

Game Theory Basics

Lecture Notes

Bernhard von Stengel

Department of Mathematics, London School of Economics,
Houghton St, London WC2A 2AE, United Kingdom
email: stengel@nash.lse.ac.uk

September 8, 2016

(minor amendments from earlier versions of
May 16, 2014, and February 15, 2011)

The PDF version of this document has *hyperlinks* for the section headings in the table of contents [[← click here](#)], and throughout the text for all numbers of sections, equations, theorems, definitions, and similar statements.

The material in this book has been adapted and developed from material originally produced for the BSc Mathematics and Economics by distance learning offered by the University of London International Programme (<http://www.londoninternational.ac.uk/>).

I am grateful to George Zouros for his valuable comments, and to Rahul Savani for his continuing interest in this project, his thorough reading of earlier drafts, and the fun we have in joint research and writing on game theory.

Preface

Game theory is the formal study of conflict and cooperation. It is concerned with situations where “players” interact, so that it matters to each player what the other players do. Game theory provides mathematical tools to model, structure and analyse such interactive scenarios. The players may be, for example, competing firms, political voters, mating animals, or buyers and sellers on the internet. The language and concepts of game theory are widely used in economics, political science, biology, and computer science, to name just a few disciplines.

Game theory helps to understand effects of interaction that seem puzzling at first. For example, the famous “prisoners’ dilemma” explains why fishers can exhaust their resources by over-fishing: they hurt themselves collectively, but each fisher on his own cannot really change this and still profits by fishing as much as possible. Other insights come from the way of looking at interactive situations. Game theory treats players equally and recommends to each player how to play well, given what the other players do. This mindset is useful in strategic questions of management, because “you put yourself in your opponent’s shoes”.

Game theory is fascinating as a topic because of its diverse applications. The ideas of game theory started with mathematicians, most notably the outstanding mathematician John von Neumann (1903–1957). In the 1950s, a group of young researchers in mathematics at Princeton developed game theory further, among them John Nash (1928–2015), Harold Kuhn (1925–2014), and Lloyd Shapley (1923–2016). I was lucky to have met these pioneers in person and that they knew me and some of my work. Most research in game theory is now done by economists and other social scientists.

My own interests are in the mathematics of games, so I see my own research in the tradition of early game theory. My research specialty is the connection of game theory to computer science. In particular, I work on methods to find equilibria of games, which make use of insights from geometry. That interest is partly reflected in the choice of topics in this book. This book has also a strong emphasis on *methods*, so you will learn a lot of “tricks” that allow you to understand games quickly. With these methods at hand, you will be in a position to analyse games that you can create for applications to management or economics.

Structure of this book

Chapter 1 on combinatorial games is on playing and winning games with perfect information defined by rules, in particular a simple game called “Nim”, which has a central role

in that theory. This chapter introduces abstract mathematics with a fun topic. You would probably not learn its content otherwise, because it does not have a typical application in economics. However, every game theorist should know the basics of combinatorial games. In fact, this can be seen as the motto for the contents of this book: *what every game theorist should know*.

Chapter 1 on combinatorial games is independent of the other material. This topic is deliberately put at the beginning because it is more mathematical and formal than the other topics, so that it can be used to test whether you can cope with the abstract parts of game theory, and with the mathematical style of this subject.

Chapters 2, 3, and 4 successively build on each other. These chapters cover the main concepts and methods of non-cooperative game theory.

With the exception of imperfect information, the fundamentals of non-cooperative game theory are laid out in Chapter 2. This part of game theory provides ways to model in detail the agents in an interactive situation, their possible actions, and their incentives. The model is called a *game* and the agents are called *players*. There are two types of games, called *game trees* and games in *strategic form*. The game tree (also called the *extensive form* of a game) describes in depth the actions that are available to the players, how these evolve over time, and what the players know or do not know about the game. (Games with imperfect information are treated in Chapter 4.) The players' incentives are modelled by payoffs that these players want to maximise, which is the sole guiding principle in non-cooperative game theory. In contrast, cooperative game theory studies, for example, how players should split their proceeds when they decide to cooperate, but leaves it open how they enforce an agreement. (A simple example of this cooperative approach is explained in Chapter 5 on bargaining.)

Chapter 3 shows that in certain games it may be useful to leave your actions uncertain. A nice example is the football penalty kick, which serves as our introduction to zero-sum games (see Figure 3.11). The striker should not always kick into the same corner, nor should the goalkeeper always jump into the same corner, even if they are better at scoring or saving a penalty there. It is better to be unpredictable! Game theory tells the players how to choose optimal *probabilities* for each of their available strategies, which are then used to *mix* these strategies randomly. With the help of mixed strategies, every game has an *equilibrium*. This is the central result of John Nash, who discovered the equilibrium concept in 1950 for general games. For zero-sum games, this was already found earlier by John von Neumann. It is easier to prove for zero-sum games that they have an equilibrium than for general games. In a logical progression of topics, in particular when starting from win/lose games like Nim, we could have treated zero-sum games before general games. However, we choose to treat zero-sum games later, as special cases of general games, because the latter are much more important in economics. In a course on game theory, one could omit zero-sum games and their special properties, which is why we treat them in the last section of Chapter 3. However, one could not omit the concept of Nash equilibrium, which is therefore given prominence early on.

Chapter 4 explains how to model the *information* that players have in a game. This is done by means of so-called information sets in game trees, introduced in 1953 by Harold Kuhn. The central result of this chapter is called Kuhn's theorem. Essentially, this result states that players can choose a "behaviour strategy", which is a way of playing

the game that is not too complicated, provided they do not forget what they knew and did earlier. This result is typically considered as technical, and given short shrift in many game theory texts. We go into great detail in explaining this result. The first reason is that it is a beautiful result of discrete mathematics, because elementary concepts like the game tree and the information sets are combined naturally to give a new result. Secondly, the result is used in other, more elaborate “dynamic” games that develop over time. For more advanced studies of game theory, it is therefore useful to have a solid understanding of game trees with imperfect information.

The final chapter 5 on bargaining is a particularly interesting application of non-cooperative theory. It is partly independent of chapters 3 and 4, so to a large extent it can be understood after Chapter 2. A first model provides conditions – called axioms – that an acceptable “solution” to bargaining situations should fulfil, and shows that these axioms lead to a unique solution of this kind. A second model of “alternating offers” is more detailed and uses, in particular, the analysis of game trees with perfect information introduced in Chapter 2. The “bargaining solution” is thereby given an incentive-based justification with a more detailed non-cooperative model.

Game theory, and this text, use only a few *prerequisites* from elementary linear algebra, probability theory, and some analysis. You should know that vectors and points have a *geometric interpretation* (for example, as either points or vectors in three-dimensional space if the dimension is three). It should be clear how to multiply matrices, and how to multiply a matrix with a vector. The required notions from probability theory are that of *expected value* of a function (that is, function values weighted with their probabilities), and that *independent* events have a probability that is the product of the probabilities of the individual events. The concepts from analysis are those of a *continuous* function, which, for example, assumes its maximum on a compact (closed and bounded) domain. None of these concepts is difficult, and you are reminded of their basic ideas in the text whenever they are needed.

We introduce and illustrate each concept with examples, many pictures, and a minimum of notation. Not every concept is defined with a formal definition (although many concepts are), but instead explained by means of an example. It should in each case be clear how the concept applies in similar settings, and in general. On the other hand, this requires some maturity in dealing with mathematical concepts, and in being able to generalise from examples.

Distinction to other textbooks

Game theory has seen an explosion in interest in recent years. Correspondingly, many new textbooks on game theory have recently appeared. The present text is complementary (or “orthogonal”) to most existing textbooks.

First of all, it is mathematical in spirit, meaning it can be used as a textbook for teaching game theory as a mathematics course. On the other hand, the mathematics in this book should be accessible enough to students of economics, management, and other social sciences.

The book is relatively swift in treating games in strategic form and game trees in a single chapter (chapter 3), which are often considered separately. I find these concepts simple enough to allow for that approach.

Mixed strategies and the best response condition are very useful for finding equilibria in games, so they are given detailed treatment. Similarly, game trees with imperfect information, and the concepts of perfect recall and behaviour strategies, and Kuhn's theorem, are also given an in-depth treatment, because this not found in many textbooks at this level.

The book does not try to be comprehensive in presenting the most relevant economic applications of game theory. Corresponding "stories" are only told for illustration of the concepts.

One book served as a starting point for the selection of the material, namely

- K. Binmore, *Fun and Games*, D. C. Heath, 1991. Its recent, significantly revised version is *Playing for Real*, Oxford University Press, 2007.

This book describes Nim as a game that one can learn to play perfectly. In turn, this led me to the question of why the binary system has such an important role, and towards the study of combinatorial games. The topic of bargaining is also motivated by Binmore's book. At the same time, the present text is very different from Binmore's book, because it tries to treat as few side topics as possible.

Two short sections can be considered as slightly non-standard side topics, namely the sections 2.9, Symmetries involving strategies, and 4.9 Perfect recall and order of moves, both marked with a star as "can be skipped at first reading". I included them as topics that I have not found elsewhere. Section 4.9 demonstrates how to reason carefully about moves in extensive games, and should help understand the perfect recall condition.

Methods, not philosophy

The general emphasis of the book is to teach *methods, not philosophy*. Game theorists tend to question and to justify the approaches they take, for example the concept of Nash equilibrium, and the assumed common knowledge of all players about the rules of the game. These questions are of course very important. In fact, in practice these are probably the very issues that make a game-theoretic analysis questionable. However, this problem is not remedied by a lengthy discussion of why one should play Nash equilibrium.

I think a student who learns about game theory should first become fluent in knowing the modelling tools (such as game trees and the strategic form) and in analysing games (finding their Nash equilibria). That toolbox will then be useful when comparing different game-theoretic models and looking at their implications. In many disciplines that are further away from mathematics, these possible implications are typically more interesting than the mathematical model itself (the model is necessarily imperfect, so there is no point in being dogmatic about the analysis).

Contents

1	Nim and combinatorial games	11
1.1	Aims of the chapter	11
1.2	Learning objectives	11
1.3	Further reading	12
1.4	What is combinatorial game theory?	12
1.5	Nim – rules	14
1.6	Combinatorial games, in particular impartial games	15
1.7	Simpler games and notation for Nim heaps	16
1.8	Sums of games	17
1.9	Equivalent games	19
1.10	Sums of Nim heaps	22
1.11	Poker Nim and the mex rule	26
1.12	Finding Nim values	27
1.13	Exercises for Chapter 1	30
2	Games as trees and in strategic form	37
2.1	Learning objectives	37
2.2	Further reading	37
2.3	Introduction	38
2.4	Definition of game trees	38
2.5	Backward induction	41
2.6	Strategies and strategy profiles	42
2.7	Games in strategic form	43
2.8	Symmetric games	44
2.9	Symmetries involving strategies*	47
2.10	Dominance and elimination of dominated strategies	48
2.11	Nash equilibrium	51
2.12	Reduced strategies	53
2.13	Subgame perfect Nash equilibrium (SPNE)	55
2.14	Commitment games	58
2.15	Exercises for Chapter 2	62
3	Mixed strategy equilibria	69
3.1	Learning objectives	69
3.2	Further reading	69

3.3	Introduction	70
3.4	Expected-utility payoffs	71
3.5	Example: Compliance inspections	72
3.6	Bimatrix games	75
3.7	Matrix notation for expected payoffs	76
3.8	Convex combinations and mixed strategy sets	78
3.9	The best response condition	79
3.10	Existence of mixed equilibria	81
3.11	Finding mixed equilibria	83
3.12	The upper envelope method	87
3.13	Degenerate games	90
3.14	Zero-sum games	95
3.15	Exercises for Chapter 3	101
4	Game trees with imperfect information	105
4.1	Learning objectives	105
4.2	Further reading	105
4.3	Introduction	106
4.4	Information sets	106
4.5	Extensive games	112
4.6	Strategies for extensive games and the strategic form	114
4.7	Reduced strategies	115
4.8	Perfect recall	116
4.9	Perfect recall and order of moves*	120
4.10	Behaviour strategies	122
4.11	Kuhn's theorem: behaviour strategies suffice	126
4.12	Subgames and subgame perfect equilibria	130
4.13	Exercises for Chapter 4	132
5	Bargaining	135
5.1	Learning objectives	135
5.2	Further reading	135
5.3	Introduction	135
5.4	Bargaining sets	136
5.5	Bargaining axioms	138
5.6	The Nash bargaining solution	140
5.7	Splitting a unit pie	144
5.8	The ultimatum game	146
5.9	Alternating offers over two rounds	150
5.10	Alternating offers over several rounds	153
5.11	Stationary strategies	156
5.12	The Nash bargaining solution via alternating offers	160
5.13	Exercises for Chapter 5	162

6	Geometric representation of equilibria	165
6.1	Learning objectives	165
6.2	Further reading	165
6.3	Introduction	166
6.4	Best response regions	166
6.5	Identifying equilibria	172
6.6	The Lemke–Howson algorithm	173
6.7	Odd number of Nash equilibria	176
6.8	Exercises for Chapter 6	178
7	Linear programming and zero-sum games	181
7.1	Learning objectives	181
7.2	Further reading	182
7.3	Linear programs and duality	182
7.4	The minimax theorem via LP duality	185
7.5	General LP duality	187
7.6	The Lemma of Farkas and proof of strong LP duality*	191
7.7	Exercises for Chapter 7	195

Chapter 1

Nim and combinatorial games

1.1 Aims of the chapter

This chapter introduces the basics of combinatorial games, and explains the central role of the game Nim. A detailed summary of the chapter is given in Section 1.4.

Furthermore, this chapter demonstrates the use of abstract mathematics in game theory. This chapter is written more formally than the other chapters, in parts in the traditional mathematical style of definitions, theorems and proofs. One reason for doing this, and why we start with combinatorial games, is that this topic and style serves as a warning shot to those who think that game theory, and this text in particular, is “easy”. If we started with the well-known “prisoner’s dilemma” (which makes its due appearance in Chapter 2), the less formally inclined student might be lulled into a false sense of familiarity and “understanding”. We therefore start deliberately with an unfamiliar topic.

This is a mathematics text, with great emphasis on rigour and clarity, and on using mathematical notions precisely. As mathematical prerequisites, game theory requires only the very basics of linear algebra, calculus and probability theory. However, game theory provides its own conceptual tools that are used to model and analyse interactive situations. This text emphasises the mathematical structure of these concepts, which belong to “discrete mathematics”. Learning a number of new mathematical concepts is exemplified by combinatorial game theory, and it will continue in the study of classical game theory in the later chapters.

1.2 Learning objectives

After studying this chapter, you should be able to:

- play Nim optimally;
- explain the concepts of game-sums, equivalent games, Nim values and the mex rule;
- apply these concepts to play other impartial games like those described in the exercises.

1.3 Further reading

Very few textbooks on game theory deal with combinatorial games. An exception is Chapter 1 of the following book:

- Mendelson, Elliot *Introducing Game Theory and Its Applications*. (Chapman & Hall / CRC, 2004) [ISBN 1584883006].

The winning strategy for the game Nim based on the binary system was first described in the following article, which is available electronically from the JSTOR archive:

- Bouton, Charles “Nim, a game with a complete mathematical theory.” *The Annals of Mathematics*, 2nd Ser., Vol. 3, No. 1/4 (1902), pp. 35–39.

The definitive text on combinatorial game theory is the set of volumes “Winning Ways” by Berlekamp, Conway and Guy. The material of this chapter appears in the first volume:

- Berlekamp, Elwyn R., John H. Conway and Richard K. Guy *Winning Ways for Your Mathematical Plays, Volume I*, second edition. (A. K. Peters, 2001) [ISBN 1568811306].

Some small pieces of that text have been copied here nearly verbatim, for example in Sections 1.6, 1.8, and 1.11 below.

The four volumes of “Winning Ways” are beautiful books. However, they are not suitable reading for a beginner, because the mathematics is hard, and the reader is confronted with a wealth of material. The introduction to combinatorial game theory given here represents a very small fraction of that body of work, but may invite you to study it further.

A very informative and entertaining mathematical tour of parlour games is

- Bewersdorff, Jörg *Logic, Luck and White Lies*. (A. K. Peters, 2005) [ISBN 1568812108].
Combinatorial games are treated in part II of that book.

1.4 What is combinatorial game theory?

This chapter is on the topic of *combinatorial games*. These are games with two players, perfect information, and no chance moves, specified by certain rules. Familiar games of this sort are Chess, Go, Checkers, Tic-Tac-Toe, Dots-and-Boxes, and Nim. Such games can be played perfectly in the sense that either one player can force a win or both can force a draw. In reality, games like Chess and Go are too complex to find an optimal strategy, and they derive their attraction from the fact that (so far) it is not known how to play them perfectly. We will, however, learn how to play Nim perfectly.

There is a “classical” game theory with applications in economics which is very different from combinatorial game theory. The games in classical game theory are typically formal models of conflict *and* cooperation which cannot only be lost or won, and in which

there is often no perfect information about past and future moves. To the economist, combinatorial games are not very interesting. Chapters 2–5 of the book are concerned with classical game theory.

Why, then, study combinatorial games at all in a text that is mostly about classical game theory, and which aims to provide an insight into the theory of games as used in economics? The reason is that combinatorial games have a rich and interesting mathematical theory. We will explain the basics of that theory, in particular the central role of the game Nim for impartial games. It is non-trivial mathematics, it is fun, and you, the student, will have learned something that you would most likely not have learned otherwise.

The first “trick” from combinatorial game theory is how to win in the game Nim, using the binary system. Historically, that winning strategy was discovered first (published by Charles Bouton in 1902). Only later did the central importance of Nim, in what is known as the Sprague–Grundy theory of impartial games, become apparent. It also revealed why the binary system is important (and not, say, the ternary system, where numbers are written in base three), and learning that is more satisfying than just learning how to use it.

In this chapter, we first define the game Nim and more general classes of games with perfect information. These are games where every player knows exactly the state of the game. We then define and study the concepts listed in the learning outcomes above, which are the concepts of game-sums, equivalent games, Nim values and the mex rule. It is best to learn these concepts by following the chapter in detail. We give a brief summary here, which will make more sense, and should be re-consulted, after a first study of the chapter (so do not despair if you do not understand this summary).

Mathematically, any game is defined by other “games” that a player can reach in his first move. These games are called the *options* of the game. This seemingly circular definition of a “game” is sound because the options are *simpler* games, which need fewer moves in total until they end. The definition is therefore not circular, but *recursive*, and the mathematical tool to argue about such games is that of mathematical *induction*, which will be used extensively (it will also recur in Chapter 2 as “backward induction” for game trees). Here, it is very helpful to be familiar with mathematical induction for proving statements about natural numbers.

We focus here on *impartial* games, where the available moves are the same no matter whether player I or player II is the player to make a move. Games are “combined” by the simple rule that a player can make a move in exactly one of the games, which defines a *sum* of these games. In a “losing game”, the first player to move loses (assuming, as always, that both players play as well as they can). An impartial game added to itself is always losing, because any move can be copied in the other game, so that the second player always has a move left. This is known as the “copycat” principle (Lemma 1.6). An important observation is that a losing game can be “added” (via the game-sum operation) to any game without changing the winning or losing properties of the original game.

In Section 1.10, the central Theorem 1.10 explains the winning strategy in Nim. The importance of Nim for impartial games is then developed in Section 1.11 via the beautiful mex rule. After the comparatively hard work of the earlier sections, we almost instantly obtain that any impartial game is equivalent to a Nim heap (Corollary 1.13).

At the end of the chapter, the sizes of these equivalent Nim heaps (called Nim values) are computed for some examples of impartial games. Many other examples are studied in the exercises.

Our exposition is distinct from the classic text “Winning Ways” in the following respects: First, we only consider impartial games, even though many aspects carry over to more general combinatorial games. Secondly, we use a precise definition of *equivalent* games (see Section 1.9), because a game where you are bound to lose against a smart opponent is not the *same* as a game where you have already lost. Two such games are merely equivalent, and the notion of equivalent games is helpful in understanding the theory. So this text is much more restricted, but to some extent more precise than “Winning Ways”, which should help make this topic accessible and enjoyable.

1.5 Nim – rules

The game Nim is played with heaps (or piles) of chips (or counters, beans, pebbles, matches). Players alternate in making a move, by removing some chips from one of the heaps (at least one chip, possibly the entire heap). The first player who cannot move any more loses the game.

The players will be called, rather unimaginatively, player I and player II, with player I to start the game.

For example, consider three heaps of size 1, 1, 2. What is a good move? Removing one of the chips from the heap with two chips will create the position 1, 1, 1, then player II must move to 1, 1, then player I to 1, and then player II takes the last chip and wins. So this is not a good opening move. The winning move is to remove all chips from the heap of size 2, to reach position 1, 1, and then player I will win. Hence we call 1, 1, 2 a *winning position*, and 1, 1 a *losing position*.

When moving in a winning position, the player to move can win by playing well, by moving to a losing position of the other player. In a losing position, the player to move will lose no matter what move she chooses, if her opponent plays well. This means that *all* moves from a losing position lead to a winning position of the opponent. In contrast, one needs only *one* good move from a winning position that goes to a losing position of the next player.

Another winning position consists of three Nim heaps of sizes 1, 1, 1. Here all moves result in the same position and player I always wins. In general, a player in a winning position must play well by picking the right move. We assume that players play well, forcing a win if they can.

Suppose Nim is played with only two heaps. If the two heaps have equal size, for example in position 4, 4, then the first player to move loses (so this is a losing position), because player II can always *copy* player I’s move by equalising the two heaps. If the two heaps have different sizes, then player I can equalise them by removing an appropriate number of chips from the larger heap, putting player II in a losing position. The rule for 2-heap Nim is therefore:

Lemma 1.1 *The Nim position m, n is winning if and only if $m \neq n$, otherwise losing, for all $m, n \geq 0$.*

This lemma applies also when $m = 0$ or $n = 0$, and thus includes the cases that one or both heap sizes are zero (meaning only one heap or no heap at all).

With three or more heaps, Nim becomes more difficult. For example, it is not immediately clear if, say, positions $1, 4, 5$ or $2, 3, 6$ are winning or losing positions.

\Rightarrow At this point, you should try Exercise 1.1(a) on page 30.

1.6 Combinatorial games, in particular impartial games

The games we study in this chapter have, like Nim, the following properties:

1. There are just two players.
2. There are several, usually finitely many, *positions*, and sometimes a particular *starting position*.
3. There are clearly defined *rules* that specify the *moves* that either player can make from a given position to the possible new positions, which are called the *options* of that position.
4. The two players move alternately, in the game as a whole.
5. In the *normal play* convention a player unable to move loses.
6. The rules are such that play will always come to an end because some player will be unable to move. This is called the *ending condition*. So there can be no games which are drawn by repetition of moves.
7. Both players know what is going on, so there is *perfect information*.
8. There are no *chance moves* such as rolling dice or shuffling cards.
9. The game is *impartial*, that is, the possible moves of a player only depend on the position but not on the player.

As a negation of condition 5, there is also the *misère play* convention where a player unable to move *wins*. In the surrealist (and unsettling) movie “Last year at Marienbad” by Alain Resnais from 1962, misère Nim is played, several times, with rows of matches of sizes $1, 3, 5, 7$. If you have a chance, try to watch that movie and spot when the other player (not the guy who brought the matches) makes a mistake! Note that this is misère Nim, not Nim, but you will be able to find out how to play it once you know how to play Nim. (For games other than Nim, normal play and misère versions are typically not so similar.)

In contrast to condition 9, games where the available moves depend on the player (as in Chess where one player can only move white pieces and the other only black pieces) are called *partisan* games. Much of combinatorial game theory is about partisan games, which we do not consider to keep matters simple.

Chess, and the somewhat simpler tic-tac-toe, also fail condition 6 because they may end in a tie or draw. The card game Poker does not have perfect information (as required in 7) and would lose all its interest if it had. The analysis of Poker, although it is also a win-or-lose game, leads to the “classical” theory of zero-sum games (with imperfect information) that we will consider later. The board game backgammon is a game with perfect information but with chance moves (violating condition 8) because dice are rolled.

We will be relatively informal in style, but our notions are precise. In condition 3 above, for example, the term *option* refers to a position that is reachable in one move from the current position; do not use “option” when you mean “move”. Similarly, we will later use the term *strategy* to define a plan of moves, one for every position that can occur in the game. Do not use “strategy” when you mean “move”. However, we will take some liberty in identifying a game with its starting position when the rules of the game are clear.

⇒ Try now exercises 1.2 and 1.3 starting on page 30.

1.7 Simpler games and notation for Nim heaps

A game, like Nim, is defined by its rules, and a particular starting position. Let G be such a particular instance of Nim, say with the starting position 1,1,2. Knowing the rules, we can identify G with its starting position. Then the options of G are 1,2, and 1,1,1, and 1,1. Here, position 1,2 is obtained by removing either the first or the second heap with one chip only, which gives the same result. Positions 1,1,1 and 1,1 are obtained by making a move in the heap of size two. It is useful to list the options systematically, considering one heap to move in at a time, so as not to overlook any option.

Each of the options of G is the starting position of another instance of Nim, defining one of the new games H, J, K , say. We can also say that G is defined by the moves to these games H, J, K , and we call these *games* also the *options* of G (by identifying them with their starting positions; recall that the term “option” has been defined in point 3 of Section 1.6).

That is, we can define a game as follows: Either the game has no move, and the player to move loses, or a game is given by one or several possible moves to new games, in which the other player makes the initial move. In our example, G is defined by the possible moves to H, J , or K . With this definition, the entire game is completely specified by listing the initial moves and what games they lead to, because all subsequent use of the rules is encoded in those games.

This is a *recursive* definition because a “game” is defined in terms of “game” itself. We have to add the *ending* condition that states that every sequence of moves in a game must eventually end, to make sure that a game cannot go on indefinitely.

This recursive condition is similar to defining the set of natural numbers as follows: (a) 0 is a natural number; (b) if n is a natural number, then so is $n + 1$; and (c) all natural numbers are obtained in this way, starting from 0. Condition (c) can be formalised by the

principle of induction that says: if a property $P(n)$ is true for $n = 0$, and if the property $P(n)$ implies $P(n + 1)$, then it is true for all natural numbers.

