, Player 1’s strategy (T, B) means “Player 1 plays *T* if of type *a* and plays *B* if of type *b*”. However, Player 1 knows that she is of type *b*: she knows what game she is playing.

²Recall that, when constructing the strategic-form game associated with an extensive-form game with chance moves, payoffs are *expected* payoffs: see Chapter 7.

Why, then, should she formulate a plan on how she would play in a counterfactual world where her type was a (that is, in a counterfactual game)? The answer is that Player 1's strategy (T, B) is best understood not as a contingent plan of action formulated by Player 1, but as a complex object that incorporates (1) the actual choice (namely to play B) made by Player 1 in the game that she knows she is playing and (2) a belief in the mind of Player 2 about what Player 1 would do in the two games that, as far as Player 2 knows, are actual possibilities.

An alternative (and easier) way to find the Bayesian Nash equilibria of the game of Figure 14.4 is to use the notion of weak sequential equilibrium for that game.

For example, to verify that $((B, T), L)$ is a Nash equilibrium we can reason as follows: If the strategy of Player 1 is (B, T) , then Bayesian updating requires Player 2 to have the following beliefs: probability $\frac{2}{3}$ on the second node from the left and probability $\frac{1}{3}$ on the third node from the left.

Given these beliefs, playing L yields him an expected payoff of $\frac{2}{3}(3) + \frac{1}{3}(3) = 3$ while playing R yields him an expected payoff of $\frac{2}{3}(0) + \frac{1}{3}(9) = 3$; thus any strategy (pure or mixed) is sequentially rational: in particular L is sequentially rational.

If Player 2 plays L then at her left node Player 1 gets 0 with T and 3 with B , so that B is sequentially rational; at her right node Player 1 gets 6 with T and 3 with B , so that T is sequentially rational. Hence $((B, T), L)$ with the stated beliefs is a weak sequential equilibrium, implying that $((B, T), L)$ is a Nash equilibrium (Theorem 11.3.1, Chapter 11).

Definition 14.2.1 In a “game of one-sided incomplete information” a pure-strategy Bayesian Nash equilibrium where the “informed” player makes the same choice at every singleton node is called a *pooling equilibrium*, while a pure-strategy Bayesian Nash equilibrium where the “informed” player makes different choices at different nodes is called a *separating equilibrium*.

Thus in the game of Figure 14.4 there are no pooling equilibria: all the pure-strategy Bayesian Nash equilibria are separating equilibria.

R Although we have only looked at the case of two players, situations of one-sided incomplete information can involve any number of players, as long as only one player is uncertain about the game being played (while all the other players are not). An example of a situation of one-sided incomplete information with three players will be given in the next chapter with reference to Selten's chain-store game (which we studied in Chapter 3, Section 3.4).

We conclude this section with an example of a situation of one-sided incomplete information involving a two-player game where each player has an infinite number of strategies. This is an incomplete-information version of Cournot's game of competition between two firms (Chapter 2, Section 2.7). This example involves the use of calculus and some readers might want to skip it.

Consider a Cournot duopoly (that is, an industry consisting of two firms, which compete in output levels) with inverse demand given by $P(Q) = 34 - Q$, where $Q = q_1 + q_2$ is total industry output (q_1 is the output of Firm 1 and q_2 is the output of Firm 2).

It is common knowledge between the two firms that Firm 1's cost function is given by: $C_1(q_1) = 6q_1$. Firm 2's cost function is $C_2(q_2) = 9q_2$. Firm 2 knows this, while Firm 1 believes that Firm 2's cost function is $C_2(q_2) = 9q_2$ with probability $\frac{1}{3}$ and $C_2(q_2) = 3q_2$ with probability $\frac{2}{3}$ (Firm 2 could be a new entrant to the industry using an old technology, or could have just invented a new technology).

Firm 1 knows that Firm 2 knows its own cost function. Everything that Firm 1 knows is common knowledge between the two firms. Thus we have the situation of one-sided incomplete information illustrated in Figure 14.6.

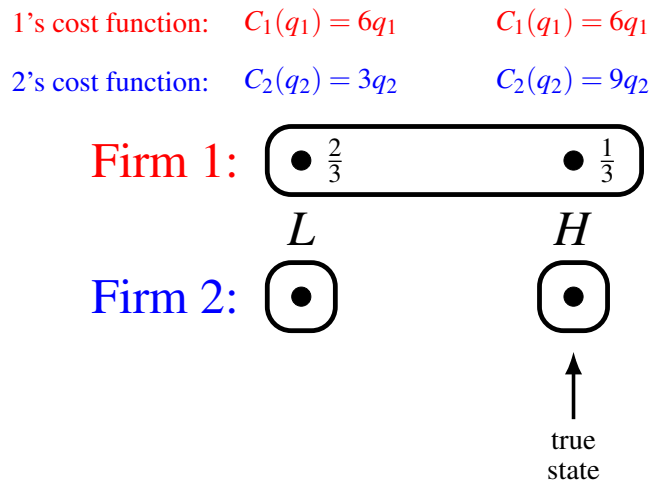


Figure 14.6: A one-sided incomplete information situation involving two firms.

Let us find a Bayesian Nash equilibrium of the game obtained by applying the Harsanyi transformation to this situation of incomplete information and compare it to the Nash equilibrium of the complete information game where Firm 2's cost function is common knowledge.

In the complete information case where it is common knowledge that Firm 2's cost function is $C_2(q_2) = 9q_2$, the profit (= payoff) functions of the two firms are given by:

$$\pi_1(q_1, q_2) = [34 - (q_1 + q_2)] q_1 - 6q_1$$

$$\pi_2(q_1, q_2) = [34 - (q_1 + q_2)] q_2 - 9q_2$$

The Nash equilibrium is given by the solution to the following pair of equations:

$$\frac{\partial \pi_1(q_1, q_2)}{\partial q_1} = 34 - 2q_1 - q_2 - 6 = 0$$

$$\frac{\partial \pi_2(q_1, q_2)}{\partial q_2} = 34 - q_1 - 2q_2 - 9 = 0$$

The solution is $q_1 = \frac{31}{3} = 10.33$ and $q_2 = \frac{22}{3} = 7.33$.

Next we go back to the incomplete-information situation illustrated in Figure 14.6. Although it is not possible to draw the extensive-form game that results from applying the Harsanyi transformation (because of the infinite number of possible output levels) we can nevertheless sketch the game as shown in Figure 14.7 (where H means “high cost” and L “low cost”):

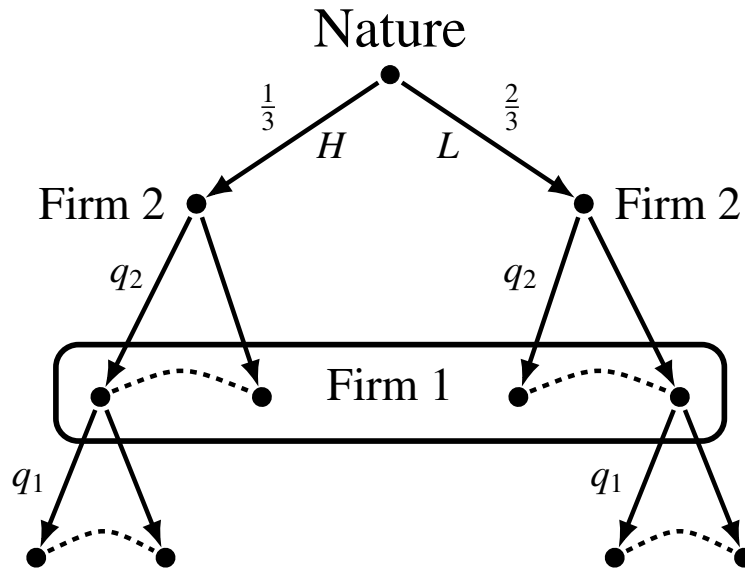


Figure 14.7: Sketch of the infinite extensive-form game obtained by applying the Harsanyi transformation to the incomplete-information situation of Figure 14.6.

To find a Bayesian Nash equilibrium it is easiest to think in terms of weak sequential equilibrium: if the strategy profile $((\hat{q}_2^H, \hat{q}_2^L), \hat{q}_1)$ is part of a weak sequential equilibrium, then – by sequential rationality – \hat{q}_2^H must maximize the expression

$$[(34 - \hat{q}_1 - q_2^H)q_2^H - 9q_2^H]$$

and \hat{q}_2^L must maximize the expression

$$[(34 - \hat{q}_1 - q_2^L)q_2^L - 3q_2^L].$$

Furthermore, by Bayesian updating, Firm 1 must assign probability $\frac{1}{3}$ to the node following choice \hat{q}_2^H and probability $\frac{2}{3}$ to the node following choice \hat{q}_2^L so that \hat{q}_1 must maximize the expression

$$\frac{1}{3} [(34 - q_1 - \hat{q}_2^H)q_1 - 6q_1] + \frac{2}{3} [(34 - q_1 - \hat{q}_2^L)q_1 - 6q_1].$$

There is a unique solution given by $\hat{q}_1 = 9$, $\hat{q}_2^H = 8$, $\hat{q}_2^L = 11$. Since the true state is where Firm 2’s cost is high, the actual output levels are $\hat{q}_1 = 9$ for Firm 1 and $\hat{q}_2^H = 8$ for Firm 2. Thus in the incomplete information case Firm 2’s output is higher than in the complete information case and Firm 1’s output is lower than in the complete information case.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 14.5.2 at the end of this chapter.

14.3 Two-sided incomplete information

Let us now consider a situation involving two players, both of whom face some uncertainty. Such a situation is called a *two-sided incomplete-information situation*. Note that for such a situation to arise, it is not necessary that both players are uncertain about some “objective” aspect of the game (such as the preferences of the opponent): one of the two players might simply be uncertain about the beliefs of the other player.

A situation of this sort is illustrated in Figure 14.8.

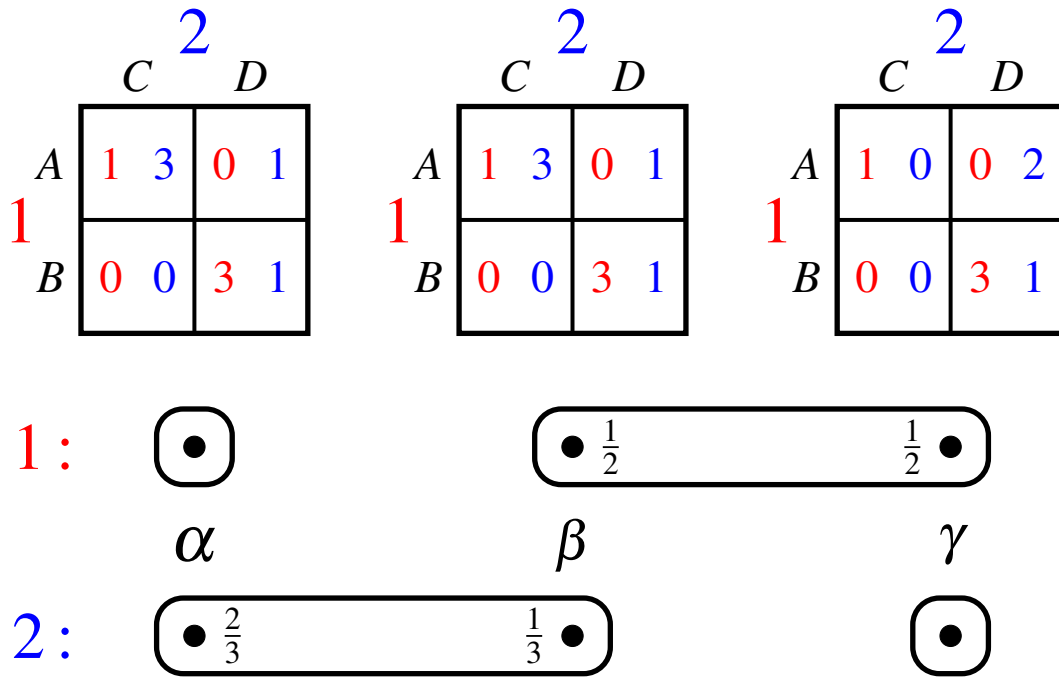


Figure 14.8: A two-sided incomplete-information situation.

Let G be the game associated with states α and β (it is the same game) and G' the game associated with state γ .

Suppose that the true state is α . Then all of the following are true at state α :

- (a) both Player 1 and Player 2 know that they are playing game G (that is, neither player has any uncertainty about the objective aspects of the game),
- (b) Player 1 knows that Player 2 knows that they are playing game G ,
- (c) Player 2 is uncertain as to whether Player 1 knows that they are playing G (which is the case if the actual state is α) or whether Player 1 is uncertain (if the actual state is β) between the possibility that they are playing game G and the possibility that they are playing game G' and considers the two possibilities equally likely; furthermore, Player 2 attaches probability $\frac{2}{3}$ to the first case (where Player 1 knows that they are playing game G) and probability $\frac{1}{3}$ to the second case (where Player 1 is uncertain between game G and game G'),

- (d) Player 1 knows the state of uncertainty of Player 2 (concerning Player 1, as described in part (c) above),
- (e) The payoffs of Player 1 are common knowledge; furthermore, it is common knowledge that Player 2 knows his own payoffs.

In principle, the Harsanyi transformation can be applied also to situations of two-sided incomplete information. However, in such cases there is an issue concerning the probabilities with which Nature chooses the states.

In the case of one-sided incomplete information we can take Nature's probabilities to be the beliefs of the uninformed player, but in the case of two-sided incomplete information we have two uninformed players and thus two sets of beliefs.

For example, if we look at state β in Figure 14.8, we have two different probabilities assigned to that state: $\frac{1}{2}$ by Player 1 and $\frac{1}{3}$ by Player 2. Which of the two should we take as Nature's probability for state β ? The answer is: neither of them. What we should take as Nature's probabilities is a probability distribution over the set of states $\{\alpha, \beta, \gamma\}$ that reflects the beliefs of *both* players. We have encountered such a notion before, in Chapter 9: it is the notion of a *common prior* (Definition 9.5.1). A *common prior* is a probability distribution over the set of states that yields the players' beliefs upon conditioning on the information represented by a cell of an information partition.

For example, in the situation illustrated in Figure 14.8 we are seeking a probability distribution

$$v : \{\alpha, \beta, \gamma\} \rightarrow [0, 1]$$

such that

$$v(\beta|\{\beta, \gamma\}) = \frac{v(\beta)}{v(\beta) + v(\gamma)} = \frac{1}{2},$$

$$v(\alpha|\{\alpha, \beta\}) = \frac{v(\alpha)}{v(\alpha) + v(\beta)} = \frac{2}{3}.$$

(and, of course, $v(\alpha) + v(\beta) + v(\gamma) = 1$).

In this case a common prior exists and is given by $v(\alpha) = \frac{2}{4}$ and $v(\beta) = v(\gamma) = \frac{1}{4}$.

Using this common prior to assign probabilities to Nature's choices we obtain the imperfect-information game shown in Figure 14.9.

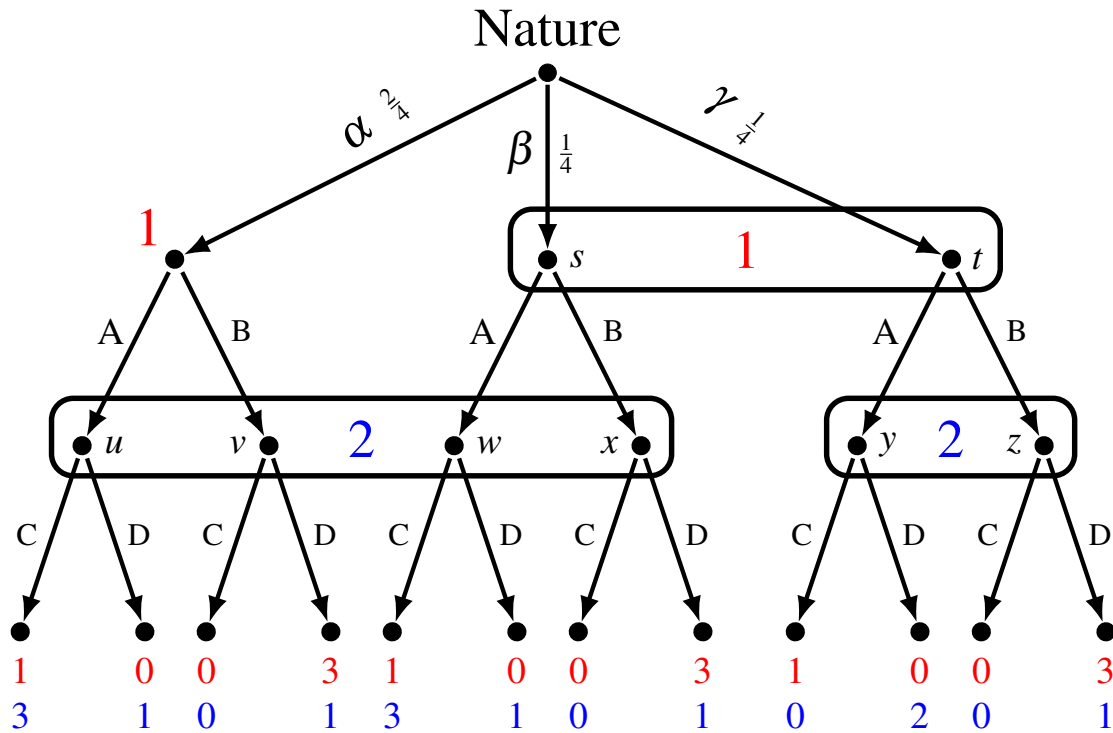


Figure 14.9: The extensive-form game obtained by applying the Harsanyi transformation to the incomplete-information situation of Figure 14.8.

The following pure-strategy profile is Bayesian Nash equilibrium of the game of Figure 14.9:

- Player 1's strategy is AB (that is, he plays A if informed that the state is α and plays B if informed that the state is either β or γ) and
- Player 2's strategy is CD (that is, she plays C at her information set on the left and D at her information set on the right).

To verify that this is a Bayesian Nash equilibrium it is easier to verify that it is a weak sequential equilibrium together with the following system of belief, obtained by using Bayesian updating (note that every information set is reached with positive probability by the strategy profile):

$$\mu = \left(\begin{array}{cc|cccc|cc} s & t & u & v & w & x & y & z \\ \frac{1}{2} & \frac{1}{2} & \frac{2}{3} & 0 & 0 & \frac{1}{3} & 0 & 1 \end{array} \right).$$

Let us check sequential rationality.

We begin with Player 1:

- ◇ at the singleton node on the left, A gives Player 1 a payoff of 1 (given Player 2's choice of C) while B gives him a payoff of 0, hence A is sequentially rational;
- ◇ at the information set on the right, given his beliefs and given Player 2's strategy CD , choosing A gives him an expected payoff of $\frac{1}{2}(1) + \frac{1}{2}(0) = \frac{1}{2}$ while B gives him an expected payoff of $\frac{1}{2}(0) + \frac{1}{2}(3) = \frac{3}{2}$, hence B is sequentially rational.

Now consider Player 2:

- ◇ at her information set on the left, given her beliefs, C gives her an expected payoff of $\frac{2}{3}(3) + \frac{1}{3}(0) = 2$, while D gives her an expected payoff of $\frac{2}{3}(1) + \frac{1}{3}(1) = 1$, hence C is sequentially rational;
- ◇ at her information set on the right, given her beliefs, C gives her a payoff of 0 while D gives her a payoff of 1, hence D is sequentially rational.

The existence of a common prior is essential in order to be able to apply the Harsanyi transformation to a situation of two-sided incomplete information. In some cases a common prior does *not* exist (see Exercise 14.6) and thus the Harsanyi transformation cannot be carried out.

R Besides the conceptual issues that arise in general with respect to the notion of Nash equilibrium, the notion of Bayesian Nash equilibrium for games with incomplete information raises the further issue of how one should understand or justify the notion of a common prior. This issue is not a trivial one and has been the object of debate in the literature.³

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 14.5.3 at the end of this chapter.

³See, for example, Bonanno and Nehring (1999), Gul (1998) and Morris (1995).

14.4 Multi-sided incomplete information

So far we have only considered strategic-form games with two players. However, the analysis extends easily to games involving more than two players. If there are $n \geq 3$ players and only one player has uncertainty about some aspects of the game, while the others have no uncertainty whatsoever, then we have a situation of one-sided incomplete information; if two or more players have uncertainty (not necessarily all about the game but some possibly about the beliefs of other players) then we have a *multi-sided incomplete-information situation* (the two-sided case being a special case).

We will consider an example of this below. Let G_1 and G_2 be the three-player games in strategic form (with cardinal payoffs) shown in Figure 14.10.

				Player 2			
		C			D		
Player 1	A	4	1	0	2	4	2
	B	4	1	0	0	0	0
				Player 3: E			
				Player 2			
		C			D		
Player 1	A	2	0	2	0	1	0
	B	0	0	0	0	2	0
				Player 3: F			

GAME G_1

				Player 2			
		C			D		
Player 1	A	4	4	0	2	1	2
	B	4	4	0	0	1	0
				Player 3: E			
				Player 2			
		C			D		
Player 1	A	2	1	2	0	0	0
	B	0	2	0	0	0	0
				Player 3: F			

GAME G_2

Figure 14.10: Two three-player games in strategic form with cardinal payoffs.

Consider the multi-sided situation of incomplete information illustrated in Figure 14.11, where with each state is associated one of the two games of Figure 14.10.

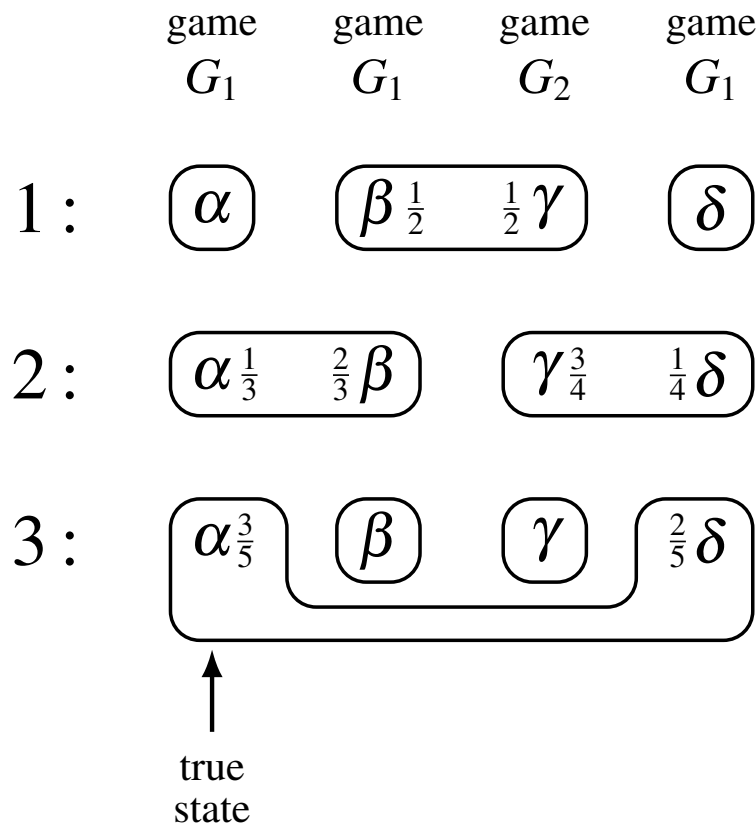


Figure 14.11: A three-sided situation of incomplete information.

At the true state α , all of the following are true:

- (a) All three players know that they are playing game G_1 ,
- (b) Player 1 knows that Players 2 and 3 know that the actual game is G_1 ;
- (c) Player 2 knows that Player 3 knows that the actual game is G_1 , but is uncertain as to whether Player 1 knows or is uncertain; furthermore, Player 2 assigns probability $\frac{2}{3}$ to Player 1 being uncertain;
- (d) Player 3 knows that Player 1 knows that the actual game is G_1 , but is uncertain as to whether Player 2 knows or is uncertain; furthermore, Player 3 assigns probability $\frac{2}{5}$ to Player 2 being uncertain;
- (e) The payoffs of Players 1 and 3 are common knowledge.

The beliefs of the players in the situation illustrated in Figure 14.11 are compatible with each other, in the sense that there exists a common prior. In fact the following is a common

$$\text{prior: } v = \begin{pmatrix} \alpha & \beta & \gamma & \delta \\ \frac{3}{17} & \frac{6}{17} & \frac{6}{17} & \frac{2}{17} \end{pmatrix}.$$

Thus we can apply the Harsanyi transformation and obtain the game shown in Figure 14.12.




Figure 14.12: The extensive-form game obtained by applying the Harsanyi transformation to the incomplete-information situation of Figure 14.11.

The task of finding Bayesian Nash equilibria for the game of Figure 14.12 is left as an exercise (see Exercise 14.7)

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 14.5.3 at the end of this chapter.

14.5 Exercises

14.5.1 Exercises for Section 14.2: One-sided incomplete information.

The answers to the following exercises are in Section 14.6 at the end of the chapter.

Exercise 14.1

Albert and Bill play a Game of Chicken^a under incomplete information. The situation is as follows.

- Choices are made simultaneously.
 - Each player can choose “swerve” or “don’t swerve.”
If a player swerves then he is a chicken.
If he does not swerve, then he is a rooster if the other player swerves, but he gets injured if the other player does not swerve.
 - Each player can be a “normal” type or a “reckless” type.
A “normal” type gets a payoff of 100 from being a rooster, 85 from being a chicken, and zero from being injured.
A “reckless” type gets a payoff of 100 from being a rooster, 50 from being injured, and zero from being a chicken.
 - As a matter of fact, both Albert and Bill are normal; it is common knowledge between them that Bill is normal, but Bill thinks that there is a 20% chance that Albert is reckless (and an 80% chance that Albert is normal).
 - Bill knows that Albert knows whether he (= Albert) is reckless or normal. Everything that Bill knows is common knowledge between Albert and Bill.
- (a) Construct an interactive knowledge-belief structure that represents the situation of incomplete information described above.
- (b) Apply the Harsanyi transformation to obtain an extensive-form game.
- (c) Construct the strategic-form associated with the extensive-form game of part (b).
- (d) Find the pure-strategy Bayesian Nash equilibria and classify them as either pooling or separating (Definition 14.2.1).

^aIn the original interpretation of this game, two drivers drive towards each other on a collision course: at least one must swerve in order to avoid a deadly crash; if one driver swerves and the other does not, the one who swerved will be called a “chicken”, meaning a coward. Hence the name “game of chicken”.

Exercise 14.2

Bill used to be Ann's boyfriend. Today it is Ann's birthday. Bill can either not give a gift to Ann or give her a nicely wrapped gift. If offered a gift, Ann can either accept or reject. This would be a pretty simple situation, if it weren't for the fact that Ann does not know if Bill is still a friend or has become an enemy. If he is a friend, she expects a nice present from him. If Bill is an enemy, she expects to find a humiliating thing in the box (he is known to have given dead frogs to his "not so nice friends" when he was in third grade!). The preferences are as follows.

- Bill's favorite outcome (payoff = 1) occurs when he offers a gift and it is accepted (in either case: if he is a friend, he enjoys seeing her unwrap a nice present, and if he is an enemy, he revels in the cruelty of insulting Ann with a humiliating "gift"). Whether he is a friend or an enemy, Bill prefers having not extended a gift (payoff = 0) to enduring the humiliation of a rejected gift (payoff = -1).
- Ann prefers accepting a gift coming from a friend (payoff = 1) to refusing a gift (payoff = 0); the worst outcome for her (payoff = -1) is one where she accepts a gift from an enemy.
- Ann attaches probability p (with $0 < p < 1$) to the event that Bill is a friend (and $1 - p$ to Bill being an enemy); however, Ann knows that Bill knows whether he is a friend or an enemy.
- Everything that Ann knows is common knowledge between Ann and Bill.
- As a matter of fact, Bill is a friend.

- (a) Construct an interactive knowledge-belief structure that represents the situation of incomplete information described above.
- (b) Apply the Harsanyi transformation to obtain an extensive-form game.
- (c) Construct the strategic-form associated with the extensive-form game.
- (d) Find all the pure-strategy Bayesian Nash equilibria and classify them as either pooling or separating.
- (e) Suppose that $p = \frac{1}{4}$. Is the outcome associated with a pure-strategy Bayesian Nash equilibrium Pareto efficient?^a

^aAn outcome is Pareto efficient if there is no other outcome that is Pareto superior to it (Definition 2.2.4, Chapter 2).