We use the following *notation for Nim heaps*. If G is a single Nim heap with n chips, $n \geq 0$, then we denote this game by $*n$. This game is completely specified by its options, and they are:

$$\text{options of } *n : \quad *0, *1, *2, \dots, *(n-1). \quad (1.1)$$

Note that $*0$ is the empty heap with no chips, which allows no moves. It is invisible when playing Nim, but it is useful to have a notation for it because it defines the most basic losing position. (In combinatorial game theory, the game with no moves, which is the empty Nim heap $*0$, is often simply denoted as 0.)

We could use (1.1) as the definition of $*n$; for example, the game $*4$ is defined by its options $*0, *1, *2, *3$. It is very important to include $*0$ in that list of options, because it means that $*4$ has a winning move. Condition (1.1) is a recursive definition of the game $*n$, because its options are also defined by reference to such games $*k$, for numbers k smaller than n . This game fulfils the ending condition because the heap gets successively smaller in any sequence of moves.

If G is a game and H is a game reachable by one or more successive moves from the starting position of G , then the game H is called *simpler* than G . We will often prove a property of games inductively, using the assumption that the property applies to all simpler games. An example is the – already stated and rather obvious – property that one of the two players can force a win. (Note that this applies to games where winning or losing are the only two outcomes for a player, as implied by the “normal play” convention in 5 above.)

Lemma 1.2 *In any game G , either the starting player I can force a win, or player II can force a win.*

Proof. When the game has no moves, player I loses and player II wins. Now assume that G does have options, which are simpler games. By inductive assumption, in each of these games one of the two players can force a win. If, in all of them, the starting player (which is player II in G) can force a win, then she will win in G by playing accordingly. Otherwise, at least one of the starting moves in G leads to a game G' where the second-moving player in G' (which is player I in G) can force a win, and by making that move, player I will force a win in G . \square

If in G , player I can force a win, its starting position is a winning position, and we call G a *winning game*. If player II can force a win, G starts with a losing position, and we call G a *losing game*.

1.8 Sums of games

We continue our discussion of Nim. Suppose the starting position has heap sizes 1, 5, 5. Then the obvious good move is to option 5, 5, which is losing.

What about Nim with four heaps of sizes 2, 2, 6, 6? This is losing, because 2, 2 and 6, 6 independently are losing positions, and any move in a heap of size 2 can be copied in the other heap of size 2, and similarly for the heaps of size 6. There is a second way of looking at this example, where it is not just two losing games put together: consider the game with heap sizes 2, 6. This is a winning game. However, two such winning games, put together to give the game 2, 6, 2, 6, result in a losing game, because any move in one of the games 2, 6, for example to 2, 4, can be copied in the other game, also to 2, 4, giving the new position 2, 4, 2, 4. So the second player, who plays “copycat”, always has a move left (the copying move) and hence cannot lose.

Definition 1.3 The *sum* of two games G and H , written $G + H$, is defined as follows: The player may move in either G or H as allowed in that game, leaving the position in the other game unchanged.

Note that $G + H$ is a notation that applies here to *games* and not to numbers, even if the games are in some way defined using numbers (for example as Nim heaps). The result is a new game.

More formally, assume that G and H are defined in terms of their options (via moves from the starting position) G_1, G_2, \dots, G_k and H_1, H_2, \dots, H_m , respectively. Then the options of $G + H$ are given as

$$\text{options of } G + H : G_1 + H, \dots, G_k + H, G + H_1, \dots, G + H_m. \quad (1.2)$$

The first list of options $G_1 + H, G_2 + H, \dots, G_k + H$ in (1.2) simply means that the player makes his move in G , the second list $G + H_1, G + H_2, \dots, G + H_m$ that he makes his move in H .

We can define the game Nim as a sum of Nim heaps, where any single Nim heap is recursively defined in terms of its options by (1.1). So the game Nim with heaps of size 1, 4, 6 is written as $*1 + *4 + *6$.

The “addition” of games with the abstract $+$ operation leads to an interesting connection of combinatorial games with abstract algebra. If you are somewhat familiar with the concept of an abstract *group*, you will enjoy this connection; if not, you do not need to worry, because this connection it is not essential for our development of the theory.

A group is a set with a binary operation $+$ that fulfils three properties:

1. The operation $+$ is *associative*, that is, $G + (J + K) = (G + J) + K$ holds for all G, J, K .
2. The operation $+$ has a *neutral element* 0, so that $G + 0 = G$ and $0 + G = G$ for all G .
3. Every element G has an *inverse* $-G$ so that $G + (-G) = 0$.

Furthermore,

4. The group is called *commutative* (or “abelian”) if $G + H = H + G$ holds for all G, H .

Familiar groups in mathematics are, for example, the set of integers with addition, or the set of positive real numbers with multiplication (where the multiplication operation is written as \cdot , the neutral element is 1, and the inverse of G is written as G^{-1}).

The games that we consider form a group as well. In the way the sum of two games G and H is defined, $G + H$ and $H + G$ define the same game, so $+$ is commutative. Moreover, when one of these games is itself a sum of games, for example $H = J + K$, then $G + H$ is $G + (J + K)$ which means the player can make a move in exactly one of the games G , J , or K . This means obviously the same as the sum of games $(G + J) + K$, that is, $+$ is associative. The sum $G + (J + K)$, which is the same as $(G + J) + K$, can therefore be written unambiguously as $G + J + K$.

An obvious neutral element is the empty Nim heap $*0$, because it is “invisible” (it allows no moves), and adding it to any game G does not change the game.

However, there is no direct way to get an inverse operation because for any game G which has some options, if one adds any other game H to it (the intention being that H is the inverse $-G$), then $G + H$ will have some options (namely at least the options of moving in G and leaving H unchanged), so that $G + H$ is not equal to the empty Nim heap.

The way out of this is to identify games that are “equivalent” in a certain sense. We will see shortly that if $G + H$ is a losing game (where the first player to move cannot force a win), then that losing game is “equivalent” to $*0$, so that H fulfils the role of an inverse of G .

1.9 Equivalent games

There is a neutral element that can be added to any game G without changing it. By definition, because it allows no moves, it is the empty Nim heap $*0$:

$$G + *0 = G. \quad (1.3)$$

However, other games can also serve as neutral elements for the addition of games. We will see that any losing game can serve that purpose, provided we consider certain games as equivalent according to the following definition.

Definition 1.4 Two games G, H are called *equivalent*, written $G \equiv H$, if and only if for any other game J , the sum $G + J$ is losing if and only if $H + J$ is losing.

In definition 1.4, we can also say that $G \equiv H$ if for any other game J , the sum $G + J$ is winning if and only if $H + J$ is winning. In other words, G is equivalent to H if, whenever G appears in a sum $G + J$ of games, then G can be replaced by H without changing whether $G + J$ is winning or losing.

One can verify easily that \equiv is indeed an equivalence relation, meaning it is reflexive ($G \equiv G$), symmetric ($G \equiv H$ implies $H \equiv G$), and transitive ($G \equiv H$ and $H \equiv K$ imply $G \equiv K$; all these conditions hold for all games G, H, K).

Using $J = *0$ in definition 1.4 and (1.3), $G \equiv H$ implies that G is losing if and only if H is losing. The converse is not quite true: just because two games are winning does not mean they are equivalent, as we will see shortly. However, any two *losing* games are equivalent, because they are all equivalent to $*0$:

Lemma 1.5 *If G is a losing game (the second player to move can force a win), then $G \equiv *0$.*

Proof. Let G be a losing game. We want to show $G \equiv *0$. By definition 1.4, this is true if and only if for any other game J , the game $G + J$ is losing if and only if $*0 + J$ is losing. According to (1.3), this holds if and only if J is losing.

So let J be any other game; we want to show that $G + J$ is losing if and only if J is losing. Intuitively, adding the losing game G to J does not change which player in J can force a win, because any intermediate move in G by his opponent is simply countered by the winning player, until the moves in G are exhausted.

Formally, we first prove by induction the simpler claim that for all games J , if J is losing, then $G + J$ is losing. (So we first ignore the “only if” part.) Our inductive assumptions for this simpler claim are: for all losing games G'' that are simpler than G , if J is losing, then $G'' + J$ is losing; and for all games J'' that are simpler than J , if J'' is losing, then $G + J''$ is losing.

So suppose that J is losing. We want to show that $G + J$ is losing. Any initial move in J leads to an option J' which is winning, which means that there is a corresponding option J'' of J' (by player II’s reply) where J'' is losing. Hence, when player I makes the corresponding initial move from $G + J$ to $G + J'$, player II can counter by moving to $G + J''$. By inductive assumption, this is losing because J'' is losing. Alternatively, player I may move from $G + J$ to $G' + J$. Because G is a losing game, there is a move by player II from G' to G'' where G'' is again a losing game, and hence $G'' + J$ is also losing, by inductive assumption, because J is losing. This completes the induction and proves the claim.

What is missing is to show that if $G + J$ is losing, so is J . If J was winning, then there would be a winning move to some option J' of J where J' is losing, but then, by our claim (the “if” part that we just proved), $G + J'$ is losing, which would be a winning option in $G + J$ for player I. But this is a contradiction. This completes the proof. \square

The preceding lemma says that any losing game Z , say, can be added to a game G without changing whether G is winning or losing (in Lemma 1.5, Z is called G). That is, extending (1.3),

$$Z \text{ losing} \implies G + Z \equiv G. \quad (1.4)$$

As an example, consider $Z = *1 + *2 + *3$, which is Nim with three heaps of sizes 1, 2, 3. To see that Z is losing, we examine the options of Z and show that all of them are winning games. Removing an entire heap leaves two unequal heaps, which is a winning position by Lemma 1.1. Any other move produces three heaps, two of which have equal size. Because two equal heaps define a losing Nim game Z , they can be ignored by (1.4), meaning that all these options are like single Nim heaps and therefore winning positions, too.

So $Z = *1 + *2 + *3$ is losing. The game $G = *4 + *5$ is clearly winning. By (1.4), the game $G + Z$ is equivalent to G and is also winning. However, verifying directly that $*1 + *2 + *3 + *4 + *5$ is winning would not be easy to see without using (1.4).

It is an easy exercise to show that in sums of games, games can be replaced by equivalent games, resulting in an equivalent sum. That is, for all games G, H, J ,

$$G \equiv H \implies G + J \equiv H + J. \quad (1.5)$$

Note that (1.5) is not merely a re-statement of definition 1.4, because equivalence of the games $G + J$ and $H + J$ means more than just that the games are either both winning or both losing (see the comments before Lemma 1.9 below).

Lemma 1.6 (The copycat principle) $G + G \equiv *0$ for any impartial game G .

Proof. Given G , we assume by induction that the claim holds for all simpler games G' . Any option of $G + G$ is of the form $G' + G$ for an option G' of G . This is winning by moving to the game $G' + G'$ which is losing, by inductive assumption. So $G + G$ is indeed a losing game, and therefore equivalent to $*0$ by Lemma 1.5. \square

We now come back to the issue of inverse elements in abstract groups, mentioned at the end of Section 1.8. If we identify equivalent games, then the addition $+$ of games defines indeed a group operation. The neutral element is $*0$, or any equivalent game (that is, a losing game).

The inverse of a game G , written as the negative $-G$, fulfils

$$G + (-G) \equiv *0. \quad (1.6)$$

Lemma 1.6 shows that for an impartial game, $-G$ is simply G itself.

Side remark: For games that are not impartial, that is, partisan games, $-G$ exists also. It is G but with the roles of the two players exchanged, so that whatever move was available to player I is now available to player II and vice versa. As an example, consider the game checkers (with the rule that whoever can no longer make a move loses), and let G be a certain configuration of pieces on the checkerboard. Then $-G$ is the same configuration with the white and black pieces interchanged. Then in the game $G + (-G)$, player II (who can move the black pieces, say), can also play “copycat”. Namely, if player I makes a move in either G or $-G$ with a white piece, then player II copies that move with a black piece on the other board ($-G$ or G , respectively). Consequently, player II always has a move available and will win the game, so that $G + (-G)$ is indeed a losing game for the starting player I, that is, $G + (-G) \equiv *0$. However, we only consider impartial games, where $-G = G$.

The following condition is very useful to prove that two games are equivalent.

Lemma 1.7 Two impartial games G, H are equivalent if and only if $G + H \equiv *0$.

Proof. If $G \equiv H$, then by (1.5) and Lemma 1.6, $G + H \equiv H + H \equiv *0$. Conversely, $G + H \equiv *0$ implies $G \equiv G + H + H \equiv *0 + H \equiv H$. \square

Sometimes, we want to prove equivalence inductively, where the following observation is useful.

Lemma 1.8 *Two games G and H are equivalent if all their options are equivalent, that is, for every option of G there is an equivalent option of H and vice versa.*

Proof. Assume that for every option of G there is an equivalent option of H and vice versa. We want to show $G + H \equiv *0$. If player I moves from $G + H$ to $G' + H$ where G' is an option in G , then there is an equivalent option H' of H , that is, $G' + H' \equiv *0$ by Lemma 1.7. Moving there defines a winning move in $G' + H$ for player II. Similarly, player II has a winning move if player I moves to $G + H'$ where H' is an option of H , namely to $G' + H'$ where G' is an option of G that is equivalent to H' . So $G + H$ is a losing game as claimed, and $G \equiv H$ by Lemma 1.5 and Lemma 1.7. \square

Note that Lemma 1.8 states only a sufficient condition for the equivalence of G and H . Games can be equivalent without that property. For example, $G + G \equiv *0$, but $*0$ has no options whereas $G + G$ has many.

We conclude this section with an important point. Equivalence of two games is a finer distinction than whether the games are both losing or both winning, because that property has to be preserved in sums of games as well. Unlike losing games, winning games are in general *not* equivalent.

Lemma 1.9 *Two Nim heaps are equivalent only if they have equal size: $*n \equiv *m \implies n = m$.*

Proof. By Lemma 1.7, $*n \equiv *m$ if and only if $*n + *m$ is a losing position. By Lemma 1.1, this implies $n = m$. \square

That is, different Nim heaps are not equivalent. In a sense, this is due to the different amount of “freedom” in making a move, depending on the size of the heap. However, all the relevant freedom in making a move in an impartial game can be captured by a Nim heap. We will later show that *any impartial game* is equivalent to *some* Nim heap.

1.10 Sums of Nim heaps

Before we show how impartial games can be represented as Nim heaps, we consider the game of Nim itself. We show in this section how any Nim game, which is a sum of Nim heaps, is equivalent to a single Nim heap. As an example, we know that $*1 + *2 + *3 \equiv *0$, so by Lemma 1.7, $*1 + *2$ is equivalent to $*3$. In general, however, the sizes of the Nim heaps cannot simply be added to obtain the equivalent Nim heap (which by Lemma 1.9 has a unique size). For example, as shown after (1.4), $*1 + *2 + *3 \equiv *0$, that is, $*1 + *2 + *3$ is a losing game and not equivalent to the Nim heap $*6$. Adding the game $*2$ to both sides of the equivalence $*1 + *2 + *3 \equiv *0$ gives $*1 + *3 \equiv *2$, and in a similar way $*2 + *3 \equiv *1$, so any two heaps from sizes 1, 2, 3 has the third size as its equivalent single heap. This rule is very useful in simplifying Nim positions with small heap sizes.

If $*k \equiv *n + *m$, we also call k the *Nim sum* of n and m , written $k = n \oplus m$. The following theorem states that for distinct powers of two, their Nim sum is the ordinary sum. For example, $1 = 2^0$ and $2 = 2^1$, so $1 \oplus 2 = 1 + 2 = 3$.

Theorem 1.10 *Let $n \geq 1$, and $n = 2^a + 2^b + 2^c + \cdots$, where $a > b > c > \cdots \geq 0$. Then*

$$*n \equiv *(2^a) + *(2^b) + *(2^c) + \cdots. \quad (1.7)$$

We first discuss the implications of this theorem, and then prove it. The right-hand side of (1.7) is a sum of games, whereas n itself is represented as a sum of powers of two. Any n is uniquely given as such a sum. This amounts to the binary representation of n , which, if $n < 2^{a+1}$, gives n as the sum of all powers of two $2^a, 2^{a-1}, 2^{a-2}, \dots, 2^0$, each power multiplied with one or zero. These ones and zeros are then the digits in the binary representation of n . For example,

$$13 = 8 + 4 + 1 = 1 \times 2^3 + 1 \times 2^2 + 0 \times 2^1 + 1 \times 2^0,$$

so that 13 in decimal is written as 1101 in binary. Theorem 1.10 uses only the powers of two $2^a, 2^b, 2^c, \dots$ that correspond to the digits “one” in the binary representation of n .

Equation (1.7) shows that $*n$ is equivalent to the game sum of many Nim heaps, all of which have a size that is a power of two. Any other Nim heap $*m$ is also a sum of such games, so that $*n + *m$ is a game sum of several heaps, where equal heaps cancel out in pairs. The remaining heap sizes are all distinct powers of two, which can be added to give the size of the single Nim heap $*k$ that is equivalent to $*n + *m$. As an example, let $n = 13 = 8 + 4 + 1$ and $m = 11 = 8 + 2 + 1$. Then $*n + *m \equiv *8 + *4 + *1 + *8 + *2 + *1 \equiv *4 + *2 \equiv *6$, which we can also write as $13 \oplus 11 = 6$. In particular, $*13 + *11 + *6$ is a losing game, which would be very laborious to show without the theorem.

One consequence of Theorem 1.10 is that the Nim sum of two numbers never exceeds their ordinary sum. Moreover, if both numbers are less than some power of two, then so is their Nim sum.

Lemma 1.11 *Let $0 \leq p, q < 2^a$. Then $*p + *q \equiv *r$ where $0 \leq r < 2^a$, that is, $r = p \oplus q < 2^a$.*

Proof. Both p and q are sums of distinct powers of two, all smaller than 2^a . By theorem 1.10, r is also a sum of such powers of two, where those that appear in both p and q cancel out, so that $r < 2^a$. \square

The following proof may be best understood by considering it along with an example, say $n = 7$.

Proof of theorem 1.10. We proceed by induction. Consider some n , and assume that the theorem holds for all smaller n . Let $n = 2^a + q$ where $q = 2^b + 2^c + \cdots$. If $q = 0$, the claim holds trivially (n is just a single power of two), so let $q > 0$. We have $q < 2^a$. By inductive assumption, $*q \equiv *(2^b) + *(2^c) + \cdots$, so all we have to prove is that $*n \equiv *(2^a) + *q$ in order to show (1.7). We show this using Lemma 1.8, that is, by showing that the options of the games $*n$ and $*(2^a) + *q$ are all equivalent. The options of $*n$ are $*0, *1, *2, \dots, *(n-1)$.

The options of $*(2^a) + *q$ are of two kinds, depending on whether the player moves in the Nim heap $*(2^a)$ or $*q$. The first kind of options are given by

$$\begin{aligned} *0 + *q &\equiv *r_0 \\ *1 + *q &\equiv *r_1 \\ &\vdots \\ *(2^a - 1) + *q &\equiv *r_{2^a-1} \end{aligned} \tag{1.8}$$

where the equivalence of $*i + *q$ with some Nim heap $*r_i$, for $0 \leq i < 2^a$, holds by inductive assumption. Moreover, by Lemma 1.11 (which is a consequence of Theorem 1.10 which can be used by inductive assumption), both i and q are less than 2^a so that also $r_i < 2^a$. On the right-hand side in (1.8), there are 2^a many Nim heaps $*r_i$ for $0 \leq i < 2^a$. We claim they are all different, so that these options form exactly the set $\{*0, *1, *2, \dots, *(2^a - 1)\}$. Namely, by adding the game $*q$ to the heap $*r_i$, (1.5) implies $*r_i + *q \equiv *i + *q + *q \equiv *i$, so that $*r_i \equiv *r_j$ implies $*r_i + *q \equiv *r_j + *q$, that is, $*i \equiv *j$ and hence $i = j$ by Lemma 1.9, for $0 \leq i, j < 2^a$.

The second kind of options of $*(2^a) + *q$ are of the form

$$\begin{aligned} *(2^a) + *0 &\equiv *(2^a + 0) \\ *(2^a) + *1 &\equiv *(2^a + 1) \\ &\vdots \\ *(2^a) + *(q - 1) &\equiv *(2^a + q - 1), \end{aligned}$$

where the heap sizes on the right-hand sides are given again by inductive assumption. These heaps form the set $\{*(2^a), *(2^a + 1), \dots, *(n - 1)\}$. Together with the first kind of options, they are exactly the options of $*n$. This shows that the options of $*n$ and of the game sum $*(2^a) + *q$ are indeed equivalent, which completes the proof. \square

The Nim sum of any set of numbers can be obtained by writing each number as a sum of distinct powers of two and then cancelling repetitions in pairs. For example,

$$6 \oplus 4 = (4 + 2) \oplus 4 = 2, \quad \text{or} \quad 11 \oplus 16 \oplus 18 = (8 + 2 + 1) \oplus 16 \oplus (16 + 2) = 8 + 1 = 9.$$

This is usually described as “write the numbers in binary and add without carrying”, which comes to the same thing. In the following tables, the top row shows the powers of two needed in the binary representation for the numbers beneath; the bottom row gives the resulting Nim sum.

	4	2	1		16	8	4	2	1
6 =	1	1	0		11 =	1	0	1	1
4 =	1	0	0		16 =	1	0	0	0
2 =	0	1	0		18 =	1	0	0	1
					9 =	0	1	0	0

However, using only the powers of two that are used and cancelling repetitions is easier to do in your head, and is less prone to error.

How does Theorem 1.10 translate into playing Nim? When the Nim sum of the heap sizes is zero, then the player is in a losing position. (Such Nim positions are sometimes called *balanced* positions.) All moves will lead to a winning position, and in practice the best advice may only be not to move to a winning position that is too obvious, like one where two heap sizes are equal, in the hope that the opponent makes a mistake.

If the Nim sum of the heap sizes is not zero, it is some sum of powers of two, say $s = 2^a + 2^b + 2^c + \dots$, like for example $11 \oplus 16 \oplus 18 = 9 = 8 + 1$ above. The winning move is then obtained as follows:

1. Identify a heap of size n which uses 2^a , the largest power of 2 in the Nim sum; at least one such heap must exist. In the example, that heap has size 11 (so it is not always the largest heap).
2. Compute $n \oplus s$. In that Nim sum, the power 2^a appears in both n and s , and it cancels out, so the result is some number m that is smaller than n . In the example, $m = 11 \oplus 9 = (8 + 2 + 1) \oplus (8 + 1) = 2$.
3. Reduce the heap of size n to size m , in the example from size 11 to size 2. The resulting heap $*m$ is equivalent to $*n + *s$, so when it replaces $*n$ in the original sum, $*s$ is added and cancels with $*s$, and the result is equivalent to $*0$, a losing position.

On paper, the binary representation may be easier to use. In step 2 above, computing $n \oplus s$ amounts to “flipping the bits” in the binary representation of the heap n whenever the corresponding bit in the binary representation of the sum s is one. In this way, a player in a winning position moves always from an “unbalanced” position (with nonzero Nim sum) to a balanced position (nim sum zero), which is losing because any move will create again an unbalanced position. This is the way Nim is usually explained. The method was discovered by Bouton who published it in 1902.

\Rightarrow You are now in a position to answer all of Exercise 1.1 on page 30.

So far, it is not fully clear why powers of two appear in the computation of Nim sums. One reason is provided by the proof of Theorem 1.10. The options of moving from $*(2^a) + *q$ in (1.8) neatly produce exactly the numbers $0, 1, \dots, 2^a - 1$, which would not work when replacing 2^a with something else.

As a second reason, the copycat principle $G + G \equiv *0$ shows that the impartial games form a group where every element is its own inverse. There is essentially only one mathematical structure that has these particular properties, namely the addition of binary vectors. In each component, such vectors are separately added modulo two, where $1 \oplus 1 = 0$. Here, the binary vectors translate into binary numbers for the sizes of the Nim heaps. The “addition without carry” of binary vectors defines exactly the winning strategy in Nim, as stated in Theorem 1.10. However, the proof of this theorem is stated directly and without any recourse to abstract algebra and groups.

A third reason uses the construction of equivalent Nim heaps for any impartial game, in particular a sum of two Nim heaps, which we explain next; see also Figure 1.2 below.

1.11 Poker Nim and the mex rule

Poker Nim is played with heaps of Poker chips. Just as in ordinary Nim, either player may reduce the size of any heap by removing some of the chips. But alternatively, a player may also *increase* the size of some heap by adding to it some of the chips he acquired in earlier moves. These two kinds of moves are the only ones allowed.

Let's suppose that there are three heaps, of sizes 3, 4, 5, and that the game has been going on for some time, so that both players have accumulated substantial reserves of chips. It's player I's turn, who moves to 1, 4, 5 because that is a good move in ordinary Nim. But now player II adds 50 chips to the heap of size 4, creating position 1, 54, 5, which seems complicated.

What should player I do? After a moment's thought, he just removes the 50 chips player II has just added to the heap, reverting to the previous position. Player II may keep adding chips, but will eventually run out of them, no matter how many she acquires in between, and then player I can proceed as in ordinary Nim.

So a player who can win a position in ordinary Nim can still win in Poker Nim. He replies to the opponent's reducing moves just as he would in ordinary Nim, and reverses the effect of any increasing move by using a reducing move to restore the heap to the same size again. Strictly speaking, the ending condition (see condition 6 in Section 1.6) is violated in Poker Nim because in theory the game could go on forever. However, a player in a winning position wants to end the game with his victory, and never has to put back any chips; then the losing player will eventually run out of chips that she can add to a heap, so that the game terminates.

Consider now an impartial game where the options of player I are games that are equivalent to the Nim heaps $*0, *1, *2, *5, *9$. This can be regarded as a rather peculiar Nim heap of size 3 which can be reduced to any of the sizes 0, 1, 2, but which can also be increased to size 5 or 9. The Poker Nim argument shows that this extra freedom is in fact of no use whatsoever.

The *mex rule* says that if the options of a game G are equivalent to Nim heaps with sizes from a set S (like $S = \{0, 1, 2, 5, 9\}$ above), then G is equivalent to a Nim heap of size m , where m is the *smallest non-negative integer not contained in S* . This number m is written $\text{mex}(S)$, where *mex* stands for “minimum excludant”. That is,

$$m = \text{mex}(S) = \min\{k \geq 0 \mid k \notin S\}. \quad (1.9)$$

For example, $\text{mex}(\{0, 1, 2, 3, 5, 6\}) = 4$, $\text{mex}(\{1, 2, 3, 4, 5\}) = 0$, and $\text{mex}(\emptyset) = 0$.

\Rightarrow Which game has the empty set \emptyset as its set of options?

Theorem 1.12 (The mex rule) *Let the impartial game G have the set of options that are equivalent to $\{*s \mid s \in S\}$ for some set S of non-negative integers (assuming S is not the set of all non-negative integers, for example if S is finite). Then $G \equiv *(\text{mex}(S))$.*

Proof. Let $m = \text{mex}(S)$. We show $G + *m \equiv *0$, which proves the theorem by Lemma 1.7. If player I moves from $G + *m$ to $G + *k$ for some $k < m$, then $k \in S$ and there is an option

K of G so that $K \equiv *k$ by assumption, so player II can counter by moving from $G + *k$ to the losing position $K + *k$. Otherwise, player I may move to $K + *m$, where K is some option K which is equivalent to $*k$ for some $k \in S$. If $k < m$, then player II counters by moving to $K + *k$. If $k > m$, then player II counters by moving to $M + *m$ where M is the option of K that is equivalent to $*m$ (because $K \equiv *k$). The case $k = m$ is excluded by the definition of $m = \text{mex}(S)$. This shows $G + *m$ is a losing game. \square

A special case of the preceding theorem is that $\text{mex}(S) = 0$, which means that all options of G are equivalent to positive Nim heaps, so they are all winning positions, or that G has no options at all. Then G is a losing game, and indeed $G \equiv *0$.

Corollary 1.13 *Any impartial game G is equivalent to some Nim heap $*n$.*

Proof. We can assume by induction that this holds for all games that are simpler than G , in particular the options of G . They are equivalent to Nim heaps whose sizes form the set S (which we assume is not the set of all non-negative integers). Theorem 1.12 then shows $G \equiv *m$ for $m = \text{mex}(S)$. \square

\Rightarrow Do Exercise 1.7 on page 33, which provides an excellent way to understand the mex rule.

1.12 Finding Nim values

By Corollary 1.13, any impartial game can be played like Nim, provided the equivalent Nim heaps of the positions of the game are known. This forms the basis of the *Sprague–Grundy* theory of impartial games, named after the independent discoveries of this principle by R. P. Sprague in 1936 and P. M. Grundy in 1939. Any sum of such games is then evaluated by taking the Nim sum of the sizes of the corresponding Nim heaps.

The Nim values of the positions can be evaluated by the mex rule (Theorem 1.12). This is illustrated in the “rook-move” game in Figure 1.1. Place a rook on a Chess board of given arbitrary size. In one move, the rook is moved either horizontally to the left or vertically upwards, for any number of squares (at least one) as long as it stays on the board. The first player who can no longer move loses, when the rook is on the top left square of the board. We number the rows and columns of the board by $0, 1, 2, \dots$ starting from the top left.

Figure 1.2 gives the Nim values for the positions of the rook on the Chess board. The top left square is equivalent to $*0$ because the rook can no longer move. The square below that allows only to reach the square with $*0$ on it, so it is equivalent to $*1$ because $\text{mex}\{0\} = 1$. The square below gets $*2$ because its options are equivalent to $*0$ and $*1$. From any square in the leftmost column in Figure 1.2, the rook can only move upwards, so any such square in row i corresponds obviously to a Nim heap $*i$. Similarly, the topmost row has entry $*j$ in column j .

In general, a position on the board is evaluated knowing all Nim values for the squares to the left and the top of it, which are the options of that position. As an example, consider




Figure 1.1 Rook move game, where the player may move the rook on the Chess board in the direction of the arrows.

the square in row 3 and column 2. To the left of that square, the entries $*3$ and $*2$ are found, and to the top the entries $*2$, $*3$, $*0$. So the square itself is equivalent to $*1$ because $\text{mex}(\{0, 2, 3\}) = 1$. The square in row 3 and column 5, where the rook is placed in Figure 1.1, gets entry $*6$.

	0	1	2	3	4	5	6	7	8	9	10
0	$*0$	$*1$	$*2$	$*3$	$*4$	$*5$	$*6$	$*7$	$*8$	$*9$	$*10$
1	$*1$	$*0$	$*3$	$*2$	$*5$	$*4$	$*7$	$*6$	$*9$	$*8$	$*11$
2	$*2$	$*3$	$*0$	$*1$	$*6$	$*7$	$*4$	$*5$	$*10$	$*11$	$*8$
3	$*3$	$*2$	$*1$	$*0$	$*7$	$*6$	$*5$	$*4$	$*11$	$*10$	$*9$
4	$*4$	$*5$	$*6$	$*7$	$*0$	$*1$	$*2$	$*3$	$*12$	$*13$	$*14$
5	$*5$	$*4$	$*7$	$*6$	$*1$	$*0$	$*3$	$*2$	$*13$	$*12$	$*15$
6	$*6$	$*7$	$*4$	$*5$	$*2$	$*3$	$*0$	$*1$	$*14$	$*15$	$*12$
7	$*7$	$*6$	$*5$	$*4$	$*3$	$*2$	$*1$	$*0$	$*15$	$*14$	$*13$
8	$*8$	$*9$	$*10$	$*11$	$*12$	$*13$	$*14$	$*15$	$*0$	$*1$	$*2$
9	$*9$	$*8$	$*11$	$*10$	$*13$	$*12$	$*15$	$*14$	$*1$	$*0$	$*3$
10	$*10$	$*11$	$*8$	$*9$	$*14$	$*15$	$*12$	$*13$	$*2$	$*3$	$*0$

Figure 1.2 Equivalent Nim heaps $*n$ for positions of the rook move game.

The astute reader will have noticed that the rook move game is just Nim with two heaps. A rook positioned in row i and column j can either diminish i by moving left, or

j by moving up. So this position is the sum of Nim heaps $*i + *j$. It is a losing position if and only if $i = j$, where the rook is on the diagonal leading to the top left square. Therefore, Figure 1.2 represents the computation of Nim heaps equivalent to $*i + *j$, or, by omitting the stars, the Nim sums $i \oplus j$, for all $i, j \geq 0$.

The Nim addition table Figure 1.2 is computed by the mex rule, and does not require Theorem 1.10. Given this Nim addition table, one can conjecture (1.7). You may find it useful to go back to the proof of Theorem 1.10 using Figure 1.2 and check the options for a position of the form $*(2^a) + *q$ for $q < 2^a$, as a square in row 2^a and column q .

Another impartial game is shown in Figure 1.3 where the rook is replaced by a queen, which may also move diagonally. The squares on the main diagonal are therefore no longer losing positions. This game can also be played with two heaps of chips where in one move, the player may either take chips from one heap as in Nim, or reduce *both* heaps by the *same* number of chips (so this is no longer a sum of two games!). In order to illustrate that we are not just interested in the winning and losing squares, we add to this game a Nim heap of size 4.




Figure 1.3 Sum of a queen move game and a Nim heap. The player may *either* move the queen in the direction of the arrows *or* take some of the 4 chips from the heap.

Figure 1.4 shows the equivalent Nim heaps for the positions of the queen move game, determined by the mex rule. The square in row 3 and column 4 occupied by the queen in Figure 1.3 has entry $*2$. So a winning move is to remove 2 chips from the Nim heap to turn it into the heap $*2$, creating the losing position $*2 + *2$.

	0	1	2	3	4	5	6
0	$*0$	$*1$	$*2$	$*3$	$*4$		
1	$*1$	$*2$	$*0$	$*4$	$*5$		
2	$*2$	$*0$	$*1$	$*5$	$*3$		
3	$*3$	$*4$	$*5$	$*6$	$*2$		

Figure 1.4 Equivalent Nim heaps for positions of the queen move game.

This concludes our introduction to combinatorial games. Further examples will be given in the exercises.

⇒ Do the remaining exercises 1.4–1.11, starting on page 31, which show how to use the mex rule and what you learned about Nim and combinatorial games.


1.13 Exercises for Chapter 1

In this chapter, which is more abstract than the others, the exercises are particularly important. Exercise 1.1 is a standard question on Nim, where part (a) can be answered even without the theory. Exercise 1.2 is an example of an impartial game, which can also be answered without much theory. Exercise 1.3 is difficult – beware not to rush into any quick and false application of Nim values here; part (c) of this exercise is particularly challenging. Exercise 1.4 tests your understanding of the queen move game. For Exercise 1.5, remember the concept of a sum of games, which applies here naturally. In Exercise 1.6, try to see how Nim heaps are hidden in the game. Exercise 1.7 is very instructive for understanding the mex rule. In Exercise 1.8, it is essential that you understand Nim values. It takes some work to investigate all the options in the game. Exercise 1.9 is another familiar game where you have to find out Nim values. In Exercise 1.10, a new game is defined that you can analyse with the mex rule. Exercise 1.11 is an impartial game that is rather different from the previous games. In the challenging part (c) of that exercise, you should first formulate a conjecture and then prove it precisely.

Exercise 1.1 Consider the game Nim with heaps of chips. The players alternately remove some chips from one of the heaps. The player to remove the last chip wins.

- (a) For all positions with three heaps, where one of the heaps has only *one* chip, describe exactly the *losing* positions. Justify your answer, for example with a proof by induction, or by theorems on Nim.
[Hint: Start with the easy cases to find the pattern.]
- (b) Determine all initial winning moves for Nim with three heaps of size 6, 10 and 15, using the theorem on Nim where the heap sizes are represented as sums of powers of two.

Exercise 1.2 The game *Dominos* is played on a board of $m \times n$ squares, where players alternately place a domino on the board which covers two adjacent squares that are free (not yet occupied by a domino), vertically or horizontally. The first player who cannot place a domino any more loses. Example play for a 2×3 board:



- (a) Who will win in 3×3 Dominos?


[Hint: Use the symmetry of the game to investigate possible moves, and remember that it suffices to find one winning strategy.]

- (b) Who will win in $m \times n$ Dominos when both m and n are even?
 (c) Who will win in $m \times n$ Dominos when m is odd and n is even?

Justify your answers.

Note (not a question): Because of the known answers from (b) and (c), this game is more interesting for “real play” on an $m \times n$ board where both m and n are odd. Play it with your friends on a 5×5 board, for example. The situation often decomposes into independent parts, like contiguous fields of 2, 3, 4, 5, 6 squares, that have a known winner, which may help you analyse the situation.

Exercise 1.3 Consider the following game chomp: A rectangular array of $m \times n$ dots is given, in m rows and n columns, like 3×4 in the next picture on the left. A dot in row i and column j is named (i, j) . A move consists in picking a dot (i, j) and removing it and *all other dots to the right and below it*, which means removing all dots (i', j') with $i' \geq i$ and $j' \geq j$, as shown for $(i, j) = (2, 3)$ in the middle picture, resulting in the picture on the right:



Player I is the first player to move, players alternate, and the last player who removes a dot *loses*.

An alternative way is to think of these dots as (real) cookies: a move is to eat a cookie and all those to the right and below it, but the top left cookie is poisoned. See also <http://www.stolaf.edu/people/molnar/games/chomp/>

- (a) Assuming optimal play, determine the winning player and a winning move for chomp of size 2×2 , size 2×3 , size $2 \times n$, and size $m \times m$, where $m \geq 3$. Justify your answers.
 (b) In the way described here, chomp is a misère game where the last player to make a move loses. Suppose we want to play the same game so that the normal play convention applies, where the last player to move wins. (This would be a boring game with the board as given, by simply taking the top left dot $(1, 1)$.) Explain how this can be done by removing one dot from the initial array of dots.
 (c) Show that when chomp is played for a game of any size $m \times n$, player I can always win.

[Hint: You only have to show that a winning move exists, but you do not have to describe that winning move.]


Exercise 1.4

- (a) Complete the entries of equivalent Nim heaps for the queen move game in columns 5 and 6, rows 0 to 3, in the table in Figure 1.4.
 (b) Describe *all* winning moves in the game-sum of the queen move game and the Nim heap in Figure 1.3.

Exercise 1.5 Consider the game Dominos from Exercise 1.2, played on a $1 \times n$ board for $n \geq 2$. Let D_n be the Nim value of that game, so that the starting position of the $1 \times n$ board is equivalent to a Nim heap of size D_n . For example, $D_2 = 1$ because the 1×2 board is equivalent to $*1$.


- (a) How is D_n computed from smaller values D_k for $k < n$?
 [Hint: Use sums of games and the mex rule. The notation for Nim sums is \oplus , where $a \oplus b = c$ if and only if $*a + *b \equiv *c$.]
- (b) Give the values of D_n up to $n = 10$ (or more, if you are ambitious – higher values come at $n = 16$, and at some point they even repeat, but before you detect that you will probably have run out of patience). For which values of n , where $1 \leq n \leq 10$, is Dominos on a $1 \times n$ board a losing game?

Exercise 1.6 Consider the following game on a rectangular board where a white and a black counter are placed in each row, like in this example:



Player I is white and starts, and player II is black. Players take turns. In a move, a player moves a counter of his colour to any other square within its row, but may not jump over the other counter. For example, in (*) above, in row 8 white may move from e8 to any of the squares c8, d8, f8, g8, or h8. The player who can no longer move loses.

- (a) Who will win in the following position?




- (b) Show that white can win in position (*) above. Give at least two winning moves from that position.

Justify your answers.

[Hint: Compare this with another game that is not impartial and that also violates the



ending condition, but that nevertheless is close to Nim and ends in finite time when played well.]

Exercise 1.7 Consider the following network (in technical terms, a directed graph or “di-graph”). Each circle, here marked with one of the letters A–P, represents a *node* of the network. Some of these nodes (here A, F, G, H, and K) have counters on them, which are allowed to share a node, like the two counters on H. In a move, one of the counters is moved to a neighbouring node in the direction of the arrow as indicated, for example from F to I (but not from F to C, nor directly from F to D, say). Players alternate, and the last player no longer able to move loses.



- Explain why this game fulfils the ending condition.
- Who is winning in the above position? If it is player I (the first player to move), describe *all* possible winning moves. Justify your answer.
- How does the answer to (b) change when the arrow from J to K is reversed so that it points from K to J instead?

Exercise 1.8 Consider the game chomp from Exercise 1.3 of size 2×4 , added to a Nim heap of size 4.



What are the winning moves of the starting player I, if any?

[Hint: Represent chomp as a game in the normal play convention (see Exercise 1.3(b), by changing the dot pattern), so that the losing player is not the player who takes the “poisoned cookie”, but the player who can no longer move. This will simplify finding the various Nim values.]

Exercise 1.9 In $m \times n$ *Dominos* (see Exercise 1.2), a rectangular board of $m \times n$ squares is given. The two players alternately place a domino either horizontally or vertically on two unoccupied adjacent squares, which then become occupied. The last player to be able to place a domino wins.


- Find the Nim-value (size of the equivalent Nim heap) of 2×3 Dominos.
- Find all winning moves, if any, for the *game-sum* of a 2×3 Domino game and a 1×4 Domino game.

Exercise 1.10 Consider the following variant of Nim called split-nim, which is played with heaps of chips as in Nim. Like in Nim, a player can remove some chips from one of the heaps, or else *split* a heap into two (not necessarily equal) new heaps. For example, a heap of size 4 can be reduced to size 3, 2, 1, or 0 as in ordinary Nim, or be split into two heaps of sizes 1 and 3, respectively, or into two heaps of sizes 2 and 2. As usual, the last player to be able to move wins.


- Find the Nim-values (size of equivalent Nim heaps) of the single split-nim heaps of size 1, 2, 3, and 4, respectively.
- Find all winning moves, if any, when split-nim is played with three heaps of sizes 1, 2, 3.
- Find all winning moves, if any, when split-nim is played with three heaps of sizes 1, 2, 4.
- Determine if the following statement is true or false: “Two heaps in split-nim are equivalent if and only if they have equal size.” Justify your answer.

Exercise 1.11 The game *Hackenbush* is played on a figure consisting of dots connected with lines (called *edges*) that are connected to the ground (the dotted line in the pictures below). A move is to remove (“chop off”) an edge, and with it all the edges that are then no longer connected to the ground. For example, in the leftmost figure below, one can remove any edge in one of the three stalks. Removing the second edge from the stalk consisting of three edges takes the topmost edge with it, leaving only a single edge. As usual, players alternate and the last player able to move wins.


- Compute the Nim values for the following three Hackenbush figures (using what you know about Nim, and the mex rule):



- Compute the Nim values for the following four Hackenbush figures:



- (c) Based on the observation in (b), give a rule how to compute the Nim value of a “tree” which is constructed by putting several “branches” with known Nim values on top of a “stem” of size n , for example $n = 2$ in the picture below, where three branches of height 3, 2, and 4 are put on top of the “stem”. Prove that rule (you may find it advantageous to glue the branches together at the bottom, as in (b)). Use this to find the Nim value of the rightmost picture.



Chapter 2

Games as trees and in strategic form

This chapter introduces the main concepts of non-cooperative game theory: game trees and games in strategic form and the ways to analyse them.

We occasionally refer to concepts discussed in Chapter 1, such as “game states”, to clarify the connection, but Chapter 1 is not a requirement to study this chapter.

2.1 Learning objectives

After studying this chapter, you should be able to:

- interpret game trees and games in strategic form;
- explain the concepts of move in a game tree, strategy (and how it differs from move), strategy profile, backward induction, symmetric games, dominance and weak dominance, dominance solvable, Nash equilibrium, reduced strategies and reduced strategic form, subgame perfect Nash equilibrium, and commitment games;
- apply these concepts to specific games.

2.2 Further reading

The presented concepts are standard in game theory. They can be found, for example, in the following book:

- Osborne, Martin J., and Ariel Rubinstein *A Course in Game Theory*. (MIT Press, 1994) [ISBN 0262650401].

Osborne and Rubinstein treat game theory as it is used in economics. Rubinstein is also a pioneer of bargaining theory. He invented the alternating-offers model treated in Chapter 5. On the other hand, the book uses some non-standard descriptions. For example, Osborne and Rubinstein define games of perfect information via “histories” and not game trees; we prefer the latter because they are less abstract. Hence, keep in mind

that the terminology used by Osborne and Rubinstein may not be standard and may differ from ours.

Another possible reference is:

- Gibbons, Robert *A Primer in Game Theory* [in the United States sold under the title *Game Theory for Applied Economists*]. (Prentice Hall / Harvester Wheatsheaf, 1992) [ISBN 0745011594].

In particular, this book looks at the commitment games of Section 2.14 from the economic perspective of “Stackelberg leadership”.

2.3 Introduction

In this chapter, we introduce several main concepts of non-cooperative game theory: *game trees* (with perfect information), which describe explicitly how a game evolves over time, and *strategies*, which describe a player’s possible “plans of action”. A game can be described in terms of strategies alone, which defines a game in *strategic form*.

Game trees can be solved by *backward induction*, where one finds optimal moves for each player, given that all future moves have already been determined. The central concept for games in strategic form is the *Nash equilibrium*, where each player chooses a strategy that is optimal given what the other players are doing. We will show that backward induction always produces a Nash equilibrium, also called “subgame perfect Nash equilibrium” or SPNE.

The main difference between game trees and games in strategic form is that in a game tree, players act sequentially, being aware of the previous moves of the other players. In contrast, players in a game in strategic form move *simultaneously*. The difference between these two descriptions becomes striking when *changing* a two-player game in strategic form to a game with *commitment*, described by a game tree where one player moves first and the other second, but which otherwise has the same payoffs.

Every game tree can be converted to strategic form, which, however, is often much larger than the tree. A general game in strategic form can only be represented by a tree by modelling “imperfect information” where a player is not aware of another player’s action. Game trees with imperfect information are treated in Chapter 4.

2.4 Definition of game trees

Figure 2.1 shows an example of a game tree. We always draw trees downwards, with the *root* at the top. (Conventions on drawing game trees vary. Sometimes trees are drawn from the bottom upwards, sometimes from left to right, and sometimes from the center with edges in any direction.)

The *nodes* of the tree denote game states. (In a combinatorial game, game states are called “positions”.) Nodes are connected by lines, called *edges*. An edge from a node u




Figure 2.1 Example of a game tree. The square node indicates a chance move. At a leaf of the tree, the top payoff is to player I, the bottom payoff to player II.

to a successor node v (where v is drawn below u) indicates a possible *move* in the game. This may be a move of a “personal” player, for example move X in Figure 2.1 of player I. Then u is also called a *decision node*. Alternatively, u is a *chance node*. A chance node is drawn here as a *square*, like the node u that follows move b of player II in Figure 2.1. The next node v is then determined by a random choice according to the probability associated with the edge leading from u to v . In Figure 2.1, these probabilities are $\frac{1}{3}$ for the left move and $\frac{2}{3}$ for the right move.

Nodes without successors in the tree are called terminal nodes or *leaves*. At such a node, every player gets a *payoff*, which is a real number (in our examples often an integer). In Figure 2.1, leaves are not explicitly drawn, but the payoffs given instead, with the top payoff to player I and the bottom payoff to player II.

The game tree, with its decision nodes, moves, chance probabilities and payoffs, is known to the players, and defines the game completely. The game is played by starting at the root. At a decision node, the respective player chooses a move which determines the next node. At a chance node, the move is made randomly according to the given probabilities. Play ends when a leaf is reached, where all players receive their payoffs.

Players are interested in maximising their own payoff. If the outcome of the game is random, then the players are assumed to be interested in maximising their *expected payoffs*. In Figure 2.1, the expected payoff to player I after the chance move is $\frac{1}{3} \times 0 + \frac{2}{3} \times 3 = 2$, and to player II it is $\frac{1}{3} \times 15 + \frac{2}{3} \times 0 = 5$. In this game, the chance node could therefore be replaced by a leaf with payoff 2 for player I and payoff 5 for player II. These expected payoffs reflect the players’ preferences for the random outcomes, including their attitude to the risk involved when facing such uncertain outcomes. This “risk-neutrality” of payoffs does not necessarily hold for actual monetary payoffs. For example, if the payoffs to player II in Figure 2.1 were millions of dollars, then in reality that player would

probably prefer receiving 5 million dollars for certain to a lottery that gave her 15 million dollars with probability $1/3$ and otherwise nothing. However, payoffs can be adjusted to reflect the player's attitude to risk, as well as representing the "utility" of an outcome like money so that one can take just the expectation. In the example, player II may consider a lottery that gives her 15 million dollars with probability $1/3$ and otherwise nothing only as valuable as getting 2 million dollars for certain. In that case, "15 million dollars" may be represented by a payoff of 6 so that the said lottery has an expected payoff 2, assuming that "2 million dollars" are represented by a payoff of 2 and "getting nothing" by a payoff of 0. This is discussed in greater detail in Section 3.4.




Figure 2.2 The left picture (a) is not a tree because the leaf where both players receive payoff 1 is reachable by two different paths. The correct tree representation is shown on the right in (b).

For the game tree, it does not matter how it is drawn, but only its "combinatorial" structure, that is, how its nodes are connected by edges. A tree can be considered as a special "directed graph", given by a set of nodes, and a set of edges, which are pairs of nodes (u, v) . A *path* in such a graph from node u to node v is a sequence of nodes $u_0 u_1 \dots u_k$ so that $u_0 = u$, $u_k = v$, and so that (u_i, u_{i+1}) is an edge for $0 \leq i < k$. The graph is a tree if it has a distinguished node, the root, so that to each node of the graph there is a unique path from the root.

Figure 2.2 demonstrates the tree property that every node is reached by a unique path from the root. This fails to hold in the left picture (a) for the middle leaf. Such a structure may make sense in a game. For example, the figure could represent two people in a pub where first player II chooses to pay for the drinks (move p) or to accept (move a) that the first round is paid by player I. In the second round, player I may then decide to either accept not to pay (move A) or to pay (P). Then the players may only care how often, but not when, they have paid a round, with the middle payoff pair $(1, 1)$ for two possible "histories" of moves. The tree property prohibits game states with more than one history, because the history is represented by the unique path from the root. Even if certain differences in the game history do not matter, these game states are distinguished, mostly as a matter of mathematical convenience. The correct representation of the above game is shown in the right picture, Figure 2.2(b).

Game trees are also called *extensive games with perfect information*. Perfect information means that a player always knows the game state and therefore the complete history

of the game up to then. Game trees can be enhanced with an additional structure that represents “imperfect information”, which is the topic of Chapter 4.

Note how game trees differ from the combinatorial games studied in Chapter 1:

- (a) A combinatorial game can be described very compactly, in particular when it is given as a sum of games. For general game trees, such sums are typically not considered.
- (b) The “rules” in a game tree are much more flexible: more than two players and chance moves are allowed, players do not have to alternate, and payoffs do not just stand for “win” or “lose”.
- (c) The flexibility of game trees comes at a price, though: The game description is much longer than a combinatorial game. For example, a simple instance of Nim may require a huge game tree. Regularities like the mex rule do not apply to general game trees.

2.5 Backward induction

Which moves should the players choose in a game tree? “Optimal” play should maximise a player’s payoff. This can be decided irrespective of other players’ actions when the player is the *last* player to move. In Figure 2.2(b), player I maximises his payoff by move *A* at both his decision nodes, because at the left node he receives 1 rather than 0 with that move, and at the right node payoff 2 rather than 1. Going backwards in time, player II has to make her move *a* or *p* at the root of the game tree, where she will receive either 1 or 0, assuming the described future behaviour of player I. Consequently, she will choose *a*.