14.5.2 Exercises for Section 14.3: Two-sided incomplete information.

The answers to the following exercises are in Section 14.6 at the end of the chapter.

Exercise 14.3

Consider the situation of two-sided incomplete information illustrated in Figure 14.13 (where the true state is α).

Use the Harsanyi transformation to represent this incomplete-information situation as an extensive-form game. Be explicit about how you calculated the probabilities for Nature's choices.




Figure 14.13: A two-sided incomplete-information situation.

Exercise 14.4

Consider the following congestion situation (a variant of the so called El Farol Bar Problem: see https://en.wikipedia.org/wiki/El_Farol_Bar_problem).

- Two students at a college are simultaneously deciding between going to a bar or going home. The bar is extremely small and it gets congested when more than one person is there.
- In principle, there are two types of students. One type, call it the b type, prefers going to the bar if he is the only customer there (in which case he gets a utility of 20) but dislikes congestion and gets a utility of -20 if he goes to the bar and is not the only customer there; furthermore, for the b type, the utility of being at home is 0.
- The other type of student, call him the not- b type, prefers being at home (in which case his utility is 20) to being at the bar; being at the bar alone gives him a utility of 0 but being at the bar with other customers is very stressful and gives him a utility of -40 . Let $G(b_1, b_2)$ be the game where both players are b types, $G(b_1, \text{not-}b_2)$ the game where Player 1 is a b type and Player 2 a not- b type, etc.
- Assume that all payoffs are von Neumann-Morgenstern payoffs.

- (a) Write the four possible strategic-form games.
- (b) (b.1) For each of the games of part (a) find the pure-strategy Nash equilibria.
(b.2) For game $G(b_1, b_2)$ find also a mixed-strategy equilibrium where each choice is made with positive probability.
- (c) Assume now that, as a matter of fact, both players are b types. However, it is a situation of incomplete information where it is common knowledge that each player knows his own type but is uncertain about the type of the other player and assigns probability $\frac{1}{5}$ to the other player being the same type as he is and probability $\frac{4}{5}$ to the other player being of the opposite type. Draw an interactive knowledge-belief structure that represents this situation of incomplete information.
- (d) Use the Harsanyi transformation to represent the above situation of incomplete information as an extensive-form game.
- (e) For the game of part (d), pick one strategy of Player 1 and explain in words what it means.
- (f) For the game of part (d), write the corresponding strategic-form game.
- (g) Find a pure-strategy Bayesian Nash equilibrium of the game of part (d).

- (h) For the Bayesian Nash equilibrium of part (g),
 - (h.1) find where the players are actually going,
 - (h.2) find the actual payoffs of the players in the game that they are actually playing (that is, at the true state),
 - (h.3) do their actual choices yield a Nash equilibrium of the game that they are actually playing?
- (i) If you didn't know what the true state was but you knew the game of part (d), what probability would you assign to the event that the players would end-up making actual choices that constitute a Nash equilibrium of the true game that they are playing?

Exercise 14.5

Consider the situation of two-sided incomplete information illustrated in Figure 14.14 (where the true state is α).

- (a) Use the Harsanyi transformation to represent this situation as an extensive-form game. Be explicit about how you calculated the probabilities for Nature.
- (b) Write all the pure strategies of Player 1 and all the pure strategies of Player 2.
- (c) Consider the following pure-strategy profile: Player 1 plays T always and Player 2 plays L always. What belief system, paired with this strategy profile, would satisfy Bayesian updating?




Figure 14.14: A two-sided situation of incomplete information.

14.5.3 Exercises for Section 14.4: Multi-sided incomplete information.

The answers to the following exercises are in Section 14.6 at the end of the chapter.

Exercise 14.6

Consider the three-sided situation of incomplete information shown in Figure 14.15 (what games G_1 and G_2 are is irrelevant to the following question). For what values of p can the Harsanyi transformation be carried out? ■





Figure 14.15: A three-sided situation of incomplete information.

Exercise 14.7

Find a pure-strategy Bayesian Nash equilibrium of the game of Figure 14.12 which is reproduced on the next page.

[Hint: it is best to think in terms of weak sequential equilibrium.] ■



Exercise 14.8 — ***Challenging Question***.

Consider the following situation. It is common knowledge between Players 1 and 2 that tomorrow one of three states will occur: a , b or c .

It is also common knowledge between them that

- if state a materializes, then Player 1 will only know that either a or b occurred and Player 2 will only know that either a or c occurred,
- if state b materializes, then Player 1 will only know that either a or b occurred and Player 2 will know that b occurred,
- if state c materializes, then Player 1 will know that c occurred while Player 2 will only know that either a or c occurred.

Tomorrow they will play the following *simultaneous* game: each will report, confidentially, one of her two possible states of information to a third party (for example, Player 1 can only report either $\{a, b\}$ or $\{c\}$).

Note that lying is a possibility: for example, if the state is a Player 1 can choose to report $\{c\}$.

Let R_1 be the report of Player 1 and R_2 the report of Player 2. The third party, who always knows the true state, then proceeds as follows:

- (1) if the reports are compatible, in the sense that

$$R_1 \cap R_2 \neq \emptyset,$$

then he gives the players the following sums of money (which depend on $R_1 \cap R_2$):

If the true state is a :

$R_1 \cap R_2 =$	a	b	c
1 gets	\$5	\$4	\$6
2 gets	\$5	\$6	\$4

If the true state is b :

$R_1 \cap R_2 =$	a	b	c
1 gets	\$5	\$4	\$4
2 gets	\$0	\$1	\$1

If the true state is c :

$R_1 \cap R_2 =$	a	b	c
1 gets	\$0	\$1	\$1
2 gets	\$5	\$4	\$4

(2) if the reports are incompatible, in the sense that

$$R_1 \cap R_2 = \emptyset,$$

then he gives the players the following sums of money:

The true state is	a	b	c
1 gets	\$5	\$4	\$1
2 gets	\$5	\$1	\$4

- (a) Represent this situation of incomplete information by means of an interactive knowledge structure (for the moment do not worry about beliefs).
- (b) Apply the Harsanyi transformation to the situation represented in part (a) to obtain an extensive-form frame (again, at this stage, do not worry about probabilities).
- (c) Suppose first that both players have no idea what the probabilities of the states are and are not willing to form subjective probabilities.

It is common knowledge that each player is selfish (i.e. only cares about how much money she herself gets) and greedy (i.e. prefers more money to less) and ranks sets of outcomes according to the worst outcome, in the sense that she is indifferent between sets X and Y if and only if the worst outcome in X is equal to the worst outcome in Y and prefers X to Y if and only if the worst outcome in X is better than the worst outcome in Y .

- (c.1) Write the normal-form (or strategic-form) of the game of part (a).
- (c.2) Find all the pure-strategy Nash equilibria of this game.
- (c.3) Among the Nash equilibria, is there one where each player tells the truth?
- (d) Suppose now that it is common knowledge between the players that there are objective probabilities for the states as follows:

$$\begin{pmatrix} a & b & c \\ \frac{2}{5} & \frac{1}{5} & \frac{2}{5} \end{pmatrix}$$

This time assume that it is common knowledge that both players are selfish, greedy and risk-neutral. (Thus ignore now the preferences described in part (c).)

- (d.1) Suppose that Player 2 expects Player 1 to report truthfully. Is it rational for Player 2 to also report truthfully?
- (d.2) Is “always lying” for each player part of a pure-strategy weak sequential equilibrium? Prove your claim.

(c) The corresponding strategic-form game is shown in Figure 14.18.

		Bill			
		s		ds	
A l b e r t	s, s	$(0.8)85 =$ 68	85	$(0.8)85 =$ 68	100
	s, ds	$(0.8)100 =$ 80	85	0	20 <small>$= (0.2)100$</small>
	ds, s	$(0.8)85 + (0.2)100 =$ 88	85	$(0.8)85 + (0.2)50 =$ 78	80 <small>$= (0.8)100$</small>
	ds, ds	100	85	$(0.2)50 =$ 10	0

Figure 14.18: The strategic form of the game of Figure 14.17.

(d) There is only one pure-strategy Bayesian Nash equilibrium, namely $((ds, ds), s)$ and it is a pooling equilibrium.

□

Solutions to Exercise 14.2

(a) The structure is shown in Figure 14.19

(g means “gift”, ng means “no gift”, a means “accept” and r means “reject”).

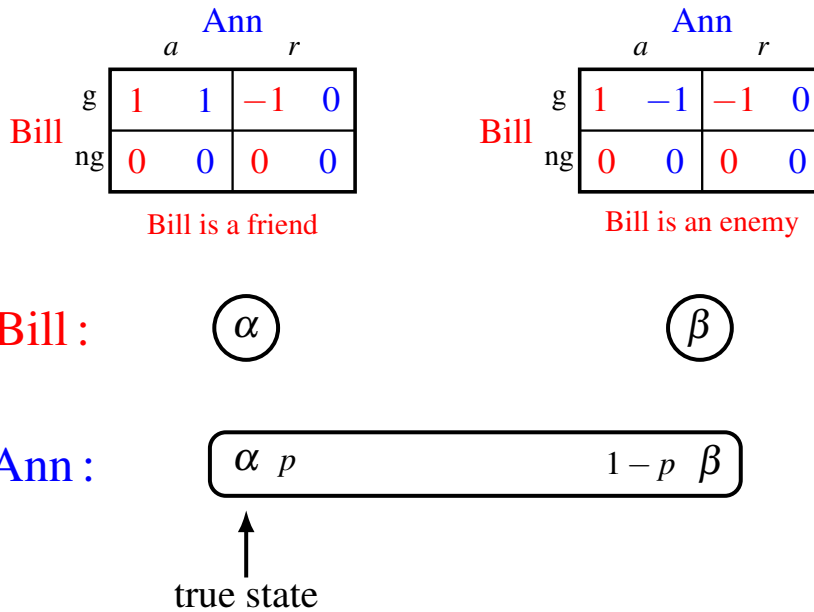


Figure 14.19: The one-sided incomplete-information structure for Exercise 14.2.

(b) The Harsanyi transformation yields the game shown in Figure 14.20.




Figure 14.20: The game obtained by applying the Harsanyi transformation to the incomplete-information situation of Figure 14.19.

(c) The corresponding strategic-form game is shown in Figure 14.21.

		Ann	
		a	r
Bill	g, g	1 $2p-1$	-1 0
	g, ng	p p	-p 0
	ng, g	1-p $p-1$	$p-1$ 0
	ng, ng	0 0	0 0

Figure 14.21: The strategic form of the game of Figure 14.20.

- (d) If $p \geq \frac{1}{2}$ then there are two pure-strategy Bayesian Nash equilibria: $((g, g), a)$ and $((ng, ng), r)$. Both of them are pooling equilibria. If $p < \frac{1}{2}$ then there is only one Bayesian -Nash equilibrium: $((ng, ng), r)$ (a pooling equilibrium).
- (e) If $p = \frac{1}{4}$ the only Bayesian-Nash equilibrium is $((ng, ng), r)$ and the outcome is that Bill does not offer a gift to Ann. This outcome is Pareto inefficient in the true game because, given that the true state of affairs is one where Bill is a friend, a Pareto superior outcome would be one where Bill offers a gift and Ann accepts.

□

Solutions to Exercise 14.3 The game is shown in Figure 14.22.




Figure 14.22: The game for exercise 14.3.

Nature's probabilities are obtained by solving the following system of equations:

$$\frac{p\alpha}{p\alpha + p\beta} = \frac{1}{3}$$

$$\frac{p\beta}{p\beta + p\gamma} = \frac{3}{4}$$

$$p\alpha + p\beta + p\gamma = 1.$$

□

Solutions to Exercise 14.4

(a) The four games are shown in Figure 14.23.

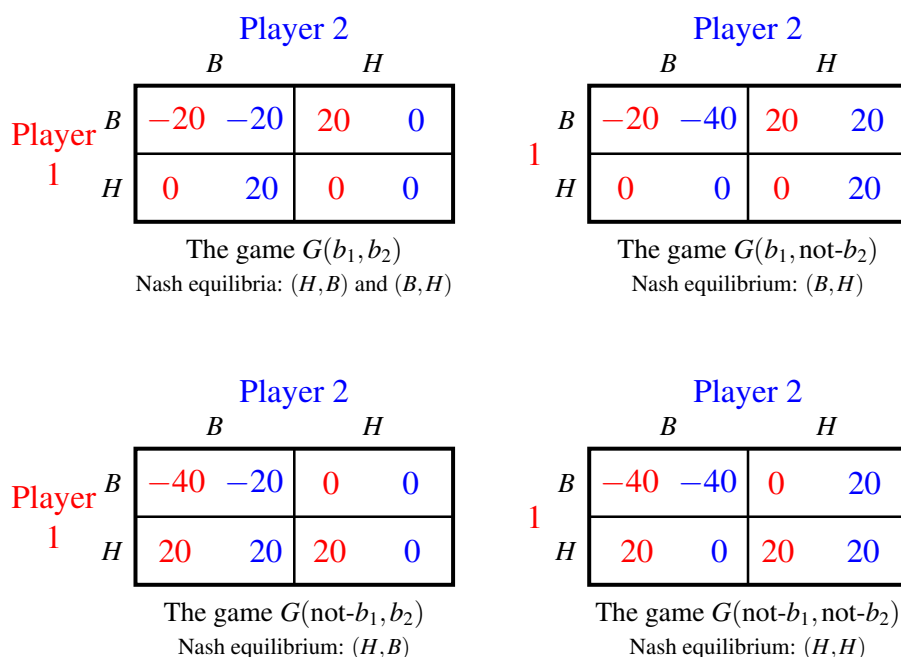


Figure 14.23: The four games of part (a) of Exercise 14.4.

(b) (b.1) The Nash equilibria are written under each figure in Figure 14.23.

(b.2) Let p be the probability that Player 1 plays B and q the probability that Player 2 plays B. Then the mixed-strategy equilibrium is given by the solution to $-20q + 20(1 - q) = 0$ and $-20p + 20(1 - p) = 0$, which is $p = q = \frac{1}{2}$.

(c) The structure is shown in Figure 14.24.




Figure 14.24: The two-sided incomplete-information structure of part (c) of Exercise 14.4

(d) First of all, note that the common prior is

$$\begin{pmatrix} (b_1, b_2) & (b_1, \text{not } -b_2) & (\text{not } -b_1, b_2) & (\text{not } -b_1, \text{not } -b_2) \\ \frac{1}{10} & \frac{4}{10} & \frac{4}{10} & \frac{1}{10} \end{pmatrix}.$$

The extensive-form game is shown in Figure 14.25.

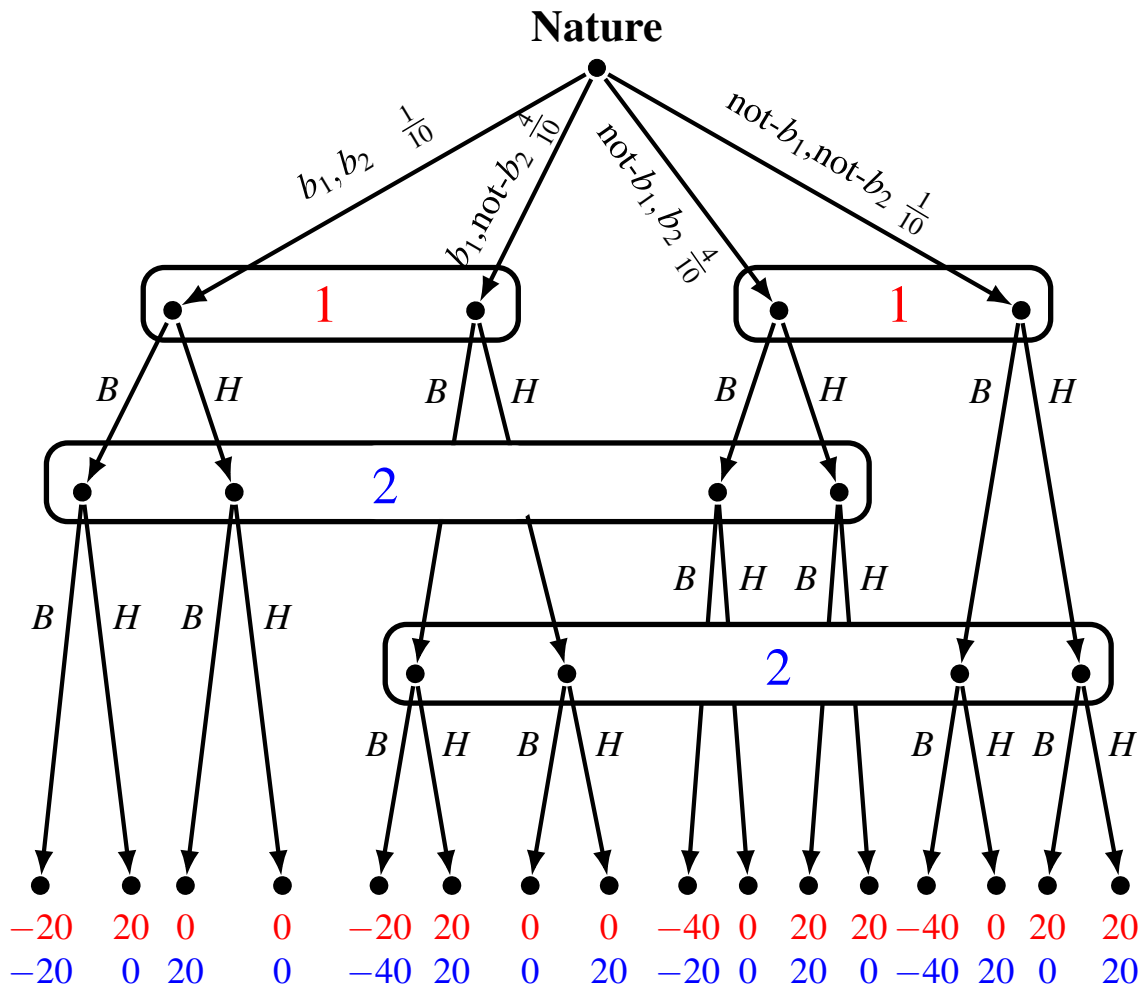


Figure 14.25: The game obtained by applying the Harsanyi transformation to the incomplete-information situation of Figure 14.24

(e) One possible strategy for Player 1 is (B, H) which means “if I am the b type then I go to the bar and if I am the $\text{not } -b$ type then I go home”.

(f) The strategic form is shown in Figure 14.26.

		Player 2			
		B, B	B, H	H, B	H, H
Player 1	B, B	-30 -30	-10 0	-10 -20	10 10
	B, H	0 -10	16 16	4 -16	20 10
	H, B	-20 -10	-16 4	-4 -4	0 10
	H, H	10 10	10 20	10 0	10 10

Figure 14.26: The strategic form of the game of Figure 14.25.

(g) There is a unique pure-strategy Nash equilibrium, namely $((B, H), (B, H))$ where each player goes to the bar if he is a b type and goes home if he is a not- b type. This can be found either the long way, by filling in all the payoffs in the above matrix, or by reasoning as follows: For each player, going home is strictly better than going to the bar if the player is of type not- b , no matter what she anticipates the other player doing (in other words, H strictly dominates B at the information set where not- b_i holds). Thus the question is what to do if you are of type b . You know that the other player is going home if he is of type not- b , thus you only need to consider his choice if he is of type b ; if his plan is to go home also in that case, then B gives you 20 for sure and H gives you 0 for sure, hence B is better; if his plan is to go to the bar, then H gives you 0 for sure while B gives you the lottery

$$\begin{pmatrix} -20 & 20 \\ \frac{1}{5} & \frac{4}{5} \end{pmatrix},$$

that is, an expected payoff of 12; hence B is better in that case too.

(h) At the true state, both players prefer going to the bar, thus

(h.1) they both end up going to the bar and

(h.2) they both get a payoff of -20 .

(h.3) (B, B) is not a Nash equilibrium of the game that they are actually playing (game $G(b_1, b_2)$).

(i) If the true game is $G(b_1, b_2)$ they end up playing (B, B) which is not a Nash equilibrium of that game,

if the true game is $G(b_1, \text{not} - b_2)$ they end up playing (B, H) which is a Nash equilibrium of that game,

if the true game is $G(\text{not} - b_1, b_2)$ they end up playing (H, B) which is a Nash equilibrium of that game

and if the true game is $G(\text{not} - b_1, \text{not} - b_2)$ they end up playing (H, H) which is a Nash equilibrium of that game.

Since the probability of $G(b_1, b_2)$ is $\frac{1}{10}$, the probability that they end up playing a Nash equilibrium of the actual game is $\frac{9}{10}$. \square

Solutions to Exercise 14.5

(a) The game is shown in Figure 14.27.




Figure 14.27: The extensive-form game for Exercise 14.5.

Nature’s probabilities are obtained solving the following system of equations:

$$\frac{p_\alpha}{p_\alpha + p_\beta} = \frac{1}{4}, \quad \frac{p_\beta}{p_\beta + p_\gamma} = \frac{1}{2}, \quad \frac{p_\gamma}{p_\gamma + p_\delta} = \frac{3}{4}, \quad p_\alpha + p_\beta + p_\gamma + p_\delta = 1.$$

(b) Player 1 has 8 strategies: *TTT*, *TTB*, *TBT*, *TBB*, *BTT*, *BTB*, *BBT*, *BBB*. Player 2 has four strategies: *LL*, *LR*, *RL*, *RR*.

(c) Player 1’s beliefs must be $\frac{1}{2}$ at the left node and $\frac{1}{2}$ at the right node of his middle information set.

Player 2’s beliefs at her information set on the left must be: $\frac{1}{4}$ at the left-most node and $\frac{3}{4}$ at the third node from the left. The same is true for the other information set of Player 2. \square

Solutions to Exercise 14.6

The Harsanyi transformation requires that there be a common prior. Thus we need a probability distribution $v : \{\alpha, \beta, \gamma, \delta\} \rightarrow (0, 1)$ such that:

(1) $\frac{v(\alpha)}{v(\alpha)+v(\beta)} = \frac{1}{5}$, (2) $\frac{v(\gamma)}{v(\gamma)+v(\delta)} = \frac{2}{3}$, (3) $\frac{v(\beta)}{v(\beta)+v(\gamma)} = \frac{1}{3}$ and (4) $\frac{v(\alpha)}{v(\alpha)+v(\delta)} = p$.

From (1) we get $v(\beta) = 4v(\alpha)$, from (2) we get $v(\gamma) = 2v(\delta)$ from (3) we get $v(\gamma) = 2v(\delta)$; these three equalities, together with $\alpha + \beta + \gamma + \delta = 1$ yield a unique

solution, namely $v = \begin{pmatrix} \alpha & \beta & \gamma & \delta \\ \frac{1}{17} & \frac{4}{17} & \frac{8}{17} & \frac{4}{17} \end{pmatrix}$. This is a common prior if and only if

$$p = \frac{v(\alpha)}{v(\alpha)+v(\delta)} = \frac{\frac{1}{17}}{\frac{1}{17} + \frac{4}{17}} = \frac{1}{5}.$$

Thus $p = \frac{1}{5}$ is the only value that makes it possible to apply the Harsanyi transformation. \square

Solutions to Exercise 14.7 The game under consideration is reproduced in Figure 14.28.




Figure 14.28: The game for Exercise 14.7.

Given the complexity of the game, it is definitely not a good idea to construct the corresponding strategic form. It is best to think in terms of weak sequential equilibrium. Consider the assessment (σ, μ) , highlighted in the Figure 14.28 by thick arrows, where $\sigma = (BBB, CC, EEE)$ (that is, σ is the pure-strategy profile where Player 1 chooses B at each of his three information sets, Player 2 chooses C at each of her two information sets and Player 3 chooses E at each of her three information sets) and μ is the following system of beliefs (shown in square brackets in Figure 14.28):

- Player 1 attaches probability $\frac{1}{2}$ to each node in his bottom information set;
- Player 2, at her information set on the left, attaches probability $\frac{1}{3}$ to the second node from the top and probability $\frac{2}{3}$ to the bottom node and, at her information set on the right, attaches probability $\frac{1}{4}$ to the second node from the top and probability $\frac{3}{4}$ to the bottom node;
- Player 3, at her top information set, attaches probability $\frac{3}{5}$ to the left-most node and probability $\frac{2}{5}$ to the second node from the right and, at her bottom-left information

set, attaches probability 1 to the third node from the left and, at her bottom-right information set, attaches probability 1 to the left-most node.

Let us verify that (σ, μ) is a weak sequential equilibrium. The beliefs described above are obtained from σ using Bayesian updating. Thus we only need to check sequential rationality.

For Player 1:

1. at the top-left node, both A and B give the same payoff (namely, 4), thus B is sequentially rational;
2. the same is true at the top-right node;
3. at the bottom information set both A and B give an expected payoff of $\frac{1}{2}(4) + \frac{1}{2}(4) = 4$, thus B is sequentially rational.

For Player 2:

1. at the left information set C gives an expected payoff of $\frac{1}{3}(1) + \frac{2}{3}(1) = 1$ while D gives $\frac{1}{3}(0) + \frac{2}{3}(0) = 0$, thus C is sequentially rational;
2. at the information on the right C gives an expected payoff of $\frac{1}{4}(1) + \frac{3}{4}(4) = \frac{13}{4}$ while D gives an expected payoff of $\frac{1}{4}(0) + \frac{3}{4}(1) = \frac{3}{4}$, thus C is sequentially rational.

For Player 3:

1. at the top information set both E and F give an expected payoff of $\frac{3}{5}(0) + \frac{2}{5}(0) = 0$, thus E is sequentially rational;
2. at the bottom-left information set both E and F give a payoff of 0, thus E is sequentially rational;
3. the same is true at the bottom-right information set. □

Solutions to Exercise 14.8

(a) The structure is shown in Figure 14.29.




Figure 14.29: The two-sided incomplete-information structure for Exercise 14.8.

(b) The game is shown in Figure 14.30.




Figure 14.30: The extensive-form game for the incomplete-information situation of Figure 14.29

(c) (c.1) The strategic form is as follows (where the strategy (x,y) for Player 1 means x if $\{a,b\}$ and y if $\{c\}$ and the strategy (z,w) for Player 2 means z if $\{a,c\}$ and w if $\{b\}$). Given inside each cell are the sets of outcomes:

		Player 2			
		$(\{a,c\}, \{a,c\})$	$(\{a,c\}, \{b\})$	$(\{b\}, \{a,c\})$	$(\{b\}, \{b\})$
PI	$(\{a,b\}, \{a,b\})$	$\{(5,5), (5,0), (0,5)\}$	$\{(5,5), (4,1), (0,5)\}$	$\{(4,6), (5,0), (1,4)\}$	$\{(4,6), (4,1), (1,4)\}$
	$(\{a,b\}, \{c\})$	$\{(5,5), (5,0), (1,4)\}$	$\{(5,5), (4,1), (1,4)\}$	$\{(4,6), (5,0), (1,4)\}$	$\{(4,6), (4,1), (1,4)\}$
1	$(\{c\}, \{a,b\})$	$\{(6,4), (4,1), (0,5)\}$	$\{(6,4), (4,1), (0,5)\}$	$\{(5,5), (4,1), (1,4)\}$	$\{(5,5), (4,1), (1,4)\}$
	$(\{c\}, \{c\})$	$\{(6,4), (4,1), (1,4)\}$	$\{(6,4), (4,1), (1,4)\}$	$\{(5,5), (4,1), (1,4)\}$	$\{(5,5), (4,1), (1,4)\}$

Taking as payoffs the smallest sum of money in each cell (for the corresponding player) the game can be written as follows:

		Player 2			
		$(\{a,c\}, \{a,c\})$	$(\{a,c\}, \{b\})$	$(\{b\}, \{a,c\})$	$(\{b\}, \{b\})$
PI	$(\{a,b\}, \{a,b\})$	0, 0	0, 1	1, 0	1, 1
	$(\{a,b\}, \{c\})$	1, 0	1, 1	1, 0	1, 1
1	$(\{c\}, \{a,b\})$	0, 1	0, 1	1, 1	1, 1
	$(\{c\}, \{c\})$	1, 1	1, 1	1, 1	1, 1

(c.2) There are 9 Nash equilibria, namely all the strategy profiles with payoffs $(1, 1)$.