This process is called *backward induction*: Starting with the decision nodes closest to the leaves, a player’s move is chosen which maximises that player’s payoff at the node. In general, a move is chosen in this way for each decision node provided all subsequent moves have already been decided. Eventually, this will determine a move for every decision node, and hence for the entire game. Backward induction is also known as *Zermelo’s algorithm*. (This is attributed to an article by Zermelo (1913) on Chess. Later, people decided that Zermelo proved something different, in fact a more complicated property, so that Zermelo’s algorithm is sometimes called “Kuhn’s algorithm”, according to a paper by Kuhn (1953), which we cite on page 106.)

The move selected by backward induction is not necessarily unique, if there is more than one move giving maximal payoff to the player. In the game in Figure 2.1, backward induction chooses either move *b* or move *c* for player II, both of which give her payoff 5 (which is an expected payoff for move *b*) that exceeds her payoff 4 for move *a*. At the right-most node, player I chooses *Q*. This determines the preceding move *d* by player II which gives her the higher payoff 3 as opposed to 2 (via move *Q*). In turn, this means that player I, when choosing between *X*, *Y*, or *Z* at the root of the game tree, will get payoff 2 for *Y* and payoff 1 for *Z*; the payoff when he chooses *X* depends on the choice of player II: if that is *b*, then player I gets 2, and can choose either *X* or *Y*, both of which give him maximal payoff 2. If player II chooses *c*, however, then the payoff to player I

is 4 when choosing X , so this is the unique optimal choice. To summarise, the possible combinations of moves that can arise in Figure 2.1 by backward induction are, simply listing the moves for each player: (XQ, bd) , (XQ, cd) , and (YQ, bd) . Note that Q and d are always chosen, but that Y can only be chosen in combination with the move b by player II.

The moves determined by backward induction are therefore, in general, not unique, and possibly interdependent.

Backward induction gives a *unique* recommendation to each player if there is always only one move that gives maximal payoff. This applies to *generic* games. A generic game is a game where the payoff parameters are real numbers that are in no special dependency of each other (like two payoffs being equal). In particular, it should be allowed to replace them with values nearby. For example, the payoffs may be given in some practical scenario which has some “noise” that effects the precise value of each payoff. Then two such payoffs are equal with probability zero, and so this case can be disregarded. In generic games, the optimal move is always unique, so that backward induction gives a unique result.

2.6 Strategies and strategy profiles

Definition 2.1 In a game tree, a *strategy* of a player specifies a move for every decision node of that player. A *strategy profile* is a tuple of strategies, with one strategy for each player of the game.

If the game has only two players, a strategy profile is therefore a pair of strategies, with one strategy for player I and one strategy for player II.

In the game tree in Figure 2.1, the possible strategies for player I are XP, XQ, YP, YQ, ZP, ZQ . The strategies for player II are ad, ae, bd, be, cd, ce . For simplicity of notation, we have thereby specified a strategy simply as a list of moves, one for each decision node of the player. When specifying a strategy in that way, this must be done with respect to a fixed order of the decision nodes in the tree, in order to identify each move uniquely. This matters when a move name appears more than once, for example in Figure 2.2(b). In that tree, the strategies of player I are AA, AP, PA, PP , with the understanding that the first move refers to the left decision node and the second move to the right decision node of player I.

Backward induction defines a move for every decision node of the game tree, and therefore for every decision node of each player, which in turn gives a strategy for each player. The result of backward induction is therefore a strategy profile.

⇒ Exercise 2.1 on page 62 studies a game tree which is not unlike the game in Figure 2.1. You can already answer part (a) of this exercise, and apply backward induction, which answers (e).

2.7 Games in strategic form

Assume a game tree is given. Consider a strategy profile, and assume that players move according to their strategies. If there are no chance moves, their play leads to a unique leaf. If there are chance moves, a strategy profile may lead to a probability distribution on the leaves of the game tree, with resulting *expected* payoffs. In general, any strategy profile defines an expected payoff to each player (which also applies to a payoff that is obtained deterministically, where the expectation is computed from a probability distribution that assigns probability one to a single leaf of the game tree).

Definition 2.2 The *strategic form* of a game is defined by specifying for each player the set of strategies, and the payoff to each player for each strategy profile.

For two players, the strategic form is conveniently represented by a table. The rows of the table represent the strategies of player I, and the columns the strategies of player II. A strategy profile is a strategy pair, that is, a row and a column, with a corresponding cell of the table that contains two payoffs, one for player I and the other for player II.

If m and n are positive integers, then an $m \times n$ game is a two-player game in strategic form with m strategies for player I (the rows of the table) and n strategies for player II (the columns of the table).




Figure 2.3 Extensive game (a) and its strategic form (b). In a cell of the strategic form, player I receives the bottom left payoff, and player II the top right payoff.

Figure 2.3 shows an extensive game in (a), and its strategic form in (b). In the strategic form, T and B are the strategies of player I given by the top and bottom row of the table, and l and r are the strategies of player II, corresponding to the left and right column of the table. The strategic form for the game in Figure 2.1 is shown in Figure 2.4.

A game in strategic form can also be given directly according to definition 2.2, without any game tree that it is derived from. Given the strategic form, the game is played as follows: Each player chooses a strategy, independently from and simultaneously with the other players, and then the players receive their payoffs as given for the resulting strategy profile.

		II					
		<i>ad</i>	<i>ae</i>	<i>bd</i>	<i>be</i>	<i>cd</i>	<i>ce</i>
I	<i>XP</i>	4 3	4 3	5 2	5 2	5 4	5 4
	<i>XQ</i>	4 3	4 3	5 2	5 2	5 4	5 4
	<i>YP</i>	3 2	3 2	3 2	3 2	3 2	3 2
	<i>YQ</i>	3 2	3 2	3 2	3 2	3 2	3 2
	<i>ZP</i>	3 1	5 0	3 1	5 0	3 1	5 0
	<i>ZQ</i>	3 1	2 4	3 1	2 4	3 1	2 4

Figure 2.4 Strategic form of the extensive game in Figure 2.1. The less redundant reduced strategic form is shown in Figure 2.13 below.

2.8 Symmetric games

Many game-theoretic concepts are based on the strategic form alone. In this section, we discuss the possible *symmetries* of such a game. We do this by presenting a number of standard games from the literature, mostly 2×2 games, which are the smallest games which are not just one-player decision problems.

		II	
		<i>c</i>	<i>d</i>
I	<i>C</i>	2 2	3 0
	<i>D</i>	0 3	1 1

		II	
		<i>c</i>	<i>d</i>
I	<i>C</i>	2 2	3 0
	<i>D</i>	0 3	1 1

Figure 2.5 Prisoner's dilemma game (a), its symmetry shown by reflection along the dotted diagonal line in (b).

Figure 2.5(a) shows the well-known “prisoner’s dilemma” game. Each player has two strategies, called *C* and *D* for player I and *c* and *d* for player II. (We use upper

case letters for player I and lower case letters for player II for easier identification of strategies and moves.) These letters stand for “cooperate” and “defect”, respectively. The story behind the name “prisoner’s dilemma” is that of two prisoners held suspect of a serious crime. There is no judicial evidence for this crime except if one of the prisoners testifies against the other. If one of them testifies, he will be rewarded with immunity from prosecution (payoff 3), whereas the other will serve a long prison sentence (payoff 0). If both testify, their punishment will be less severe (payoff 1 for each). However, if they both “cooperate” with each other by not testifying at all, they will only be imprisoned briefly for some minor charge that can be held against them (payoff 2 for each). The “defection” from that mutually beneficial outcome is to testify, which gives a higher payoff no matter what the other prisoner does (which makes “defect” a dominating strategy, as discussed in the next section). However, the resulting payoff is lower to both. This constitutes their “dilemma”.

The prisoner’s dilemma game is *symmetric* in the sense that if one exchanges player I with player II, and strategy *C* with *c*, and *D* with *d*, then the game is unchanged. This is demonstrated in Figure 2.5(b) with the dotted line that connects the top left of the table with the bottom right: the payoffs remain the same when reflected along that dotted line. In order to illustrate the symmetry in this visual way, payoffs have to be written in diagonally opposite corners in a cell of the table.

This representation of payoffs in different corners of cells of the table is due to Thomas Schelling, which he used, for example, in his book *The Strategy of Conflict* (1961). Schelling modestly (and not quite seriously) calls the staggered payoff matrices his most important contribution to game theory, despite that fact that this book was most influential in applications of game theory to social science. Schelling received the 2005 Nobel prize in economics for his contribution to the understanding of conflict and cooperation using game theory.

The payoff to player I has to appear in the bottom left corner, the payoff to player II at the top right. These are also the natural positions for the payoffs that leave no ambiguity about which player each payoff belongs to, because the table is read from top to bottom and from left to right.

For this symmetry, the order of the strategies matters (which it does not for other aspects of the game), so that, when exchanging the two players, the first row is exchanged with the first column, the second row with the second column, and so on. A non-symmetric game can sometimes be made symmetric when re-ordering the strategies of one player, as illustrated in Figure 2.7 below. Obviously, in a symmetric game, both players must have the same number of strategies.

The game of “chicken” is another symmetric game, shown in Figure 2.6. The two strategies are *A* and *C* for player I and *a* and *c* for player II, which may be termed “aggressive” and “cautious” behaviour, respectively. The aggressive strategy is only advantageous (with payoff 2 rather than 0) if the other player is cautious, whereas a cautious strategy always gives payoff 1 to the player using it.

The game known as the “battle of sexes” is shown in Figure 2.7(a). In this scenario, player I and player II are a couple each deciding whether to go to a concert (strategies *C*

		II	
		<i>a</i>	<i>c</i>
I	<i>A</i>	0, 0	1, 2
	<i>C</i>	2, 1	1, 1

Figure 2.6 The game “chicken”, its symmetry indicated by the diagonal dotted line.

(a)		II	
		<i>c</i>	<i>s</i>
I	<i>C</i>	2, 1	0, 0
	<i>S</i>	0, 0	1, 2

(b)		II	
		<i>s</i>	<i>c</i>
I	<i>C</i>	0, 0	2, 1
	<i>S</i>	1, 2	0, 0

Figure 2.7 The “battle of sexes” game (a), which is symmetric if the strategies of one player are exchanged, as shown in (b).

and *c*) or to a sports event (strategies *S* and *s*). The players have different payoffs arising from which event they go to, but that payoff is zero if they have to attend the event alone.

This game is not symmetric when written as in Figure 2.7(a), where strategy *C* of player I would be exchanged with strategy *c* of player II, and correspondingly *S* with *s*, because the payoffs for the strategy pairs (*C, c*) and (*S, s*) on the diagonal are not the same for both players, which is clearly necessary for symmetry. However, changing the order of the strategies of one player, for example of player II as shown in Figure 2.7(b), makes this a symmetric game.

		II		
		<i>r</i>	<i>s</i>	<i>p</i>
I	<i>R</i>	0, 0	-1, 1	1, -1
	<i>S</i>	1, -1	0, 0	-1, 1
	<i>P</i>	-1, 1	1, -1	0, 0
	<i>P</i>	1, -1	-1, 1	0, 0

Figure 2.8 The “rock-scissors-paper” game.

Figure 2.8 shows a 3×3 game known as “rock-scissors-paper”. Both players choose simultaneously one of their three strategies, where rock beats scissors, scissors beat paper, and paper beats rock, and it is a draw otherwise. This is a *zero-sum* game because the payoffs in any cell of the table sum to zero. The game is symmetric, and because it is zero-sum, the pairs of strategies on the diagonal, here (R, r) , (S, s) , and (P, p) , must give payoff zero to both players: otherwise the payoffs for these strategy pairs would not be the same for both players.

2.9 Symmetries involving strategies*

In this section,¹ we give a general definition of symmetry that may apply to any game in strategic form, with possibly more than two players. This definition allows also for symmetries among strategies of a player, or across players. As an example, the rock-scissors-paper game also has a symmetry with respect to its three *strategies*: When cyclically replacing R by S , S by P , and P by R , and r by s , s by p , and p by r , which amounts to moving the first row and column in Figure 2.8 into last place, respectively, then the game remains the same. (In addition, the players may also be exchanged as described.) The formal definition of a symmetry of the game is as follows.

Definition 2.3 Consider a game in strategic form with N as the finite set of players, and strategy set Σ_i for player $i \in N$, where $\Sigma_i \cap \Sigma_j = \emptyset$ for $i \neq j$. The payoff to player i is given by $a_i((s_j)_{j \in N})$ for each strategy profile $(s_j)_{j \in N}$, where $s_j \in \Sigma_j$. Then a *symmetry* of the game is given by a bijection $\phi: N \rightarrow N$ and a bijection $\psi_i: \Sigma_i \rightarrow \Sigma_{\phi(i)}$ for each $i \in N$ so that

$$a_i((s_j)_{j \in N}) = a_{\phi(i)}((\psi_j(s_j))_{j \in N}) \quad (2.1)$$

for each player $i \in N$ and each strategy profile $(s_j)_{j \in N}$.

Essentially, a symmetry is a way of re-naming players and strategies that leaves the game unchanged. In definition 2.3, the strategy sets Σ_i of the players are assumed to be pairwise disjoint. This assumption removes any ambiguity when looking at a strategy profile $(s_j)_{j \in N}$, which defines one strategy s_j in Σ_j for each player j , because then the particular order of these strategies does not matter, because every strategy belongs to exactly one player. In our examples, strategy sets have been disjoint because we used upper-case letters for player I and lower-case letters for player II.

A symmetry as in definition 2.3 is, first, given by a re-naming of the players with the bijection ϕ , where player i is re-named as $\phi(i)$, which is either another player in the game or the same player. Secondly, a bijection ψ_i re-names the strategies s_i of player i as strategies $\psi_i(s_i)$ of player $\phi(i)$. Any strategy profile $(s_j)_{j \in N}$ thereby becomes a re-named strategy profile $(\psi_j(s_j))_{j \in N}$. With this re-naming, player $\phi(i)$ has taken the role of player i , so one should evaluate his payoff, normally given by the payoff function a_i , as $a_{\phi(i)}((\psi_j(s_j))_{j \in N})$ when applied to the re-named strategy profile. If this produces in

¹This is an additional section which can be omitted at first reading, and has non-examinable material, as indicated by the star sign * following the section heading.

all cases the same payoffs, as stated in (2.1), then the re-naming has produced the same game in strategic form, and is therefore called a symmetry of the game.

For two players, the reflection along the diagonal is obviously simpler to state and to observe. In this symmetry among the two players, with player set $N = \{I, II\}$ and with strategy sets Σ_I and Σ_{II} , we used $\phi(I) = II$ and thus $\phi(II) = I$, and a *single* bijection $\psi_I : \Sigma_I \rightarrow \Sigma_{II}$ for re-naming the strategies, where ψ_{II} is the inverse of ψ_I . This last condition does not have to apply in definition 2.3, which makes this definition more general.

		II	
		<i>h</i>	<i>t</i>
I	<i>H</i>	-1 1	1 -1
	<i>T</i>	1 -1	-1 1

Figure 2.9 The “matching pennies” game.

The game of “matching pennies” in Figure 2.9 illustrates the greater generality of definition 2.3. This zero-sum game is played with two players that each have a penny, which the player can choose to show as heads or tails, which is the strategy H or T for player I, and h or t for player II. If the pennies match, for the strategy pairs (H, h) and (T, t) , then player II has to give her penny to player I; if they do not, for the strategy pairs (H, t) and (T, h) , then player I loses his penny to player II. This game has an obvious symmetry in strategies where heads are exchanged with tails for both players, but where the players are not re-named, so that ϕ is the identity function on the player set. However, the game cannot be written so that it remains the same when reflected along the diagonal, because the diagonal payoffs would have to be zero, as in the rock-scissors-paper game, so it is not symmetric among players in the sense of any of the earlier examples. Definition 2.3 does capture the symmetry between the two players as follows (for example): ϕ exchanges I and II, and $\psi_I(H) = h$, $\psi_I(T) = t$, but $\psi_{II}(h) = T$, $\psi_{II}(t) = H$. That is, the sides of the penny keep their meaning for player I when he is re-named as player II, but heads and tails change their role for player II. In effect, this exchanges the players’ preference for matching versus non-matching pennies, as required for the symmetry.

Please note: In the following discussion, when we call a game “symmetric”, we always mean the simpler symmetry of a two-player game described in the previous Section 2.8 that can be seen by reflecting the game along the diagonal.

2.10 Dominance and elimination of dominated strategies

This section discusses the concepts of strict and weak dominance, which apply to strategies and therefore to games in strategic form.

Definition 2.4 Consider a game in strategic form, and let s and t be two strategies of some player P of that game. Then s *dominates* (or *strictly dominates*) t if for any fixed strategies of the other players, the payoff to player P when using s is higher than his payoff when using t . Strategy s *weakly dominates* t if for any fixed strategies of the other players, the payoff to player P when using s is at least as high as when using t , and in at least one case strictly higher.

In a two-player game, player P in definition 2.4 is either player I or player II. For player I, his strategy s dominates t if row s of the payoffs to player I is in each component larger than row t . If one denotes the payoff to player I for row i and column j by $a(i, j)$, then s dominates t if $a(s, j) > a(t, j)$ for each strategy j of player II. Similarly, if the payoff to player II is denoted by $b(i, j)$, and s and t are two strategies of player II, which are columns of the payoff table, then s dominates t if $b(i, s) > b(i, t)$ for each strategy i of player I.

In the prisoner's dilemma game in Figure 2.5, strategy D of player I dominates strategy C , because, with the notation of the preceding paragraph, $a(D, c) = 3 > 2 = a(C, c)$ and $a(D, d) = 1 > 0 = a(C, d)$. Because the game is symmetric, strategy d of player II also dominates c .

It would be inaccurate to define dominance by saying that a strategy s dominates strategy t if s is "always" better than t . This is only correct if "always" means "given the *same* strategies of the other players". Even if s dominates t , it may happen that *some* payoff when playing s is worse than some other payoff when playing t . For example, $a(C, c) = 2 > 1 = a(D, d)$ in the prisoner's dilemma game, so this is a case where the dominating strategy D gives a lower payoff than C . However, the strategy used by the other player is necessarily different in that case. When s dominates t , then strategy s is better than t when considering any arbitrary (but same) fixed strategies of the other players.

Dominance is sometimes called *strict dominance* in order to distinguish it from *weak dominance*. Consider a two-player game with payoffs $a(i, j)$ to player I and $b(i, j)$ to player II when they play row i and column j . According to definition 2.4, strategy s of player I weakly dominates t if $a(s, j) \geq a(t, j)$ for all j and $a(s, j) > a(t, j)$ for at least one column j . The latter condition ensures that if the two rows s and t of payoffs to player I are equal, s is not said to weakly dominate t (because for the same reason, t could also be said to dominate s). Similarly, if s and t are strategies of player II, then s weakly dominates t if the column s of payoffs $b(i, s)$ to player II is in each component i at least as large as column t , and strictly larger, with $b(i, s) > b(i, t)$, in at least one row i . An example of such a strategy is l in Figure 2.3(b), which is weakly dominated by r .

When a strategy s of player P dominates his strategy t , player P can always improve his payoff by playing s rather than t . This follows from the way a game in strategic form is played, where the players choose their strategies simultaneously, and the game is played only once.² Then player P may consider the strategies of the other players as fixed, and

²The game in strategic form is considered as a "one-shot" game. Many studies concern the possible emergence of cooperation in the prisoner's dilemma when the game is repeated, which is a different context.

his own strategy choice cannot be observed and should not influence the choice of the others, so he is better off playing s rather than t .

In consequence, one may study the game where all dominated strategies are *eliminated*. If some strategies are eliminated in this way, one then obtains a game that is simpler because some players have fewer strategies. In the prisoner's dilemma game, elimination of the dominated strategies C and c results in a game that has only one strategy per player, D for player I and d for player II. This strategy profile (D, d) may therefore be considered as a "solution" of the game, a recommendation of a strategy for each player.

		II	
		l	r
I	T	2, 2	1, 0
	B	0, 3	1, 1

Figure 2.10 The "quality game", with T and B as good or bad quality offered by player I, and l and r as buying or refusing to buy the product as strategies of player II.

If one accepts that a player will never play a dominated strategy, one may eliminate it from the game and continue eliminating strategies that are dominated in the resulting game. This is called *iterated elimination of dominated strategies*. If this process ends in a unique strategy profile, the game is said to be *dominance solvable*, with the final strategy profile as its solution.

The "quality game" in Figure 2.10 is a game that is dominance solvable in this way. The game is nearly identical to the prisoner's dilemma game in Figure 2.5, except for the payoff to player II for the top right cell of the table, which is changed from 3 to 1. The game may describe the situation of, say, player I as a restaurant owner, who can provide food of good quality (strategy T) or bad quality (B), and a potential customer, player II, who may decide to eat there (l) or not (r). The customer prefers l to r only if the quality is good. However, whatever player II does, player I is better off by choosing B , which therefore dominates T . After eliminating T from the game, player II's two strategies l and r remain, but in this smaller game r dominates l , and l is therefore eliminated in a second iteration. The resulting strategy pair is (B, r) .

When eliminating dominated strategies, the order of elimination does not matter, because if s dominates t , then s still dominates t in the game where some strategies (other than t) are eliminated. In contrast, when eliminating weakly dominated strategies, the order of elimination may matter. Moreover, a weakly dominated strategy, such as strategy l of player II in Figure 2.3, may be just as good as the strategy that weakly dominates it, if the other player chooses some strategy, like T in Figure 2.3, where the two strategies have equal payoff. Hence, there are no strong reasons for eliminating a weakly dominated strategy in the first place.

⇒ Exercise 2.5 on page 64 studies weakly dominated strategies, and what can happen when these are eliminated. You can answer this exercise except for (d) which you should do after having understood the concept of Nash equilibrium, treated in the next section.

2.11 Nash equilibrium

Not every game is dominance solvable, as the games of chicken and the battle of sexes demonstrate. The central concept of non-cooperative game theory is that of *equilibrium*, often called *Nash equilibrium* after John Nash, who introduced this concept in 1950 for general games in strategic form (the equivalent concept for zero-sum games was considered earlier). An equilibrium is a strategy profile where each player's strategy is a "best response" to the remaining strategies.

Definition 2.5 Consider a game in strategic form and a strategy profile given by a strategy s_j for each player j . Then for player i , his strategy s_i is a *best response* to the strategies of the remaining players if no other strategy gives player i a higher payoff, when the strategies of the other players are unchanged. An *equilibrium* of the game, also called *Nash equilibrium*, is a strategy profile where the strategy of each player is a best response to the other strategies.

In other words, a Nash equilibrium is a strategy profile so that no player can gain by changing his strategy *unilaterally*, that is, with the remaining strategies kept fixed.




Figure 2.11 Indicating best responses in the quality game with arrows (a), or by putting best response payoffs in boxes (b).

Figure 2.11(a) shows the quality game of Figure 2.10 where the best responses are indicated with arrows: The downward arrow on the left shows that if player II plays her left strategy l , then player I has best response B ; the downward arrow on the right shows that the best response to r is also B ; the leftward arrow at the top shows that player II's best response to T is l , and the rightward arrow at the bottom that the best response to B is r .

This works well for 2×2 games, and shows that a Nash equilibrium is where the arrows "flow together". For the quality game, player I's arrows both point downwards,

which shows that B dominates T , and the bottom arrow points right, so that starting from any cell and following the arrows, one arrives at the equilibrium (B, r) .

In Figure 2.11(b), best responses are indicated by boxes drawn around the payoffs for each strategy that is a best response. For player I, a best response is considered against a strategy of player II. That strategy is a column of the game, so the best response payoff is the maximum of each column (which may occur more than once) of the payoffs of player I. Similarly, a best response payoff for player II is the maximum in each row of the payoffs to player II. Unlike arrows, putting best response payoffs into boxes can be done easily even when a player has more than two strategies. A Nash equilibrium is then given by a cell of the table where both payoffs are boxed.




Figure 2.12 Best responses and Nash equilibria in the game of chicken (a), and in the battle of sexes game (b).

The game of chicken in Figure 2.12(a) has two equilibria, (A, c) and (a, C) . If a game is symmetric, like in this case, then an equilibrium is symmetric if it does not change under the symmetry (that is, when exchanging the two players). Neither of the two equilibria (A, c) and (a, C) is symmetric, but they map to each other under the symmetry (any non-symmetric equilibrium must have a symmetric counterpart; only symmetric equilibria map to themselves).

As Figure 2.12(b) shows, the battle of sexes game has two Nash equilibria, (C, c) and (S, s) . (When writing the battle of sexes game symmetrically as in Figure 2.7(b), its equilibria are not symmetric either.) The prisoner's dilemma game has one equilibrium (D, d) , which is symmetric.

It is clear that a dominated strategy is never a best response, and hence cannot be part of an equilibrium. Consequently, one can eliminate any dominated strategy from a game, and not lose any equilibrium. Moreover, this elimination cannot create additional equilibria because any best response in the absence of a dominated strategy remains a best response when adding the dominated strategy back to the game. By repeating this argument when considering the iterated elimination of dominated strategies, we obtain the following proposition.

Proposition 2.6 *If a game in strategic form is dominance solvable, its solution is the only Nash equilibrium of the game.*

⇒ In order to understand the concept of Nash equilibrium in games with more than two players, Exercise 2.6 on page 65 is very instructive.

We have seen that a Nash equilibrium may not be unique. Another drawback is illustrated by the rock-scissors-paper game in Figure 2.8, and the game of matching pennies in Figure 2.9, namely that the game may have no equilibrium “in pure strategies”, that is, when the players may only use exactly one of their given strategies in a deterministic way. This drawback is overcome by allowing each player to use a “mixed strategy”, which means that the player chooses his strategy randomly according to a certain probability distribution. Mixed strategies are the topic of the next chapter.

In the following section, we return to game trees, which do have equilibria even when considering only non-randomised or “pure” strategies.

2.12 Reduced strategies

The remainder of this chapter is concerned with the connection of game trees and their strategic form, and the Nash equilibria of the game.

We first consider a simplification of the strategic form. Recall Figure 2.4, which shows the strategic form of the extensive game in Figure 2.1. The two rows of the strategies XP and XQ in that table have identical payoffs for both players, as do the two rows YP and YQ . This is not surprising, because after move X or Y of player I at the root of the game tree, the decision node where player I can decide between P or Q cannot be reached, because that node is preceded by move Z , which is excluded by choosing X or Y . Because player I makes that move himself, it makes sense to replace both strategies XP and XQ by a less specific “plan of action” X^* that prescribes only move X . We *always* use the star “ $*$ ” as a placeholder. It stands for an unspecified move at the respective unreachable decision node, to identify the node with its unspecified move in case the game has many decision nodes.

Leaving moves at unreachable nodes unspecified in this manner defines a *reduced strategy* according to the following definition. Because the resulting expected payoffs remain uniquely defined, tabulating these reduced strategies and the payoff for the resulting reduced strategy profiles gives the *reduced strategic form* of the game. The reduced strategic form of the game tree of Figure 2.1 is shown in Figure 2.13.

Definition 2.7 In a game tree, a *reduced strategy* of a player specifies a move for every decision node of that player, except for those decision nodes that are unreachable due to an earlier own move. A *reduced strategy profile* is a tuple of reduced strategies, one for each player of the game. The *reduced strategic form* of a game tree lists all reduced strategies for each player, and tabulates the expected payoff to each player for each reduced strategy profile.

The preceding definition generalises definitions 2.1 and 2.2. It is important that the only moves that are left unspecified are at decision nodes which are unreachable due to an earlier *own* move of the player. A reduced strategy must not disregard a move because

		II					
I		<i>ad</i>	<i>ae</i>	<i>bd</i>	<i>be</i>	<i>cd</i>	<i>ce</i>
X^*		4	4	5	5	5	5
		3	3	2	2	4	4
Y^*		3	3	3	3	3	3
		2	2	2	2	2	2
ZP		3	5	3	5	3	5
		1	0	1	0	1	0
ZQ		3	2	3	2	3	2
		1	4	1	4	1	4

Figure 2.13 Reduced strategic form of the extensive game in Figure 2.1. The star * stands for an arbitrary move at the second decision node of player I, which is not reachable after move X or Y .

another player may not move there, because that possibility cannot be excluded by looking only at the player's own moves. In the game tree in Figure 2.1, for example, no reduction is possible for the strategies of player II, because neither of her moves at one decision node precludes a move at her other decision node. Therefore, the reduced strategies of player II in that game are the same as her (unreduced, original) strategies.

In definition 2.7, reduced strategies are defined without any reference to payoffs, only by looking at the structure of the game tree. One could derive the reduced strategic form in Figure 2.13 from the strategic form in Figure 2.4 by identifying the strategies that have identical payoffs for both players and replacing them by a single reduced strategy. Some authors define the reduced strategic form in terms of this elimination of duplicate strategies in the strategic form. (Sometimes, dominated strategies are also eliminated when defining the reduced strategic form.) We define the reduced strategic form without any reference to payoffs because strategies may have identical payoffs by accident, due to special payoffs at the leaves of the game (this does not occur in generic games). Moreover, we prefer that reduced strategies refer to the structure of the game tree, not to some relationship of the payoffs, which is a different aspect of the game.

Strategies are combinations of moves, so for every additional decision node of a player, each move at that node can be combined with move combinations already considered. The number of move combinations therefore grows exponentially with the number of decision nodes of the player, because it is the product of the numbers of moves at each node. In the game in Figure 2.14, for example, where m possible initial moves of player I are followed each time by two possible moves of player II, the number of strategies of player II is 2^m . Moreover, this is also the number of reduced strategies of player II because no reduction is possible. This shows that even the reduced strategic form can be exponentially large in the size of the game tree. If a player's move is preceded by an own




Figure 2.14 Extensive game with m moves C_1, \dots, C_m of player I at the root of the game tree. To each move C_i , player II may respond with two possible moves l_i or r_i . Player II has 2^m strategies.

earlier move, the reduced strategic form is smaller because then that move can be left unspecified if the preceding own move is not made.

\Rightarrow Exercise 2.4 on page 64 tests your understanding of reduced strategies.

In the reduced strategic form, a Nash equilibrium is defined, as in definition 2.5, as a profile of reduced strategies, each of which is a best response to the others.

In this context, we will for brevity sometimes refer to reduced strategies simply as “strategies”. This is justified because, when looking at the reduced strategic form, the concepts of dominance and equilibrium can be applied directly to the reduced strategic form, for example to the table defining a two-player game. Then “strategy” means simply a row or column of that table. The term “reduced strategy” is only relevant when referring to the extensive form.

The Nash equilibria in Figure 2.13 are identified as those pairs of (reduced) strategies that are best responses to each other, with both payoffs surrounded by a box in the respective cell of the table. These Nash equilibria in reduced strategies are (X^*, bd) , (X^*, cd) , (X^*, ce) , and (Y^*, bd) .

\Rightarrow You are now in a position answer (b)–(e) of Exercise 2.1 on page 62.

2.13 Subgame perfect Nash equilibrium (SPNE)

In this section, we consider the relationship between Nash equilibria of games in extensive form and backward induction.

The 2×2 game in Figure 2.3(b) has two Nash equilibria, (T, l) and (B, r) . The strategy pair (B, r) is also obtained, uniquely, by backward induction. We will prove shortly that backward induction always defines a Nash equilibrium.

The equilibrium (T, l) in Figure 2.3(b) is not obtained by backward induction, because it prescribes the non-optimal move l for player II at her only decision node. Moreover, l is a weakly dominated strategy. Nevertheless, this is a Nash equilibrium because T is the best response to l , and, moreover, l is a best response to T because against T , the

payoff to player II is no worse than when choosing r . This can also be seen when considering the game tree in Figure 2.3(a): Move l is a best response to T because player II never has to make that move when player I chooses T , so player II's move does not affect her payoff, and she is indifferent between l and r . On the other hand, only when she makes move l is it optimal for player I to respond by choosing T , because against r player I would get more by choosing B .

The game in Figure 2.3(a) is also called a “threat game” because it has a Nash equilibrium (T, l) where player II “threatens” to make the move l that is bad for both players, against which player I chooses T , which is then advantageous for player II compared to the backward induction outcome when the players choose (B, r) . The threat works only because player II never has to execute it, given that player I acts rationally and chooses T .

The concept of Nash equilibrium is based on the strategic form because it applies to a strategy profile. When applied to a game tree, the strategies in that profile are assumed to be chosen by the players *before* the game starts, and the concept of best response applies to this given expectation of what the other players will do.

With the game tree as the given specification of the game, it is often desirable to keep its sequential interpretation. The strategies chosen by the players should therefore also express some “sequential rationality” as expressed by backward induction. That is, the moves in a strategy profile should be optimal for any part of the game, including subtrees that cannot be reached due to earlier moves, possibly of other players, like the tree starting at the decision node of player II in the threat game in Figure 2.3(a).

In a game tree, a *subgame* is any subtree of the game, given by a node of the tree as the root of the subtree and all its descendants. (Note: This definition of a subgame applies only to games with perfect information, which are the game trees considered so far. In extensive games with imperfect information, which we consider later, the subtree is a subgame only if all players *know* that they are in that subtree.) A strategy profile that defines a Nash equilibrium for every subgame is called a *subgame perfect Nash equilibrium* or SPNE.

⇒ You can now answer the final question (e) of Exercise 2.1 on page 62, and Exercise 2.2 on page 63.

Theorem 2.8 *Backward induction defines an SPNE.*

Proof. Recall the process of backward induction: Starting with the nodes closest to the leaves, consider a decision node u , say, with the assumption that all moves after u have already been selected. Among the moves at u , select a move that *maximises* the expected payoff to the player that moves at u . (Expected payoffs must be regarded if there are chance moves after u .) In that manner, a move is selected for every decision node, which determines an entire strategy profile.

We prove the theorem inductively: Consider a non-terminal node u of the game tree, which may be a decision node (as in the backward induction process), or a chance node. Suppose that the moves at u are c_1, c_2, \dots, c_k , which lead to subtrees T_1, T_2, \dots, T_k of the game tree, and assume, as inductive hypothesis, that the moves selected so far define an

SPNE in each of these trees. (As the “base case” of the induction, this is certainly true if each of T_1, T_2, \dots, T_k is just a leaf of the tree, so that u is a “last” decision node considered first in the backward induction process.) The induction step is completed if one shows that, by selecting the move at u , one obtains an SPNE for the subgame with root u .

First, suppose that u is a chance node, so that the next node is chosen according to the fixed probabilities specified in the game tree. Then backward induction does not select a move for u , and the inductive step holds trivially: For every player, the payoff in the subgame starting at u is the expectation of the payoffs for each subgame T_i (weighted with the probability for move c_i), and if a player could improve on that payoff, she would have to do so by changing her moves within at least one subtree T_i , which, by inductive hypothesis, she cannot.

Secondly, suppose that u is a decision node, and consider a player *other* than the player to move at u . Again, for that player, the moves in the subgame starting at u are completely specified, and, irrespective of what move is selected for u , she cannot improve her payoff because that would mean she could improve her payoff already in some subtree T_i .

Finally, consider the player to move at u . Backward induction selects a move for u that is best for that player, given the remaining moves. Suppose the player could improve his payoff by choosing some move c_i and additionally change his moves in the subtree T_i . But the resulting improved payoff would *only* be the improved payoff in T_i , that is, the player could already get a better payoff in T_i itself, contradicting the inductive assumption that the moves selected so far defined an SPNE for T_i . This completes the induction. \square

This theorem has two important consequences. In backward induction, each move can be chosen deterministically, so that backward induction determines a profile of pure strategies. Theorem 2.8 therefore implies that game trees have Nash equilibria, and it is not necessary to consider randomised strategies. Secondly, subgame perfect Nash equilibria exist. For game trees, we can use “SPNE” synonymously with “strategy profile obtained by backward induction”.

Corollary 2.9 *Every game tree with perfect information has a pure-strategy Nash equilibrium.*

Corollary 2.10 *Every game tree has an SPNE.*

In order to describe an SPNE, it is important to consider *unreduced* (that is, fully specified) strategies. Recall that in the game in Figure 2.1, the SPNE are (XQ, bd) , (XQ, cd) , and (YQ, bd) . The reduced strategic form of the game in Figure 2.13 shows the Nash equilibria $(X*, bd)$, $(X*, cd)$, and $(Y*, bd)$. However, we cannot call any of these an SPNE because they leave the second move of player I unspecified, as indicated by the $*$ symbol. In this case, replacing $*$ by Q results in all cases in an SPNE. The full strategies are necessary to determine if they define a Nash equilibrium in every subgame. As seen in Figure 2.13, the game has, in addition, the Nash equilibrium $(X*, ce)$. This is not subgame perfect because it prescribes the move e , which is not part of a Nash equilibrium in the subgame starting with the decision node where player II chooses between

d and e : When replacing $*$ by P , this would not be the best response by player I to e in that subgame, and when replacing $*$ by Q , player II's best response would not be e . This property can only be seen in the game tree, not even in the unreduced strategic form in Figure 2.4, because it concerns an unreachable part of the game tree due to the earlier move X of player I.

⇒ For understanding subgame perfection, Exercise 2.3 on page 64 is very instructive. Exercise 2.7 on page 66 studies a three-player game.

2.14 Commitment games

The main difference between game trees and games in strategic form is that in a game tree, players act sequentially, being aware of the previous moves of the other players. In contrast, players in a game in strategic form move *simultaneously*. The difference between these two descriptions becomes striking when *changing* a two-player game in strategic form to a *commitment game*, as described in this section.

We start with a two-player game in strategic form, say an $m \times n$ game. The corresponding commitment game is an extensive game with perfect information. One of the players (we always choose player I) moves *first*, by choosing one of his m strategies. Then, unlike in the strategic form, player II is *informed* about the move of player I. For each move of player I, the possible moves of player II are the n strategies of the original strategic form game, with resulting payoffs to both players as specified in that game.




Figure 2.15 The extensive game (a) is the *commitment game* of the quality game in Figure 2.10. Player I moves first, and player II can *react* to that choice. The arrows show the backward induction outcome. Its strategic form with best-response payoffs is shown in (b).

An example is Figure 2.15, which is the commitment game for the quality game in Figure 2.10. The original strategies T and B of player I become his moves T and B . In general, if player I has m strategies in the given game, he also has m strategies in the commitment game. In the commitment game, each strategy l or r of player II becomes a move, which depends on the preceding move of player I, so we call them l_T and r_T when following move T , and l_B and r_B when following move B . In general, when player I has m moves in the commitment game and player II has n responses, then each combination

of moves defines a pure strategy of player II, so she has n^m strategies in the commitment game. Figure 2.15(b) shows the resulting strategic form of the commitment game in the example. It shows that the given game and the commitment game are very different.

We then look for subgame perfect equilibria of the commitment game, that is, apply backward induction, which we can do because the game has perfect information. Figure 2.15(a) shows that the resulting unique SPNE is $(T, l_T r_B)$, which means that (going backwards) player II, when offered high quality T , chooses l_T , and r_B when offered low quality. In consequence, player I will then choose T because that gives him the higher payoff 2 rather than just 1 with B . Note that it is important that player I has the power to *commit* to T and cannot change later back to B after seeing that player II chose l ; without that commitment power, the game would not be accurately reflected by the extensive game as described.

The commitment game in Figure 2.15 also demonstrates the difference between an equilibrium, which is a strategy profile, and the *equilibrium path*, which is described by the moves that are actually played in the game when players use their equilibrium strategies. If the game has no chance moves, the equilibrium path is just a certain path in the game tree. In contrast, a strategy profile defines a move in every part of the game tree. Here, the SPNE $(T, l_T r_B)$ specifies moves l_T and r_B for both decision points of player II. Player I chooses T , and then only l_T is played. The equilibrium path is then given by move T followed by move l_T . However, it is *not sufficient* to simply call this equilibrium (T, l_T) or (T, l) , because player II's move r_B must be specified to know that T is player I's best response.

We only consider SPNE when analysing commitment games. For example, the Nash equilibrium (B, r) of the original quality game in Figure 2.10 can also be found in Figure 2.15, where player II *always* chooses the move r corresponding to her equilibrium strategy in the original game, which is the strategy $r_T r_B$ in the commitment game. This defines the Nash equilibrium $(B, r_T r_B)$ in the commitment game, because B is a best response to $r_T r_B$, and $r_T r_B$ is a best response to B (which clearly holds as a general argument, starting with any Nash equilibrium of the original game). However, this Nash equilibrium of the commitment game is not subgame perfect, because it prescribes the suboptimal move r_T off the equilibrium path; this does not affect the Nash equilibrium property because T is not chosen by player I. In order to compare a strategic-form game with its commitment game in an interesting way, we consider only SPNE of the commitment game.

A practical application of a game-theoretic analysis may be to reveal the potential effects of changing the “rules” of the game. This is illustrated by changing the quality game to its commitment version.

⇒ This is a good point to try Exercise 2.8 on page 66.

Games in strategic form, when converted to a commitment game, typically have a *first-mover advantage*. A player in a game becomes a first mover or “leader” when he can commit to a strategy as described, that is, choose a strategy irrevocably and inform the other players about it; this is a change of the “rules of the game”. The first-mover advantage states that a player who can become a leader is not worse off than in the original

game where the players act simultaneously. In other words, if one of the players has the power to commit, he or she should do so.

This statement must be interpreted carefully. For example, if more than one player has the power to commit, then it is not necessarily best to go first. For example, consider changing the game in Figure 2.10 so that player II can commit to her strategy, and player I moves second. Then player I will always respond by choosing B because this is his dominant strategy in Figure 2.10. Backward induction would then amount to player II choosing r , and player I choosing $B_l B_r$, with the low payoff 1 to both. Then player II is not worse off than in the simultaneous-choice game, as asserted by the first-mover advantage, but does not gain anything either. In contrast, making player I the first mover as in Figure 2.15 is beneficial to both.

The first-mover advantage is also known as *Stackelberg leadership*, after the economist Heinrich von Stackelberg, who formulated this concept for the structure of markets in 1934. The classic application is to the duopoly model by Cournot, which dates back to 1838.

		II			
		h	m	l	n
I	H	0 0	8 12	9 18	0 36
	M	12 8	16 16	15 20	0 32
	L	18 9	20 15	18 18	0 27
	N	36 0	32 0	27 0	0 0

Figure 2.16 Duopoly game between two chip manufacturers who can decide between high, medium, low, or no production, denoted by H, M, L, N for player I and h, m, l, n for player II. Payoffs denote profits in millions of dollars.

As an example, suppose that the market for a certain type of memory chip is dominated by two producers. The players can choose to produce a certain quantity of chips, say either high, medium, low, or none at all, denoted by H, M, L, N for player I and h, m, l, n for player II. The market price of the memory chips decreases with increasing total quantity produced by both companies. In particular, if both choose a high quantity of production, the price collapses so that profits drop to zero. The players know how increased production lowers the chip price and their profits. Figure 2.16 shows the game in strategic form, where both players choose their output level simultaneously. The symmetric payoffs are derived from Cournot's model, explained below.

This game is dominance solvable (see Section 2.10 above). Clearly, no production is dominated by low or medium production, so that row N and column n in Figure 2.16 can be eliminated. Then, high production is dominated by medium production, so that row H and column h can be omitted. At this point, only medium and low production remain. Then, regardless of whether the opponent produces medium or low, it is always better for each player to produce a medium quantity, eliminating L and l . Only the strategy pair (M, m) remains. Therefore, the Nash equilibrium of the game is (M, m) , where both players make a profit of \$16 million.

Consider now the commitment version of the game, with a game tree corresponding to Figure 2.16 just as Figure 2.15 is obtained from Figure 2.10. We omit this tree to save space, but backward induction is easily done using the strategic form in Figure 2.16: For each row H , M , L , or N , representing a possible commitment of player I, consider the best response of player II, as shown by the boxed payoff entry for player II *only* (the best responses of player I are irrelevant). The respective best responses are unique, defining the backward induction strategy $l_H m_M m_L h_N$ of player II, with corresponding payoffs 18, 16, 15, 0 to player I when choosing H , M , L , or N , respectively. Player I gets the maximum of these payoffs when he chooses H , to which player II will respond with l_H . That is, among the anticipated responses by player II, player I does best by announcing H , a high level of production. The backward induction outcome is thus that player I makes a profit \$18 million, as opposed to only \$16 million in the simultaneous-choice game. When player II must play the role of the follower, her profits fall from \$16 million to \$9 million.

The first-mover advantage comes from the ability of player I to credibly commit himself. After player I has chosen H , and player II replies with l , player I would like to be able switch to M , improving profits even further from \$18 million to \$20 million. However, once player I is producing M , player II would change to m . This logic demonstrates why, when the players choose their quantities simultaneously, the strategy combination (H, l) is not an equilibrium. The commitment power of player I, and player II's appreciation of this fact, is crucial.

The payoffs in Figure 2.16 are derived from the following simple model due to Cournot. The high, medium, low, and zero production numbers are 6, 4, 3, and 0 million memory chips, respectively. The profit per chip is $12 - Q$ dollars, where Q is the total quantity (in millions of chips) on the market. The entire production is sold. As an example, the strategy combination (H, l) yields $Q = 6 + 3 = 9$, with a profit of \$3 per chip. This yields the payoffs of \$18 million and \$9 million for players I and II in the (H, l) cell in Figure 2.16. Another example is player I acting as a monopolist (player II choosing n), with a high production level H of 6 million chips sold at a profit of \$6 each.

In this model, a monopolist would produce a quantity of 6 million even if numbers other than 6, 4, 3, or 0 were allowed, which gives the maximum profit of \$36 million. The two players could cooperate and split that amount by producing 3 million each, corresponding to the strategy combination (L, l) in Figure 2.16. The equilibrium quantities, however, are 4 million for each player, where both players receive less. The central four cells in Figure 2.16, with low and medium production in place of “cooperate” and “defect”, have the structure of a prisoner's dilemma game (Figure 2.5), which arises here in

a natural economic context. The optimal commitment of a first mover is to produce a quantity of 6 million, with the follower choosing 3 million.

These numbers, and the equilibrium (“Cournot”) quantity of 4 million, apply even when arbitrary quantities are allowed. That is, suppose x and y are the quantities of production for player I and player II, respectively. The payoffs $a(x, y)$ to player I and $b(x, y)$ to player II are defined as

$$\begin{aligned} a(x, y) &= x \cdot (12 - x - y), \\ b(x, y) &= y \cdot (12 - x - y). \end{aligned} \tag{2.2}$$


Then it is easy to see that player I’s best response $x(y)$ to y is given by $6 - y/2$, and player II’s best response $y(x)$ to x is given by $6 - x/2$. The pair of linear equations $x(y) = x$ and $y(x) = y$ has the solution $x = y = 4$, which is the above Nash equilibrium (M, m) of Figure 2.16. In the commitment game, player I maximises his payoff, assuming the unique best response $y(x)$ of player II in the SPNE, by maximising $a(x, y(x))$, which is $x \cdot (12 - x - 6 + x/2) = x \cdot (12 - x)/2$. That maximum is given for $x = 6$, which happens to be the strategy H (high production) in Figure 2.16. The best response of player II to that commitment is $6 - 6/2 = 3$, which we have named strategy l for low production.

⇒ Exercise 2.9 on page 67 studies a game with infinite strategy sets like (2.2), and makes interesting observations, in particular in (d), concerning the payoffs to the two players in the commitment game compared to the original simultaneous game.

2.15 Exercises for Chapter 2

Exercises 2.1 and 2.2 ask you to apply the main concepts of this chapter to a simple game tree. A more detailed study of the concept of SPNE is Exercise 2.3. Exercise 2.4 is on counting strategies and reduced strategies. Exercise 2.5 shows that iterated elimination of *weakly* dominated strategies is a problematic concept. Exercises 2.6 and 2.7 concern three-player games, which are very important to understand because the concepts of dominance and equilibrium require to keep the strategies of *all other* players fixed; this is not like just having another player in a two-player game. Exercise 2.8 is an instructive exercise on commitment games. Exercise 2.9 studies commitment games where in the original game both players have infinitely many strategies.

Exercise 2.1 Consider the following game tree. At a leaf, the top payoffs are for player I, the bottom payoffs are for player II.



- What is the number of strategies of player I and of player II?
- How many reduced strategies do they have?
- Give the reduced strategic form of the game.
- What are the Nash equilibria of the game in reduced strategies?
- What are the subgame perfect Nash equilibria of the game?

Exercise 2.2 Consider the game tree in Figure 2.17. At a leaf, the top payoff is for player I, the bottom payoff for player II.




Figure 2.17 Game tree for Exercise 2.2.

- What is the number of strategies of player I and of player II? How many reduced strategies does each of the players have?
- Give the reduced strategic form of the game.

- (c) What are the Nash equilibria of the game in reduced strategies? What are the subgame perfect equilibria of the game?
- (d) Identify every pair of reduced strategies where one strategy weakly or strictly dominates the other, and indicate if the dominance is weak or strict.

Exercise 2.3 Consider the game trees in Figure 2.18.

- (a) For the game tree in Figure 2.18(a), find all Nash equilibria (in pure strategies). Which of these are subgame perfect?
- (b) In the game tree in Figure 2.18(b), the payoffs a, b, c, d are positive real numbers. For each of the following statements (i), (ii), (iii), decide if it is true or false, justifying your answer with an argument or counterexample; you may refer to any standard results. For any $a, b, c, d > 0$,
- the game always has a subgame perfect Nash equilibrium (SPNE);
 - the payoff to player II in any SPNE is always at least as high as her payoff in any Nash equilibrium;
 - the payoff to player I in any SPNE is always at least as high as his payoff in any Nash equilibrium.




Figure 2.18 Game trees for Exercise 2.3

Exercise 2.4 Consider the two game trees (a) and (b) in Figure 2.19. In each case, how many strategies does each player have? How many reduced strategies?

Exercise 2.5 Consider the 3×3 game in Figure 2.20.

- (a) Identify all pairs of strategies where one strategy weakly dominates the other.
- (b) Assume you are allowed to remove a weakly dominated strategy of some player. Do so, and repeat this process (of iterated elimination of weakly dominated strategies) until you find a single strategy pair of the original game. (Remember that two strategies with identical payoffs do *not* weakly dominate each other!)




Figure 2.19 Game trees for Exercise 2.4. Payoffs have been omitted because they are not relevant for the question.

		II		
		<i>l</i>	<i>c</i>	<i>r</i>
I	<i>T</i>	0 1	1 3	1 1
	<i>M</i>	1 1	0 3	1 0
	<i>B</i>	2 2	3 3	2 0

Figure 2.20 3×3 game for Exercise 2.5.

- (c) Find such an iterated elimination of weakly dominated strategies that results in a strategy pair other than the one found in (b), where *both* strategies, and the payoffs to the players, are different.
- (d) What are the Nash equilibria (in pure strategies) of the game?

Exercise 2.6 Consider the three-player game in strategic form in Figure 2.21. Each player has two strategies: Player I has the strategies *T* and *B* (the top and bottom row below), player II has the strategies *l* and *r* (left or right column in each 2×2 panel), and player III has the strategies *L* and *R* (the right or left panel). The payoffs to the players in each cell are given as triples of numbers to players I, II, III.

- (a) Identify all pairs of strategies where one strategy strictly, or weakly, dominates the other.
[Hint: Make sure you understand what dominance means for more than two players. Be careful to consider the correct payoffs.]
- (b) Apply iterated elimination of strictly dominated strategies to this game. What are the Nash equilibria of the game?

		II		
		<i>l</i>	<i>r</i>	
I	<i>T</i>	3, 4, 4	1, 3, 3	
	<i>B</i>	8, 1, 4	2, 0, 6	
		III: <i>L</i>		

		II		
		<i>l</i>	<i>r</i>	
I	<i>T</i>	4, 0, 5	0, 1, 6	
	<i>B</i>	5, 1, 3	1, 2, 5	
		III: <i>R</i>		

Figure 2.21 Three-player game for Exercise 2.6.