(c.3) Truth telling is represented by the following strategy profile and it is one of the Nash equilibria: $((\{a,b\}, \{c\}) (\{a,c\}, \{b\}))$.

(d) (d.1) No. If the state is b then it is a good idea for Player 2 to report truthfully because $\{a,c\}$ yields her 0 while $\{b\}$ yields her 1.

But if the state is either a or c then, by Bayesian updating, Player 2 must assign probability $\frac{1}{2}$ to the left-most node and probability $\frac{1}{2}$ to the right-most node of her bottom information set;

thus her expected payoff from reporting $\{a,c\}$ is $\frac{1}{2}(5) + \frac{1}{2}(4) = 4.5$

while the expected payoff from reporting $\{b\}$ is $\frac{1}{2}(6) + \frac{1}{2}(4) = 5$.

(d.2) Yes. "Always lie" corresponds to the strategy profile $((\{c\}, \{a,b\}) (\{b\}, \{a,c\}))$.

By Bayesian updating the corresponding beliefs must be: for Player 1 $(\frac{2}{3}, \frac{1}{3})$ and for Player 2 $(0, \frac{1}{2}, \frac{1}{2}, 0)$ at the bottom information set and $(0,1)$ at the top information set.

Sequential rationality is then satisfied at every information set: for Player 1 at the top information set $\{c\}$ gives an expected payoff of $\frac{2}{3}(5) + \frac{1}{3}(4) = \frac{14}{3}$

while $\{a,b\}$ gives $\frac{2}{3}(4) + \frac{1}{3}(5) = \frac{13}{3}$

and at the singleton node on the right $\{a,b\}$ gives 1 and so does $\{c\}$.

For Player 2 at the bottom information set $\{b\}$ gives an expected payoff of $\frac{1}{2}(5) + \frac{1}{2}(4) = 4.5$ and $\{a,c\}$ gives $\frac{1}{2}(4) + \frac{1}{2}(5) = 4.5$,

and at the top information set both $\{a,c\}$ and $\{b\}$ give 1. □

15. Dynamic Games

15.1 One-sided incomplete information

At the conceptual level, situations of incomplete information involving dynamic (or extensive-form) games are essentially the same as situations involving static games: the only difference in the representation is that one would associate with every state a dynamic game instead of a static game.¹

As in the case of static games, we will distinguish between one-sided and multi-sided incomplete information. In this section we will go through two examples of the former, while the latter will be discussed in Section 15.2.

Recall that a situation of incomplete information is said to be one-sided if only one of the players has some uncertainty about what game is being played. In Chapter 14 we restricted attention to the case where the uncertainty is about the payoffs and we will continue to do so in this chapter. The player who is uncertain about the game is often referred to as the “uninformed” player, whereas the other players are referred to as the “informed” players. It is common knowledge among all the players that the informed players know what game they are playing; furthermore, the beliefs of the uninformed player are common knowledge among all the players.

¹The reader might have noticed that in Exercise 14.2 (Chapter 14) we “sneaked in” a dynamic game (where Bill moved first and decided whether or not to offer a gift and Ann – if offered a gift – decided whether or not to accept it). However, the game can also be written as a simultaneous game, where Ann decides whether or not to accept before knowing whether Bill will offer a gift (without knowing Ann’s decision).

Consider the perfect-information frame shown in Figure 15.1, where the question marks in place of Player 2's payoffs indicate that her payoffs are not common knowledge between the two players.




Figure 15.1: A perfect-information game-frame where the question marks indicate uncertainty about Player 2's payoffs.

Suppose that Player 1's payoffs are common knowledge between the two players. Furthermore, suppose that Player 1 is uncertain between the two possibilities shown in Figure 15.2 and assigns probability $\frac{1}{3}$ to the one on the left and probability $\frac{2}{3}$ to the one on the right. The thick edges represent the backward-induction solutions of the two games. Thus, if Player 1 knew that he was playing the game on the left, he would choose *B*, while if he knew that he was playing the game on the right, he would choose *T*.




Figure 15.2: The two possibilities in the mind of Player 1.

If we assume that Player 1's beliefs are common knowledge between the players, then we have the situation of one-sided incomplete information shown in Figure 15.3. Let us take state α to be the true state.

Using the Harsanyi transformation we can convert the situation illustrated in Figure 15.3 into the extensive-form game shown in Figure 15.4.




Figure 15.3: A one-sided situation of incomplete information involving the two games of Figure 15.2.




Figure 15.4: The game obtained by applying the Harsanyi transformation to the incomplete-information situation of Figure 15.3.

The solution concept used for the case of static games was Bayesian Nash equilibrium (that is, Nash equilibrium of the imperfect-information game obtained by applying the Harsanyi transformation to the situation of incomplete information).

In the case of dynamic games this is no longer an appropriate solution concept, since – as we know from Chapter 4 – it allows a player to “choose” a strictly dominated action at an unreached information set. To see this, consider the strategic-form game associated with the game of Figure 15.4, shown in Figure 15.5.

		Player 2			
		<i>DD</i>	<i>DA</i>	<i>AD</i>	<i>AA</i>
Player 1	<i>B</i>	1 0	1 0	1 0	1 0
	<i>T</i>	0 4/3	4/3 2	2/3 1	2 5/3

Figure 15.5: The strategic form of the game of Figure 15.4.

The Nash equilibria of the strategic-form game shown in Figure 15.5 (and thus the Bayesian Nash equilibria of the game shown in Figure 15.4) are:

$$(B, DD), \quad (B, AD) \quad \text{and} \quad (T, DA).$$

Of these, only (T, DA) is a subgame-perfect equilibrium. From now on we shall use either the notion of subgame-perfect equilibrium or the notion of weak sequential equilibrium (in the game of Figure 15.4 the two notions coincide). Thus we will take (T, DA) to be the solution of the game of Figure 15.4. Since we postulated that the true game was the one associated with state α , as a matter of fact the players end up playing (T, D) which is not the backward induction solution of the true game. As pointed out in Chapter 14, this is not surprising given the uncertainty in the mind of Player 1 as to which game she is playing.

Next we consider a more complex example, built on Selten’s Chain-Store game analyzed in Chapter 3 (Section 4). Recall that the story is as follows (we will consider the special case where m , the number of towns and thus of potential entrants, is 2). A chain store is a monopolist in an industry. It owns stores in two different towns. In each town the chain store makes \$5 million if left to enjoy its privileged position undisturbed. In each town there is a businesswoman who could enter the industry in that town, but earns \$1 million in an alternative investment if she chooses not to enter; if she decides to enter, then the monopolist can either fight the entrant, leading to zero profits for both the chain store and the entrant in that town, or it can accommodate entry and share the market with the entrant, in which case both players make \$1.5 million in that town. Thus, in each town the interaction between the incumbent monopolist and the potential entrant is as illustrated in Figure 15.6.

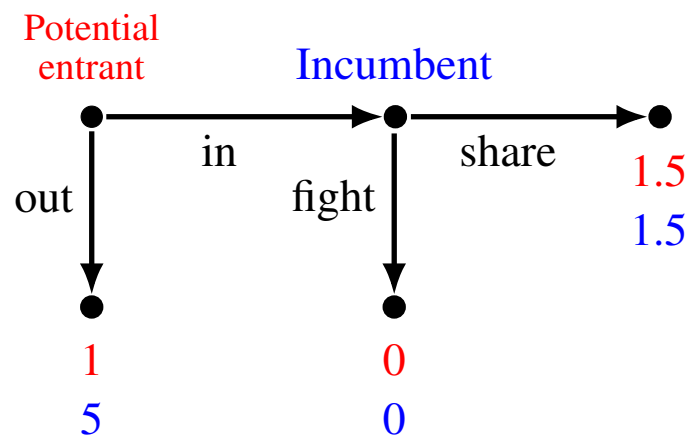


Figure 15.6: The game with only one potential entrant.

Decisions are made sequentially, as follows. At date t ($t = 1, 2$) the businesswoman in town t decides whether or not to enter and, if she enters, then the chain store decides whether or not to fight in that town. What happens in town t (at date t) becomes known to everybody. Thus the businesswoman in town 2 at date 2 knows what happened in town 1 at date 1 (either that there was no entry or that entry was met with a fight or that entry was accommodated).

Intuition suggests that in this game the threat by the Incumbent to fight early entrants might be credible, for the following reason. The Incumbent could tell Businesswoman 1 the following:

“It is true that, if you enter and I fight you, I will make zero profits, while by accommodating your entry I would make \$1.5 million and thus it would seem that it cannot be in my interest to fight you. However, somebody else is watching us, namely Businesswoman 2. If she sees that I have fought your entry then she might fear that I would do the same with her and decide to stay out, in which case, in town 2, I would make \$5 million, so that my total profits in towns 1 and 2 would be $\$(0 + 5) = \5 million. On the other hand, if I accommodate your entry, then she will be encouraged to enter herself and I will make \$1.5 million in each town, for a total profit of \$3 million. Hence, as you can see, it is indeed in my interest to fight you and thus you should stay out.”

We showed in Chapter 3 that the notion of backward induction does not capture this intuition. In the game depicting the entire interaction (Figure 3.11, Chapter 3) there was a unique backward-induction solution whose corresponding outcome was that both businesswomen entered and the Incumbent accommodated entry in both towns. The reason why the backward-induction solution did not capture the “reputation” argument outlined above was explained in Chapter 3. We remarked there that, in order to capture the reputation effect, one would need to allow for some uncertainty in the mind of some of the players. This is what we now show.

Suppose that, in principle, there are two types of incumbent monopolists: the rational type and the hotheaded type. The payoffs of a rational type are shown in Figure 15.7 (a rational incumbent prefers 'share' to 'fight' if there is entry), while the payoffs of a hotheaded type are shown in Figure 15.8 (a hotheaded type prefers 'fight' to 'share' if there is entry). On the other hand, there is only one type of potential entrant.

		Incumbent	
		fight	share
Potential entrant	in	0 0	1.5 1.5
	out	1 5	1 5

Figure 15.7: The payoffs of a **rational incumbent**.

		Incumbent	
		fight	share
Potential entrant	in	0 2	1.5 1.5
	out	1 5	1 5

Figure 15.8: The payoffs of a **hotheaded incumbent**.

Consider the following situation of one-sided incomplete information.

- As a matter of fact, the Incumbent is rational and this fact is common knowledge between the Incumbent and Potential Entrant 1 (from now on denoted by *PE-1*).
- Potential Entrant 2 (*PE-2*) is uncertain whether the Incumbent is rational or hotheaded and attaches probability p to the latter case. The beliefs of *PE-2* are common knowledge as are the payoffs of *PE-1* and *PE-2*.

This situation can be illustrated with the knowledge-belief structure of Figure 15.9.

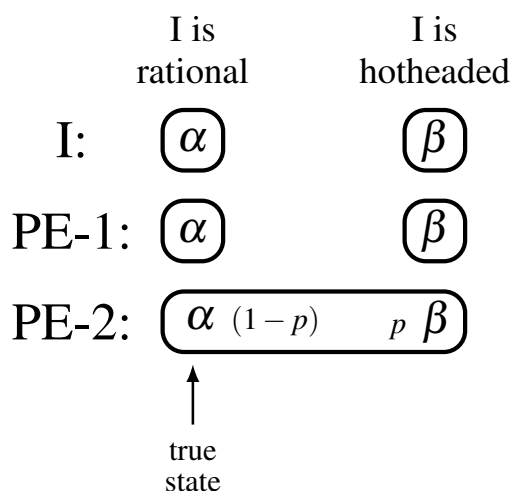


Figure 15.9: The one-sided incomplete-information situation.

Applying the Harsanyi transformation to the situation depicted in Figure 15.9 yields the extensive-form game shown in Figure 15.10. We want to see if this situation of one-sided incomplete information can indeed capture the reputation effect discussed above; in particular, we want to check if there is a weak sequential equilibrium of the game of Figure 15.10 where the Incumbent's strategy would be to fight $PE-1$'s entry in order to scare off $PE-2$ and, as a consequence, $PE-1$ decides to stay out and so does $PE-2$.

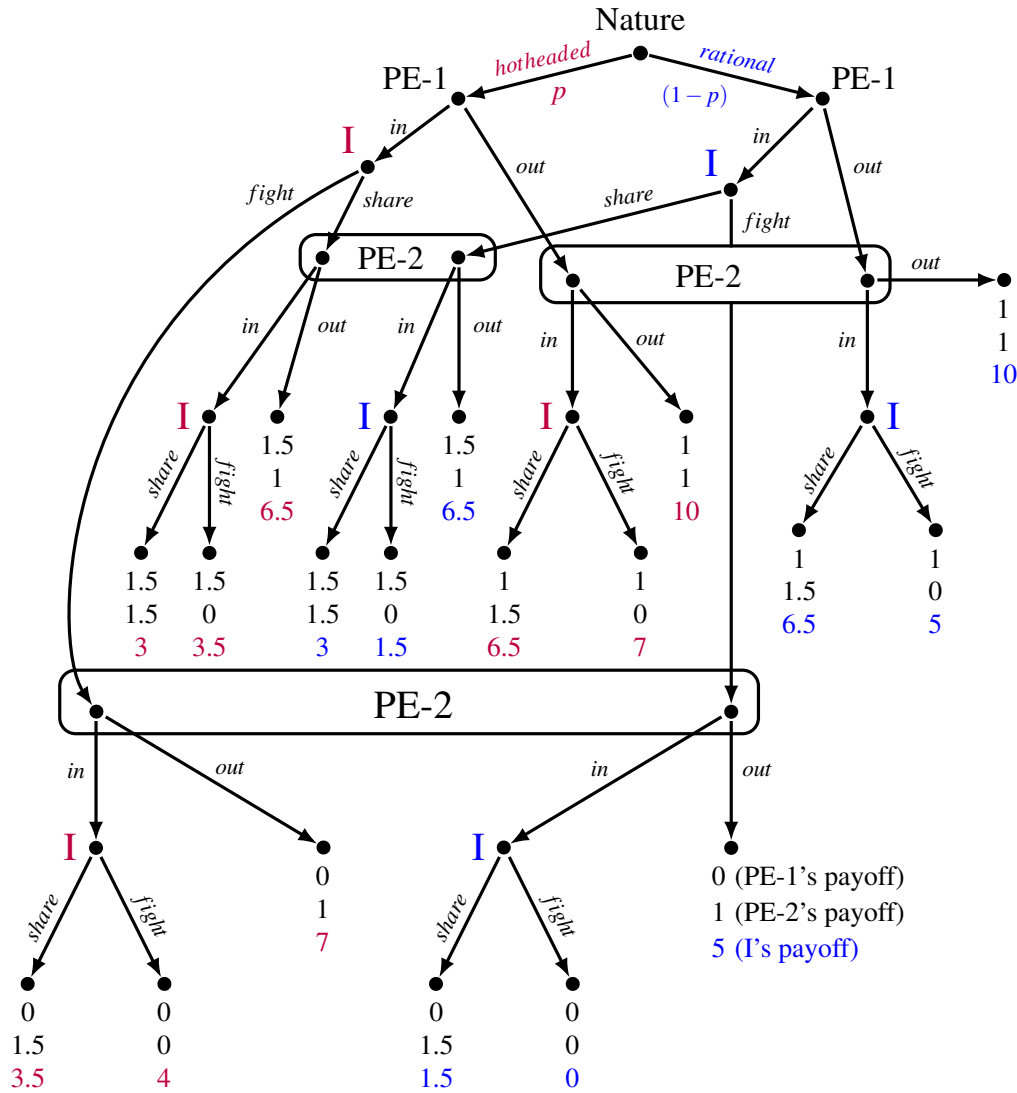


Figure 15.10: The game obtained by applying the Harsanyi transformation to the one-sided situation of incomplete information of Figure 15.9.

The notion of weak sequential equilibrium allows us to simplify the game, by selecting the strictly dominant choice for the Incumbent at each of his singleton nodes followed only by terminal nodes. It is straightforward to check that at such nodes a hotheaded Incumbent would choose “fight” while a rational Incumbent would choose “share”. Thus we can delete those decision nodes and replace them with the payoff vectors associated with the optimal choice.

The simplified game is shown in Figure 15.11.

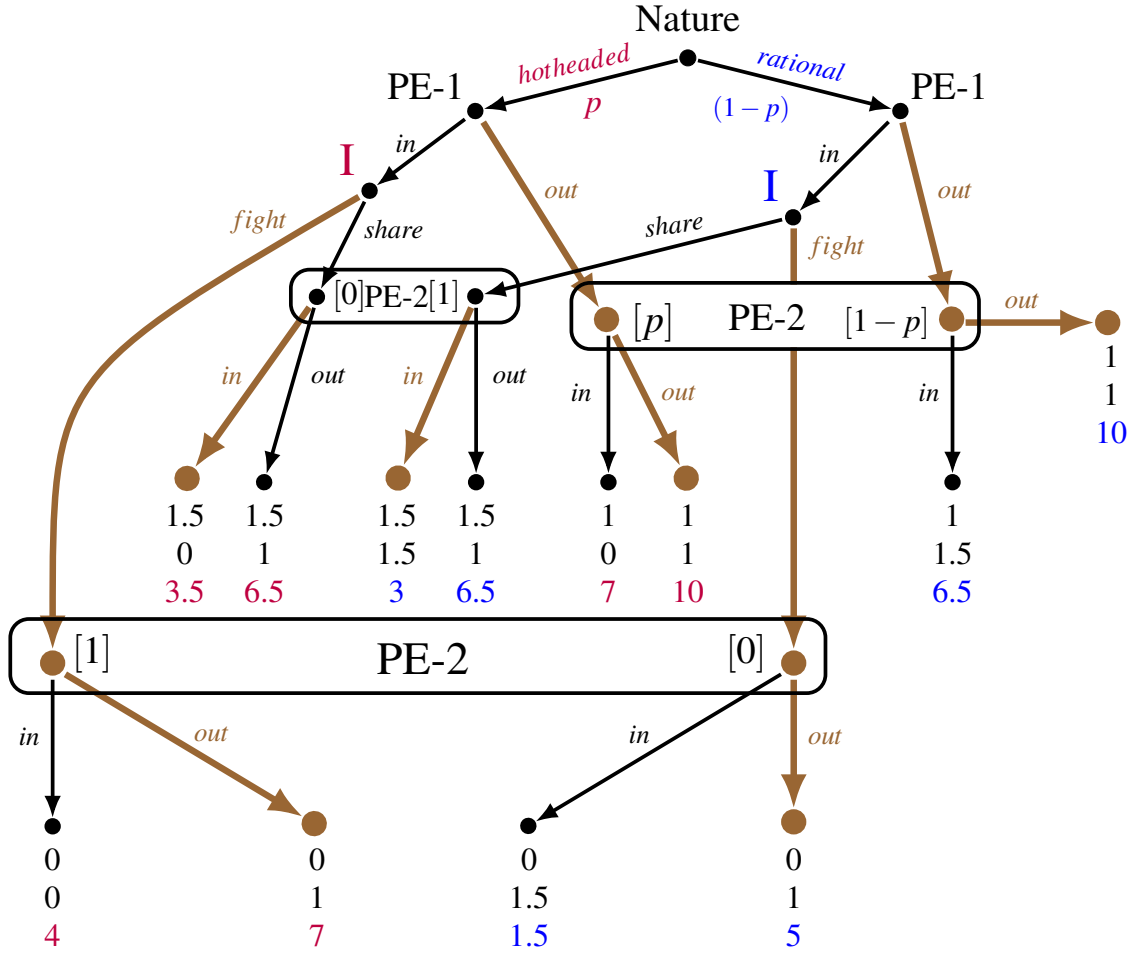


Figure 15.11: The reduced game obtained from Figure 15.10 by eliminating strictly dominated choices.

Consider the following pure-strategy profile, call it σ , for the simplified game of Figure 15.11 (it is highlighted by thick arrows in Figure 15.11):

1. $PE-1$'s strategy is "out" at both nodes,
2. $PE-2$'s strategy is
 - "out" at the top information set (after having observed that $PE-1$ stayed out),
 - "in" at the middle information set (after having observed that $PE-1$'s entry was followed by the Incumbent sharing the market with $PE-1$)
 - "out" at the bottom information set (after having observed that $PE-1$'s entry was followed by the Incumbent fighting against $PE-1$),
3. the Incumbent's strategy is to fight entry of $PE-1$ in any case (that is, whether the Incumbent himself is hotheaded or rational).

We want to show that σ , together with the system of beliefs μ shown in square brackets in Figure 15.11 (at her top information $PE-2$ assigns probability p to the Incumbent being hotheaded, at her middle information set she assigns probability 1 to the Incumbent being rational and at her bottom information set she assigns probability 1 to the Incumbent being hotheaded) constitutes a weak sequential equilibrium for any value of $p \geq \frac{1}{3}$.

Bayesian updating is satisfied at $PE-2$'s top information set (the only information set reached by σ), and at the other two information sets of $PE-2$ any beliefs are allowed by the notion of weak sequential equilibrium.

Sequential rationality also holds:

- For $PE-1$, at the left node “in” yields 0 and “out” yields 1, so that “out” is sequentially rational; the same is true at the right node.
- For the Incumbent,
 - at the left node “fight” yields 7 and “share” yields 3.5, so that “fight” is sequentially rational;
 - at the right node “fight” yields 5 and “share” yields 3, so that “fight” is sequentially rational.
- For $PE-2$,
 - at the top information set “in” yields an expected payoff of $p(0) + (1 - p)(1.5)$ and “out” yields 1, so that “out” is sequentially rational as long as $p \geq \frac{1}{3}$;
 - at the middle information set, given her belief that I is rational, “in” yields 1.5 and “out” yields 1, so that “in” is sequentially rational;
 - at the bottom information set, given her belief that I is hotheaded, “in” yields 0 and “out” yields 1, so that “out” is sequentially rational.

Thus the equilibrium described above captures the intuition suggested in Chapter 3, namely that – even though it is common knowledge between the Incumbent and $PE-1$ that the Incumbent is rational and thus would suffer a loss of 1.5 by fighting $PE-1$'s entry – it is still credible for the Incumbent to threaten to fight $PE-1$'s entry because it would influence the beliefs of $PE-2$ and induce her to stay out; understanding the credibility of this threat, it is optimal for $PE-1$ to stay out.

The above argument exploits the fact that the notion of weak sequential equilibrium allows for any beliefs whatsoever at unreached information sets. However, the beliefs postulated for $PE-2$ at her middle and bottom information sets seem highly reasonable. Indeed, as shown in Exercise 15.2, the assessment described above is also a sequential equilibrium.

We now turn to an example that deals with the issue of labor-management negotiations and the inefficiency of strikes. It is not uncommon to observe a firm and a union engaging in unsuccessful negotiations leading to a strike by the workers, followed by renewed negotiations and a final agreement. Strikes are costly for the workers, in terms of lost wages, and for the firm, in terms of lost production. Why, then, don't the parties reach the agreement at the very beginning, thus avoiding the inefficiency of a strike? The answer in many cases has to do with the fact that there is incomplete information on the side of the labor union and enduring a strike is the only credible way for the firm to convince the union to reduce its demands. We shall illustrate this in a simple example of one-sided incomplete information.

Consider the following game between a new firm and a labor union. The union requests a wage (either high, w_H , or low, w_L) and the firm can either accept or reject. If the union's request is accepted, then a contract is signed and production starts immediately. If the request is rejected, then the union goes on strike for one period and at the end of that period makes a second, and last, request to the firm, which the firm can accept or reject. If the firm rejects, then it cannot enter the industry. When no agreement is signed, both parties get a payoff of 0.

Both the firm and the union have a discount factor of δ with $0 < \delta < 1$, which means that \$1 accrued one period into the future is considered to be equivalent to δ at the present time (this captures the cost of waiting and thus the desirability of avoiding a strike).

The extensive-form game is shown in Figure 15.12, where π denotes the firm's profit (gross of labor costs).

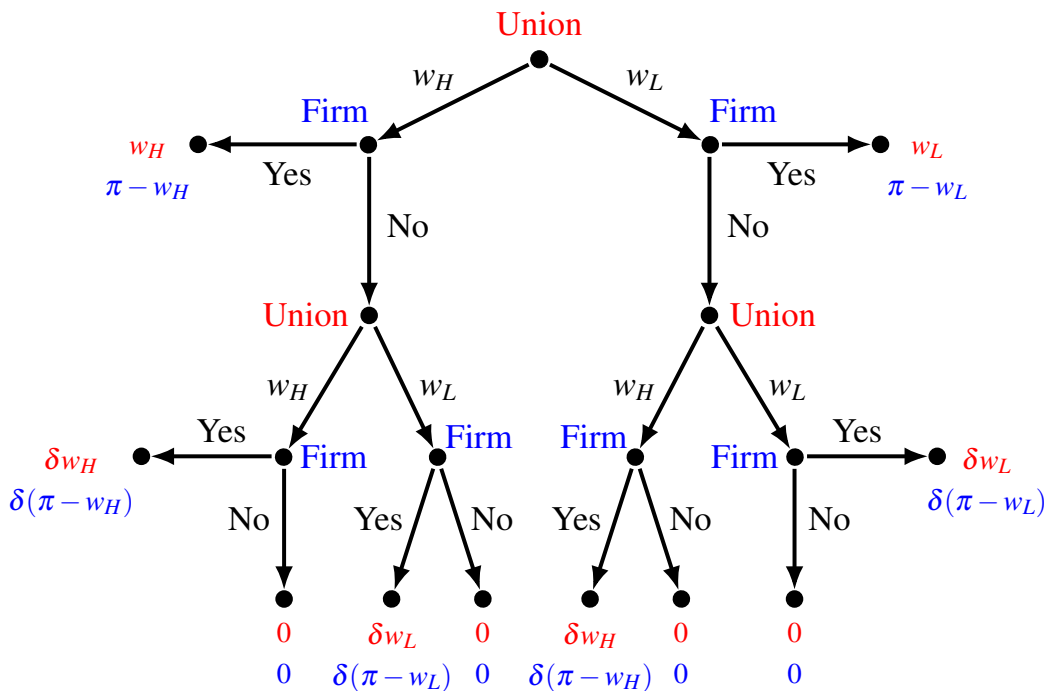


Figure 15.12: The structure of the wage bargaining game.

Suppose, however, that the union does not know the firm's expected profit π (gross of labor costs). It is not unlikely that the management will know what profits to expect, while the union will not (because, for example, the union does not have enough information on the intensity of competition, the firm's non-labor costs, etc.).

Suppose that we have a one-sided incomplete information situation where the union believes that π can have two values: π_H (high) and π_L (low) and assigns probability α to π_H and $(1-\alpha)$ to π_L .

This situation of one-sided incomplete information is illustrated in Figure 15.13.




Figure 15.13: The one-sided situation of incomplete information involving the wage bargaining game of Figure 15.12.

Let $\pi_H > \pi_L > 0$ and $w_H > w_L > 0$ (so that H means “high” and L means “low”).

We also assume that

- $\pi_H - w_H > 0$ (so that the high-profit firm could in fact afford to pay a high wage),
- $\pi_L - w_L > 0$ (that is, the low wage is sufficiently low for the low-profit firm to be able to pay it), and
- $\pi_L - w_H < 0$ (that is, the low-profit firm cannot afford to pay a high wage).

Finally, we assume that the true state is the one where $\pi = \pi_L$, that is, the firm's potential profits are in fact low. These assumptions imply that it is in the interest of both the firm and the union to sign a contract (even if the firm's profits are low).

Using the Harsanyi transformation we can convert the situation shown in Figure 15.13 into the extensive-form game shown in Figure 15.14.