Figure 2.22 Game tree with three players for Exercise 2.7. At a leaf, the topmost payoff is to player I, the middle payoff is to player II, and the bottom payoff is to player III.

Exercise 2.7 Consider the three-player game tree in Figure 2.22.

- How many strategy *profiles* does this game have?
- Identify all pairs of strategies where one strategy strictly, or weakly, dominates the other.
- Find all Nash equilibria in pure strategies. Which of these are subgame perfect?

Exercise 2.8 Consider a game G in strategic form. Recall that the *commitment game* derived from G is defined by letting player I choose one of his pure strategies x which is then *announced* to player II, who can then in each case choose one of her strategies in G as a response to x . The resulting payoffs are as in the original game G .

- If G is an $m \times n$ game, how many strategies do player I and player II have, respectively, in the commitment game?

For each of the following statements, determine whether they are true or false; justify your answer by a short argument or counterexample.

- (b) In an SPNE of the commitment game, player I never commits to a strategy that is strictly dominated in G .
- (c) In an SPNE of the commitment game, player II never chooses a move that is a strictly dominated strategy in G .

Exercise 2.9 Let G be the following game: player I chooses a (not necessarily integer) non-negative number x , and player II in the same way a non-negative number y . The resulting (symmetric) payoffs are

$$\begin{aligned} &x \cdot (4 + y - x) \text{ for player I,} \\ &y \cdot (4 + x - y) \text{ for player II.} \end{aligned}$$

- (a) Given x , determine player II's best response $y(x)$ (which is a function of x), and player I's best response $x(y)$ to y . Find a Nash equilibrium, and give the payoffs to the two players.
- (b) Find an SPNE of the commitment game, and give the payoffs to the two players.
- (c) Are the equilibria in (a) and (b) unique?
- (d) Let G be a game where the best response $y(x)$ of player II to any strategy x of player I is always unique. Show that in any SPNE of the commitment game, the payoff to player I is at least as large as his payoff in any Nash equilibrium of the original game G .

Chapter 3

Mixed strategy equilibria

This chapter studies mixed strategy equilibria, extending the considerations of Chapter 2.

3.1 Learning objectives

After studying this chapter, you should be able to:

- explain what mixed strategies and mixed equilibria are;
- state the main theorems about mixed strategies: the best response condition, Nash's theorem on the existence of mixed equilibria, and the minimax theorem;
- find mixed equilibria of small bimatrix games, using geometric tools like the upper envelope method where necessary;
- give the definition of a degenerate game, and find all equilibria of small degenerate games;
- explain the concept of max-min strategies and their importance for zero-sum games, and apply this to given games.

3.2 Further reading

The strong emphasis on geometric concepts for understanding mixed strategies is not found to this extent in other books. For zero-sum games, which are the subject of Section 3.14, a detailed treatment is given in the following book:

- Mendelson, Elliot *Introducing Game Theory and Its Applications*. (Chapman & Hall / CRC, 2004) [ISBN 1584883006].

Later sections and appendices in that book treat, more briefly, also the non-zero-sum games studied here.

In Section 3.5, an “inspection game” is used to motivate mixed strategies. A book on inspection games is

- Avenhaus, Rudolf, and Morton J. Canty *Compliance Quantified*. (Cambridge University Press, 1996) [ISBN 0521019192].

Nash's central paper is

- Nash, John F. "Non-cooperative games." *Annals of Mathematics*, Vol. 54 (1951), pp. 286–295.

This paper is easy to obtain. It should be relatively accessible after a study of this chapter.

3.3 Introduction

A game in strategic form does not always have a Nash equilibrium in which each player deterministically chooses one of his strategies. However, players may instead randomly select from among these *pure* strategies with certain probabilities. Randomising one's own choice in this way is called a *mixed* strategy. A profile of mixed strategies is called a *mixed equilibrium* if no player can gain on average by unilateral deviation. Nash showed in 1951 that any finite strategic-form game has a mixed equilibrium.

In this chapter, we first discuss how a player's payoffs represent his preference for random outcomes via an *expected-utility* function. This is illustrated first with a single-player decision problem and then with a game whether to comply or not with a legal requirement. The game is known as an *inspection game* between a player II called inspectee and an inspector as player I. This game has no equilibrium using pure (that is, deterministic) strategies. However, active randomisation comes up naturally in this game, because the inspector can and will choose to inspect only some of the time. The mixed equilibrium in this game demonstrates many aspects of general mixed equilibria, in particular that a player's randomisation probabilities depend on the payoffs of the *other* player.

We then turn to general two-player games in strategic form, called bimatrix games. The payoff matrices for the two players are convenient to represent expected payoffs via multiplication with vectors of probabilities which are the players' mixed strategies. These mixed strategies, as vectors of probabilities, have a geometric interpretation, which we study in detail. Mixing with probabilities is geometrically equivalent to taking *convex combinations*.

The *best response condition* gives a convenient, finite condition when a mixed strategy is a best response to the mixed strategy of the other player (or against the mixed strategies of all other players in a game with more than two players, which we do not consider). Namely, only the *pure* strategies that are best responses are allowed to have positive probability. These pure best responses must have maximal, and hence equal, expected payoff. This is stated as an "obvious fact" in Nash's paper, and it is not hard to prove. However, the best response condition is central for understanding and computing mixed-strategy equilibria.

In Section 3.10, we state and prove Nash's theorem that every game has an equilibrium. We follow Nash's original proof, and state and briefly motivate *Brouwer's fixed point theorem* used in that proof, but do not prove the fixed point theorem itself.

Sections 3.11–3.13 show how to find equilibria in bimatrix games, in particular when one player has only two strategies.

The special and interesting case of zero-sum games is treated in Section 3.14.

3.4 Expected-utility payoffs

The payoffs in a game represent a player's preferences, in that higher payoffs correspond to more desirable outcomes for the player. In this sense, the payoffs have the role of an “ordinal” utility function, meaning that only the order of preference is important. For example, the order among payoffs 0, 8, and 10 can equally be represented by the numbers 0, 2, and 10.

As soon as the outcome is random, the payoffs also represent a “cardinal” utility, which means that the relative size of the numbers also matters. The reason is that average (that is, *expected*) payoffs are considered to represent the preference for the random outcomes. For example, the payoffs 0, 8, and 10 do not represent the same preferences as 0, 2 and 10 when the player has to decide between a coin-toss that gives 0 or 10 with probability $1/2$ each, which is a risky choice, or taking the “middle” payoff. The expected payoff of the coin-toss is 5, which is larger than 2 but smaller than 8, so it matters whether the middle payoff is 2 or 8. In a game, a player's payoffs always represent his “expected utility” function in the sense that the payoffs can be weighted with probabilities in order to represent the player's preference for a random outcome.




Figure 3.1 One-player decision problem to decide between “comply” and “cheat”. This problem demonstrates expected-utility payoffs. With these payoffs, the player is indifferent.

As a concrete example, Figure 3.1 shows a game with a single player who can decide to comply with a regulation, like to buy a parking permit, or to cheat otherwise. The payoff when she chooses to comply is 0. Cheating involves a 10 percent chance of getting caught and having to pay a penalty, stated as the negative payoff -90 , and otherwise a 90 percent chance of gaining a payoff of 10. With these numbers, cheating leads to a random outcome with an expected payoff of $0.9 \times 10 + 0.1 \times (-90)$, which is zero, so that the player is exactly *indifferent* (prefers neither one nor the other) between her two available moves.

If the payoffs are monetary amounts, each payoff unit standing for a dollar, say, one would not necessarily assume such a *risk neutrality* on the part of the player. In practice, decision-makers are typically *risk averse*, meaning they prefer the safe payoff of 0 to the gamble with an expectation of 0.

In a game-theoretic model with random outcomes, as in the game above, the payoff is not necessarily to be interpreted as money. Rather, the player's attitude towards risk is incorporated into the payoff figure as well. To take our example, the player faces a punishment or reward when cheating, depending on whether she is caught or not. Suppose that the player's decision only depends on the probability of being caught, which is 0.1 in Figure 3.1, so that she would cheat if that probability was zero. Moreover, set the reward for cheating arbitrarily to 10 units, as in the figure above, and suppose that being caught has clearly defined consequences for the player, like regret and losing money and time. Then there must be a certain probability of getting caught where the player in the above game is indifferent, say 4 percent. This determines the utility $-u$, say, for "getting caught" by the equation

$$0 = 0.96 \times 10 + 0.04 \times (-u)$$

which states equal expected utility for the choices "comply" and "cheat". This equation is equivalent to $u = 9.6/0.04 = 240$. That is, in the above game, the negative utility -90 would have to be replaced by -240 to reflect the player's attitude towards the risk of getting caught. With that payoff, she will now *prefer* to comply if the probability of getting caught stays at 0.1.

The point of this consideration is to show that payoffs exist, and can be constructed, that represent a player's preference for a risky outcome, as measured by the resulting *expected payoff*. These payoffs do not have to represent money. The existence of such expected-utility payoffs depends on a certain consistency of the player when facing choices with random outcomes. This can be formalised, but the respective theory, known as the *von Neumann–Morgenstern* axioms for expected utility, is omitted here for brevity.

In practice, the risk attitude of a player may not be known. A game-theoretic analysis should be carried out for different choices of the payoff parameters in order to test how much they influence the results. Often, these parameters represent the "political" features of a game-theoretic model, those most sensitive to subjective judgement, compared to the more "technical" part of a solution. In particular, there are more involved variants of the inspection game discussed in the next section. In those more complicated models, the technical part often concerns the optimal usage of limited inspection resources, like maximising the probability of catching a player who wants to cheat. A separate step afterwards is the "political decision" when to declare that the inspectee has actually cheated. Such models and practical issues are discussed in the book by Avenhaus and Canty, *Compliance Quantified* (see Section 3.2).

3.5 Example: Compliance inspections

Suppose a player, whom we simply call "inspectee", must fulfil some legal obligation (such as buying a transport ticket, or paying tax). The inspectee has an incentive to violate

this obligation. Another player, the “inspector”, would like to verify that the inspectee is abiding by the rules, but doing so requires inspections which are costly. If the inspector does inspect and catches the inspectee cheating, the inspector can demand a large penalty payment for the noncompliance.

		II comply → cheat	
I Don't inspect ↑		0 10	
		0 -10	
Inspect ↓		0 -90	
		-1 -6	
		←	

Figure 3.2 Inspection game between an inspector (player I) and inspectee (player II).

Figure 3.2 shows possible payoffs for such an inspection game. The standard outcome, which defines the reference payoff zero to both inspector (player I) and inspectee (player II), is that the inspector chooses “Don’t inspect” and the inspectee chooses to comply. Without inspection, the inspectee prefers to cheat because that gives her payoff 10, with resulting negative payoff -10 to the inspector. The inspector may also decide to inspect. If the inspectee complies, inspection leaves her payoff 0 unchanged, while the inspector incurs a cost resulting in a negative payoff -1 . If the inspectee cheats, however, inspection will result in a heavy penalty (payoff -90 for player II) and still create a certain amount of hassle for player I (payoff -6).

In all cases, player I would strongly prefer if player II complied, but this is outside of player I’s control. However, the inspector prefers to inspect if the inspectee cheats (because -6 is better than -10), indicated by the downward arrow on the right in Figure 3.2. If the inspector always preferred “Don’t inspect”, then this would be a dominating strategy and be part of a (unique) equilibrium where the inspectee cheats.

The circular arrow structure in Figure 3.2 shows that this game has no equilibrium in pure strategies. If any of the players settles on a deterministic choice (like “Don’t inspect” by player I), the best response of the other player would be unique (here to cheat by player II), to which the original choice would not be a best response (player I prefers to inspect when the other player cheats, against which player II in turn prefers to comply). The strategies in a Nash equilibrium must be best responses to each other, so in this game this fails to hold for any pure strategy profile.

What should the players do in the game of Figure 3.2? One possibility is that they prepare for the worst, that is, choose a max-min strategy. A max-min strategy maximises the player’s worst payoff against all possible choices of the opponent. The max-min strategy (as a pure strategy) for player I is to inspect (where the inspector guarantees

himself payoff -6), and for player II it is to comply (which guarantees her payoff 0). However, this is not a Nash equilibrium and hence not a stable recommendation to the two players, because player I could switch his strategy and improve his payoff.

A *mixed strategy* of player I in this game is to inspect only with a certain probability. In the context of inspections, randomising is also a practical approach that reduces costs. Even if an inspection is not certain, a sufficiently high chance of being caught should deter from cheating, at least to some extent.

The following considerations show how to find the probability of inspection that will lead to an equilibrium. If the probability of inspection is very low, for example one percent, then player II receives (irrespective of that probability) payoff 0 for complying, and payoff $0.99 \times 10 + 0.01 \times (-90) = 9$, which is bigger than zero, for cheating. Hence, player II will still cheat, just as in the absence of inspection.

If the probability of inspection is much higher, for example 0.2, then the expected payoff for “cheat” is $0.8 \times 10 + 0.2 \times (-90) = -10$, which is less than zero, so that player II prefers to comply. If the inspection probability is either too low or too high, then player II has a unique best response. As shown above, such a pure strategy cannot be part of an equilibrium.

Hence, the only case where player II herself could possibly randomise between her strategies is if both strategies give her the same payoff, that is, if she is *indifferent*. As stated and proved formally in Theorem 3.1 below, it is never optimal for a player to assign a positive probability to a pure strategy that is inferior to other pure strategies, given what the other players are doing. It is not hard to see that player II is indifferent if and only if player I inspects with probability 0.1, because then the expected payoff for cheating is $0.9 \times 10 + 0.1 \times (-90) = 0$, which is then the same as the payoff for complying.

With this mixed strategy of player I (don’t inspect with probability 0.9 and inspect with probability 0.1), player II is indifferent between her strategies. Hence, she can *mix* them (that is, play them randomly) without losing payoff. The only case where, in turn, the original mixed strategy of player I is a best response is if player I is indifferent. According to the payoffs in Figure 3.2, this requires player II to comply with probability 0.8 and to cheat with probability 0.2. The expected payoffs to player I are then for “Don’t inspect” $0.8 \times 0 + 0.2 \times (-10) = -2$, and for “Inspect” $0.8 \times (-1) + 0.2 \times (-6) = -2$, so that player I is indeed indifferent, and his mixed strategy is a best response to the mixed strategy of player II.

This defines the only Nash equilibrium of the game. It uses mixed strategies and is therefore called a *mixed equilibrium*. The resulting expected payoffs are -2 for player I and 0 for player II.

The preceding analysis shows that the game in Figure 3.2 has a mixed equilibrium, where the players choose their pure strategies according to certain probabilities. These probabilities have several noteworthy features.

First, the equilibrium probability of 0.1 for inspecting makes player II indifferent between complying and cheating. As explained in Section 3.4 above, this requires payoffs to be *expected utilities*.

Secondly, mixing seems paradoxical when the player is indifferent in equilibrium. If player II, for example, can equally well comply or cheat, why should she gamble? In particular, she could comply and get payoff zero for certain, which is simpler and safer. The answer is that precisely because there is no incentive to choose one strategy over the other, a player can mix, and only in that case there can be an equilibrium. If player II would comply for certain, then the only optimal choice of player I is not to inspect, making the choice of complying not optimal, so this is not an equilibrium.

The least intuitive aspect of mixed equilibrium is that the probabilities depend on the *opponent's payoffs* and not on the player's own payoffs (as long as the qualitative preference structure, represented by the arrows, remains intact). For example, one would expect that raising the penalty -90 in Figure 3.2 for being caught lowers the probability of cheating in equilibrium. In fact, it does not. What does change is the probability of inspection, which is reduced until the inspectee is indifferent.

3.6 Bimatrix games

In the following, we discuss mixed equilibria for general games in strategic form. We always assume that each player has only a finite number of given pure strategies. In order to simplify notation, we consider the case of two players. Many definitions and results carry over without difficulty to the case of more than two players.

Recall that a game in strategic form is specified by a finite set of “pure” strategies for each player, and a payoff for each player for each *strategy profile*, which is a tuple of strategies, one for each player. The game is played by each player independently and simultaneously choosing one strategy, whereupon the players receive their respective payoffs.

For two players, a game in strategic form is also called a *bimatrix game* (A, B) . Here, A and B are two payoff matrices. By definition, they have equal dimensions, that is, they are both $m \times n$ matrices, having m rows and n columns. The m rows are the pure strategies i of player I and the n columns are the pure strategies j of player II. For a row i , where $1 \leq i \leq m$, and column j , where $1 \leq j \leq n$, the matrix entry of A is a_{ij} as payoff to player I, and the matrix entry of B is b_{ij} as payoff to player II.

Usually, we depict such a game as a table with m rows and n columns, so that each cell of the table corresponds to a pure strategy pair (i, j) , and we enter both payoffs a_{ij} and b_{ij} in that cell, a_{ij} in the lower-left corner, preferably written in red if we have colours at hand, and b_{ij} in the upper-right corner of the cell, displayed in blue. The “red” numbers are then the entries of the matrix A , the “blue” numbers those of the matrix B . It does not matter if we take two matrices A and B , or a single table where each cell has two entries (the respective components of A and B).

A *mixed strategy* is a randomised strategy of a player. It is defined as a probability distribution on the set of pure strategies of that player. This is played as an “active randomisation”: Using a lottery device with the given probabilities, the player picks each pure strategy according to its probability. When a player plays according to a mixed

strategy, the other player is not supposed to know the outcome of the lottery. Rather, it is assumed that the opponent knows that the strategy chosen by the player is a random event, and bases his or her decision on the resulting distribution of payoffs. The payoffs are then “weighted with their probabilities” to determine the *expected payoff*, which represents the player’s preference, as explained in Section 3.4.

A *pure strategy* is a *special mixed strategy*. Namely, consider a pure strategy i of player I. Then the mixed strategy x that selects i with probability one and any other pure strategy with probability zero is effectively the same as the pure strategy i , because x chooses i with certainty. The resulting expected payoff is the same as the pure strategy payoff, because any unplayed strategy has probability zero and hence does not affect the expected payoff, and the pure strategy i is weighted with probability one.

3.7 Matrix notation for expected payoffs

Unless specified otherwise, we assume that in the two-player game under consideration, player I has m strategies and player II has n strategies. The pure strategies of player I, which are the m rows of the bimatrix game, are denoted by $i = 1, \dots, m$, and the pure strategies of player II, which are the n columns of the bimatrix game, are denoted by $j = 1, \dots, n$.

A mixed strategy is determined by the probabilities that it assigns to the player’s pure strategies. For player I, a mixed strategy x can therefore be identified with the m -tuple of probabilities (x_1, x_2, \dots, x_m) that it assigns to the pure strategies $1, 2, \dots, m$ of player I. We can therefore consider x as an element of m -space (written \mathbb{R}^m). We assume that the vector x with m components is a *row vector*, that is, a $1 \times m$ matrix with a single row and m columns. This will allow us to write expected payoffs in a short way.

A mixed strategy y of player II is an n -tuple of probabilities y_j for playing the pure strategies $j = 1, \dots, n$. That is, y is an element of \mathbb{R}^n . We write y as a *column vector*, as $(y_1, y_2, \dots, y_n)^\top$, that is, the row vector (y_1, y_2, \dots, y_n) *transposed*. Transposition in general applies to any matrix. The transpose B^\top of the payoff matrix B , for example, is the $n \times m$ matrix where the entry in row j and column i is b_{ij} , because transposition means exchanging rows and columns. A column vector with n components is therefore considered as an $n \times 1$ matrix; transposition gives a row vector, a $1 \times n$ matrix.

Normally, all vectors are considered as column vectors, so \mathbb{R}^n is equal to $\mathbb{R}^{n \times 1}$, the set of all $n \times 1$ matrices with n rows and one column. We have made an exception in defining a mixed strategy x of player I as a row vector. Whether we mean row or column vectors will be clear from the context.

Suppose that player I uses the mixed strategy x and that player II uses the mixed strategy y . With these conventions, we can now succinctly express the expected payoff to player I as xAy , and the expected payoff to player II as xBy .

In order to see this, recall that the *matrix product* CD of two matrices C and D is defined when the number of columns of C is equal to the number of rows of D . That is, C is a $p \times q$ matrix, and D is a $q \times r$ matrix. The product CD is then a $p \times r$ matrix with

entry $\sum_{k=1}^q c_{ik}d_{kj}$ in row i and column j , where c_{ik} and d_{kj} are the respective entries of C and D . Matrix multiplication is associative, that is, for another $r \times s$ matrix E the matrix product CDE is a $p \times s$ matrix, which can be computed either as $(CD)E$ or as $C(DE)$.

For mixed strategies x and y , we read xAy and xB_y as matrix products. This works because x , considered as a matrix, is of dimension $1 \times m$, both A and B are of dimension $m \times n$, and y is of dimension $n \times 1$. The result is a 1×1 matrix, that is, a single real number.

It is best to think of xAy being computed as $x(Ay)$, that is, as the product of a row vector x that has m components with a column vector Ay that has m components. (The matrix product of two such vectors is also known as the *scalar product* of these two vectors.) The column vector Ay has m rows. We denote the entry of Ay in row i by $(Ay)_i$ for each row i . It is given by

$$(Ay)_i = \sum_{j=1}^n a_{ij}y_j \quad \text{for } 1 \leq i \leq m. \quad (3.1)$$

That is, the entries a_{ij} of row i of player I's payoff matrix A are multiplied with the probabilities y_j of their columns, so $(Ay)_i$ is the *expected payoff to player I when playing row i* . One can also think of y_j as a linear coefficient of the j th column of the matrix A . That is, Ay is the linear combination of the column vectors of A , each multiplied with its probability under y . This linear combination Ay is a vector of expected payoffs, with one expected payoff $(Ay)_i$ for each row i .

Furthermore, xAy is the expected payoff to player I when the players use x and y , because

$$xAy = \sum_{i=1}^m x_i(Ay)_i = \sum_{i=1}^m x_i \sum_{j=1}^n a_{ij}y_j = \sum_{i=1}^m \sum_{j=1}^n (x_iy_j) a_{ij}. \quad (3.2)$$

Because the players choose their pure strategies i and j independently, the probability that they choose the pure strategy pair (i, j) is the product x_iy_j of these probabilities, which is the coefficient of the payoff a_{ij} in (3.2).

Analogously, xB_y is the expected payoff to player II when the players use the mixed strategies x and y . Here, it is best to read this as $(xB)y$. The vector xB , as the product of a $1 \times m$ with an $m \times n$ matrix, is a $1 \times n$ matrix, that is, a row vector. Each column of that row corresponds to a strategy j of player II, for $1 \leq j \leq n$. We denote the respective column entry by $(xB)_j$. It is given by $\sum_{i=1}^m x_i b_{ij}$, which is the scalar product of x with the j th column of B . That is, $(xB)_j$ is the expected payoff to player II when player I plays x and player II plays the pure strategy j . If these numbers are multiplied with the column probabilities y_j and added up, then the result is the expected payoff to player II, which in analogy to (3.2) is given by

$$(xB)y = \sum_{j=1}^n (xB)_j y_j = \sum_{j=1}^n \left(\sum_{i=1}^m x_i b_{ij} \right) y_j = \sum_{j=1}^n \sum_{i=1}^m (x_iy_j) b_{ij}. \quad (3.3)$$

3.8 Convex combinations and mixed strategy sets

It is useful to regard mixed strategy vectors as geometric objects. A mixed strategy x of player I assigns probabilities x_i to the pure strategies i . The pure strategies, in turn, are special mixed strategies, namely the unit vectors in \mathbb{R}^m , for example $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$ if $m = 3$. The mixed strategy (x_1, x_2, x_3) is then a linear combination of the pure strategies, namely $x_1 \cdot (1, 0, 0) + x_2 \cdot (0, 1, 0) + x_3 \cdot (0, 0, 1)$, where the linear coefficients are just the probabilities. Such a linear combination is called a *convex* combination because the coefficients sum to one and are non-negative.




Figure 3.3 The line through the points x and y is given by the points $x + p(y - x)$ where $p \in \mathbb{R}$. Examples are point a for $p = 0.6$, point b for $p = 1.5$, and point c when $p = -0.4$. The line segment connecting x and y results when p is restricted to $0 \leq p \leq 1$.

Figure 3.3 shows two points x and y , here in the plane, but the picture may also be regarded as a suitable view of the situation in a higher-dimensional space. The line that goes through the points x and y is obtained by adding to the point x , regarded as a vector, any multiple of the difference $y - x$. The resulting vector $x + p \cdot (y - x)$, for $p \in \mathbb{R}$, gives x when $p = 0$, and y when $p = 1$. Figure 3.3 gives some examples a , b , c of other points. When $0 \leq p \leq 1$, as for point a , the resulting points give the *line segment* joining x and y . If $p > 1$, then one obtains points on the line through x and y on the other side of y relative to x , like the point b in Figure 3.3. For $p < 0$, the corresponding point, like c in Figure 3.3, is on that line but on the other side of x relative to y .

The expression $x + p(y - x)$ can be re-written as $(1 - p)x + py$, where the given points x and y appear only once. This expression (with $1 - p$ as the coefficient of the first vector and p of the second) shows how the line segment joining x to y corresponds to the real interval $[0, 1]$ for the possible values of p , with the endpoints 0 and 1 of the interval corresponding to the endpoints x and y , respectively, of the line segment.

In general, a *convex combination* of points z_1, z_2, \dots, z_k in some space is given as any linear combination $p_1 \cdot z_1 + p_2 \cdot z_2 + \dots + p_k \cdot z_k$ where the linear coefficients p_1, \dots, p_k are non-negative and sum to one. The previously discussed case corresponds to $z_1 = x$, $z_2 = y$, $p_1 = 1 - p$, and $p_2 = p \in [0, 1]$.

A set of points is called *convex* if it contains with any points z_1, z_2, \dots, z_k also every convex combination of these points. Equivalently, one can show that a set is convex if it

contains with any two points also the line segment joining these two points; one can then obtain combinations of k points for $k > 2$ by iterating convex combinations of only two points.

The coefficients in a convex combination can also be regarded as probabilities, and conversely, a probability distribution on a finite set can be seen as a convex combination of the unit vectors.

In a two-player game with m pure strategies for player I and n pure strategies for player II, we denote the sets of mixed strategies of the two players by X and Y , respectively:

$$\begin{aligned} X &= \{(x_1, \dots, x_m) \mid x_i \geq 0 \text{ for } 1 \leq i \leq m, \sum_{i=1}^m x_i = 1\}, \\ Y &= \{(y_1, \dots, y_n)^\top \mid y_j \geq 0 \text{ for } 1 \leq j \leq n, \sum_{j=1}^n y_j = 1\}. \end{aligned} \quad (3.4)$$

For consistency with Section 3.7, we assume that X contains row vectors and Y column vectors, but this is not an important concern.




Figure 3.4 Examples of player I's mixed strategy set X when $m = 2$ (left) and $m = 3$ (right), as the set of convex combinations of the unit vectors.

Examples of X are shown in Figure 3.4. When $m = 2$, then X is just the line segment joining $(1,0)$ to $(0,1)$. If $m = 3$, then X is a triangle, given as the set of convex combinations of the unit vectors, which are the vertices of the triangle.

It is easily verified that in general X and Y are convex sets.

3.9 The best response condition

A mixed strategy equilibrium is a profile of mixed strategies such that no player can improve his expected payoff by unilaterally changing his own strategy. In a two-player game, an equilibrium is a pair (x,y) of mixed strategies such that x is a best response to y and vice versa. That is, player I cannot get a better expected payoff than xAy by

choosing any other strategy than x , and player II cannot improve her expected payoff xBy by changing y .

It does not seem easy to decide if x is a best response to y among all possible mixed strategies, that is, if x maximises xAy for all x in X , because X is an infinite set. However, the following theorem, known as the *best response condition*, shows how to recognise this. This theorem is not difficult but it is important to understand. We discuss it afterwards.

Theorem 3.1 (Best response condition) *Let x and y be mixed strategies of player I and II, respectively. Then x is a best response to y if and only if for all pure strategies i of player I,*

$$x_i > 0 \implies (Ay)_i = \max\{(Ay)_k \mid 1 \leq k \leq m\}. \quad (3.5)$$

Proof. Recall that $(Ay)_i$ is the i th component of Ay , which is the expected payoff to player I when playing row i , according to (3.1). Let $u = \max\{(Ay)_k \mid 1 \leq k \leq m\}$, which is the maximum of these expected payoffs for the *pure* strategies of player I. Then

$$\begin{aligned} xAy &= \sum_{i=1}^m x_i (Ay)_i = \sum_{i=1}^m x_i (u - (u - (Ay)_i)) = \sum_{i=1}^m x_i u - \sum_{i=1}^m x_i (u - (Ay)_i) \\ &= u - \sum_{i=1}^m x_i (u - (Ay)_i). \end{aligned} \quad (3.6)$$

Because, for any pure strategy i , both x_i and the difference of the maximum payoff u and the payoff $(Ay)_i$ for row i is non-negative, the sum $\sum_{i=1}^m x_i (u - (Ay)_i)$ is also non-negative, so that $xAy \leq u$. The expected payoff xAy achieves the maximum u if and only if that sum is zero, that is, if $x_i > 0$ implies $(Ay)_i = u$, as claimed. \square

Consider the phrase “ x is a best response to y ” in the preceding theorem. This means that among all mixed strategies in X of player I, x gives maximum expected payoff to player I. However, the pure best responses to y in (3.5) only deal with the pure strategies of player I. Each such pure strategy corresponds to a row i of the payoff matrix. In that row, the payoffs a_{ij} are multiplied with the column probabilities y_j , and the sum over all columns gives the expected payoff $(Ay)_i$ for the pure strategy i according to (3.1). This pure strategy is a best response if and only if no other row gives a higher payoff.

The first point of the theorem is that the condition whether a pure strategy is a best response is very easy to check, as one only has to compute the m expected payoffs $(Ay)_i$ for $i = 1, \dots, m$. For example, if player I has three pure strategies ($m = 3$), and the expected payoffs in (3.1) are $(Ay)_1 = 4$, $(Ay)_2 = 4$, and $(Ay)_3 = 3$, then only the first two strategies are pure best responses. If these expected payoffs are 3, 5, and 3, then only the second strategy is a best response. Clearly, at least one pure best response exists, because the numbers $(Ay)_k$ in (3.5) have their maximum u for at least one k . The theorem states that only pure best responses i may have positive probability x_i if x is to be a best response to y .

A second consequence of Theorem 3.1, used also in its proof, is that a mixed strategy can never give a higher payoff than the best pure strategy. This is intuitive because “mixing” amounts to *averaging*, which is an average weighted with the probabilities, in

the way that the overall expected payoff xAy in (3.6) is obtained from those in (3.1) by multiplying (weighting) each row i with weight x_i and summing over all rows i , as shown in (3.2). Consequently, any pure best response i to y is also a mixed best response, so the maximum of xAy for $x \in X$ is the same as when x is restricted to the unit vectors in \mathbb{R}^m that represent the pure strategies of player I.

\Rightarrow Try now Exercise 3.1 on page 101, with the help of Theorem 3.1. This will help you appreciate methods for finding mixed equilibria that you will learn in later sections. You can also answer Exercise 3.2 on page 101, which concerns a game with three players.

3.10 Existence of mixed equilibria

In this section, we give the original proof of John Nash from 1951 that shows that any game with a finite number of players, and finitely many strategies per player, has a mixed equilibrium. This proof uses the following theorem about continuous functions. The theorem concerns compact sets; here, a set is compact if it is closed (containing all points near the set) and bounded.

Theorem 3.2 (Brouwer's fixed point theorem) *Let S be a subset of \mathbb{R}^N that is convex and compact, and let f be a continuous function from S to S . Then f has at least one fixed point, that is, a point s in S so that $f(s) = s$.*

We do not prove Theorem 3.2, but instead demonstrate its assumptions with examples where not all assumptions hold and the conclusion fails. Consider the function $f: \mathbb{R} \rightarrow \mathbb{R}$ defined by $f(x) = x + 1$. This is clearly a continuous function, and the set \mathbb{R} of all real numbers is closed and convex, but f has no fixed point. Some assumption of the fixed point theorem must be violated, and in this case it is compactness, because the set \mathbb{R} is not bounded. Consider another function $f: [0, 1] \rightarrow [0, 1]$, given by $f(x) = x^2$. The fixed points of this function are 0 and 1. If we consider the function $f(x) = x^2$ as a function on the open interval $(0, 1)$ (which is $[0, 1]$ without its endpoints), then this function has no longer any fixed points. In this case, the missing condition is that the function is not defined on a closed set, which is therefore not compact. Another function is $f(x) = 1 - x$ for $x \in \{0, 1\}$, where the domain of this function has just two elements, so this is a compact set. This function has no fixed points which is possible because its domain is not convex. Finally, the function on $[0, 1]$ defined by $f(x) = 1$ for $0 \leq x \leq 1/2$ and by $f(x) = 0$ for $1/2 < x \leq 1$ has no fixed point, which is possible in this case because the function is not continuous. These simple examples demonstrate why the assumptions of Theorem 3.2 are necessary.

Theorem 3.3 (Nash [1951]) *Every finite game has at least one equilibrium in mixed strategies.*

Proof. We will give the proof for two players, to simplify notation. It extends in the same manner to any finite number of players. The set S that is used in the present context is the product of the sets of mixed strategies of the players. Let X and Y be the sets of mixed strategies of player I and player II as in (3.4), and let $S = X \times Y$.

Then the function $f: S \rightarrow S$ that we are going to construct maps a pair of mixed strategies (x, y) to another pair $f(x, y) = (\bar{x}, \bar{y})$. Intuitively, a mixed strategy probability x_i (of player I, similarly y_j of player II) is changed to \bar{x}_i , such that it will decrease if the pure strategy i does worse than the average of all pure strategies. In equilibrium, all pure strategies of a player that have positive probability do equally well, so no sub-optimal pure strategy can have a probability that is reduced further. This means that the mixed strategies do not change, so this is indeed equivalent to the fixed point property $(x, y) = (\bar{x}, \bar{y}) = f(x, y)$.

In order to define f as described, consider the following functions $\chi: X \times Y \rightarrow \mathbb{R}^m$ and $\psi: X \times Y \rightarrow \mathbb{R}^n$ (we do not worry whether these vectors are row or column vectors; it suffices that \mathbb{R}^m contains m -tuples of real numbers, and similarly \mathbb{R}^n contains n -tuples). For each pure strategy i of player I, let $\chi_i(x, y)$ be the i th component of $\chi(x, y)$, and for each pure strategy j of player II, let $\psi_j(x, y)$ be the j th component of $\psi(x, y)$. The functions χ and ψ are defined by

$$\chi_i(x, y) = \max\{0, (Ay)_i - xAy\}, \quad \psi_j(x, y) = \max\{0, (xB)_j - xBy\},$$

for $1 \leq i \leq m$ and $1 \leq j \leq n$. Recall that $(Ay)_i$ is the expected payoff to player I against y when he uses the pure strategy i , and that $(xB)_j$ is the expected payoff to player II against x when she uses the pure strategy j . Moreover, xAy and xBy are the overall expected payoffs to player I and player II, respectively. So the difference $(Ay)_i - xAy$ is positive if the pure strategy i gives more than the average xAy against y , zero if it gives the same payoff, and negative if it gives less. The term $\chi_i(x, y)$ is this difference, except that it is replaced by zero if the difference is negative. The term $\psi_j(x, y)$ is defined analogously. Thus, $\chi(x, y)$ is a non-negative vector in \mathbb{R}^m , and $\psi(x, y)$ is a non-negative vector in \mathbb{R}^n . The functions χ and ψ are continuous.

The pair of vectors (x, y) is now changed by replacing x by $x + \chi(x, y)$ in order to get \bar{x} , and y by $y + \psi(x, y)$ to get \bar{y} . Both sums are non-negative. The only problem is that in general these new vectors are no longer probabilities because their components do not sum to one. For that purpose, they are “re-normalised” by the following functions $r: \mathbb{R}^m \rightarrow \mathbb{R}^m$ and $s: \mathbb{R}^n \rightarrow \mathbb{R}^n$, defined by their components r_i and s_j , that is, $r(x) = (r_1(x), \dots, r_m(x))$, and $s(y) = (s_1(y), \dots, s_n(y))$:

$$r_i(x_1, \dots, x_m) = \frac{x_i}{\sum_{k=1}^m x_k}, \quad s_j(y_1, \dots, y_n) = \frac{y_j}{\sum_{k=1}^n y_k}.$$

Clearly, if $x_i \geq 0$ for $1 \leq i \leq m$ and $\sum_{k=1}^m x_k > 0$, then $r(x)$ is defined and is a probability distribution, that is, an element of the mixed strategy set X . Analogously, $s(y) \in Y$.

The function $f: S \rightarrow S$ is now defined by

$$f(x, y) = (r(x + \chi(x, y)), s(y + \psi(x, y))).$$

What is a fixed point (x, y) of that function, so that $f(x, y) = (x, y)$? Among all strategies i where $x_i > 0$, consider the smallest pure strategy payoff $(Ay)_i$ against y , that is, $(Ay)_i = \min\{(Ay)_k \mid x_k > 0\}$. Then $(Ay)_i \leq xAy$, which is proved analogously to (3.6), so the component $\chi_i(x, y)$ of $\chi(x, y)$ is zero. This means that the respective term $x_i + \chi_i(x, y)$

is equal to x_i . Conversely, consider any other pure strategy l of player I that gets the maximum payoff $(Ay)_l = \max_k (Ay)_k$. If that payoff is better than the average xAy , then clearly $\chi_l(x, y) > 0$, so that $x_l + \chi_l(x, y) > x_l$. Because $\chi_k(x, y) \geq 0$ for all k , this implies $\sum_{k=1}^m (x_k + \chi_k(x, y)) > 1$, which is the denominator in the re-normalisation with r in $r(x + \chi(x, y))$. This re-normalisation will now *decrease* the value of x_i for the pure strategy i with $(Ay)_i \leq xAy$, so the relative weight x_i of the pure strategy i decreases. (It would stay unchanged, that is, not decrease, if $x_i = 0$.) But $(Ay)_l > xAy$ can only occur if there is some sub-optimal strategy i (with $(Ay)_i \leq xAy < (Ay)_l$) that has positive probability x_i . In that case, $r(x + \chi(x, y))$ is *not equal* to x , so that $f(x, y) \neq (x, y)$.

Analogously, if $\psi(x, y)$ has some component that is positive, then the respective pure strategy of player II has a better payoff than xB_y , so $y \neq s(y + \psi(x, y))$ and (x, y) is not a fixed point of f . In that case, y is also not a best response to x .

Hence, the function f has a fixed point (x, y) if and only if both $\chi(x, y)$ and $\psi(x, y)$ are zero in all components. But that means that xAy is the maximum possible payoff $\max_i (Ay)_i$ against y , and xB_y is the maximum possible payoff $\max_j (xB)_j$ against x , that is, x and y are mutual best responses. The fixed points (x, y) are therefore exactly the Nash equilibria of the game. \square

3.11 Finding mixed equilibria

How can we find all Nash equilibria of a two-player game in strategic form? It is easy to determine the Nash equilibria where both players use a pure strategy, because these are the cells of the payoff table where both payoffs are best-response payoffs, shown by a box around that payoff. We now describe how to find mixed strategy equilibria, using the best response condition Theorem 3.1.

We first consider 2×2 games. Our first example above, the inspection game in Figure 3.2, has no pure-strategy equilibrium. There we have determined the mixed strategy probabilities of a player so as to make the other player indifferent between his or her pure strategies, because only then that player will mix between these strategies. This is a consequence of Theorem 3.1: Only pure strategies that get maximum, and hence equal, expected payoff can be played with positive probability in equilibrium.

A mixed strategy equilibrium can exist also in 2×2 games that have pure-strategy equilibria. As an example, consider the battle of sexes game, shown on the left in Figure 3.5, which has the pure strategy equilibria (C, c) and (S, s) . As indicated by the boxes, the best response of the other player depends on the player's own strategy: If player I plays C then player II's best response is c , and the best response to S is s . Suppose now that player I plays a mixed strategy, playing C with probability $1 - p$ and S with probability p . Clearly, when p is close to zero so that player I almost certainly chooses C , the best response will still be c , whereas if p is close to one the best response will be s . Consequently, there is some probability so that player II is indifferent. This probability p is found as follows. Given p , which defines the mixed strategy of player I, the expected payoff to player II when she plays c is $2(1 - p)$, and when she plays s that expected payoff

		II	
		<i>c</i>	<i>s</i>
I	<i>C</i>	<div>2</div> 0	0
	<i>S</i>	0 <div>1</div>	0

		II	
		<i>c</i>	<i>s</i>
I	<i>C</i>	<div>2</div> 0	0
	<i>S</i>	0	<div>1</div>
	<i>B</i>	<div>4</div> 1	0

Figure 3.5 Battle of sexes game (left) and a 3×2 game (right) which is the same game with an extra strategy B for player I.

is simply p . She is indifferent between c and s if and only if these expected payoffs are equal, that is, $2(1 - p) = p$ or $2 = 3p$, that is, $p = 2/3$. This mixed strategy of player I, where he plays C and S with probabilities $1/3$ and $2/3$, respectively, gives player II for both her strategies, the two columns, the same expected payoff $2/3$. Then player II can mix between c and s , and a similar calculation shows that player I is indifferent between his two strategies if player II uses the mixed strategy $(2/3, 1/3)$, which is the vector of probabilities for her two strategies c, s . Then and only then player I gets the same expected payoff for both rows, which is $2/3$. This describes the mixed equilibrium of the game, which we write as $((1/3, 2/3), (2/3, 1/3))$, as a pair (x, y) of mixed strategies x and y , which are probability vectors.

In short, the rule to find a mixed strategy equilibrium in a 2×2 game is to make the other player indifferent, because only in that case can the other player mix. This indifference means the expected payoffs for the two opponent strategies have to be equal, and the resulting equation determines the player's own probability.

The mixed strategy probabilities in the battle of sexes game can be seen relatively quickly by looking at the payoff matrices: For example, player II must give twice as much weight to strategy c compared to s , because player I's strategy C only gets one payoff unit from c whereas his strategy S gets two payoff units when player II chooses s . Because the strategy pairs (C, s) and (S, c) give payoff zero, the two rows C and S give the same expected payoff only when c has probability $2/3$ and s has probability $1/3$.

We now explain a quick method, the “difference trick”, to find the probabilities in a mixed strategy equilibrium of a general 2×2 game. Consider the game on the left in Figure 3.6. When player II plays T , then l is a best response, and when player II plays B , then r is a best response. Consequently, there must be a way to mix T and B so that player II is indifferent between l and r . We consider now the difference Δ in payoffs to the other player for both rows: When player I plays T , then the difference between the two payoffs to player II is $\Delta = 2 - 1 = 1$, and when player I plays B , then that difference, in absolute value, is $\Delta = 7 - 3 = 4$, as shown on the side of the game. (Note that these differences, when considered as differences between the payoffs for l




Figure 3.6 The “difference trick” to find equilibrium mixed strategy probabilities in a 2×2 game. The left figure shows the game and the difference in payoffs to the other player for each strategy. As shown on the right, these differences are assigned to the respective other own strategy and are re-normalised to become probabilities. The fractions $11/5$ and $38/10$ are the resulting equal expected payoffs to player II and player I, respectively.

versus r , have opposite sign because l is preferred to r against T , and the other way around against B ; otherwise, player II would always prefer the same strategy and could not be made indifferent.) Now the payoff difference Δ is assigned as a probability weight to the respective *other* strategy of the row player, meaning T is given weight 4 (which is the Δ computed for B), and B is given weight 1 (which is the Δ for T). The probabilities for T and B are then chosen proportional to these weights, so one has to divide each weight by 5 (which is the sum of the weights) in order to obtain probabilities. This is shown on the right in Figure 3.6. The expected payoffs to player II are also shown there, at the bottom, and are for l given by $(4 \cdot 2 + 1 \cdot 3)/5 = 11/5$, and for r by $(4 \cdot 1 + 1 \cdot 7)/5 = 11/5$, so they are indeed equal as claimed. Similarly, the two payoff differences for player I in columns l and r are 3 and 7, respectively, so l and r should be played with probabilities that are proportional to 7 and 3, respectively. With the resulting probabilities $7/10$ and $3/10$, the two rows T and B get the same expected payoff $38/10$.

\Rightarrow A proof that this “difference trick” always works is the subject of Exercise 3.3 on page 102.

Next, we consider games where one player has two and the other more than two strategies. The 3×2 game on the right in Figure 3.5 is like the battle of sexes game, except that player I has an additional strategy B . Essentially, such a game can be analysed by considering the 2×2 games obtained by restricting both players to two strategies only, and checking if the resulting equilibria carry over to the whole game. First, the pure-strategy equilibrium (C, c) of the original battle of sexes game is also an equilibrium of the larger game, but (S, s) is not because S is no longer a best response to s because B gives a larger payoff to player I. Second, consider the mixed strategy equilibrium of the smaller game where player I chooses C and S with probabilities $1/3$ and $2/3$ (and B with probability zero), so that player II is indifferent between c and s . Hence, the mixed strategy $(2/3, 1/3)$ of player II is still a best response, against which C and S both give the same expected payoff $2/3$. However, this is not enough to guarantee an equilibrium,

because C and S have to be best responses, that is, their payoff must be at least as large as that for the additional strategy B . That payoff is $(-1) \cdot 2/3 + 3 \cdot 1/3 = 1/3$, so it is indeed not larger than the payoff for C and S . In other words, the mixed equilibrium of the smaller game is also a mixed equilibrium of the larger game, given by the pair of mixed strategies $((1/3, 2/3, 0), (2/3, 1/3))$.

Other “restricted” 2×2 games are obtained by letting player I play only C and B , or only S and B . In the first case, the difference trick gives the mixed strategy $(3/5, 2/5)$ of player II (for playing c, s) so that player I is indifferent between C and B , where both strategies receive expected payoff $3/5$. However, the expected payoff to the third strategy S is then $4/5$, which is higher, so that C and B are not best responses, which means we cannot have an equilibrium where only C and B are played with positive probability by player I. Another reason why no equilibrium strategy of player II mixes C and B is that against both C and B , player II’s best response is always c , so that player I could not make player II indifferent by playing in that way.

Finally, consider S and B as pure strategies that player I mixes in a possible equilibrium. This requires, via the difference trick, the mixed strategy $(0, 3/4, 1/4)$ of player I (as probability vector for playing his three pure strategies C, S, B) so that player II is indifferent between c and s , both strategies receiving payoff 1 in that case. Then if player II uses the mixed strategy $(1/2, 1/2)$ (in order to make player I indifferent between S and B), the expected payoffs for the three rows C, S, B are $1/2, 1$, and 1 , respectively, so that indeed player I only uses best responses with positive probability, which gives a third Nash equilibrium of the game.

Why is there no mixed equilibrium of the 3×2 game in Figure 3.5 where player I mixes between his three pure strategies? The reason is that player II has only a single probability at her disposal (which determines the complementary probability for the other pure strategy), which does not give enough freedom to satisfy two equations, namely indifference between C and S as well as between S and B (the third indifference between C and B would then hold automatically). We have already computed the probabilities $(2/3, 1/3)$ for c, s that are needed to give the same expected payoff for C and B , against which the payoff for row B is different, namely $1/3$. This alone suffices to show that it is not possible to make player I indifferent between all three pure strategies, so there is no equilibrium where player I mixes between all of them. In certain games, which are called “degenerate” and which are treated in Section 3.13, it is indeed the case that player II, say, mixes only between two pure strategies, but where “by accident” three strategies of player I have the same optimal payoff. This is normally not the case, and leads to complications when determining equilibria, which will be discussed in the later section on degenerate games.

\Rightarrow You can now attempt Exercise 3.1 on page 101, if you have not done this earlier. Compare your solution with the approach for the 3×2 game in Figure 3.5.

3.12 The upper envelope method

The topic of this section is the “upper-envelope” diagram that simplifies finding mixed equilibria, in particular of games where one player has only two strategies.

The left graph in Figure 3.7 shows the upper-envelope diagram for the left game in Figure 3.5. This diagram is a plot of the expected payoff of one player against the mixed strategy of the other player. We use this typically when that mixed strategy is determined by a single probability, that is, we consider a player who has only two pure strategies, in this example player II. The horizontal axis represents the probability for playing the second pure strategy s of player II, so this is $q = \text{prob}(s)$. In the vertical direction, we plot the resulting expected payoffs to the other player, here player I, for his pure strategies. In this game, this is the expected payoff $1 - q$ when he plays row C , and the expected payoff $2q$ when he plays row S . The plots are lines because expected payoffs are linear functions of the mixed strategy probabilities.

These lines are particularly easy to obtain graphically. To see this, consider row C of player I. When player II plays her left strategy c (where $q = 0$), row C gives payoff 1 (according to the table of payoffs to player I). When player II plays her right strategy s (where $q = 1$), row C gives payoff 0. The probabilities $q = 0$ or $q = 1$ and corresponding payoffs define the two endpoints with co-ordinates $(0, 1)$ and $(1, 0)$ of a line that describes the expected payoff $1 - q$ as a function of q , which consists of the points $(q, 1 - q)$ for $q \in [0, 1]$. As explained in Section 3.8 above, this line is just the set of convex combinations $(1 - q) \cdot (0, 1) + q \cdot (1, 0)$ of the two endpoints. The first component of such a convex combination is always q and second component $(1 - q) \cdot 1 + q \cdot 0$ is the expected payoff. For row S , this expected payoff is $(1 - q) \cdot 0 + q \cdot 2$, so the line of expected payoffs for S connects the left endpoint $(0, 0)$ with the right endpoint $(1, 2)$.




Figure 3.7 Upper envelope of expected payoffs to player I for the two games in Figure 3.5.

The upper-envelope diagram can also be described as the “goalpost” method. First, identify a player who has only two pure strategies (here player II). Second, plot the

probability (here q) for the second of these pure strategies horizontally, along the interval $[0, 1]$. Third, erect a vertical line (a “goalpost”) at the left endpoint 0 and at the right endpoint 1 of that interval. Fourth, do the following for each pure strategy of the other player (which are here rows C and S for player I): Mark the payoff (as given in the game description) against the left strategy (of player II, that is, against $q = 0$) as a “height” on the left goalpost, and against the right strategy ($q = 1$) as a height on the right goalpost, and connect the two heights by a line, which gives the expected payoff as a function of q for $0 \leq q \leq 1$.

Note that for these goalposts it is very useful to consider player II’s mixed strategy as the vector of probabilities $(1 - q, q)$ (according to the second step above), because then the left goalpost corresponds to the left strategy of player II, and the right goalpost to the right strategy of player II. (This would not be the case if q was taken as the probability for the left strategy, writing player II’s mixed strategy as $(q, 1 - q)$, because then $q = 0$ would be the left goalpost but represent the right pure strategy of player II, creating an unnecessary source of confusion.)

In the left diagram in Figure 3.7, the two lines of expected payoffs for rows C and S obviously intersect exactly where player I is indifferent between these two rows, which happens when $q = 1/3$. However, the diagram not only gives information about this indifference, but about player I’s preference in general: Any pure strategy that is a *best response* of player I is obviously the *topmost* line segment of all the lines describing the expected payoffs. Here, row C is a best response whenever $0 \leq q \leq 1/3$, and row S is a best response whenever $1/3 \leq q \leq 1$, with indifference between C and S exactly when $q = 1/3$. The *upper envelope* of expected payoffs is the *maximum* of the lines that describe the expected payoffs for the different pure strategies, and it is indicated by the bold line segments in the diagram. Marking this upper envelope in bold is therefore the fifth step of the “goalpost” method.