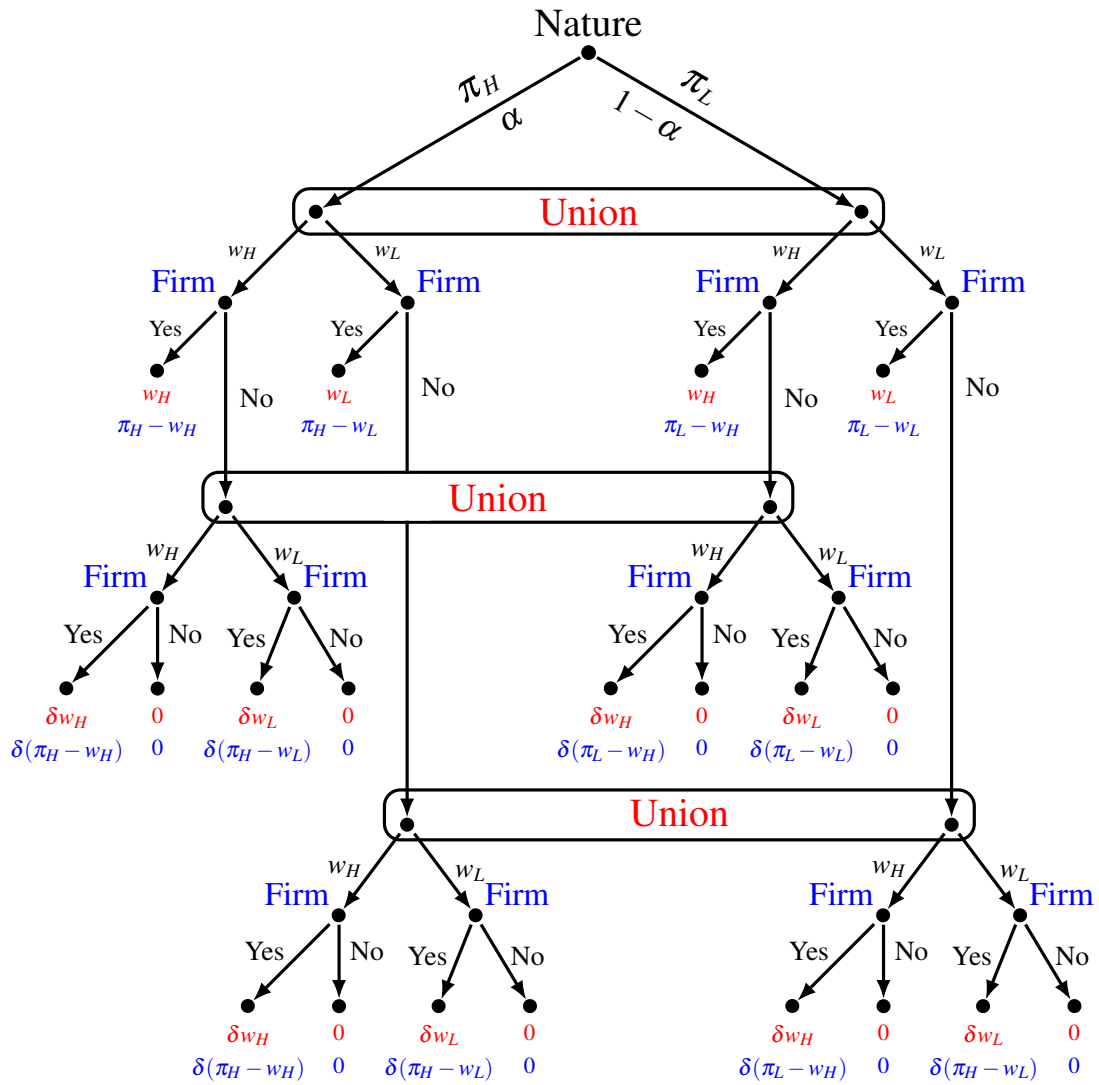


Figure 15.14: The game obtained by applying the Harsanyi transformation to the situation of incomplete information of Figure 15.13.

Let us find conditions under which the following strategy profile is part of a separating weak sequential equilibrium:

1. The union requests a high wage in the first period.
2. The high-profit firm accepts while the low-profit firm rejects.
3. After rejection of the first-period high-wage offer the union requests a low wage and the firm accepts.

That is, we want an equilibrium where the low-profit firm endures a strike to signal to the union that its expected profits are low and cannot afford to pay a high wage. The union reacts to the signal by lowering its demand.

Of course, we need to worry about the fact that the high-profit firm might want to masquerade as a low-profit firm by rejecting the first-period offer (that is, by sending the same signal as the low-profit firm).

First of all, we can simplify the game by replacing the second-period choice of the firm with the outcome that follows the best choice for the firm (which is “accept w_L ” for both the π_H and the π_L firm, “accept w_H ” for the π_H firm and “reject w_H ” for the π_L firm).

This space left blank
to allow for proper
alignment of text and figures

The simplified game is shown in Figure 15.15.

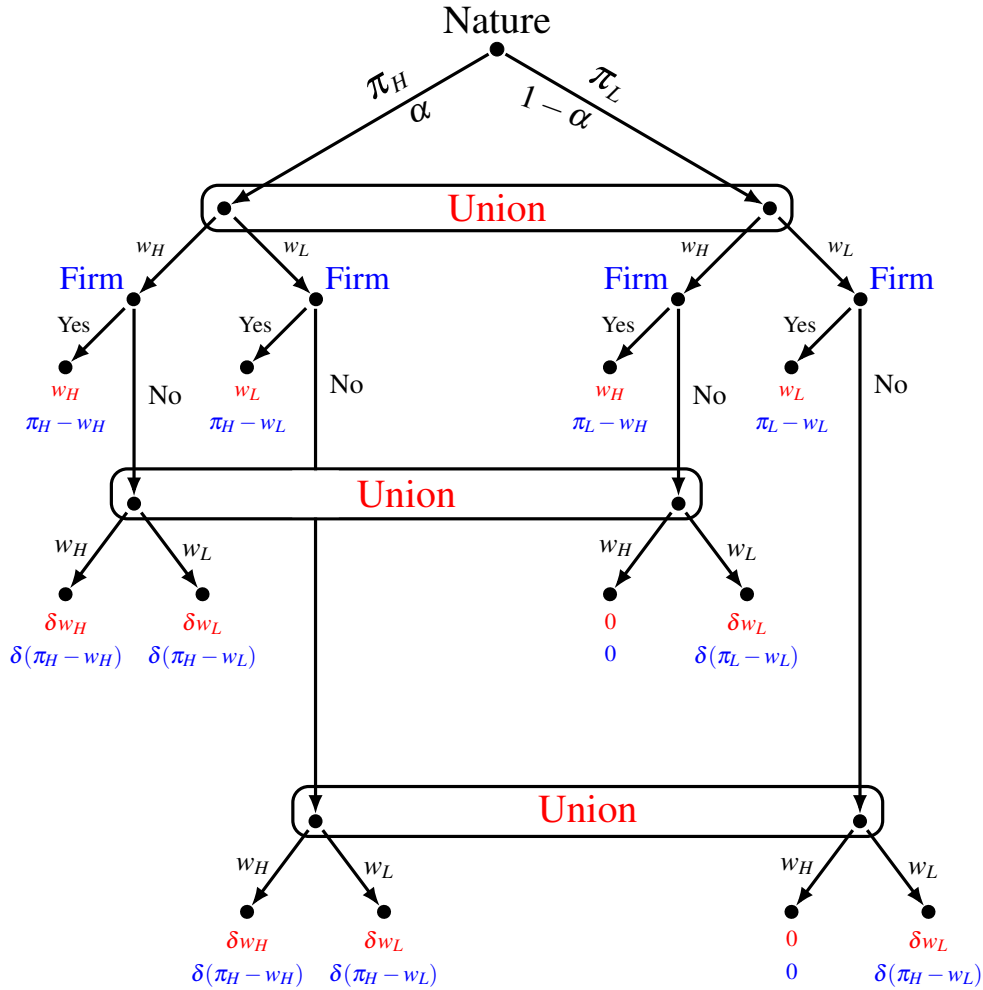


Figure 15.15: The simplified game after eliminating the second-period choice of the firm.

In order for a low-wage offer to be optimal for the union in Period 2 it is necessary for the union to assign sufficiently high probability to the firm being a low-profit one: if p is this probability, then we need the following (for example, $p = 1$ would be fine):

$$\delta w_L \geq (1 - p)\delta w_H, \quad \text{that is,} \quad p \geq 1 - \frac{w_L}{w_H}$$

For the high-profit firm not to have an incentive to send the same signal as the low-profit firm it is necessary that

$$\pi_H - w_H \geq \delta (\pi_H - w_L), \quad \text{that is,} \quad \delta \leq \frac{\pi_H - w_H}{\pi_H - w_L}.$$

If this condition is satisfied, then the high-profit firm will accept w_H in period 1, the low-profit firm will reject and Bayesian updating requires the union to assign probability 1 to the right node of its middle information set (that is, to the firm being a low-profit type), in which case adjusting its demand to w_L is sequentially rational.

On the other hand, requesting a low wage in period 1 will lead to both types of the firm accepting immediately. So in order for a high-wage offer to be optimal for the union in period 1 it is necessary that $w_L \leq \alpha w_H + (1 - \alpha)\delta w_L$, which will be true if α is sufficiently large, that is, if

$$\alpha \geq \frac{(1 - \delta)w_L}{w_H - \delta w_L}.$$

Finally, given the above equilibrium choices, the notion of weak sequential equilibrium imposes no restrictions on the beliefs (and hence choice) of the union at the bottom information set.

For example, all the above inequalities are satisfied if:

$$\pi_H = 100, \quad \pi_L = 55, \quad w_H = 60, \quad w_L = 50, \quad \alpha = 0.7, \quad \delta = 0.6.$$

Given the above values, $\pi_H - w_H = 40 > \delta(\pi_H - w_L) = 0.6(50) = 30$
and $w_L = 50 < \alpha w_H + (1 - \alpha)\delta w_L = (0.7)60 + 0.3(0.6)50 = 51$.

We conclude this section with one more example that has to do with the effectiveness of truth-in-advertising laws.

Consider the case of a seller and a buyer. The seller knows the quality x of his product, while the buyer does not, although she knows that he knows. The buyer can thus ask the seller to reveal the information he has. Suppose that there is a truth-in-advertising law which imposes harsh penalties for false claims. This is not quite enough, because the seller can tell the truth without necessarily revealing all the information.

For example, if x is the octane content of gasoline and $x = 89$, then the following are all true statements:

- “the octane content of this gasoline is at least 70”
- “the octane content of this gasoline is at least 85”
- “the octane content of this gasoline is at most 89”
- “the octane content of this gasoline is exactly 89”, etc.

Other examples are:

- the fat content of food (the label on a package of ground meat might read “not more than 30% fat”),
- fuel consumption for cars (a car might be advertised as yielding “at least 35 miles per gallon”),
- the label on a mixed-nut package might read “not more than 40% peanuts”, etc.

An interesting question is: Would the seller reveal all the information he has or would he try to be as vague as possible (while being truthful)? In the latter case, will the buyer be in a worse position than she would be in the case of complete information?

Milgrom and Roberts (1986) consider the following game (resulting from applying the Harsanyi transformation to the situation of one-sided incomplete information described above).

1. First Nature selects the value of x , representing the seller's information, from a finite set X ; as usual, the probabilities with which Nature chooses reflect the buyer's beliefs.
2. The seller "observes" Nature's choice and makes an assertion A to the buyer; A is a subset of X . The seller is restricted to make true assertions, that is, we require that $x \in A$.
3. The buyer observes the seller's claim A and then selects a quantity $q \geq 0$ to buy.

Milgrom and Roberts show that (under reasonable hypotheses), if the buyer adopts a skeptical view concerning the seller's claim, that is, she always interprets the seller's claim in a way which is least favorable to the seller (for example "not more than 30% fat" is interpreted as "exactly 30% fat"), then there is an equilibrium where the outcome is the same as it would be in the case of complete information. We shall illustrate this phenomenon by means of a simple example.

Suppose that there are three possible quality levels for the good under consideration: low (l), medium (m) and high (h); thus $X = \{l, m, h\}$. The buyer believes that the three quality levels are equally likely. The buyer has to choose whether to buy one unit or two units. The seller can only make truthful claims. For example, if the quality is l , the seller's possible claims are $\{l\}$ (full revelation) or vague claims such as $\{l, m\}$, $\{l, h\}$, $\{l, m, h\}$. The extensive-form game is shown in Figure 15.16.

It is straightforward to check that the following is a weak sequential equilibrium:

- The seller claims $\{l\}$ if Nature chooses l , $\{m\}$ if Nature chooses m and $\{h\}$ if Nature chooses h (that is, the seller reveals the whole truth by choosing not to make vague claims).
- The buyer buys one unit if told $\{l\}$, two units if told $\{m\}$ or if told $\{h\}$; furthermore, the buyer adopts beliefs which are least favorable to the seller:
 - if told $\{l, m\}$ or $\{l, h\}$ or $\{l, m, h\}$ she will believe l with probability one and buy one unit,
 - if told $\{m, h\}$ she will believe m with probability one and buy two units.

All these beliefs are admissible because they concern information sets that are not reached in equilibrium and therefore Bayesian updating does not apply.

On the other hand, as shown in Exercise 15.4, if the buyer is naïve then there is an equilibrium where the seller chooses to make vague claims.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 15.3.1 at the end of this chapter.




Figure 15.16: The buyer-seller game representing a situation of one-sided incomplete information.

15.2 Multi-sided incomplete information

The case of situations of multi-sided incomplete information involving dynamic games is conceptually the same as the case of multi-sided situations of incomplete information involving static games. In this section we shall go through one example.

Consider the following situation of two-sided incomplete information. A seller (player S) owns an item that a buyer (player B) would like to purchase. The seller's reservation price is s (that is, she is willing to sell if and only if the price paid by the buyer is at least s) and the buyer's reservation price is b (that is, he is willing to buy if and only if the price is less than or equal to b). It is common knowledge between the two that

- both b and s belong to the set $\{1, 2, \dots, n\}$,
- the buyer knows the value of b and the seller knows the value of s ,
- both the buyer and the seller attach equal probability to all the possibilities among which they are uncertain.

Buyer and Seller play the following game. First the buyer makes an offer of a price $p \in \{1, \dots, n\}$ to the seller. If $p = n$ the game ends and the object is exchanged for $\$p$.

If $p < n$ then the seller either accepts (in which case the game ends and the object is exchanged for p) or makes a counter-offer of $p' > p$, in which case either the buyer accepts (and the game ends and the object is exchanged for p') or the buyer rejects, in which case the game ends without an exchange.

Payoffs are as follows:

$$\pi_{seller} = \begin{cases} 0 & \text{if there is no exchange} \\ x - s & \text{if exchange takes place at price } \$x \end{cases}$$

$$\pi_{buyer} = \begin{cases} 0 & \text{if there is no exchange} \\ b - p & \text{if exchange takes place at price } \$p \text{ (the initial offer)} \\ b - p' - \varepsilon & \text{if exchange takes place at price } \$p' \text{ (the counter-offer)} \end{cases}$$

where $\varepsilon > 0$ is a measure of the buyer's "hurt feelings" for seeing his initial offer rejected.

These are von Neumann-Morgenstern payoffs.

Let us start by focusing on the case $n = 2$. First we represent the situation described above by means of an interactive knowledge-belief structure. A possible state can be written as a pair (b, s) , where b is the reservation price of the buyer and s that of the seller.

Thus when $n = 2$ the possible states are $(1,1)$, $(1,2)$, $(2,1)$ and $(2,2)$. Figure 15.17 represents this two-sided situation of incomplete information.

For each state there is a corresponding extensive-form game. The four possible games are shown in Figure 15.18 (the top number is the Buyer's payoff and the bottom number the Seller's payoff).




Figure 15.17: The two-sided situation of incomplete information between a buyer and a seller.

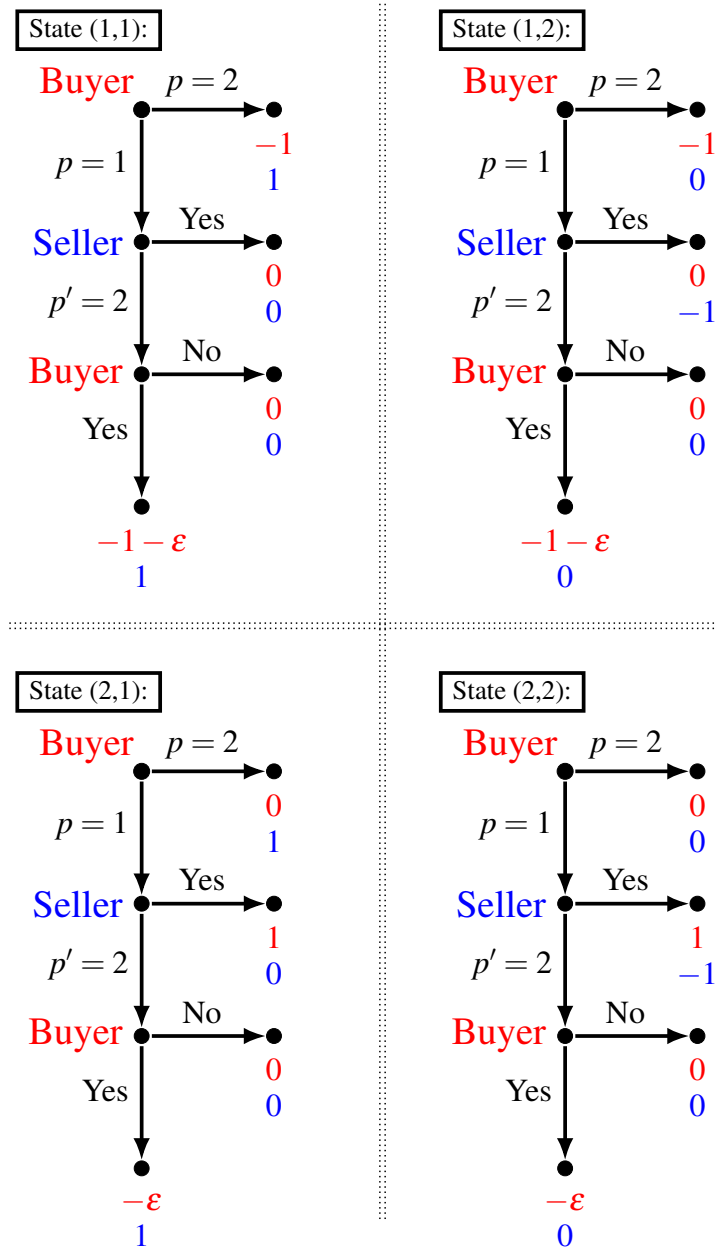


Figure 15.18: The four games corresponding to the four states of Figure 15.17.

The extensive-form game that results from applying the Harsanyi transformation to the situation illustrated in Figure 15.17 is shown in Figure 15.19.

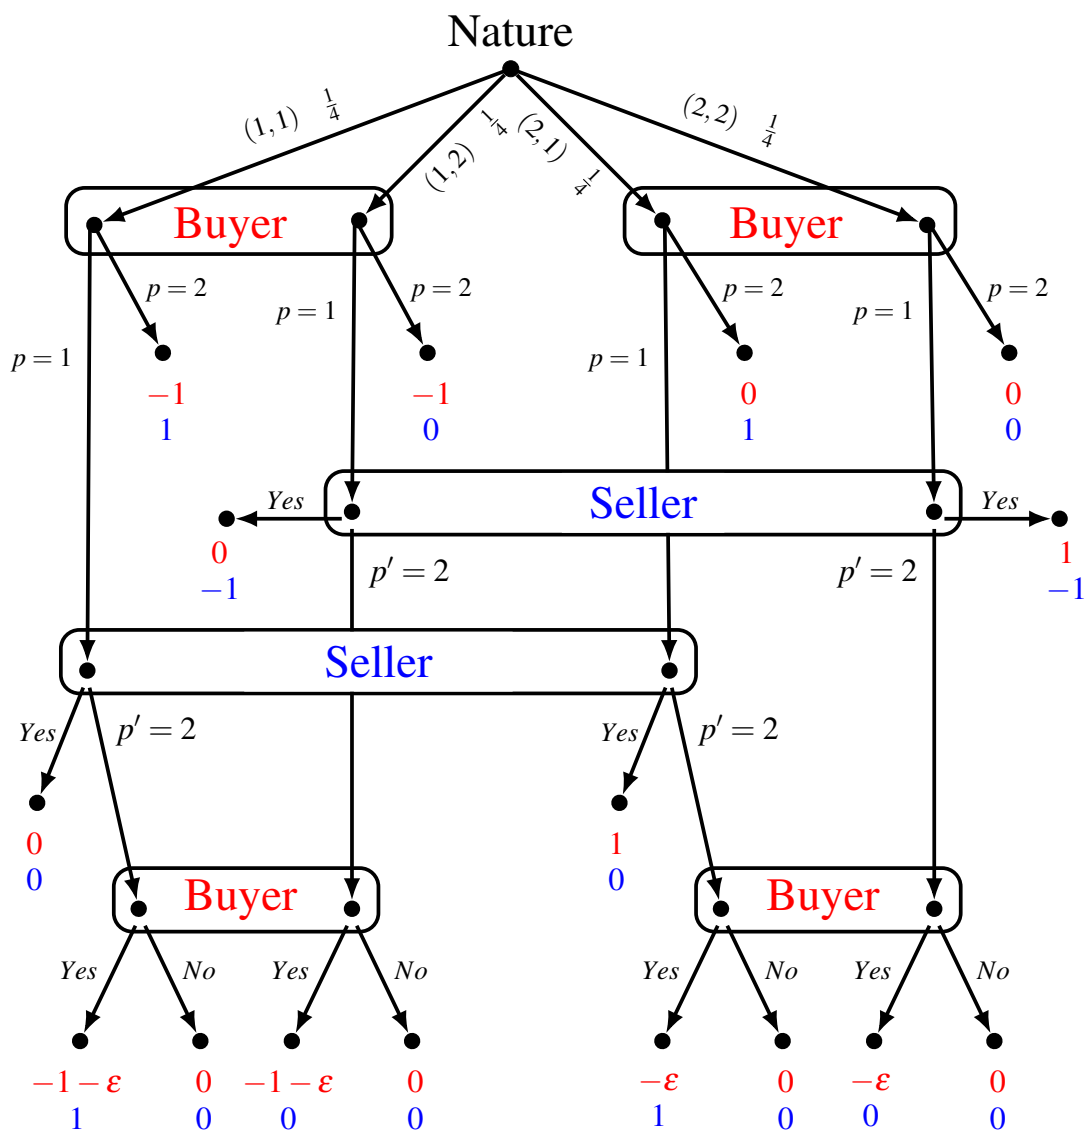


Figure 15.19: The extensive-form game that results from applying the Harsanyi transformation to the situation illustrated in Figure 15.17.

Let us find all the pure-strategy weak sequential equilibria of the game of Figure 15.19. First of all, note that at the bottom information sets of the Buyer, “Yes” is strictly dominated by “No” and thus a weak sequential equilibrium must select “No”.

Hence the game simplifies to the one shown in Figure 15.20.

In the game of Figure 15.20, at the middle information of the Seller, making a counteroffer of $p' = 2$ strictly dominates “Yes”.

Thus the game can be further simplified as shown in Figure 15.21.




Figure 15.20: The reduced game after eliminating the strictly dominated choices.

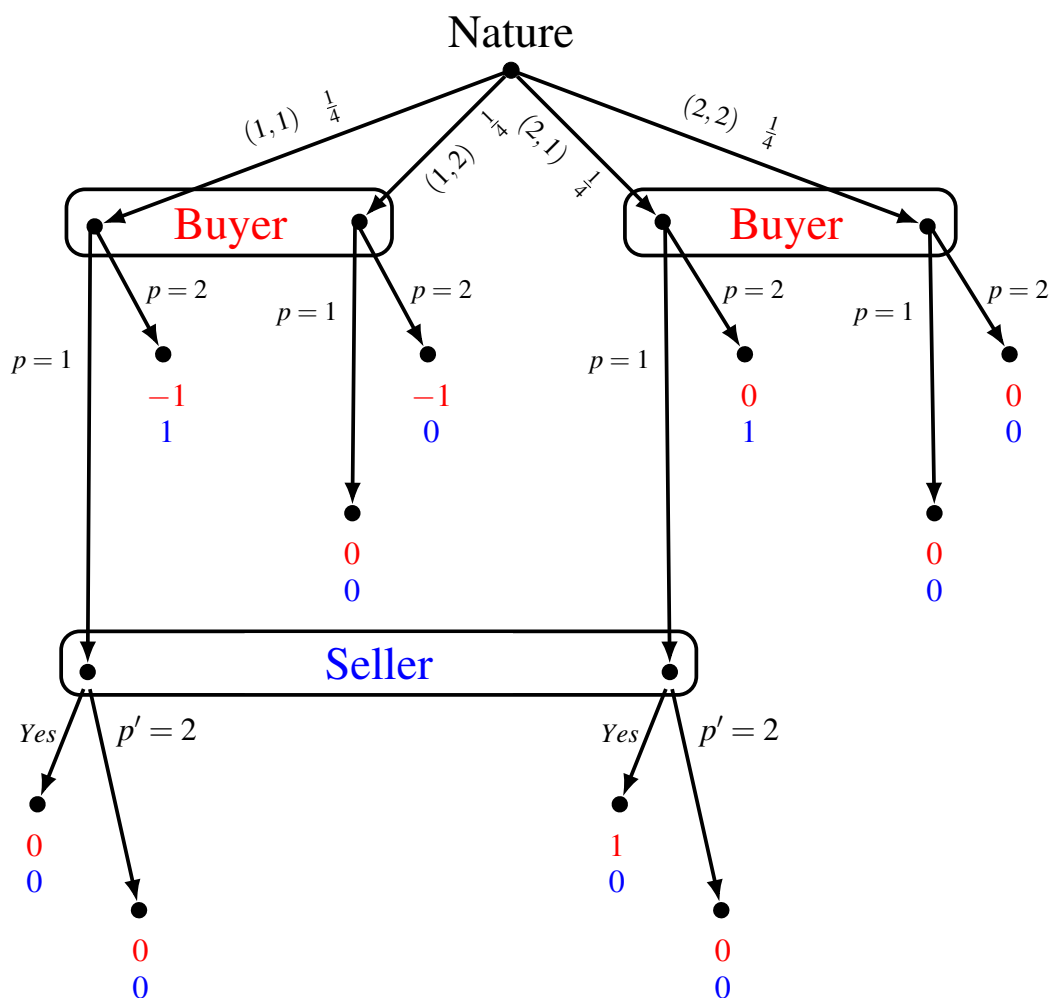


Figure 15.21: The further reduced game after eliminating strictly dominated choices from the game of Figure 15.20.

In the game of Figure 15.21 at the left information set of the Buyer $p = 1$ strictly dominates $p = 2$. At the right information set the Buyer's beliefs must be $\frac{1}{2}$ on each node so that if the Seller's strategy is to say "Yes", then $p = 1$ is the only sequentially rational choice there, otherwise both $p = 1$ and $p = 2$ are sequentially rational.

Thus the pure-strategy weak sequential equilibria of the reduced game shown in Figure 15.21 are as follows:

1. $((p = 1, p = 1), \text{Yes})$ with beliefs given by probability $\frac{1}{2}$ on each node at every information set.
2. $((p = 1, p = 1), p' = 2)$ with beliefs given by probability $\frac{1}{2}$ on each node at every information set.
3. $((p = 1, p = 2), p' = 2)$ with beliefs given by (i) probability $\frac{1}{2}$ on each node at both information sets of the Buyer and (ii) probability 1 on the left node at the information

set of the Seller.

These equilibria can be extended to the original game of Figure 15.19 as follows:

1. $((p = 1, p = 1, No, No), (p' = 2, Yes))$ with beliefs given by (i) probability $\frac{1}{2}$ on each node at both information sets of the Buyer at the top and at both information sets of the Seller and (ii) probability 1 on the right node at both information sets of Buyer at the bottom. The corresponding payoffs are: $\frac{1}{4}$ for the Buyer and 0 for the Seller.
2. $((p = 1, p = 1, No, No), (p' = 2, p' = 2))$ with beliefs given by probability $\frac{1}{2}$ on each node at every information set. The corresponding payoffs are 0 for both Buyer and Seller.
3. $((p = 1, p = 2, No, No), (p' = 2, p' = 2))$ with beliefs given by
 - (i) probability $\frac{1}{2}$ on each node at both information sets of the Buyer at the top and at the lower left information set of the Buyer,
 - (ii) any beliefs at the lower right information set of the Buyer and
 - (iii) probability 1 on the left node at each information set of the Seller.
 The corresponding payoffs are: 0 for the Buyer and $\frac{1}{4}$ for the Seller.

Now let us consider the case $n = 100$. Drawing the interactive knowledge-belief structure and the corresponding extensive-form game (obtained by applying the Harsanyi transformation) is clearly not practical. However, one can still reason about what could be a pure-strategy Bayesian Nash equilibrium of that game. As a matter of fact, there are many Bayesian Nash equilibria.

One of them is the following:

- each type of the Buyer offers a price equal to his reservation price,
- the Seller accepts if, and only if, that price is greater than or equal to her reservation price,
- at information sets of the Seller that are not reached, the Seller rejects and counteroffers \$100,
- at information sets of the Buyer that are not reached the Buyer says “No”.

The reader should convince himself/herself that this is indeed a Nash equilibrium.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 15.3.2 at the end of this chapter.

15.3 Exercises

15.3.1 Exercises for Section 15.1: One-sided incomplete information

The answers to the following exercises are in Section 15.4 at the end of this chapter.




Figure 15.22: The extensive form for Exercise 15.1.

Exercise 15.1

Consider the following situation of one-sided incomplete information. Players 1 and 2 are playing the extensive-form game shown in Figure 15.22 (where z_1, z_2 , etc. are outcomes and the question mark stands for either outcome z_5 or outcome z_6). The outcome that is behind the question mark is actually outcome z_5 and Player 1 knows this, but Player 2 does not know. Player 2 thinks that the outcome behind the question mark is either z_5 or z_6 and assigns probability 25% to it being z_5 and probability 75% to it being z_6 . Player 2 also thinks that whatever the outcome is, Player 1 knows (that is, if it is z_5 , then Player 1 knows that it is z_5 , and if it is z_6 then Player 1 knows that it is z_6). The beliefs of Player 2 are common knowledge between the two players.

- Represent this situation of incomplete information using an interactive knowledge-belief structure.
- Apply the Harsanyi transformation to transform the situation represented in part (a) into an extensive-form frame. [Don't worry about payoffs for the moment.]

From now on assume that the following is common knowledge:

1. Both players have von Neumann-Morgenstern preferences,
2. The ranking of Player 1 is

$$\left(\begin{array}{c|c|c} \textit{best} & \textit{second} & \textit{worst} \\ \hline z_4, z_6 & z_1 & z_2, z_3, z_5 \end{array} \right)$$

and he is indifferent between z_1 for sure and the lottery $\begin{pmatrix} z_6 & z_5 \\ 0.5 & 0.5 \end{pmatrix}$.

3. The ranking of player 2 is

$$\left(\begin{array}{c|c|c|c} \textit{best} & \textit{second} & \textit{third} & \textit{worst} \\ \hline z_6 & z_4 & z_2, z_5 & z_1, z_3 \end{array} \right)$$

and she is indifferent between z_4 for sure and the lottery $\begin{pmatrix} z_6 & z_3 \\ 0.5 & 0.5 \end{pmatrix}$

and she is also indifferent between z_2 for sure and the lottery $\begin{pmatrix} z_6 & z_3 \\ 0.25 & 0.75 \end{pmatrix}$

- (c) Calculate the von Neumann-Morgenstern normalized utility functions for the two players.
- (d) Is there a weak sequential equilibrium of the game of part (b) where Player 1 always plays A (thus a pooling equilibrium)?
- (e) Is there a weak sequential equilibrium of the game of part (b) where Player 1 always plays C (thus a pooling equilibrium)?
- (f) Is there a pure-strategy weak sequential equilibrium of the game of part (b) where Player 1 does not always choose the same action (thus a separating equilibrium)?



Exercise 15.2

Show that the assessment highlighted in the Figure 15.23 (which reproduces Figure 15.11) is a sequential equilibrium as long as $p \geq \frac{1}{3}$. [The assessment is fully described on page 528.]

Exercise 15.3

Consider again the game of Figure 15.23 (which reproduces Figure 15.11).

- (a) For the case where $p \leq \frac{1}{3}$, find a pure-strategy weak sequential equilibrium where PE-1 stays out but PE-2 enters; furthermore, the Incumbent would fight PE-1 if she entered. Prove that what you propose is a weak sequential equilibrium.
- (b) Either prove that the weak sequential equilibrium of part (a) is a sequential equilibrium or prove that it is not a sequential equilibrium.

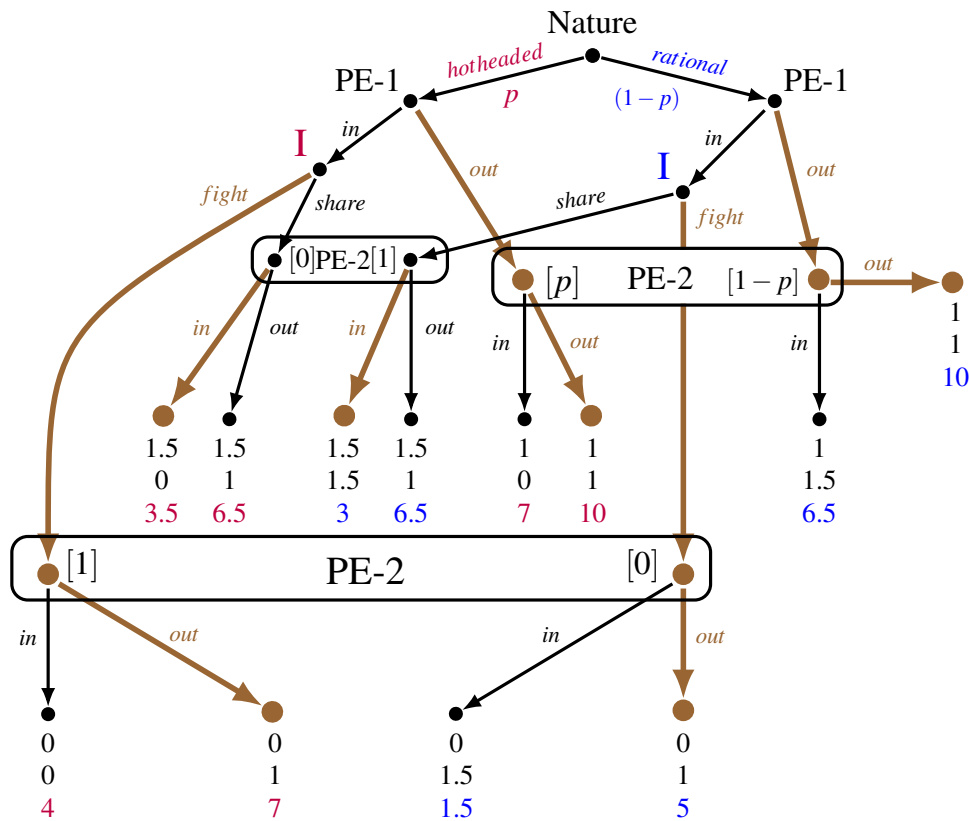
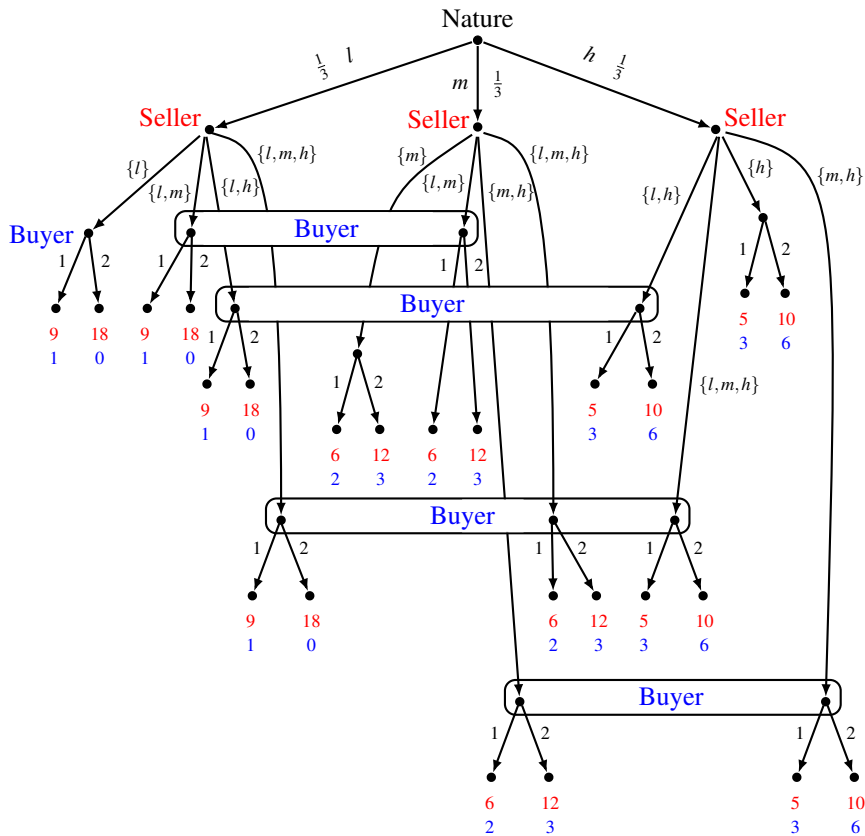


Figure 15.23: Copy of Figure 15.11

Exercise 15.4

Show that in the “truth-in-advertising” game of Figure 15.16, (which is reproduced on the next page), there is a weak sequential equilibrium where the seller makes vague claims (that is, does not reveal the whole truth).



Exercise 15.5

Consider a simpler version of the “truth-in-advertising” game, where there are only two quality levels: L and H . The payoffs are as follows:

- If the quality is L and the buyer buys one unit, the seller’s payoff is 9 and the buyer’s payoff is 1,
- If the quality is L and the buyer buys two units, the seller’s payoff is 18 and the buyer’s payoff is 0,
- If the quality is H and the buyer buys one unit, the seller’s payoff is 6 and the buyer’s payoff is 2,
- If the quality is H and the buyer buys two units, the seller’s payoff is 12 and the buyer’s payoff is 3,

Let the buyer’s initial beliefs be as follows: the good is of quality L with probability $p \in (0, 1)$ (and of quality H with probability $1 - p$).

- Draw the extensive-form game that results from applying the Harsanyi transformation to this one-sided situation of incomplete information.
- Find the pure-strategy subgame-perfect equilibria of the game of part (a) for every possible value of $p \in (0, 1)$.



15.3.2 Exercises for Section 15.2: Multi-sided incomplete information

The answers to the following exercises are in Section 15.4 at the end of this chapter.

Exercise 15.6

Consider the following situation of two-sided incomplete information. Players 1 and 2 are having dinner with friends and during dinner Player 2 insinuates that Player 1 is guilty of unethical behavior.

- Player 1 can either demand an apology (D) or ignore (I) Player 2's remark;
- if Player 1 demands an apology, then Player 2 can either apologize (A) or refuse to apologize (not- A);
- if Player 2 refuses to apologize, Player 1 can either concede (C) or start a fight (F).

Thus the sequence of moves is as shown in Figure 15.24 (where z_1, \dots, z_4 are the possible outcomes).

Let U_i be the von Neumann-Morgenstern utility function of Player i ($i = 1, 2$). The following is common knowledge between Players 1 and 2:

- $U_1(z_1) = 4, \quad U_1(z_2) = 2, \quad U_1(z_3) = 6$
- $U_2(z_1) = 6, \quad U_2(z_2) = 8, \quad U_2(z_3) = 2, \quad U_2(z_4) = 0.$
- The only uncertainty concerns the value of $U_1(z_4)$. As a matter of fact, $U_1(z_4) = 0$; Player 1 knows this, but Player 2 does not. It is common knowledge between Players 1 and 2 that Player 2 thinks that either $U_1(z_4) = 0$ or $U_1(z_4) = 3$.


Furthermore,

- Player 2 assigns probability $\frac{1}{4}$ to $U_1(z_4) = 0$ and probability $\frac{3}{4}$ to $U_1(z_4) = 3$;
- Player 1 is uncertain as to whether Player 2's beliefs are

$$\begin{pmatrix} U_1(z_4) = 0 & U_1(z_4) = 3 \\ \frac{1}{4} & \frac{3}{4} \end{pmatrix} \quad \text{or} \quad \begin{pmatrix} U_1(z_4) = 0 \\ 1 \end{pmatrix}$$

and he attaches probability $\frac{1}{2}$ to each.

- (a) Construct a three-state interactive knowledge-belief structure that captures all of the above.
- (b) Draw the extensive-form game that results from applying the Harsanyi transformation to the incomplete-information situation of part (a).
- (c) Find the pure-strategy weak sequential equilibria of the game of part (b).



I = ignore D = demand apology
 A = apologize $\text{not-}A$ = not apologize
 C = concede F = fight

Figure 15.24: The extensive form for Exercise 15.6.

Exercise 15.7 — *Challenging Question***.**

There are two parties to a potential lawsuit: the owner of a chemical plant and a supplier of safety equipment. The chemical plant owner, from now on called the plaintiff, alleges that the supplier, from now on called the defendant, was negligent in providing the safety equipment. The defendant knows whether or not he was negligent, while the plaintiff does not know; the plaintiff believes that there was negligence with probability q . These beliefs are common knowledge between the parties.

The plaintiff has to decide whether or not to sue. If she does not sue then nothing happens and both parties get a payoff of 0. If the plaintiff sues then the defendant can either offer an out-of-court settlement of $\$S$ or resist. If the defendant offers a settlement, the plaintiff can either accept (in which case her payoff is S and the defendant's payoff is $-S$) or go to trial. If the defendant resists then the plaintiff can either drop the case (in which case both parties get a payoff of 0) or go to trial. If the case goes to trial then legal costs are created in the amount of $\$P$ for the plaintiff and $\$D$ for the defendant.

Furthermore (if the case goes to trial), the judge is able to determine if there was negligence and, if there was, requires the defendant to pay $\$W$ to the plaintiff (and each party has to pay its own legal costs), while if there was no negligence the judge will drop the case without imposing any payments to either party (but each party still has to pay its own legal costs). It is common knowledge that each party is “selfish and greedy” (that is, only cares about its own wealth and prefers more money to less) and is risk neutral.

Assume the following about the parameters:

$$0 < q < 1, \quad 0 < D < S, \quad 0 < P < S < W - P.$$

- (a) Represent this situation of incomplete information by means of an interactive knowledge-belief structure (the only two players are the plaintiff and the defendant).
- (b) Apply the Harsanyi transformation to represent the situation in part (a) as an extensive-form game. [Don't forget to subtract the legal expenses from each party's payoff if the case goes to trial.]
- (c) Write all the pure strategies of the plaintiff.
- (d) Prove that there is no pure-strategy weak sequential equilibrium which (1) is a separating equilibrium and (2) involves suing.
- (e) For what values of the parameters (q, S, P, W, D) are there pure-strategy weak sequential equilibria which (1) are pooling equilibria and (2) involve suing? Consider all types of pooling equilibria and prove your claim.
- (f) Now drop the assumption that $S < W - P$. For the case where $q = \frac{1}{12}$, $P = 70$, $S = 80$ and $W = 100$ find all the pure-strategy weak sequential equilibria which (1) are pooling equilibria and (2) involve suing.

15.4 Solutions to Exercises

Solutions to Exercise 15.1

(a) Let G_1 be the game with outcome z_5 and G_2 the game with outcome z_6 . Then the structure is as shown in Figure 15.25, with $q = \frac{1}{4}$. The true state is α .




Figure 15.25: The one-sided incomplete-information situation of Exercise 15.1.

(b) The extensive form is shown in Figure 15.26.




Figure 15.26: The extensive form obtained by applying the Harsanyi transformation to the incomplete-information situation shown in Figure 15.25.

(c) The von Neumann-Morgenstern utility functions are as follows:

$$\begin{pmatrix} z_1 & z_2 & z_3 & z_4 & z_5 & z_6 \\ U_1: & 0.5 & 0 & 0 & 1 & 0 & 1 \\ U_2: & 0 & 0.25 & 0 & 0.5 & 0.25 & 1 \end{pmatrix}.$$

Adding these payoffs to the extensive form we obtain the game shown in Figure 15.27.




Figure 15.27: The game obtained by adding the payoffs to the game-frame of Figure 15.26.

- (d) No, because at the right node of Player 1, C gives a payoff of 1 no matter what Player 2 does and thus Player 1 would choose C rather than A (which only gives him a payoff of 0.5).
- (e) If Player 1 always plays C (that is, his strategy is CC), then – by Bayesian updating – Player 2 should assign probability $\frac{1}{4}$ to node y and probability $\frac{3}{4}$ to node z , in which case D gives her a payoff of 0.5 while E gives her a payoff of $\frac{1}{4}(0.25) + \frac{3}{4}(1) = 0.8125$; hence she must choose E . But then at his left node Player 1 with C gets 0 while with A he gets 0.5. Hence choosing C at the left node is not sequentially rational.
- (f) Yes, the following is a weak sequential equilibrium:
 $\sigma = (AC, E)$ with beliefs $\mu = \begin{pmatrix} x & y & w & z \\ 0 & 0 & 0 & 1 \end{pmatrix}$. It is straightforward to verify that sequential rationality and Bayesian updating are satisfied. □

Solutions to Exercise 15.2 Sequential rationality was verified in Section 15.1. Thus we only need to show that the highlighted pure-strategy profile together with the following system of beliefs $\mu = \left(\begin{array}{cc|cc|cc} s & t & u & v & x & y \\ p & 1-p & 0 & 1 & 1 & 0 \end{array} \right)$ (for the names of the nodes refer to Figure 15.28) constitutes a consistent assessment (Definition 12.1.1, Chapter 12).




Figure 15.28: The game for Exercise 15.2.

Let $\langle \sigma_n \rangle_{n=1,2,\dots}$ be the sequence of completely mixed strategies shown in Figure 15.28. It is clear that $\lim_{n \rightarrow \infty} \sigma_n = \sigma$. Let us calculate the system of beliefs μ_n obtained from σ_n by using Bayesian updating:

$$\begin{aligned} \mu_n(s) &= \frac{p(1-\frac{1}{n})}{p(1-\frac{1}{n}) + (1-p)(1-\frac{1}{n^2})} \quad \text{thus} \quad \lim_{n \rightarrow \infty} \mu_n(s) = \frac{p}{p+(1-p)} = p; \\ \mu_n(u) &= \frac{p(\frac{1}{n})(\frac{1}{n^3})}{p(\frac{1}{n})(\frac{1}{n^3}) + (1-p)(\frac{1}{n^2})(\frac{1}{n})} = \frac{p}{p+n(1-p)} \quad \text{thus} \quad \lim_{n \rightarrow \infty} \mu_n(u) = 0; \\ \mu_n(x) &= \frac{p(\frac{1}{n})(1-\frac{1}{n^3})}{p(\frac{1}{n})(1-\frac{1}{n^3}) + (1-p)(\frac{1}{n^2})(1-\frac{1}{n})} = \frac{p(1-\frac{1}{n^3})}{p(1-\frac{1}{n^3}) + (1-p)(\frac{1}{n})(1-\frac{1}{n})}. \end{aligned}$$

Thus, $\lim_{n \rightarrow \infty} \mu_n(x) = 1$ (the numerator tends to $p(1) = p$ and the denominator tends to $p(1) + (1-p)(0)(1) = p$). Hence (σ, μ) is consistent and sequentially rational and therefore it is a sequential equilibrium. \square

Solutions to Exercise 15.3

(a) Consider the following pure-strategy profile, call it σ , which is highlighted by thick arrows in Figure 15.29:

- $PE-1$'s strategy is "out" at both nodes,
- $PE-2$'s strategy is
 - "in" at the top information set (after having observed that $PE-1$ stayed out),
 - "in" at the middle information set (after having observed that $PE-1$'s entry was followed by the Incumbent sharing the market with $PE-1$),
 - "out" at the bottom information set (after having observed that $PE-1$'s entry was followed by the Incumbent fighting against $PE-1$).
- The Incumbent's strategy is to fight entry of $PE-1$ in any case (that is, whether the Incumbent himself is hotheaded or rational).

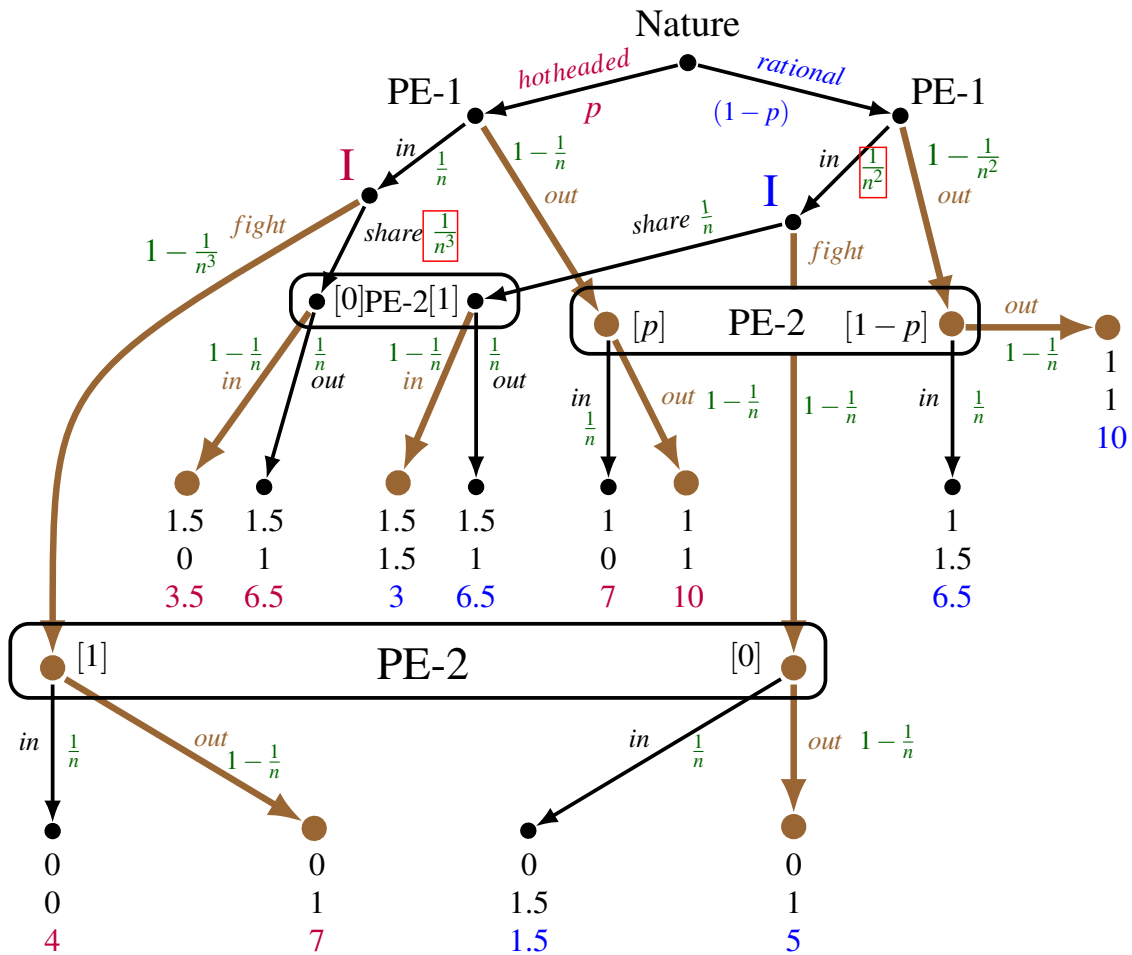


Figure 15.29: The game for Exercise 15.3.

We want to show that σ , together with the system of beliefs μ shown in Figure 15.29 (where at her top information $PE-2$ assigns probability p to the Incumbent being hotheaded, at her middle information set she assigns probability 1 to the Incumbent being rational and at her bottom information set she assigns probability 1 to the Incumbent being hotheaded) constitutes a weak sequential equilibrium for any value of $p \leq \frac{1}{3}$.


Bayesian updating is satisfied at $PE-2$'s top information set (the only non-singleton information set reached by σ), and at the other two information sets of $PE-2$ any beliefs are allowed by the notion of weak sequential equilibrium.

Sequential rationality also holds:

- For $PE-1$, at the left node “in” yields 0 and “out” yields 1, so that “out” is sequentially rational; the same is true at the right node.
 - For the Incumbent,
 - at the left node “fight” yields 7 and “share” yields 3.5, so that “fight” is sequentially rational;
 - at the right node “fight” yields 5 and “share” yields 3, so that “fight” is sequentially rational.
 - For $PE-2$,
 - at the top information set “in” yields an expected payoff of $p(0) + (1 - p)(1.5) = (1 - p)(1.5)$ and “out” yields 1, so that “in” is sequentially rational as long as $p \leq \frac{1}{3}$;
 - at the middle information set “in” yields 1.5 and “out” yields 1, so that “in” is sequentially rational;
 - at the bottom information set “in” yields 0 and “out” yields 1, so that “out” is sequentially rational.
- (b) The assessment described in part (a) is in fact a sequential equilibrium. This can be shown using the sequence of completely mixed strategies marked in Figure 15.29, which coincides with the sequence of mixed strategies considered in Exercise 15.2; thus the calculations to show consistency are identical to the ones carried out in Exercise 15.2.

□

Solutions to Exercise 15.4 The game under consideration is the following (which reproduces Figure 15.16):



Let the buyer be naïve, in the sense that at every unreached information set she assigns equal probability to each node; furthermore, let the buyer's strategy be as follows: if told $\{1\}$, buy one unit, in every other case buy two units.

Let the seller's strategy be as follows: if Nature chooses l or m , claim $\{l, m\}$ and if Nature chooses h , then claim $\{h\}$.

Let us verify that this assessment constitutes a weak sequential equilibrium. The only non-singleton information set that is reached by the strategy profile is the top information set (where the buyer hears the claim $\{l, m\}$) and Bayesian updating requires the buyer to assign equal probability to each node. Every other non-singleton information set is not reached and thus the notion of weak sequential equilibrium allows any beliefs: in particular beliefs that assign probability $\frac{1}{2}$ to each node.

Now we check sequential rationality.

Let us start with the buyer.

- At the singleton node following claim $\{l\}$, buying one unit is optimal and at the singleton nodes following claims $\{m\}$ and $\{h\}$ buying two units is optimal.
- At the top information set (after the claim $\{l, m\}$) choosing one unit gives an expected payoff of $\frac{1}{2}(1) + \frac{1}{2}(2) = 1.5$ and choosing two units yields $\frac{1}{2}(0) + \frac{1}{2}(3) = 1.5$, thus buying two units is sequentially rational.
- At the information set following claim $\{l, h\}$ choosing one unit gives an expected payoff of $\frac{1}{2}(1) + \frac{1}{2}(3) = 2$ while choosing two units yields $\frac{1}{2}(0) + \frac{1}{2}(6) = 3$, thus buying two units is sequentially rational.
- At the information set following claim $\{l, m, h\}$ choosing one unit gives an expected payoff of $\frac{1}{3}(1) + \frac{1}{3}(2) + \frac{1}{3}(3) = 2$ while choosing two units yields $\frac{1}{3}(0) + \frac{1}{3}(3) + \frac{1}{3}(6) = 3$, thus buying two units is sequentially rational.
- Finally, at the information set following claim $\{m, h\}$ choosing one unit gives an expected payoff of $\frac{1}{2}(2) + \frac{1}{2}(3) = 2.5$ while choosing two units yields $\frac{1}{2}(3) + \frac{1}{2}(6) = 4.5$, thus buying two units is sequentially rational.

Now let us check sequential rationality for the seller's strategy.

- o At the left node (after Nature chooses l) claiming $\{l\}$ yields a payoff of 9, while every other claim yields a payoff of 18. Thus $\{l, m\}$ is sequentially rational.
- o At the other two nodes, the seller is indifferent among all the claims, because they yield the same payoff (12 at the node after Nature chooses m and 10 at the node after Nature chooses h) thus the postulated choices are sequentially rational. \square

Solutions to Exercise 15.5

(a) The extensive form is shown in Figure 15.30.

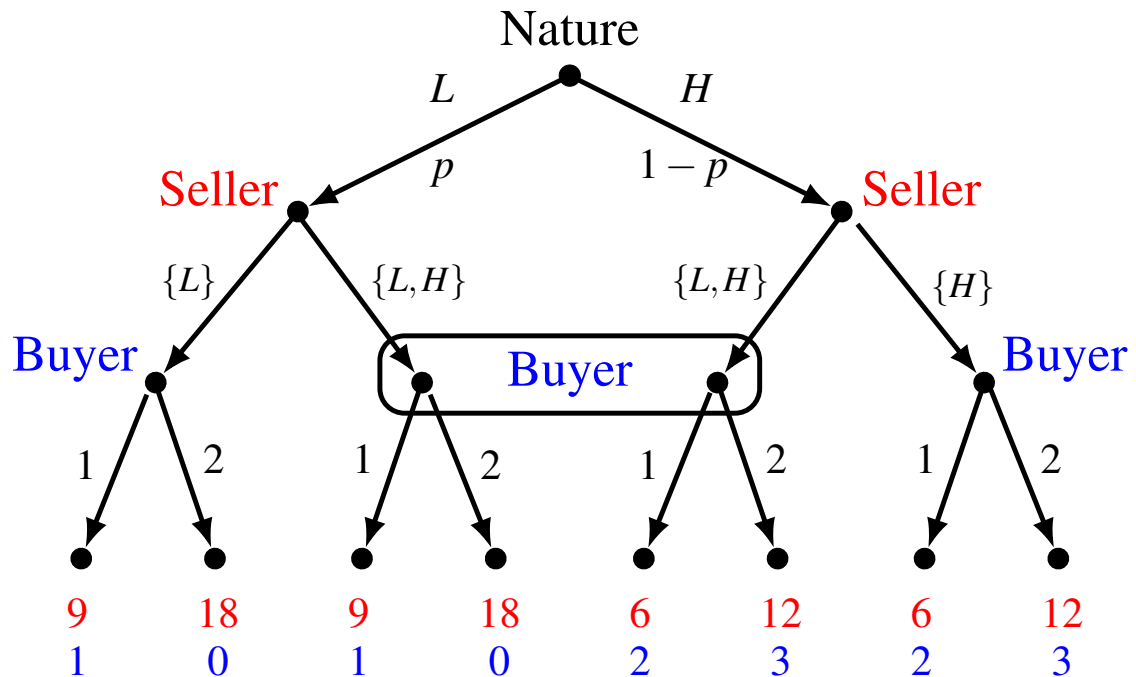


Figure 15.30: The game for Exercise 15.5.

(b) At any subgame-perfect equilibrium, the buyer buys one unit at the singleton information set on the left and two units at the singleton information set on the right. Thus we can simplify the game by replacing the buyer's node on the left with the payoff vector $(9,1)$ and the buyer's node on the right with the payoff vector $(12,3)$.

The strategic form corresponding to the simplified game is shown in Figure 15.31. The buyer’s strategy refers to the buyer’s only information set that has remained, where he is told $\{L, H\}$.

		BUYER			
		Buy 1 unit		Buy 2 units	
SELLER	L, LH	$9p + 6(1 - p)$	$p + 2(1 - p)$	$9p + 12(1 - p)$	$p + 3(1 - p)$
	L, H	$9p + 12(1 - p)$	$p + 3(1 - p)$	$9p + 12(1 - p)$	$p + 3(1 - p)$
	LH, LH	$9p + 6(1 - p)$	$p + 2(1 - p)$	$18p + 12(1 - p)$	$3(1 - p)$
	LH, H	$9p + 12(1 - p)$	$p + 3(1 - p)$	$18p + 12(1 - p)$	$3(1 - p)$

simplifying:

		BUYER			
		Buy 1 unit		Buy 2 units	
SELLER	L, LH	$6 + 3p$	$2 - p$	$12 - 3p$	$3 - 2p$
	L, H	$12 - 3p$	$3 - 2p$	$12 - 3p$	$3 - 2p$
	LH, LH	$6 + 3p$	$2 - p$	$12 + 6p$	$3 - 3p$
	LH, H	$12 - 3p$	$3 - 2p$	$12 + 6p$	$3 - 3p$

Figure 15.31: The strategic form of the simplified game.

Since $p < 1$, $12 - 3p > 6 + 3p$. Thus $((LH, H), 1)$ and $((L, H), 1)$ are always Nash equilibria for every value of p . Note that $2 - p > 3 - 3p$ if and only if $p > \frac{1}{2}$; thus if $p > \frac{1}{2}$ then there are no other pure-strategy Nash equilibria.

On the other hand, if $p \leq \frac{1}{2}$ then $((LH, LH), 2)$ is also a Nash equilibrium.

To sum up, the pure-strategy Nash equilibria of this normal form, which correspond to the pure-strategy subgame-perfect equilibria of the reduced extensive form, are as follows:

- If $p > \frac{1}{2}$, $((LH, H), 1)$ and $((L, H), 1)$;
- If $p \leq \frac{1}{2}$, $((LH, H), 1)$, $((L, H), 1)$ and $((LH, LH), 2)$.

□

Solutions to Exercise 15.6

(a) Let G_1 and G_2 be the perfect-information games (whose backward induction solution has been highlighted by means of thick arrows) shown in Figure 15.32.

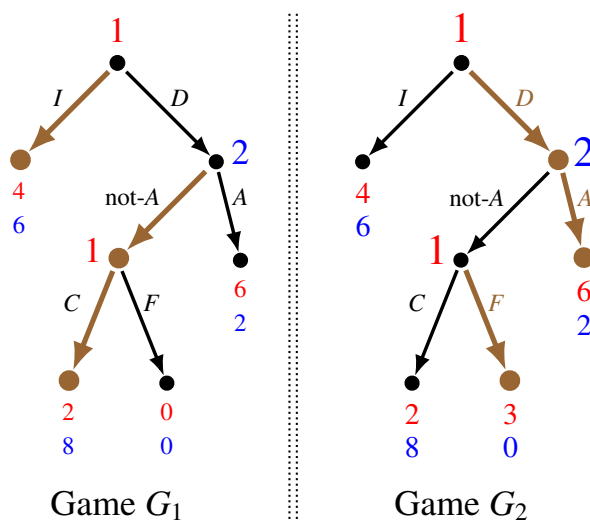


Figure 15.32: The possible perfect information games for Exercise 15.6.

The interactive knowledge-belief structure is shown in Figure 15.33




Figure 15.33: The two-sided incomplete-information situation.

(b) The common prior is given by $\begin{pmatrix} \alpha & \beta & \gamma \\ \frac{1}{5} & \frac{3}{5} & \frac{1}{5} \end{pmatrix}$.

The Harsanyi transformation yields the game shown in Figure 15.34.

- (c) - At the bottom information set of Player 1, C strictly dominates F and thus we can replace the two nodes in that information set with the payoff vector $(2,8)$.
- At the bottom singleton node of Player 1, F strictly dominates C and thus we can replace that node with the payoff vector $(3,0)$.

Thus – by appealing to sequential rationality – the game can be reduced to the game shown in Figure 15.35.

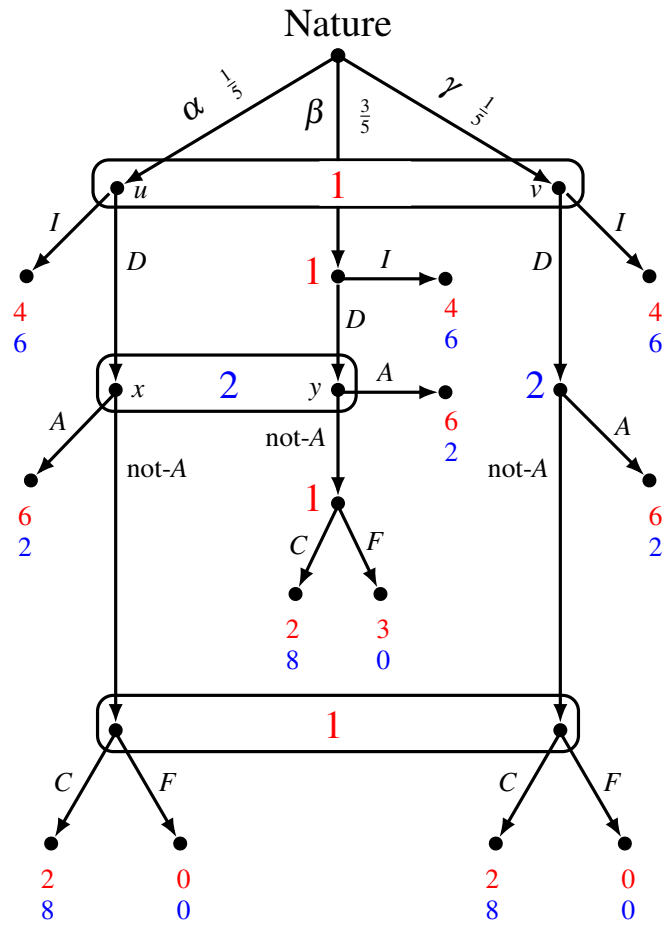


Figure 15.34: The game obtained by applying the Harsanyi transformation to Figure 15.33.




Figure 15.35: The reduced game.

Now at the bottom-right node of Player 2, n - A strictly dominates A and thus we can replace that node with the payoff vector $(2,8)$ and further simplify the game as shown in Figure 15.36.




Figure 15.36: The further reduced game.

The following are the pure-strategy weak sequential equilibria of the reduced game (which can be extended into weak sequential equilibria of the original game by adding the choices that were selected during the simplification process):

- $((I, I), n\text{-}A)$ with beliefs $\mu = \left(\begin{array}{cc|cc} u & v & x & y \\ \frac{1}{2} & \frac{1}{2} & p & 1-p \end{array} \right)$ for any $p \geq \frac{1}{4}$.
- $((I, D), A)$ with beliefs $\mu = \left(\begin{array}{cc|cc} u & v & x & y \\ \frac{1}{2} & \frac{1}{2} & 0 & 1 \end{array} \right)$.
- $((D, D), A)$ with beliefs $\mu = \left(\begin{array}{cc|cc} u & v & x & y \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{4} & \frac{3}{4} \end{array} \right)$. □

Solutions to Exercise 15.7

(a) Let G_1 and G_2 be the games shown in Figure 15.37 (in G_1 the defendant is negligent and in G_2 he is not).

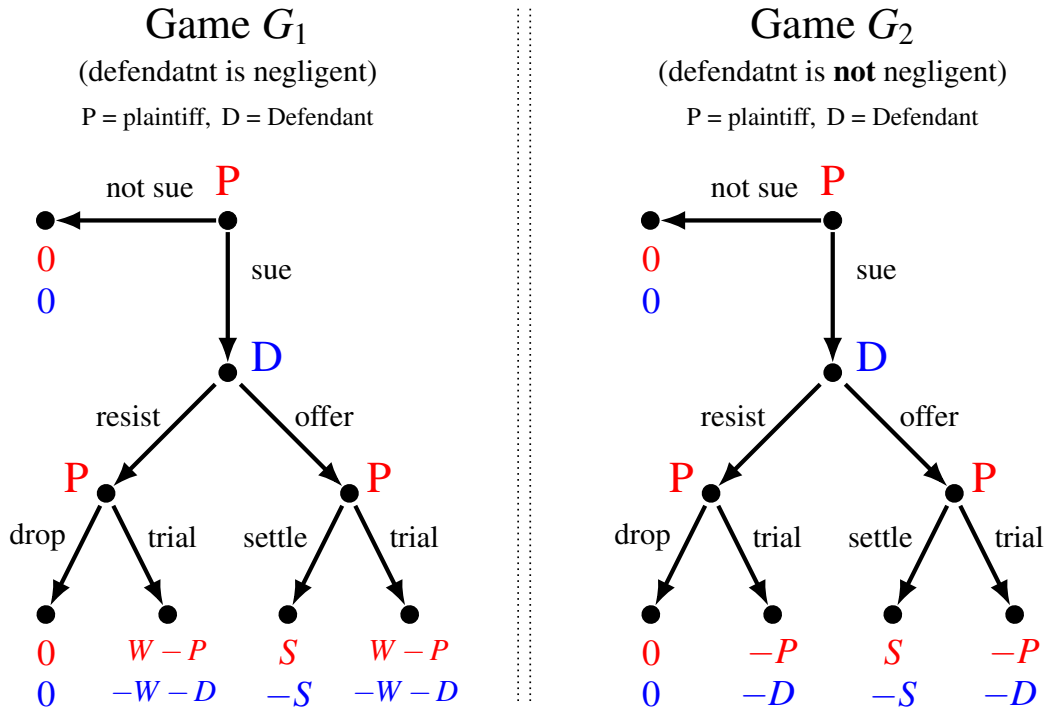


Figure 15.37: The two possible games.

Then the situation can be represented as shown in Figure 15.38 (at state α is the defendant is negligent and at state β the defendant is not negligent).




Figure 15.38: The one-sided incomplete-information situation.

(b) The extensive-form game is shown in Figure 15.39.

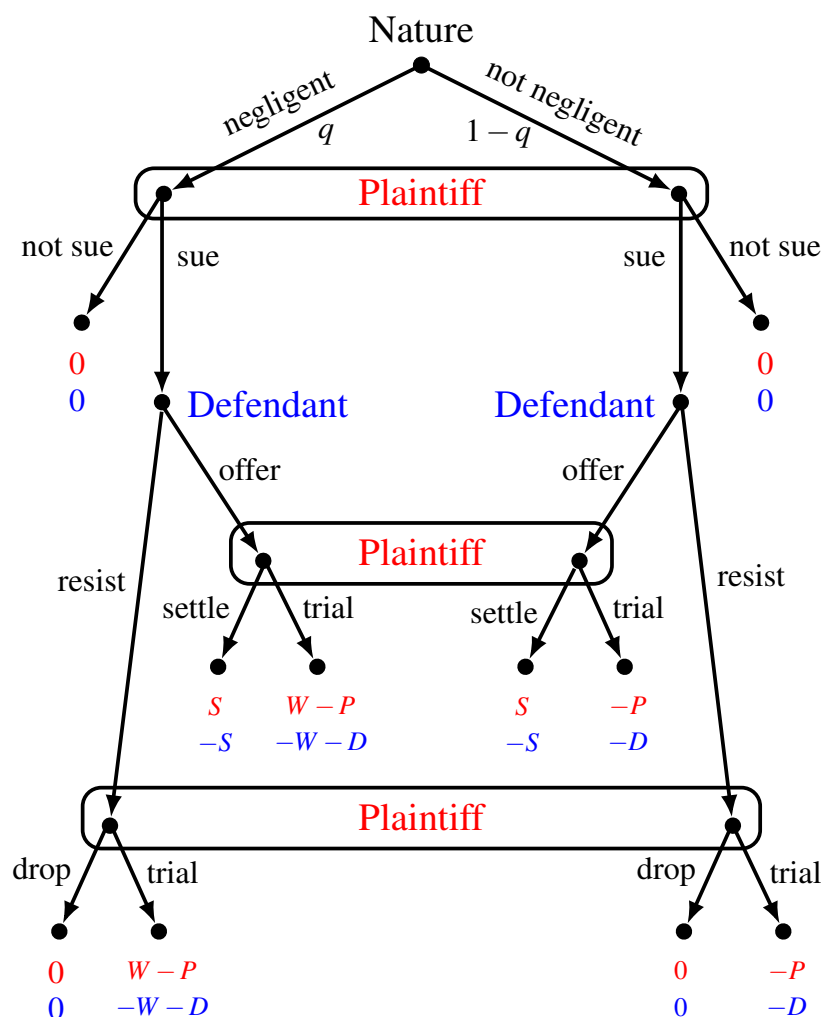


Figure 15.39:

(c) Plaintiff's strategies:

1. (sue; if offer settle; if resist drop),
2. (sue; if offer settle; if resist go to trial),
3. (sue; if offer go to trial; if resist drop),
4. (sue; if offer go to trial; if resist go to trial),
5. (not sue; if offer settle; if resist drop),
6. (not sue; if offer settle; if resist go to trial),
7. (not sue; if offer go to trial; if resist drop),
8. (not sue; if offer go to trial; if resist go to trial).

(d) There are two possibilities for a separating equilibrium:

Case S1: the defendant's strategy is "resist if negligent and offer if not negligent",

Case S2: the defendant's strategy is "offer if negligent and resist if not negligent".

In both cases we assume that the plaintiff's strategy involves suing with probability 1.

Consider Case S1 first. By Bayesian updating, at the bottom information set the plaintiff must attach probability 1 to the negligent type and thus, by sequential rationality, must choose “trial” (because $W - P > 0$).

Similarly, by Bayesian updating, at the middle information set the plaintiff must attach probability 1 to the non-negligent type and thus by sequential rationality must choose “settle”. But then the negligent type of the defendant gets $-(W + D)$ by resisting and would get $-S$ by offering.

Since, by assumption, $S < W (< W + D)$, the choice of resisting is not sequentially rational.

Now consider Case S2. By Bayesian updating, at the bottom information set the plaintiff must attach probability 1 to the non-negligent type and thus by sequential rationality must choose “drop”.

But then the negligent type of the defendant gets a negative payoff by offering, while he would get 0 by resisting. Hence the choice of offering is not sequentially rational.

(e) There are two candidates for a pure-strategy pooling equilibrium:

Case P1: both types of the defendant choose “offer” and

Case P2: both types of the defendant choose “resist”.

Consider Case P1 first. (both types of the defendant choose “offer”). In order for “offer” to be sequentially rational for the non-negligent type, it cannot be that the plaintiff’s strategy involves “settle” at the middle information set (the non-negligent type would get either 0 or $-D$ by resisting and both payoffs are greater than $-S$) and/or “drop” at the bottom information set. That is, it must be that the plaintiff chooses “trial” at both information sets.

By Bayesian updating, at the *middle* information set the plaintiff must attach probability q to the negligent type and probability $(1 - q)$ to the non-negligent type.

Hence at the middle information set “trial” is sequentially rational if and only if $qW - P \geq S$, that is, $q \geq \frac{S+P}{W}$.

In order for “trial” to be sequentially rational at the *bottom* information set, the plaintiff must attach sufficiently high probability (namely $p \geq \frac{P}{W}$) to the negligent type. This is allowed by weak sequential equilibrium because the bottom information set is not reached.

Finally, in order for “sue” to be sequentially rational it must be that $qW - P \geq 0$, that is, $q \geq \frac{P}{W}$, which is implied by $q \geq \frac{S+P}{W}$.

Thus there is a pooling equilibrium with $((\text{sue}, \text{trial}, \text{trial}), (\text{offer}, \text{offer}))$

if and only if $q \geq \frac{S+P}{W}$.

Now consider Case P2. (both types of the defendant choose “resist”). [Note: since in part (f) below the restriction $S < W - P$ does not hold, we will carry out the analysis at first without imposing the restriction.]

If the plaintiff’s strategy involves “drop” at the bottom information set, then it is indeed sequentially rational for both types of the defendant to choose “resist”.

Furthermore, “drop” is sequentially rational in this case if, and only if, $qW - P \leq 0$ that is, $q \leq \frac{P}{W}$.

Then “sue” is also sequentially rational, since the Plaintiff’s payoff is 0 no matter whether he sues or does not sue.

Thus there is a pooling equilibrium with ((sue, x , drop),(resist,resist))
if and only if $q \leq \frac{P}{W}$

and appropriate beliefs as follows (p is the probability on the left node of the unreached middle information set):

- $x = \text{settle}$ and either any p if $W \leq S + P$ or $p \leq \frac{S+P}{W}$ if $W > S + P$, or
- $x = \text{trial}$ and $p \geq \frac{S+P}{W}$, which requires $W \geq S + P$ (since $p \leq 1$).

Since it is assumed that $W > S + P$, we can conclude that

((sue,settle,drop),(resist,resist)) is an equilibrium if and only if $q \leq \frac{P}{W}$ with $p \leq \frac{S+P}{W}$

((sue,trial,drop),(resist,resist)) is an equilibrium if and only if $q \leq \frac{P}{W}$ with $p \geq \frac{S+P}{W}$

If, on the other hand, $q \geq \frac{P}{W}$, then “trial” is sequentially rational at the bottom information set. Then, in order for the non-negligent type of the defendant to choose “resist” it must be that the plaintiff’s strategy involves “trial” also at the middle information set, for which we need him to assign probability $p \geq \frac{S+P}{W}$ to the negligent type (which is possible, since the middle information set is not reached); of course, this requires $W \geq S + P$. Thus,

((sue,trial,trial),(resist,resist)) is an equilibrium if and only if $q \geq \frac{P}{W}$ with $p \geq \frac{S+P}{W}$.

(f) Note that here the restriction $W - P > S$ does not hold.

In this case, $q < \frac{S+P}{W} = \frac{80+70}{100} = \frac{3}{2}$ and thus, by the previous analysis, there is no pooling equilibrium of type **P1**, that is, ((sue,trial,trial),(offer,offer)) is not an equilibrium.

As for pooling equilibria of type **P2**, ((sue,trial,trial),(resist,resist)) is not an equilibrium because $S + P > W$; ((sue,trial,drop),(resist,resist)) is not an equilibrium either. However, there is a pooling equilibrium of type **P2** with ((sue,settle,drop),(resist,resist)) with any beliefs at the middle information set, since “settle” strictly dominates “trial” there (and, of course, belief $q = \frac{1}{12}$ on the left node of the bottom information set). \square

16. The Type-Space Approach

16.1 Types of players

As noted in Chapter 14, the theory of “games of incomplete information”¹ was pioneered by John Harsanyi (1967-68). Harsanyi’s approach was developed using a different approach from the one we employed in Chapters 14 and 15, which is based on the interactive knowledge-belief structures introduced in Chapters 8 and 9. The interactive knowledge structures of Chapter 9 are in fact a special case of the more general notion of *interactive Kripke structure*, named after the philosopher and logician Saul Kripke, whose work on this goes back to 1959 and was written while he was still an undergraduate.² Although well known among logicians and computer scientists, these structures were not known to game theorists. Perhaps, if Harsanyi had been aware of Kripke’s work he might have developed his theory using those structures. We find the interactive Kripke structures more natural and elegant and thus chose to explain the “theory of games of incomplete information” using those structures. In this chapter we will explain the “type-space” approach developed by Harsanyi and show that the two approaches are equivalent. We will limit ourselves to situations of incomplete information involving *static* games. We will begin in the next section with a simple special case, which is often all that one finds in game-theory textbooks, and then explain the general case in Section 16.3.

¹We use quotation marks because, strictly speaking, there is no such thing as a game of incomplete information. There are situations of incomplete information involving the playing of a game and what Harsanyi did was to suggest a way of transforming such situations into extensive-form games with imperfect information. Once the so-called “Harsanyi transformation” has been applied, the resulting game is a game (of complete information). Thus the “theory of games of incomplete information” is a theory on how to represent a situation of incomplete information concerning the playing of a game and how to transform it into a dynamic game with imperfect (but complete) information.

²Kripke (1959, 1963). For further details the reader is referred to van Ditmarsch et al. (2015).

16.2 Types that know their payoffs

We take as a starting point the case of incomplete information concerning a strategic-form game-frame.³ In the special case considered in this section, any uncertainty Player i has (if any) concerns either the von Neumann-Morgenstern utility function $U_j : O \rightarrow \mathbb{R}$ ($j \neq i$) of another player (or several other players) or the beliefs of the other player(s) (or both). Within the approach of Chapter 14, we would represent such a situation with an interactive knowledge-belief structure by associating with every state ω a game based on the given game-frame.⁴ Let $U_{i,\omega} : O \rightarrow \mathbb{R}$ be the utility function of Player i at state ω .

We say that every player knows her own payoffs if the following condition is satisfied, for every Player i (recall that $I_i(\omega)$ is the information set of Player i 's partition that contains state ω):

$$\text{if } \omega_1, \omega_2 \in \Omega \text{ and } \omega_2 \in I_i(\omega_1) \text{ then } U_{i,\omega_2} = U_{i,\omega_1}.$$

An example of such a situation of incomplete information is given in Figure 16.1, which reproduces Figure 14.3 of Chapter 14.⁵

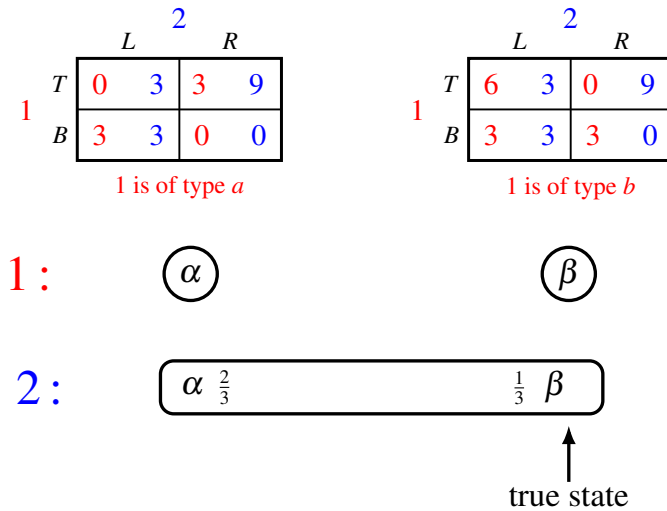


Figure 16.1: A situation of incomplete information represented by means of an interactive knowledge-belief structure.

³Recall (Definition 2.1.1, Chapter 2) that the elements of a strategic-form game-frame $\langle I, (S_1, \dots, S_n), O, f \rangle$ are as follows: $I = \{1, \dots, n\}$ is a set of players, S_i is the set of strategies of Player $i \in I$ (and $S = S_1 \times \dots \times S_n$ is the set of strategy profiles), O is a set of outcomes and $f : S \rightarrow O$ is a function that associates with every strategy profile an outcome.

⁴Recall that, given a game-frame, a game based on it is obtained by specifying, for every Player i , a von Neumann-Morgenstern utility function $U_i : O \rightarrow \mathbb{R}$ on the set of outcomes.

⁵In this case the game-frame is given by: $I = \{1, 2\}, S_1 = \{T, B\}, S_2 = \{L, R\}$ (so that $S = \{(T, L), (T, R), (B, L), (B, R)\}$), $O = \{o_1, o_2, o_3, o_4\}$, $f(T, L) = o_1, f(T, R) = o_2, f(B, L) = o_3, f(B, R) = o_4$. The (state-dependent) utility functions are given by $U_{1,\alpha} = \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ 0 & 3 & 3 & 0 \end{pmatrix}$,

$$U_{1,\beta} = \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ 6 & 0 & 3 & 3 \end{pmatrix} \text{ and } U_{2,\alpha} = U_{2,\beta} = \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ 3 & 9 & 3 & 0 \end{pmatrix}.$$

In the “type-space” approach the situation illustrated in Figure 16.1 would be represented using “types of players”: each type of Player i represents a utility function of Player i as well as Player i ’s beliefs about the types of the other players. The formal definition is as follows.

Definition 16.2.1 A static Bayesian game of incomplete information with knowledge of one’s own payoffs consists of the following elements:

- a set $I = \{1, \dots, n\}$ of players;
- for every Player $i \in I$, a set S_i of strategies (as usual, we denote by S the set of strategy profiles);
- for every Player $i \in I$, a set T_i of possible types; we denote by $T = T_1 \times \dots \times T_n$ the set of profiles of types and by $T_{-i} = T_1 \times \dots \times T_{i-1} \times T_{i+1} \times \dots \times T_n$ the set of profiles of types for the players other than i ;
- for every Player i and for every type $t_i \in T_i$ of Player i , a probability distribution $P_{i,t_i} : T_{-i} \rightarrow [0, 1]$ representing the beliefs of type t_i about the types of the other players.

The beliefs of all the types are said to be *Harsanyi consistent* if there exists a *common prior*, that is, a probability distribution $P : T \rightarrow [0, 1]$ such that, for every Player i and for every type $t_i \in T_i$ of Player i , P_{i,t_i} coincides with the probability distribution obtained from P by conditioning on the event $\{t_i\}$, that is (denoting a profile $t \in T$ by (t_{-i}, t_i)),

$$P_{i,t_i}(t_{-i}) = \frac{P(t_{-i}, t_i)}{\sum_{t'_{-i} \in T_{-i}} P(t'_{-i}, t_i)}.$$

Let us recast the situation illustrated in Figure 16.1 in the terminology of Definition 16.2.1.

First of all, we have that

$$I = \{1, 2\}, \quad S_1 = \{T, B\}, \quad S_2 = \{L, R\} \quad (\text{so that } S = \{(T, L), (T, R), (B, L), (B, R)\}).$$

Furthermore, there are two types of Player 1 and only one type of Player 2:

$$T_1 = \{t_1^a, t_1^b\}, \quad T_2 = \{t_2\}, \quad \text{so that } T = \{(t_1^a, t_2), (t_1^b, t_2)\}.$$

The utility functions are given by:

$$U_{1,t_1^a} = \begin{pmatrix} (T, L) & (T, R) & (B, L) & (B, R) \\ 0 & 3 & 3 & 0 \end{pmatrix} \quad U_{1,t_1^b} = \begin{pmatrix} (T, L) & (T, R) & (B, L) & (B, R) \\ 6 & 0 & 3 & 3 \end{pmatrix}$$

$$\text{and} \quad U_{2,t_2} = \begin{pmatrix} (T, L) & (T, R) & (B, L) & (B, R) \\ 3 & 9 & 3 & 0 \end{pmatrix}.$$

$$\text{The beliefs are given by} \quad P_{1,t_1^a} = P_{1,t_1^b} = \begin{pmatrix} t_2 \\ 1 \end{pmatrix}, \quad P_{2,t_2} = \begin{pmatrix} t_1^a & t_1^b \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}$$

$$\text{and the common prior is } P = \begin{pmatrix} (t_1^a, t_2) & (t_1^b, t_2) \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}.$$

Before we explain in detail how to transform a “state-space” structure into a “type-space” structure and *vice versa*, we give one more example, this time with double-sided incomplete information. Consider the situation illustrated in Figure 16.2, which reproduces Figure 14.8 of Chapter 14.




Figure 16.2: A situation with two-sided incomplete information.

In this case we have that each player has two types: we can identify a type of a player with a cell of the player’s information partition.

Thus we have that

$$T_1 = \{t_1^a, t_1^b\}, \quad T_2 = \{t_2^a, t_2^b\}, \quad U_{1,t_1^a} = U_{1,t_1^b} = \begin{pmatrix} (A,C) & (A,D) & (B,C) & (B,D) \\ 1 & 0 & 0 & 3 \end{pmatrix},$$

$$P_{1,t_1^a} = \begin{pmatrix} t_2^a \\ 1 \end{pmatrix}, \quad P_{1,t_1^b} = \begin{pmatrix} t_2^a & t_2^b \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}, \quad U_{2,t_2^a} = \begin{pmatrix} (A,C) & (A,D) & (B,C) & (B,D) \\ 3 & 1 & 0 & 1 \end{pmatrix},$$

$$U_{2,t_2^b} = \begin{pmatrix} (A,C) & (A,D) & (B,C) & (B,D) \\ 0 & 2 & 0 & 1 \end{pmatrix}, \quad P_{2,t_2^a} = \begin{pmatrix} t_1^a & t_1^b \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}, \quad P_{2,t_2^b} = \begin{pmatrix} t_1^b \\ 1 \end{pmatrix}.$$

The common prior is given by $P = \begin{pmatrix} (t_1^a, t_2^a) & (t_1^a, t_2^b) & (t_1^b, t_2^a) & (t_1^b, t_2^b) \\ \frac{2}{4} & 0 & \frac{1}{4} & \frac{1}{4} \end{pmatrix}.$

Thus the two types of Player 1 have the same utility function but different beliefs about the types of Player 2, while the two types of Player 2 differ both in terms of utility function and in terms of beliefs about the types of Player 1.

From these examples it should be clear how to transform a “state-space” model into a “type-space” model.

- First of all, for every Player i , create one type for every cell of Player i 's partition, making sure that different cells are associated with different types. In this way we have identified each state with a profile of types.
- Since there is a probability distribution over each information set of Player i , that probability distribution will give us a probability distribution for the associated type of Player i over some set of type-profiles for the other players.
- Finally, since our assumption (to be relaxed in the next section) is that each player knows her own utility function (that is, the utility function of a player does not vary from state to state within the same information set of that player), with each type t_i of Player i is associated a unique utility function U_{i,t_i} .

Conversely, we can convert a “type-space” structure into a “state-space” structure as follows.

- Let the set of states be the set T of profiles of types. For every Player i and for every two states $t, t' \in T$, let t and t' belong to the same information set of Player i (that is, to the same cell of Player i 's partition) if and only if Player i 's type is the same in t and t' : $t' \in I_i(t)$ if and only if $t_i = t'_i$.
- The beliefs of each type of Player i then yield a probability distribution over the information set of Player i corresponding to that type. An example of this transformation is given in Exercise 16.2.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 16.4.1 at the end of this chapter.

16.3 The general case

As pointed out in Chapter 14, it may very well be the case that a rational player does not know her own payoffs (for example, because she is uncertain about what outcomes might occur if she chooses a particular action: see Exercise 14.2, Chapter 14). Uncertainty about a player's own payoffs is compatible with the player knowing her own preferences (that is, how she ranks the outcomes that she considers possible). Definition 16.2.1 is not general enough to encompass such possibilities.

The following, more general, definition allows the utility function of a player to depend not only on the player's own type but also on the types of the other players. Definition 16.3.1 is identical to Definition 16.2.1, except for the starred items and the boldface part in the last item.

Definition 16.3.1 A static Bayesian game of incomplete information consists of the following elements:

- a set $I = \{1, \dots, n\}$ of players;
- for every Player $i \in I$, a set S_i of strategies (as usual, we denote by S the set of strategy profiles);
- for every Player $i \in I$, a set T_i of possible types; we denote by $T = T_1 \times \dots \times T_n$ the set of profiles of types and by $T_{-i} = T_1 \times \dots \times T_{i-1} \times T_{i+1} \times \dots \times T_n$ the set of profiles of types for the players other than i ;
- ★ a set $Y \subseteq T$ of relevant profiles of types;
- ★ for every Player i and for every profile of types $t \in Y$, a utility (or payoff) function $U_{i,t} : S \rightarrow \mathbb{R}$;
- for every Player i and for every type $t_i \in T_i$ of Player i , a probability distribution $P_{i,t_i} : T_{-i} \rightarrow [0, 1]$ representing the beliefs of type t_i about the types of the other players **satisfying the restriction that if $P_{i,t_i}(t_{-i}) > 0$ then $(t_i, t_{-i}) \in Y$.**

As in the special case considered in the previous section, also in the general case one can transform a “state-space” structure into a “type-space” structure and vice versa.

Given a “state-space” structure, for every Player i we identify the cells of Player i 's partition with the types of Player i (one type for every information set). Since there is a probability distribution over each information set of Player i , that probability distribution will yield the probability distribution for the associated type of Player i over some set of type-profiles for the other players. Finally, having identified each state with a profile of types (since each state belongs to one and only one information set of each player), we can assign to the corresponding type of Player i the utility function of Player i at that state.

We shall illustrate this conversion using the “state-space” structure shown in Figure 16.3 (taken from Exercise 14.2, Chapter 14).

The corresponding “type-space” structure (static Bayesian game of incomplete information) is as follows:

- $I = \{1, 2\}$ (letting Bill be Player 1 and Ann Player 2),
- $S_1 = \{g, ng\}$ $S_2 = \{a, r\}$ (so that $S = \{(g, a), (g, r), (ng, a), (ng, r)\}$).
- There are two types of Player 1 and only one type of Player 2:
 $T_1 = \{t_1^f, t_1^e\}$ (f stands for ‘friend’ and e for ‘enemy’),
 $T_2 = \{t_2\}$ so that $T = \{(t_1^f, t_2), (t_1^e, t_2)\}$.
- Thus in this case we have that $Y = T$.

- o The utility functions are given by:

$$U_{1,(t_1^f,t_2)} = U_{1,(t_1^e,t_2)} = \begin{pmatrix} (g,a) & (g,r) & (ng,a) & (ng,r) \\ 1 & -1 & 0 & 0 \end{pmatrix}$$

$$U_{2,(t_1^f,t_2)} = \begin{pmatrix} (g,a) & (g,r) & (ng,a) & (ng,r) \\ 1 & 0 & 0 & 0 \end{pmatrix}$$

$$U_{2,(t_1^e,t_2)} = \begin{pmatrix} (g,a) & (g,r) & (ng,a) & (ng,r) \\ -1 & 0 & 0 & 0 \end{pmatrix}$$

- o The beliefs are given by $P_{1,t_1^f} = P_{1,t_1^e} = \begin{pmatrix} t_2 \\ 1 \end{pmatrix}$, $P_{2,t_2} = \begin{pmatrix} t_1^f & t_1^e \\ p & 1-p \end{pmatrix}$

and the common prior is $P = \begin{pmatrix} (t_1^f,t_2) & (t_1^e,t_2) \\ p & 1-p \end{pmatrix}$.

Conversely, we can convert a “type-space” structure into a “state-space” structure as follows:

- Let the set of states be the set Y of relevant profiles of types.
- For every Player i and for every two states $t, t' \in Y$, let t and t' belong to the same information set of Player i (that is, to the same cell of Player i 's partition) if and only if Player i 's type is the same in t and t' : $t' \in I_i(t)$ if and only if $t_i = t'_i$.
- The beliefs of each type of Player i then yield a probability distribution over the information set of Player i corresponding to that type.

An example of this transformation is given in Exercise 16.3.

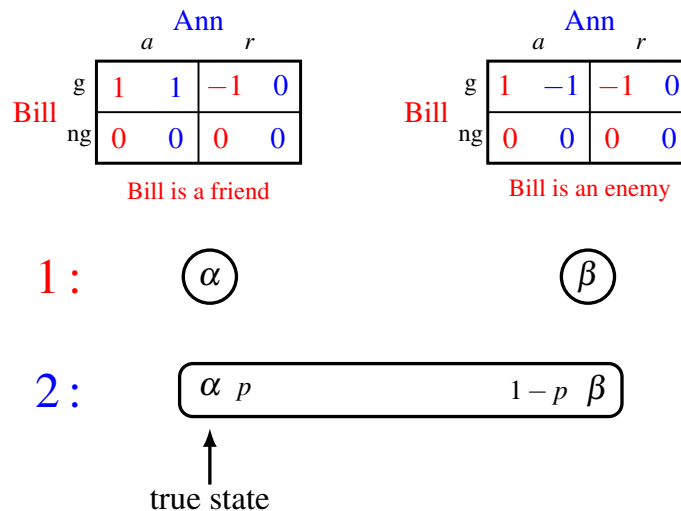


Figure 16.3: The situation of incomplete information described in Exercise 14.2 (Chapter 14) where it is not true that each player knows his/her own payoffs.

Test your understanding of the concepts introduced in this section, by going through the exercises in Section 16.4.2 at the end of this chapter.

16.4 Exercises

The answers to the following exercises are in Section 16.5 at the end of this chapter.

16.4.1 Exercises for Section 16.2: Types that know their own payoffs

Exercise 16.1 Transform the situation of incomplete information shown in Figure 16.4, where G_1 and G_2 are the games shown in Figure 16.5, into a “type-space” structure. ■




Figure 16.4: A situation of incomplete information. The games G_1 and G_2 are shown in Figure 16.5.

Exercise 16.2 Consider the following Bayesian game of incomplete information:

$$I = \{1, 2\} \quad S_1 = \{A, B\} \quad S_2 = \{C, D\} \quad T_1 = \{t_1^a, t_1^b\} \quad T_2 = \{t_2^a, t_2^b, t_2^c\}$$

$$U_{1,t_1^a} = \begin{pmatrix} AC & AD & BC & BD \\ 4 & 1 & 0 & 2 \end{pmatrix} \quad U_{1,t_1^b} = \begin{pmatrix} AC & AD & BC & BD \\ 0 & 0 & 2 & 1 \end{pmatrix}$$

$$U_{2,t_2^a} = U_{2,t_2^b} = \begin{pmatrix} AC & AD & BC & BD \\ 2 & 1 & 2 & 3 \end{pmatrix} \quad U_{2,t_2^c} = \begin{pmatrix} AC & AD & BC & BD \\ 0 & 2 & 2 & 0 \end{pmatrix}$$

$$P_{1,t_1^a} = \begin{pmatrix} t_2^a & t_2^b & t_2^c \\ \frac{1}{4} & \frac{3}{4} & 0 \end{pmatrix}, \quad P_{1,t_1^b} = \begin{pmatrix} t_2^a & t_2^b & t_2^c \\ \frac{4}{13} & \frac{3}{13} & \frac{6}{13} \end{pmatrix}, \quad P_{2,t_2^a} = \begin{pmatrix} t_1^a & t_1^b \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix},$$

$$P_{2,t_2^b} = \begin{pmatrix} t_1^a & t_1^b \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}, \quad P_{2,t_2^c} = \begin{pmatrix} t_1^a & t_1^b \\ 0 & 1 \end{pmatrix}.$$

- (a) Are the beliefs of the types consistent (that is, is there a common prior)?
- (b) Transform this type-space structure into an interactive knowledge-belief structure. ■

		Player 2					
		<i>C</i>			<i>D</i>		
Player 1	<i>A</i>	4	1	0	2	4	2
	<i>B</i>	4	1	0	0	0	0

Player 3: *E*

		Player 2					
		<i>C</i>			<i>D</i>		
Player 1	<i>A</i>	2	0	2	0	1	0
	<i>B</i>	0	0	0	0	2	0

Player 3: *F*

GAME G_1

.....

		Player 2					
		<i>C</i>			<i>D</i>		
Player 1	<i>A</i>	4	4	0	2	1	2
	<i>B</i>	4	4	0	0	1	0

Player 3: *E*

		Player 2					
		<i>C</i>			<i>D</i>		
Player 1	<i>A</i>	2	1	2	0	0	0
	<i>B</i>	0	2	0	0	0	0

Player 3: *F*

GAME G_2

Figure 16.5: The games for the situation of incomplete information of Figure 16.4.

16.4.2 Exercises for Section 16.3: The general case

Exercise 16.3 Consider the following Bayesian game of incomplete information:

$$I = \{1, 2, 3\} \quad S_1 = \{A, B\} \quad S_2 = \{C, D\} \quad S_3 = \{E, F\}$$

$$T_1 = \{t_1^a, t_1^b\} \quad T_2 = \{t_2^a, t_2^b\} \quad T_3 = \{t_3^a, t_3^b\}$$

$$Y = \{(t_1^a, t_2^a, t_3^a), (t_1^b, t_2^b, t_3^a), (t_1^b, t_2^a, t_3^b), (t_1^b, t_2^b, t_3^b)\}$$

$$U_{1, (t_1^a, t_2^a, t_3^a)} = \begin{pmatrix} ACE & ADE & BCE & BDE & ACF & ADF & BCF & BDF \\ 2 & 2 & 3 & 0 & 0 & 0 & 1 & 1 \end{pmatrix}$$

$$U_{1, (t_1^b, t_2^b, t_3^a)} = U_{1, (t_1^b, t_2^b, t_3^b)} = \begin{pmatrix} ACE & ADE & BCE & BDE & ACF & ADF & BCF & BDF \\ 1 & 2 & 0 & 0 & 1 & 0 & 0 & 1 \end{pmatrix}$$

$$U_{1, (t_1^b, t_2^a, t_3^b)} = \begin{pmatrix} ACE & ADE & BCE & BDE & ACF & ADF & BCF & BDF \\ 0 & 0 & 2 & 1 & 2 & 0 & 0 & 1 \end{pmatrix}$$

$$U_{2, (t_1^a, t_2^a, t_3^a)} = \begin{pmatrix} ACE & ADE & BCE & BDE & ACF & ADF & BCF & BDF \\ 2 & 4 & 2 & 0 & 0 & 1 & 0 & 2 \end{pmatrix}$$

$$U_{2, (t_1^b, t_2^b, t_3^a)} = U_{2, (t_1^b, t_2^b, t_3^b)} = \begin{pmatrix} ACE & ADE & BCE & BDE & ACF & ADF & BCF & BDF \\ 4 & 1 & 4 & 1 & 1 & 0 & 2 & 0 \end{pmatrix}$$

$$U_{2, (t_1^b, t_2^a, t_3^b)} = \begin{pmatrix} ACE & ADE & BCE & BDE & ACF & ADF & BCF & BDF \\ 4 & 3 & 1 & 2 & 0 & 1 & 0 & 2 \end{pmatrix}$$

$$U_{3, (t_1^a, t_2^a, t_3^a)} = \begin{pmatrix} ACE & ADE & BCE & BDE & ACF & ADF & BCF & BDF \\ 0 & 2 & 0 & 0 & 2 & 0 & 0 & 0 \end{pmatrix}$$

$$U_{3, (t_1^b, t_2^b, t_3^a)} = U_{3, (t_1^b, t_2^b, t_3^b)} = \begin{pmatrix} ACE & ADE & BCE & BDE & ACF & ADF & BCF & BDF \\ 3 & 2 & 0 & 1 & 1 & 2 & 2 & 0 \end{pmatrix}$$

$$U_{3, (t_1^b, t_2^a, t_3^b)} = \begin{pmatrix} ACE & ADE & BCE & BDE & ACF & ADF & BCF & BDF \\ 1 & 2 & 0 & 1 & 2 & 0 & 0 & 1 \end{pmatrix}$$

$$P_{1, t_1^a} = \begin{pmatrix} (t_2^a, t_3^a) \\ 1 \end{pmatrix} \quad P_{1, t_1^b} = \begin{pmatrix} (t_2^b, t_3^a) & (t_2^a, t_3^b) & (t_2^b, t_3^b) \\ \frac{2}{5} & \frac{2}{5} & \frac{1}{5} \end{pmatrix}$$

$$P_{2, t_2^a} = \begin{pmatrix} (t_1^a, t_3^a) & (t_1^b, t_3^b) \\ \frac{5}{7} & \frac{2}{7} \end{pmatrix} \quad P_{2, t_2^b} = \begin{pmatrix} (t_1^b, t_3^a) & (t_1^b, t_3^b) \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}$$

$$P_{3, t_3^a} = \begin{pmatrix} (t_1^a, t_2^a) & (t_1^b, t_2^b) \\ \frac{5}{7} & \frac{2}{7} \end{pmatrix} \quad P_{3, t_3^b} = \begin{pmatrix} (t_1^b, t_2^a) & (t_1^b, t_2^b) \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}$$

- (a) Are the beliefs of the types consistent (that is, is there a common prior)?
 (b) Transform this type-space structure into a knowledge-belief structure.

Exercise 16.4 — *** Challenging Question ***

Consider the following two-player Bayesian game of incomplete information:

$$I = \{1, 2\} \quad S_1 = \{T, B\} \quad S_2 = \{L, R\} \quad T_1 = \{t_1^A, t_1^B\} \quad T_2 = \{t_2^a, t_2^b\},$$

$$U_{1,(t_1^A,t_2^a)} = U_{1,(t_1^A,t_2^b)} = \begin{pmatrix} (T,L) & (T,R) & (B,L) & (B,R) \\ 6 & 0 & 3 & 3 \end{pmatrix}$$

$$U_{1,(t_1^B,t_2^a)} = U_{1,(t_1^B,t_2^b)} = \begin{pmatrix} (T,L) & (T,R) & (B,L) & (B,R) \\ 0 & 3 & 3 & 0 \end{pmatrix}$$

$$U_{2,(t_1^A,t_2^a)} = U_{2,(t_1^A,t_2^b)} = U_{2,(t_1^B,t_2^a)} = U_{2,(t_1^B,t_2^b)} = \begin{pmatrix} (T,L) & (T,R) & (B,L) & (B,R) \\ 3 & 9 & 3 & 0 \end{pmatrix}$$

$$P_{1,t_1^A} = \begin{pmatrix} t_2^a & t_2^b \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} \quad P_{1,t_1^B} = \begin{pmatrix} t_2^a & t_2^b \\ \frac{3}{5} & \frac{2}{5} \end{pmatrix}$$

$$P_{2,t_2^a} = \begin{pmatrix} t_1^A & t_1^B \\ \frac{1}{4} & \frac{3}{4} \end{pmatrix} \quad P_{2,t_2^b} = \begin{pmatrix} t_1^A & t_1^B \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}.$$

- Transform the situation described above into an interactive knowledge-belief structure. Assume that the true state is where Player 1 is of type t_1^A and Player 2 is of type t_2^a .
- Apply the Harsanyi transformation to obtain an extensive-form game.
- Write the strategic-form game corresponding to the game of part (b).
- Find all the pure-strategy Bayesian Nash equilibria. Does any of these yield a Nash equilibrium in the true game being played?
- Of the pure-strategy Bayesian Nash equilibria select one where Player 1 is uncertain about Player 2's choice of action.

For this equilibrium complete the structure of part (a) by turning it into a model (that is, associate with each state an action – not a strategy – for each player: see Chapter 10) and verify that at the true state there is common knowledge of rationality.



16.5 Solutions to Exercises

Solutions to Exercise 16.1 The structure is as follows (the elements are given as listed in Definition 16.2.1).

$$I = \{1, 2, 3\} \quad S_1 = \{A, B\} \quad S_2 = \{C, D\} \quad S_3 = \{E, F\}$$

$$T_1 = \{t_1^a, t_1^b, t_1^c\} \quad T_2 = \{t_2^a, t_2^b\} \quad T_3 = \{t_3^a, t_3^b\}$$

$$U_{1,t_1^a} = U_{1,t_1^b} = U_{1,t_1^c} = \begin{pmatrix} ACE & ADE & BCE & BDE & ACF & ADF & BCF & BDF \\ 4 & 2 & 4 & 0 & 2 & 0 & 0 & 0 \end{pmatrix}$$

$$U_{2,t_2^a} = \begin{pmatrix} ACE & ADE & BCE & BDE & ACF & ADF & BCF & BDF \\ 1 & 4 & 1 & 0 & 0 & 1 & 0 & 2 \end{pmatrix}$$

$$U_{2,t_2^b} = \begin{pmatrix} ACE & ADE & BCE & BDE & ACF & ADF & BCF & BDF \\ 4 & 1 & 4 & 1 & 1 & 0 & 2 & 0 \end{pmatrix}$$

$$U_{3,t_3^a} = U_{3,t_3^b} = \begin{pmatrix} ACE & ADE & BCE & BDE & ACF & ADF & BCF & BDF \\ 0 & 2 & 0 & 0 & 2 & 0 & 0 & 0 \end{pmatrix}$$

$$P_{1,t_1^a} = \begin{pmatrix} (t_2^a, t_3^a) \\ 1 \end{pmatrix} \quad P_{1,t_1^b} = \begin{pmatrix} (t_2^b, t_3^a) & (t_2^a, t_3^b) \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} \quad P_{1,t_1^c} = \begin{pmatrix} (t_2^b, t_3^b) \\ 1 \end{pmatrix}$$

$$P_{2,t_2^a} = \begin{pmatrix} (t_1^a, t_3^a) & (t_1^b, t_3^b) \\ \frac{7}{9} & \frac{2}{9} \end{pmatrix} \quad P_{2,t_2^b} = \begin{pmatrix} (t_1^b, t_3^a) & (t_1^c, t_3^b) \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}$$

$$P_{3,t_3^a} = \begin{pmatrix} (t_1^a, t_2^a) & (t_1^b, t_2^b) \\ \frac{7}{9} & \frac{2}{9} \end{pmatrix} \quad P_{3,t_3^b} = \begin{pmatrix} (t_1^b, t_2^a) & (t_1^c, t_2^b) \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}$$

The common prior is given by:

$$P = \begin{pmatrix} t_1^a t_2^a t_3^a & t_1^a t_2^a t_3^b & t_1^a t_2^b t_3^a & t_1^a t_2^b t_3^b & t_1^b t_2^a t_3^a & t_1^b t_2^a t_3^b \\ \frac{7}{12} & 0 & 0 & 0 & 0 & \frac{2}{12} \\ & t_1^b t_2^b t_3^a & t_1^b t_2^b t_3^b & t_1^c t_2^a t_3^a & t_1^c t_2^a t_3^b & t_1^c t_2^b t_3^a & t_1^c t_2^b t_3^b \\ & \frac{2}{12} & 0 & 0 & 0 & 0 & \frac{1}{12} \end{pmatrix}$$

□

Solutions to Exercise 16.2

(a) Yes, the following is a common prior:

$$P = \begin{pmatrix} (t_1^a, t_2^a) & (t_1^a, t_2^b) & (t_1^a, t_2^c) & (t_1^b, t_2^a) & (t_1^b, t_2^b) & (t_1^b, t_2^c) \\ \frac{2}{21} & \frac{6}{21} & 0 & \frac{4}{21} & \frac{3}{21} & \frac{6}{21} \end{pmatrix}$$

(b) The knowledge-belief structure is shown in Figure 16.6. The set of states is

$$\Omega = \{\alpha, \beta, \gamma, \delta, \varepsilon\}, \text{ where}$$

$$\alpha = (t_1^a, t_2^a) \quad \beta = (t_1^a, t_2^b) \quad \gamma = (t_1^b, t_2^a) \quad \delta = (t_1^b, t_2^b), \quad \varepsilon = (t_1^b, t_2^c).$$

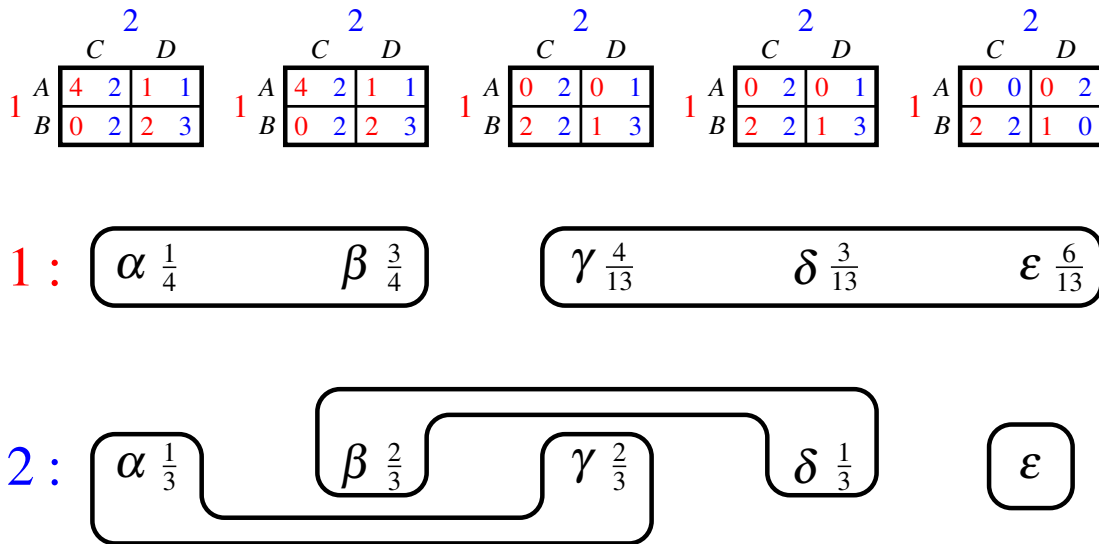


Figure 16.6: The interactive knowledge-belief structure representing the Bayesian game of incomplete information of Exercise 16.2.

□

Solutions to Exercise 16.3

(a) Yes: the following is a common prior:

$$P = \left(\begin{array}{cccc} t_1^a t_2^a t_3^a & t_1^b t_2^a t_3^b & t_1^b t_2^b t_3^a & t_1^b t_2^b t_3^b \\ \frac{1}{2} & \frac{1}{5} & \frac{1}{5} & \frac{1}{10} \end{array} \right) \text{ and } P(t) = 0 \text{ for every other } t \in T.$$

(b) The interactive knowledge-belief structure is shown in Figure 16.7, where

$$\alpha = (t_1^a, t_2^a, t_3^a) \quad \beta = (t_1^b, t_2^b, t_3^a) \quad \gamma = (t_1^b, t_2^a, t_3^b) \quad \delta = (t_1^b, t_2^b, t_3^b)$$

and the games G_1 , G_2 and G_3 are as shown in Figure 16.8. □

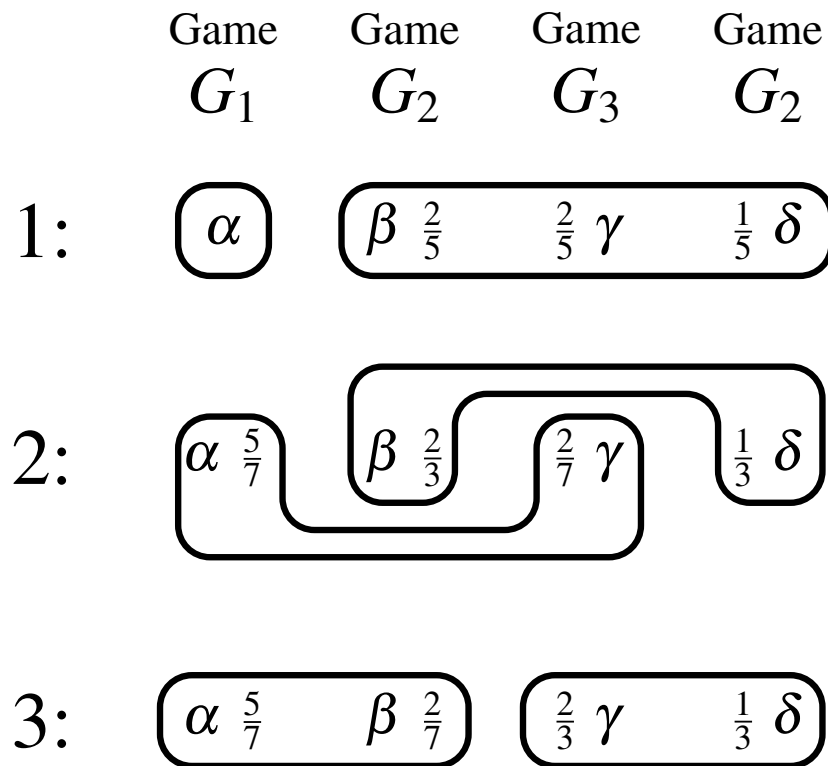


Figure 16.7: The knowledge-belief structure for Exercise 16.3; the games G_1 , G_2 and G_3 are as shown in Figure 16.8.

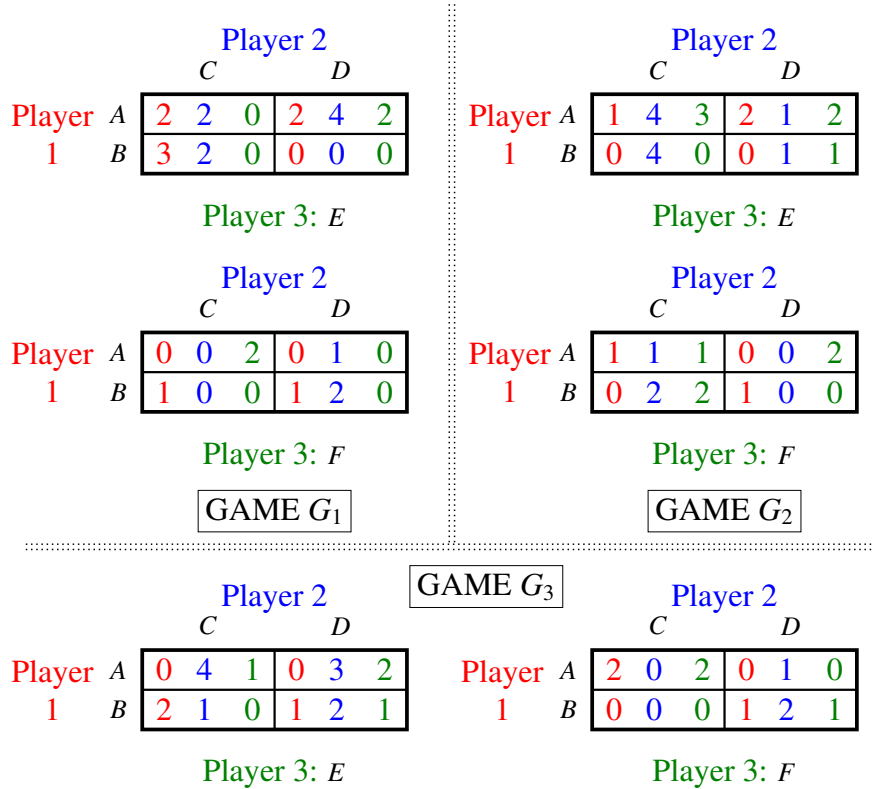


Figure 16.8: The games in the situation of incomplete information shown in Figure 16.7.

Solutions to Exercise 16.4

(a) The interactive knowledge-belief structure is represented in Figure 16.9 where

$$\alpha = (t_1^A, t_2^a) \quad \beta = (t_1^A, t_2^b) \quad \gamma = (t_1^B, t_2^a) \quad \delta = (t_1^B, t_2^b).$$




Figure 16.9: The interactive knowledge-belief structure for Exercise 16.4.

- (b) First of all, there is a common prior: $\begin{pmatrix} \alpha & \beta & \gamma & \delta \\ \frac{1}{7} & \frac{1}{7} & \frac{3}{7} & \frac{2}{7} \end{pmatrix}$.

The extensive-form game is shown in Figure 16.10.

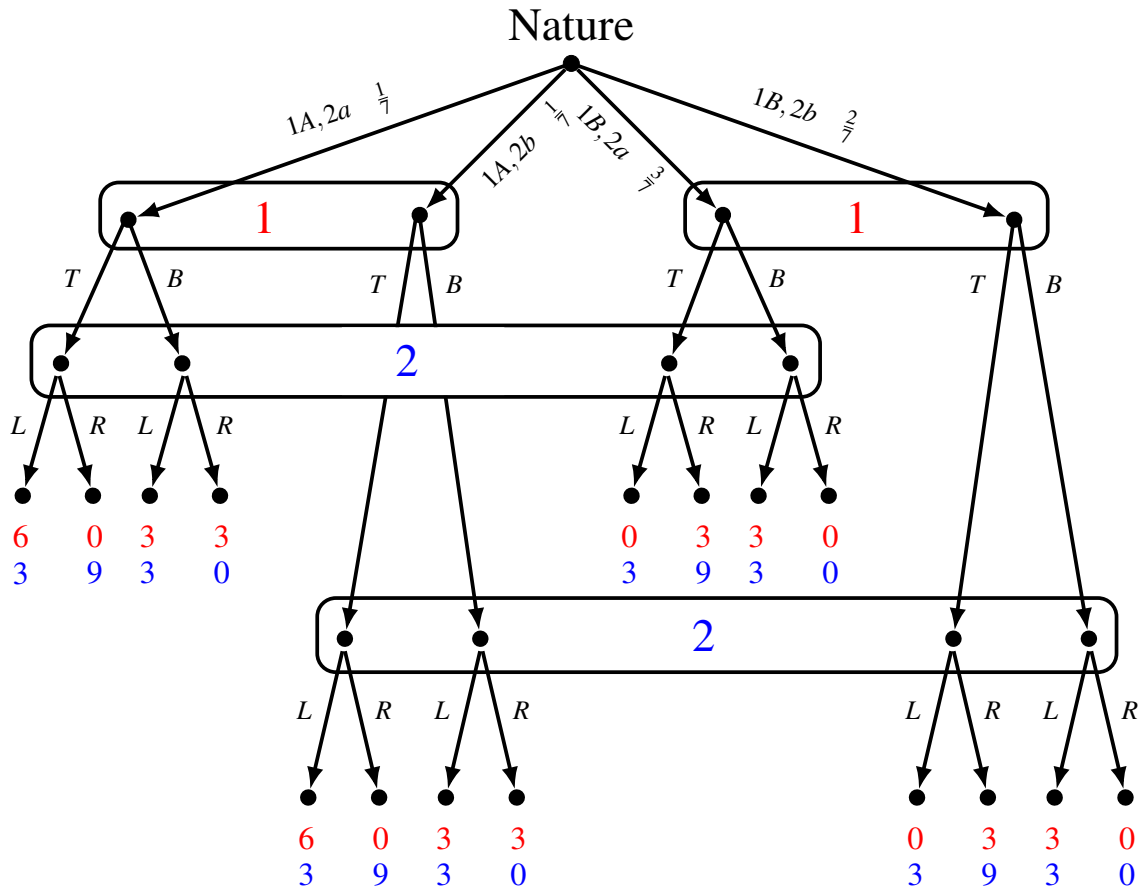


Figure 16.10: The game obtained by applying the Harsanyi transformation to the situation of incomplete information of Figure 16.9.

- (c) The strategy profiles give rise to the lotteries shown in Figure 16.11. Thus the strategic form is as shown in Figure 16.12.
- (d) The pure-strategy Nash equilibria are: (TB, LL) , (TB, LR) and (BT, RR) . None of them yields a Nash equilibrium of the game associated with state α , since neither (T, L) nor (B, R) are Nash equilibria of that game.

		Player 2			
		L if type 2a L if type 2b	L if type 2a R if type 2b	R if type 2a L if type 2b	R if type 2a R if type 2b
Player 1	T if type 1A T if type 1B	$\frac{1}{7}$ $\frac{1}{7}$ $\frac{3}{7}$ $\frac{2}{7}$ TL TL TL TL 6 6 0 0 3 3 3 3	$\frac{1}{7}$ $\frac{1}{7}$ $\frac{3}{7}$ $\frac{2}{7}$ TL TR TL TR 6 0 0 3 3 9 3 9	$\frac{1}{7}$ $\frac{1}{7}$ $\frac{3}{7}$ $\frac{2}{7}$ TR TL TR TL 0 6 3 0 9 3 9 3	$\frac{1}{7}$ $\frac{1}{7}$ $\frac{3}{7}$ $\frac{2}{7}$ TR TR TR TR 0 0 3 3 9 9 9 9
	T if type 1A B if type 1B	$\frac{1}{7}$ $\frac{1}{7}$ $\frac{3}{7}$ $\frac{2}{7}$ TL TL BL BL 6 6 3 3 3 3 3 3	$\frac{1}{7}$ $\frac{1}{7}$ $\frac{3}{7}$ $\frac{2}{7}$ TL TR BL BR 6 0 3 0 3 9 3 0	$\frac{1}{7}$ $\frac{1}{7}$ $\frac{3}{7}$ $\frac{2}{7}$ TR TL BR BL 0 6 0 3 9 3 0 3	$\frac{1}{7}$ $\frac{1}{7}$ $\frac{3}{7}$ $\frac{2}{7}$ TR TR BR BR 0 0 0 0 9 9 0 0
	B if type 1A T if type 1B	$\frac{1}{7}$ $\frac{1}{7}$ $\frac{3}{7}$ $\frac{2}{7}$ BL BL TL TL 3 3 0 0 3 3 3 3	$\frac{1}{7}$ $\frac{1}{7}$ $\frac{3}{7}$ $\frac{2}{7}$ BL BR TL TR 3 3 0 3 3 0 3 9	$\frac{1}{7}$ $\frac{1}{7}$ $\frac{3}{7}$ $\frac{2}{7}$ BR BL TR TL 3 3 3 0 0 3 9 3	$\frac{1}{7}$ $\frac{1}{7}$ $\frac{3}{7}$ $\frac{2}{7}$ BR BR TR TR 3 3 3 3 0 0 9 9
	B if type 1A B if type 1B	$\frac{1}{7}$ $\frac{1}{7}$ $\frac{3}{7}$ $\frac{2}{7}$ BL BL BL BL 3 3 3 3 3 3 3 3	$\frac{1}{7}$ $\frac{1}{7}$ $\frac{3}{7}$ $\frac{2}{7}$ BL BR BL BR 3 3 3 0 3 0 3 0	$\frac{1}{7}$ $\frac{1}{7}$ $\frac{3}{7}$ $\frac{2}{7}$ BR BL BR BL 3 3 0 3 0 3 0 3	$\frac{1}{7}$ $\frac{1}{7}$ $\frac{3}{7}$ $\frac{2}{7}$ BR BR BR BR 3 3 0 0 0 0 0 0

Figure 16.11: The lotteries associated with the strategy profiles for the game of Figure 16.10.

		Player 2			
		L if type 2a L if type 2b	L if type 2a R if type 2b	R if type 2a L if type 2b	R if type 2a R if type 2b
Player 1	T if type 1A T if type 1B	$\frac{12}{7}$ 3	$\frac{12}{7}$ $\frac{39}{7}$	$\frac{15}{7}$ $\frac{45}{7}$	$\frac{15}{7}$ 9
	T if type 1A B if type 1B	$\frac{27}{7}$ 3	$\frac{15}{7}$ 3	$\frac{12}{7}$ $\frac{18}{7}$	0 $\frac{18}{7}$
	B if type 1A T if type 1B	$\frac{6}{7}$ 3	$\frac{12}{7}$ $\frac{30}{7}$	$\frac{15}{7}$ $\frac{36}{7}$	3 $\frac{45}{7}$
	B if type 1A B if type 1B	3 3	$\frac{15}{7}$ $\frac{12}{7}$	$\frac{12}{7}$ $\frac{9}{7}$	$\frac{6}{7}$ 0

Figure 16.12: The strategic form of the game of Figure 16.10.

(e) The only pure-strategy equilibrium where Player 1 is uncertain of Player 2's choice is (TB, LR) . The corresponding model is shown in Figure 16.13.

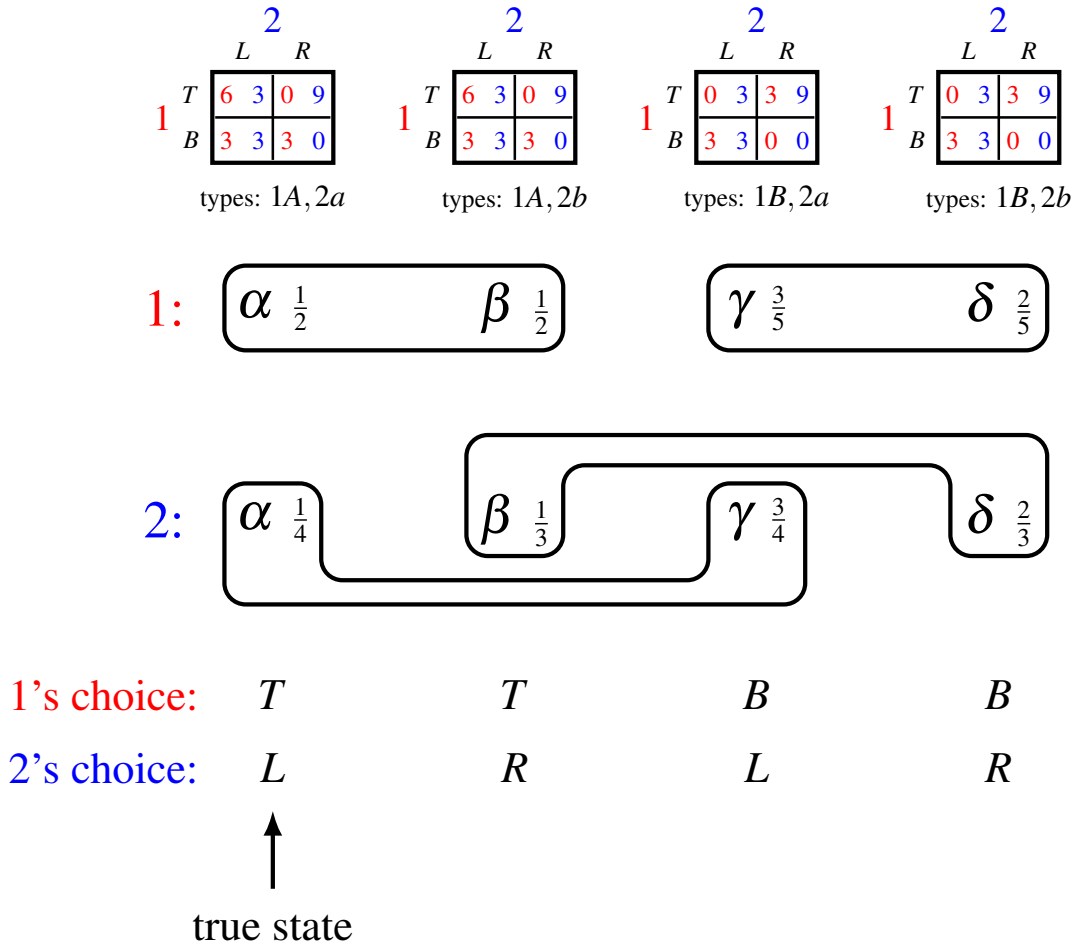


Figure 16.13: A model of the game of Figure 16.10.

There is common knowledge of rationality because at every state both players are rational:

- At states α and β Player 1 has an expected payoff of 3 from both T and B (thus T is a best reply) and at states γ and δ Player 1 has an expected payoff of $\frac{9}{5}$ from B and $\frac{6}{5}$ from T (thus B is a best reply).
- At states α and γ Player 2 has an expected payoff of 3 from L and $\frac{9}{4}$ from R (thus L is a best reply) and at states β and δ Player 2 has an expected payoff of 3 from both L and R (thus R is a best reply). \square

References

- Alchourrón, Carlos, Peter Gärdenfors, and David Makinson. On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic*, 1985, 50: 510–530.
- Aumann, Robert. Agreeing to disagree, *Annals of Statistics*, 1976, 4: 1236-1239.
- Aumann, Robert. What is game theory trying to accomplish? In: Kenneth Arrow and Seppo Honkapohja, eds, *Frontiers in economics*, Basil Blackwell, 1985, 28–76.
- Battigalli, Pierpaolo and Giacomo Bonanno. Recent results on belief, knowledge and the epistemic foundations of game theory. *Research in Economics*, 1999, 53: 149-225.
- Bonanno, Giacomo. Rational choice and AGM belief revision. *Artificial Intelligence*, 2009, 173: 1194-1203.
- Bonanno, Giacomo. AGM belief revision in dynamic games, in: K. R. Apt (ed.), *Proceedings of the 13th conference on theoretical aspects of rationality and knowledge (TARK XIII)*, ACM, New York, 2011, 37-45.
- Bonanno, Giacomo. A doxastic behavioral characterization of generalized backward induction, *Games and Economic Behavior*, 2014, 88: 221-241.
- Bonanno, Giacomo. Epistemic foundations of game theory, in: H. van Ditmarsch, J.Y. Halpern, W. van der Hoek and B. Kooi (eds), *Handbook of Logics for Knowledge and Belief*, College Publications, 2015, pp. 411–450.
- Bonanno, Giacomo. AGM-consistency and perfect Bayesian equilibrium. Part I: definition and properties, *International Journal of Game Theory*, 2013, 42: 567-592.
- Bonanno, Giacomo. AGM-consistency and perfect Bayesian equilibrium. Part II: from PBE to sequential equilibrium, *International Journal of Game Theory*,

2016,45:1071-1094.

- Bonanno, Giacomo. Reasoning about strategies and rational play in dynamic games, in: J. van Benthem, S. Ghosh, and R. Verbrugge (eds.), *Models of Strategic Reasoning: Logics, Games and Communities*, LNCS 8972. Springer, 2015.
- Bonanno, Giacomo and Klaus Nehring. How to make sense of the common prior assumption under incomplete information, *International Journal of Game Theory*, 1999, 28: 409-434.
- Camerer, Colin. *Behavioral game theory: experiments in strategic interaction*. Princeton University Press, 2003.
- Clarke, Edward H. Multipart pricing of public goods. *Public Choice*, 1971, 11: 17–33.
- Geanakoplos, John and Heraklis Polemarchakis. We can't disagree forever. *Journal of Economic Theory*, 1982, 28: 192-200.
- Gerstung, Moritz, Hani Nakhoul, and Niko Beerenwinkel. Evolutionary games with affine fitness functions: applications to cancer. *Dynamic Games and Applications*, 1, 2011, 370–385.
- Grove, Adam. Two modellings for theory change. *Journal of Philosophical Logic*, 1988, 17: 157-170.
- Gul, Faruk. A comment on Aumann's Bayesian view, *Econometrica*, 1998, 66: 923-927.
- Harsanyi, John. Games with incomplete information played by "Bayesian player", Parts I-III, *Management Science*, 8, 1967-68, pp. 159-182, 320-334, 486-502.
- Hendon, Ebbe, Jørgen Jacobsen and Brigitte Sloth, The one-shot-deviation principle for sequential rationality, *Games and Economic Behavior*, 1996, 12: 274-282.
- Kreps, David and Robert Wilson. Sequential equilibrium. *Econometrica*, 1982, 50: 863–894.
- Kripke, Saul. A Completeness Theorem in Modal Logic, *Journal of Symbolic Logic*, 1959, 24: 1–14.
- Kripke, Saul. Semantical analysis of modal logic I. Normal modal propositional calculi, *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 1963, 9: 67–96.
- Milgrom, Paul and John Roberts. Relying on the information of interested parties, *The RAND Journal of Economics*, 1986, 17: 18-32.
- Morris, Stephen. The common prior assumption in economic theory. *Economics and Philosophy*, 1995, 11: 227-53.
- Osborne, Martin J. and Ariel Rubinstein (1994). *A course in game theory*. MIT Press.
- Perea, Andrés. A note on the one-deviation property in extensive form games, *Games and Economic Behavior*, 2002, 40: 322-338.

- Perea, Andrés. *Epistemic game theory: reasoning and choice*, Cambridge University Press, 2012.
- Selten, Reinhard. Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragerträgeit, *Zeitschrift für die gesamte Staatswissenschaft* 121, 1965, Teil I, pp. 301–324, Teil II, pp. 667–689.
- Selten, Reinhard. The chain store paradox, *Theory and Decision*, 1978, 9: 127–159
- Selten, Reinhard. Re-examination of the perfectness concept for equilibrium points in extensive games, *International Journal of Game Theory*, 1975, 4: 25-55.
- Shapley, Lloyd. A Value for n-person Games. In: H.W. Kuhn and A.W. Tucker, eds., *Contributions to the Theory of Games*, Volume II, 1953.
- Shapley, Lloyd. Stochastic Games. *Proceedings of National Academy of Sciences*, Vol. 39, 1953, 1095–1100.
- Shoham, Yoav and Kevin Leyton-Brown. *Multiagent systems: algorithmic, game theoretic, and logical foundations*. Cambridge University Press, 2008.
- Smith, John Maynard. *Evolution and the theory of games*. Cambridge University Press, 1982.
- Nash, John. Equilibrium Points in N-person Games. *Proceedings of the National Academy of Sciences*, 1950, 36: 48–49.
- Nash, John. Non-cooperative Games. *Annals of Mathematics*, 1951, 54: 286–95.
- van Damme, Eric. *Stability and perfection of Nash equilibria*. Springer, 2002.
- van Ditmarsch, Hand, Joseph Halpern, Wiebe van der Hoek and Barteld Kooi (eds), *Handbook of Logics for Knowledge and Belief*, College Publications, 2015.
- Vickrey, William, Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 1961, 16: 8–37.
- von Neumann, John (1928a). Sur la theorie des jeux. *Comptes Rendus de l'Academie des Sciences*, 1928, 186.25: 1689–1691.
- von Neumann, John (1928b). Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100, 1928: 295–320.
- von Neumann, John and Oskar Morgenstern. *Theory of games and economic behavior*. Princeton University Press, 1944.
- Zermelo, Ernst. Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels, *Proc. Fifth Int. Congress of Mathematicians, National University of Singapore*, 1913, pp. 501–504. The English translation “On an application of set theory to the theory of the game of chess” can be found in Eric Rasmusen, ed., *Readings in games and information*, Wiley-Blackwell, 2001, 79-82.

Index

A

advertising game	532
affine transformation	172
AGM belief revision	301
AGM-consistency of beliefs	427
agreeing to disagree	310
Allais paradox	177
assessment	365
auction	
second-price	26
Aumann's Agreement Theorem	311

B

backward induction	
algorithm	76
reasoning	73
solution	76
bargaining game	526
basic outcome	167
Bayes' rule	296–298
Bayesian consistency of beliefs	430
Bayesian game of incomplete information	
565, 568	
Bayesian Nash equilibrium	483
Bayesian updating	299, 367, 372
behavioral strategy	225

belief

AGM revision function	303
revision	299
revision function	301
updating	299

beliefs

probabilistic	294
---------------	-----

benevolent preferences	14
------------------------	----

best reply	38
------------	----

best response	38
---------------	----

C

cardinal payoffs	190
------------------	-----

certainty	295
-----------	-----

chain-store game	83, 520
------------------	---------

chance	130
--------	-----

choice measurability	443
----------------------	-----

Clarke mechanism	29
------------------	----

Clarke, Edward H.	32
-------------------	----

common knowledge	267
------------------	-----

operator	271
----------	-----

common prior	307, 489, 565
--------------	---------------

consistent assessment	401
-----------------------	-----

Cournot game	41, 486
--------------	---------

D

- dominance
 - strict 21, 202
 - weak 21
- dominant strategy profile
 - strict 23
 - weak 23

E

- Ellsberg paradox 178
- equilibrium
 - Bayesian Nash 483
 - Nash 36, 82, 379
 - perfect Bayesian 432, 444
 - pooling 485
 - separating 485
 - sequential 404, 444
 - subgame perfect . . 124, 233, 406, 433
 - weak sequential 374, 433
- event 258, 293
- expected utility 168
- expected value 132
- extensive form
 - game 73, 229
 - game frame 72, 115

F

- fair-minded preferences 14

G

- game
 - in extensive form 73, 118, 229
 - in strategic form 18, 190
 - mixed-strategy extension 194
 - ordinal 18
 - reduced form 18, 190
 - with cardinal payoffs 229
- game frame
 - in extensive form 72, 115
 - in strategic form 15, 190
- Golden Balls game 13
- greedy preferences 14

H

- Harsanyi consistency of beliefs . . 304, 565
- Harsanyi transformation 483

I

- IDS DS 32
 - cardinal 203
- IDWDS 34
- imperfect information 113, 116
- incomplete information 480
 - dynamic games 517
 - multi-sided 492, 534
 - one-sided 481, 517
 - two-sided 488
- incredible threat 82
- information partition 258
- information set 258, 270
- iterated deletion 32
 - IDS DS 32, 203
 - IDWDS 34
 - strict dominance 32, 203
 - weak dominance 34

K

- knowledge
 - common 267
 - individual 258
 - interactive 260
 - operator 259
- Kuhn's theorem 228
- Kuhn, Harold 228

L

- like mindedness 305
- lottery
 - compound 173
 - simple 167

M

- mechanism
 - pivotal 29
- mixed strategy 192, 226
- model

of a game 344
 state-space 567
 type-space 567
 money lottery 132
 Morgenstern, Oskar 7, 168

N

Nash equilibrium 36, 38, 82
 in mixed strategies 194
 Nash's theorem 195
 Nash, John 195
 Nature 130
 node
 decision 71
 predecessor 114
 successor 114
 terminal 71
 normalized utility function 171

O

ordinal game
 in extensive form 73
 in strategic form 18
 outcome 15

P

Pareto ranking 25
 partition 114, 258
 common knowledge 269
 payoff
 cardinal 190
 von Neumann Morgenstern 190
 payoff function 18, 190
 perfect Bayesian equilibrium 432
 strongly independent 440
 weakly independent 438
 perfect information 72, 73
 perfect recall 115, 228
 pivotal mechanism 29
 plausibility order 302, 427
 pooling equilibrium 485
 pre-order 426
 preference
 indifference 15

relation 15
 strict 15
 weak 15
 preferences
 benevolent 14
 complete 16, 173
 fair-minded 14
 greedy 14
 selfish 14
 transitive 16, 173
 Prisoner's Dilemma 25
 probability
 Bayes' rule 296
 conditional 296
 distribution 293
 measure 293
 pure strategy 192

R

rationality 346
 common knowledge of 348
 rationalizable strategy 203
 risk attitude 166
 risk neutrality 132
 rooted tree 71

S

second-price auction 26
 self-enforcing agreement 37
 selfish preferences 14
 separating equilibrium 485
 sequential equilibrium 404
 sequential rationality 367, 369
 sets
 complement 291
 De Morgan's Laws 292
 intersection 292
 subset 291
 union 292
 Split or Steal game 13
 state 258
 state-space model 567
 strategy 15
 behavioral 225
 equivalent 21

in a perfect-information game 79
 in an imperfect-information game 120
 mixed 192, 226
 profile 15
 pure 192
 rationalizable 203
 strictly dominant 22
 weakly dominant 22
 winning 87
 strict dominance 21, 202
 subgame
 minimal 122
 proper 122
 subgame perfect
 algorithm 125
 equilibrium 124, 233
 subgame perfect equilibrium 406
 system of beliefs 365

T

total pre-order 426
 tree 71
 type-space model 567

U

updating 299
 utility function
 normalized 171
 ordinal 17
 von Neumann-Morgenstern 168

V

Vickrey auction 26
 Vickrey, William 29
 von Neumann, John 7, 168
 von Neumann-Morgenstern
 ranking 168
 utility function 168

W

weak dominance 21
 weak sequential equilibrium 374
 winning strategy 87

Z

Zermelo, Ernst 7