The upper envelope is of particular use when a player has more than two pure strategies, because mixed equilibria of a 2×2 are anyhow quickly determined with the “difference trick”. Consider the upper-envelope diagram on the right of Figure 3.7, for the 3×2 game in Figure 3.5. The extra strategy B of player I gives the line connecting height -1 on the left goalpost to height 3 on the right goalpost, which shows that B is a best response for all q so that $1/2 \leq q \leq 1$, with indifference between S and B when $q = 1/2$. The upper envelope consists of three line segments for the three pure strategies C , S , and B of player I. Moreover, there are only two points where player I is indifferent between any two such strategies, namely between C and S when $q = 1/3$, and between S and B when $q = 1/2$. These two points, indicated by small circles on the upper envelope, give the only two mixed strategy probabilities of player II where player I can mix between two pure strategies (the small circle for $q = 0$ corresponds to the pure equilibrium strategy c). There is a third indifference between C and B , for $q = 2/5$, where the two lines for C and B intersect, but this intersection point is below the line for S and hence not on the upper envelope. Consequently, it does not have to be investigated as a possible mixed strategy equilibrium strategy where player I mixes between the two rows C and B because they are not best responses.

Given these candidates of mixtures for player I, we can quickly find if and how the two respective pure strategies can be mixed in order to make player II indifferent using the difference trick, as described in the previous section.

So the upper envelope restricts the pairs of strategies that need to be tested as possible strategies that are mixed in equilibrium. Exercise 3.4, for example, gives a 2×5 game, where the upper envelope may consist of up to five line segments (there may be fewer line segments if some pure strategy of player II is never a best response). Consequently, there are up to four points where these line segments join up, so that player II is indifferent between two strategies and can mix. Without the upper envelope, there would not be four but ten (which are the different ways to pick two out of five) pure strategy pairs to be checked as possible pure strategies that player II mixes in equilibrium, which would be much more laborious to investigate.

⇒ The familiar Exercise 3.1 on page 101 can also be solved with the upper envelope method. This method is of particular help for larger games as in Exercise 3.4 on page 102.

In principle, the upper-envelope diagram is also defined for the expected payoffs against a mixed strategy that mixes between more than two pure strategies. This diagram is much harder to draw and visualise in this case because it requires higher dimensions. We demonstrate this for the game on the right in Figure 3.5, although we do not suggest using this as a method to find all equilibria of such a game.

Consider the set of mixed strategies of player I for the 3×2 game of Figure 3.5, which is a triangle. Then the payoff to the other player II needs an extra “vertical” dimension. A perspective drawing of the resulting three-dimensional picture is shown in Figure 3.8, in several stages. First (top left), the mixed strategy triangle is shown with goalposts erected on its corners $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 0, 1)$. Next (top right), we mark the heights 2, 0, and 4 on these posts which are the payoffs to player II for her first strategy c . The resulting expected payoffs define a plane above the triangle through these height markings. The third picture (bottom left) does the same for the second strategy s of player II, which has payoffs 0, 1, and 1 as heights in the three corners of the triangle. The final picture (bottom right) gives the upper envelope as the maximum of the two expected-payoff planes (which in the lower dimension as in Figure 3.7 is drawn in bold). The small circles indicate the three mixed strategies of player I which are his strategies, respectively, in the three Nash equilibria $((1, 0, 0), (1, 0))$, $((1/3, 2/3, 0), (2/3, 1/3))$, and $((0, 3/4, 1/4), (1/2, 1/2))$. The first of these three equilibria is the pure strategy equilibrium (C, c) , and the last two involve mixing both columns c and s of player II. Hence, player II must be indifferent between both c and s , so player I must choose a mixed strategy where the planes of expected payoffs for these two strategies c and s of player II meet. (These two planes meet in a line, which is given by any convex combination of the two mixed strategies $(1/3, 2/3, 0)$ and $(0, 3/4, 1/4)$ of player I; however, player I must choose one of these two endpoints because he can only mix either between C and S or between S and B as described earlier, depending on the mixed strategy of player II.)

⇒ The simple Exercise 3.5 on page 103 helps to understand dominated strategies with the upper-envelope diagram.

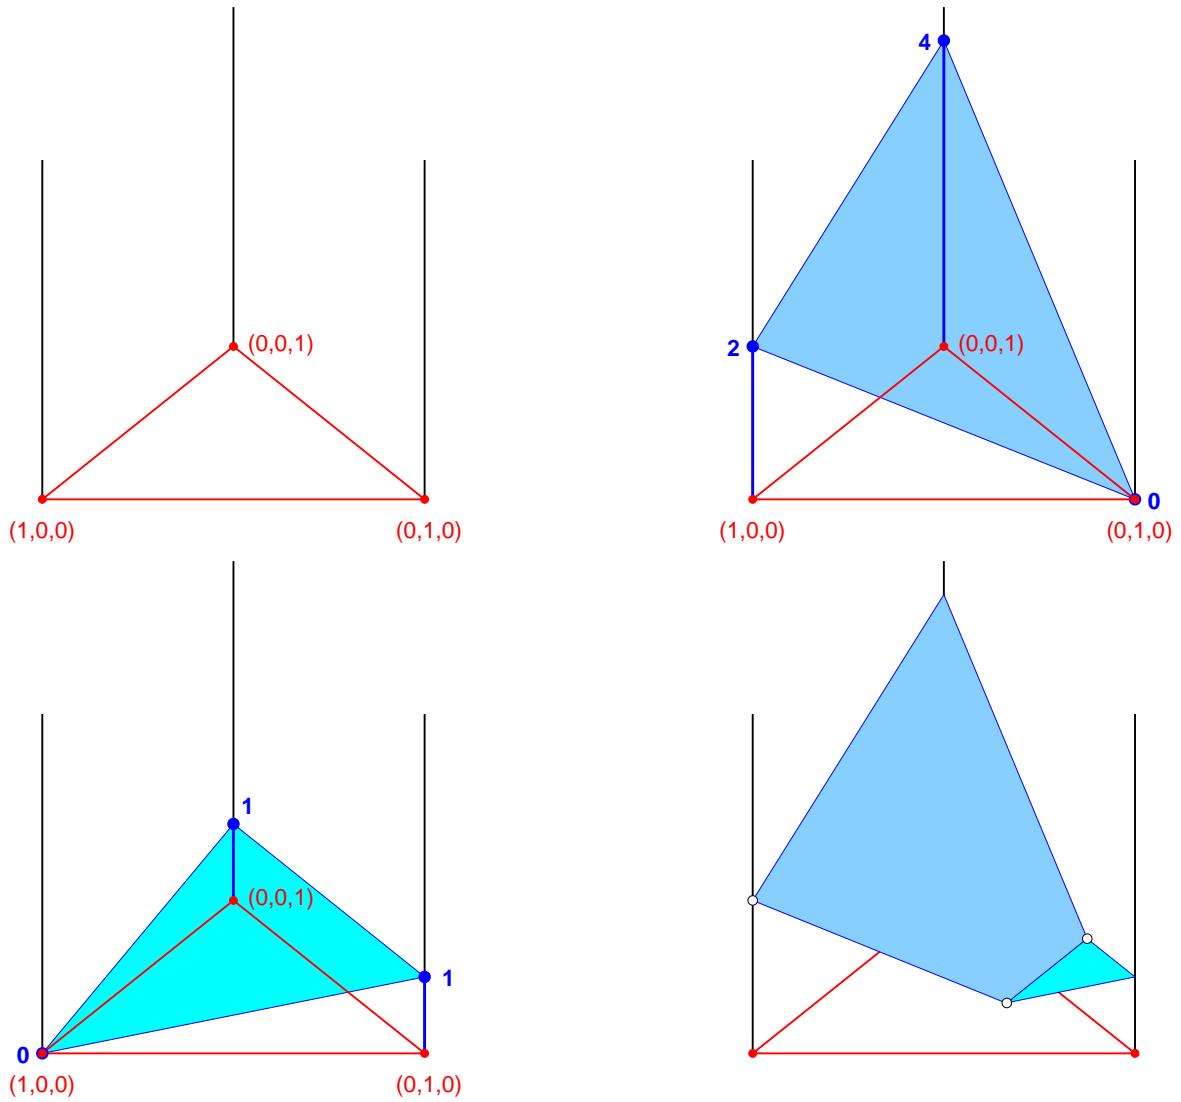


Figure 3.8 Expected payoffs to player II as a function of the mixed strategy (x_1, x_2, x_3) of player I for the 3×2 game in Figure 3.5. Top left: goalposts for the three pure strategies $C = (1, 0, 0)$, $S = (0, 1, 0)$, $B = (0, 0, 1)$; top right: plane of expected payoffs for the left column c , with payoffs 2, 0, 4 against C, S, B ; bottom left: plane of expected payoffs for the right column s , with payoffs 0, 1, 1; bottom right: *upper envelope* (maximum) of expected payoffs, with small white circles indicating equilibrium strategies, which match the right picture of Figure 3.7.

3.13 Degenerate games

In this section, we treat games that have to be analysed more carefully in order to find all their Nash equilibria. We completely analyse an example, and then give the definition of a *degenerate game* that applies to 2×2 games like in this example. We then discuss larger games, and give the general definition 3.5 when an $m \times n$ game is degenerate.

Figure 3.9 is the strategic form, with best response payoffs, of the “threat game” of Figure 2.3. It is similar to the battle of sexes game in that it has two pure Nash equilibria, here (T, l) and (B, r) . As in the battle of sexes game, we should expect an additional mixed strategy equilibrium, which indeed exists. Assume this equilibrium is given by the mixed strategy $(1 - p, p)$ for player I and $(1 - q, q)$ for player II. The difference trick gives $q = 1/2$ so that player I is indifferent between T and B and can mix, and for player I the difference trick gives $p = 0$ so that player II is indifferent between l and r .

In more detail, if player I plays T and B with probabilities $1 - p$ and p , then the expected payoff to player II is $3(1 - p)$ for her strategy l , and $3(1 - p) + 2p$ for r , so that she is indifferent between l and r only when $p = 0$. Indeed, because l and r give the same payoff when player I plays T , but the payoff for r is larger when player I plays B , any positive probability for B would make r the unique best response of player II, against which B is the unique best response of player I. Hence, the only Nash equilibrium where B is played with positive probability is the pure-strategy equilibrium (B, r) .

		II	
		l	r
I	T	<div style="border: 1px solid black; padding: 2px; display: inline-block;">3</div>	<div style="border: 1px solid black; padding: 2px; display: inline-block;">3</div>
	B	0	<div style="border: 1px solid black; padding: 2px; display: inline-block;">2</div>
		<div style="border: 1px solid black; padding: 2px; display: inline-block;">1</div>	<div style="border: 1px solid black; padding: 2px; display: inline-block;">2</div>

Figure 3.9 Strategic form of the threat game, which is a degenerate game.

Consequently, a mixed-strategy equilibrium of the game is $((1, 0), (1/2, 1/2))$. However, only player II uses a properly mixed strategy in this equilibrium, whereas the strategy of player I is the pure strategy T . The best response condition requires that player II is indifferent between her pure strategies l and r because both are played with positive probability, and this implies that T is played with probability one, as just described. For player I, however, the best response condition requires that T has maximal expected payoff, which does *not* have to be the same as the expected payoff for B because B does not have positive probability! That is to say, player I’s indifference between T and B , which we have used to determine player II’s mixed strategy $(1 - q, q)$ as $q = 1/2$, is too strong a requirement. All we need is that T is a best response to this mixed strategy, so the expected payoff 1 to player I when he plays T has to be *at least as large* as his expected payoff $2q$ when he plays B , that is, $1 \geq 2q$. In other words, the equation $1 = 2q$ has to be replaced by an *inequality*.

The inequality $1 \geq 2q$ is equivalent to $0 \leq q \leq 1/2$. With these possible values for q , we obtain an infinite set of equilibria $((1, 0), (1 - q, q))$. The two extreme cases for q , namely $q = 0$ and $q = 1/2$, give the pure-strategy equilibrium (T, l) and the mixed equilibrium $((1, 0), (1/2, 1/2))$ that we found earlier.

Note: An equilibrium where at least one player uses a mixed strategy that is not a pure strategy is always called a mixed equilibrium. Also, when we want to find “all

mixed equilibria” of a game, we usually include the pure-strategy equilibria among them (which are at any rate easy to find).

The described set of equilibria (where player II uses a mixed strategy probability q from some interval) has an intuition in terms of the possible “threat” in this game that player II plays l (with an undesirable outcome after player I has chosen B , in the game tree in Figure 2.3(a)). First, r weakly dominates l , so if there is any positive probability that player I ignores the threat and plays B , then the threat of playing l at all is not sustainable as a best response, so any equilibrium where player II does play l requires that player I’s probability p for playing B is zero. Second, if player I plays T , then player II is indifferent between l and r and she can in principle play anything, but in order to maintain the threat T must be a best response, which requires that the probability $1 - q$ for l is sufficiently high. Player I is indifferent between T and B (and therefore T still optimal) when $1 - q = 1/2$, but whenever $1 - q \geq 1/2$ the threat works as well. In other words, “threats” can work even when the threatened action with the undesirable outcome is uncertain, as long as its probability is sufficiently high.

The complications in this game arise because the game in Figure 3.9 is “degenerate” according to the following definition.

Definition 3.4 A 2×2 game is called *degenerate* if some player has a pure strategy with two pure best responses of the other player.

A degenerate game is complicated to analyse because one player can use a pure strategy (like T in the threat game) so that the other player can mix between her two best responses, but the mixed strategy probabilities are not constrained by the equation that the other player has to be indifferent between his strategies. Instead of that equation, it suffices to fulfil the *inequality* that the first player’s pure strategy is a best response. Note that the mentioned equation, as well as the inequality, is a consequence of the best response condition, Theorem 3.1.

\Rightarrow Answer Exercise 3.6 on page 103.




Figure 3.10 Degenerate 2×3 game and upper envelope of expected payoffs to player II.

Degeneracy occurs for larger games as well, which requires a more general definition. Consider the 2×3 game in Figure 3.10. Each pure strategy has a unique best response, so the condition in definition 3.4 does not hold. This game does not have a pure-strategy

equilibrium, so both players have to use mixed strategies. We apply the upper envelope method, plotting the payoff to player II as a function of the probability that player I plays B , shown on the right in Figure 3.10. Here, the lines of expected payoffs cross in a single point, and all three pure strategies of player II are best responses when player I uses the mixed strategy $(1/2, 1/2)$.

How should we analyse this game? Clearly, we only have an equilibrium if both players mix, which requires the mixed strategy $(1/2, 1/2)$ for player I. Then player I has to be made indifferent, and in order to achieve this, player II can mix between any of her three best responses l , c , and r . We now have the same problem as earlier in the threat game, namely too much freedom: Three probabilities for player II, call them y_l , y_c , and y_r , have to be chosen so that $y_l + y_c + y_r = 1$, and so that player I is indifferent between his pure strategies T and B , which gives the equation $3y_c + 2y_r = y_l + 2y_c$. These are three unknowns subject to two equations, with an underdetermined solution.

We give a method to find out all possible probabilities, which can be generalised to the situation that three probabilities are subject to only two linear equations, like here. First, sort the coefficients of the probabilities in the equation $3y_c + 2y_r = y_l + 2y_c$, which describes the indifference of the expected payoffs to the other player, so that each probability appears only once and so that its coefficient is *positive*, which here gives the equation $y_c + 2y_r = y_l$. Of the three probabilities, normally two appear on one side of the equation and one on the other side. The “extreme solutions” of this equation are obtained by setting either of the probabilities on the side of the equation with two probabilities to zero, here either $y_c = 0$ or $y_r = 0$. In the former case, this gives the solution $(2/3, 0, 1/3)$ for (y_l, y_c, y_r) , in the latter case the solution $(1/2, 1/2, 0)$. In general, we have $0 \leq y_c \leq 1/2$, which determines the remaining probabilities as $y_l = y_c + 2y_r = y_c + 2(1 - y_l - y_c) = 2 - 2y_l - y_c$ or $y_l = 2/3 - y_c/3$, and $y_r = 1 - y_l - y_c = 1 - (2/3 - y_c/3) - y_c = 1/3 - 2y_c/3$.

Another way to find the extreme solutions to the underdetermined mixed strategy probabilities of player II is to simply ignore one of the three best responses and apply the difference trick: Assume that player II uses only his best responses l and c . Then the difference trick, to make player I indifferent, gives probabilities $(1/2, 1/2, 0)$ for l, c, r . If player II uses only his best responses l and r , the difference trick gives the mixed strategy $(2/3, 0, 1/3)$. These are exactly the two mixed strategies just described. If player II uses only his best responses c and r , then the difference trick does not work because player I’s best response to both c and r is always T .

The following is a definition of degeneracy for general $m \times n$ games.

Definition 3.5 (Degenerate game) A two-player game is called *degenerate* if some player has a mixed strategy that assigns positive probability to exactly k pure strategies so that the other player has more than k pure best responses to that mixed strategy.

In a degenerate game, a mixed strategy with “too many” best responses creates the difficulty that we have described with the above examples. Consider, in contrast, the “normal” situation that the game is non-degenerate.

Proposition 3.6 *In any equilibrium of a two-player game that is not degenerate, both players use mixed strategies that mix the same number of pure strategies.*

⇒ You are asked to give the easy and instructive proof of Proposition 3.6 in Exercise 3.7 on page 103.

In other words, in a non-degenerate game we have pure-strategy equilibria (both players use exactly one pure strategy), or mixed-strategy equilibria where both players mix between exactly two strategies, or equilibria where both player mix between exactly three strategies, and so on. In general, both players mix between exactly k pure strategies. Each of these strategies has a probability, and these probabilities are subject to k equations in order to make the other player indifferent between his or her pure strategies: One of these equations states that the probabilities sum to one, and the other $k - 1$ equations state the equality between the first and second, second and third, etc., up to the $(k - 1)$ st and k th strategy of the other player. These equations have unique solutions. (It can be shown that non-unique solutions can occur only in a degenerate game.) These solutions have to be checked for the equilibrium property: The resulting probabilities have to be non-negative, and the unused pure strategies of the other player must not have a higher payoff. Finding mixed strategy probabilities by equating the expected payoffs to the other player no longer gives unique solutions in degenerate games, as we have demonstrated.

Should we or should we not care about degenerate games? In a certain sense, degenerate games can be ignored when considering games that are given in strategic form, and where each cell of the payoff table arises from a different circumstance of the interactive situation that is modelled by the game. In that case, it is “unlikely” that two payoffs are identical, which is necessary to get two pure best responses to a pure strategy, which is the only case where a 2×2 game can be degenerate. A small change of the payoff will result in a unique preference of the other player. Similarly, it is unlikely that more than two lines defining the upper envelope of expected payoffs cross in one point, or that more than three planes of expected payoffs against a mixed strategy that mixes three pure strategies (like in the two planes in Figure 3.8) cross in one point, and so on. In other words, “generic” (or “almost all”) games in strategic form are not degenerate. On the other hand, degeneracy is very likely when looking at a game tree, because a payoff of the game tree will occur repeatedly in the strategic form, as demonstrated by the threat game.

We conclude with a useful proposition that can be used when looking for all equilibria of a degenerate game. In the threat game in Figure 3.9, there are two equilibria $((1,0), (1,0))$ and $((1,0), (1/2, 1/2))$ that are obtained, essentially, by “ignoring” the degeneracy. These two equilibria have the same strategy $(1,0)$ of player I. Similarly, the game in Figure 3.10 has two mixed equilibria $((1/2, 1/2), (1/2, 1/2, 0))$ and $((1/2, 1/2), (2/3, 0, 1/3))$, which again share the same strategy of player I. The following proposition states that two mixed equilibria that have the same strategy of one player (here stated for player I, but it holds similarly for the other player) can be combined by convex combinations to obtain further equilibria, as it is the case in the above examples.

Proposition 3.7 *Consider a bimatrix game (A, B) with X as the set of mixed strategies of player I, and Y as the set of mixed strategies of player II. Suppose that (x, y) and (x', y) are equilibria of the game, where $x, x' \in X$ and $y \in Y$. Then $((1 - p)x + px', y)$ is also an equilibrium of the game, for any $p \in [0, 1]$.*

Proof. Let $\bar{x} = (1-p)x + px'$. Clearly, $\bar{x} \in X$ because $\bar{x}_i \geq 0$ for all pure strategies i of player I, and

$$\sum_{i=1}^m \bar{x}_i = \sum_{i=1}^m ((1-p)x_i + px'_i) = (1-p) \sum_{i=1}^m x_i + p \sum_{i=1}^m x'_i = (1-p) + p = 1.$$

The pair (\bar{x}, y) is an equilibrium if \bar{x} is a best response to y and vice versa. For any mixed strategy \tilde{x} of player I, we have $xAy \geq \tilde{x}Ay$, and $x'Ay \geq \tilde{x}'Ay$, and consequently $\bar{x}Ay = ((1-p)x + px')Ay = (1-p)xAy + px'Ay \geq (1-p)\tilde{x}Ay + p\tilde{x}'Ay = \tilde{x}Ay$, which shows that \bar{x} is a best response to y . In other words, these inequalities hold because they are preserved under convex combinations. In a similar way, y is a best response to x and x' , that is, for any $\tilde{y} \in Y$ we have $xBy \geq xB\tilde{y}$ and $x'B\tilde{y} \geq x'B\tilde{y}$. Again, taking convex combinations shows $\bar{x}By \geq \bar{x}B\tilde{y}$. \square

We come back to the example of Figure 3.10 to illustrate Proposition 3.7: This game has two equilibria (x, y) and (x, y') where $x = (1/2, 1/2)$, $y = (1/2, 1/2, 0)$, and $y' = (2/3, 0, 1/3)$, so this is the situation of Proposition 3.7, except that the two equilibria have same mixed strategy of player I rather than player II. Consequently, for any \bar{y} which is a convex combination of y and y' we have a Nash equilibrium (x, \bar{y}) . The equilibria (x, y) and (x, y') are “extreme” in the sense that y and y' have as many zero probabilities as possible, which means that y and y' are the endpoints of the line segment consisting of equilibrium strategies \bar{y} (that is, (x, \bar{y}) is an equilibrium), where $\bar{y} = (1-p)y + py'$.

\Rightarrow Express the set of Nash equilibria for the threat game in Figure 3.9 with the help of Proposition 3.7.

3.14 Zero-sum games

Zero-sum games are games of two players where the interests of the two players are directly opposed: One player's loss is the other player's gain. Competitions between two players in sports or in parlor games can be thought of as zero-sum games. The combinatorial games studied in Chapter 1 are also zero-sum games.

Consider a zero-sum game in strategic form. If the row player chooses row i and the column player chooses column j , then the payoff $a(i, j)$ to the row player can be considered as a cost to the column player which she tries to minimise. That is, her payoff is given by $b(i, j) = -a(i, j)$, that is, $a(i, j) + b(i, j) = 0$, which explains the term “zero-sum”. The zero-sum game is completely specified by the matrix A with entries $a(i, j)$ of payoffs to the row player, and the game is therefore often called a *matrix game*. In principle, such a matrix game is a special case of a bimatrix game (A, B) , where the column player's payoff matrix is $B = -A$. Consequently, we can analyse matrix games just like bimatrix games. However, zero-sum games have additional strong properties, which are the subject of this section, that do not hold for general bimatrix games.

Figure 3.11 shows the football penalty kick as a zero-sum game between striker and goalkeeper. For simplicity, assume that the striker's possible strategies are L and R , that




Figure 3.11 Strategic options for striker and goalkeeper in a football penalty.

is, to kick into the left or right corner, and that the goalkeeper's strategies are l , w , r , where l and r mean that he jumps immediately into the left or right corner, respectively, and w that he waits and sees where the ball goes and jumps afterwards. On the left in Figure 3.12, we give an example of resulting probabilities that the striker scores a goal, depending on the choices of the two players. Clearly, the striker (player I) tries to maximise and the goalkeeper (player II) to minimise that probability, and we only need this payoff to player I to specify the game. The two players in a zero-sum game are therefore often called "Max" and "min", respectively (where we keep our tradition of using lower case letters for the column player). In reality, one would expect higher scoring probabilities when the goalkeeper chooses w , but the given numbers lead to a more interesting equilibrium. In this case, the striker does not score particularly well, especially when he kicks into the right corner.




Figure 3.12 Left: the football penalty as a zero-sum game with scoring probabilities as payoffs to the striker (row player), which the goalkeeper (column player) tries to minimise. Right: lower envelope of best response costs to the minimiser.

In the payoff table in Figure 3.12, we have indicated the pure best responses with little boxes, in the lower left for the row player and in the upper right for the column player, similar to the boxes put around the payoffs in a bimatrix game when each cell has

two numbers. (Other conventions are to underline a best response cost for the minimiser, and to put a line above the best response payoff for the maximiser, or to put a circle or box, respectively, around the number.) When considering these best responses, remember that the column player minimises.

The best responses show that the game has no pure-strategy equilibrium. This is immediate from the sports situation: obviously, the goalkeeper's best response to a kick into the right or left corner is to jump into that corner, whereas the striker would rather kick into the opposite corner chosen by the goalkeeper. We therefore have to find a mixed equilibrium, and active randomisation is obviously advantageous for both players. Conveniently, the resulting expected scoring probability is just a probability, so the players are risk-neutral with respect to the numbers in the game matrix, which therefore represent the expected-utility function (as discussed in Section 3.4) for player I, and cost for player II.

We first find the equilibria of the game as we would do this for a 2×3 game, except that we only use the payoffs to player I. Of course, this game must be solved with the goalpost method, because it is about football! The right picture in Figure 3.12 gives the expected payoffs to player I as a function of his probability, say x , of choosing R , depending on the responses l , w , or r of player II. These payoffs to player I are costs to player II, so the best-response costs to player II are given by the minimum of these lines, which defines the *lower envelope* shown by bold line segments. There are two intersection points on the lower envelope, of the lines for l and w when $x = 1/3$, and of the lines for w and r when $x = 3/5$, which are easily found with the difference trick (where it does not matter if one considers costs or payoffs, because the differences are the same in absolute value). So player II's best response to the mixed strategy $(1 - x, x)$ is l for $0 \leq x \leq 1/3$, and w for $1/3 \leq x \leq 3/5$, and r for $3/5 \leq x \leq 1$. Only for $x = 1/3$ and for $x = 3/5$ can player II mix, and only in the first case, when l and w are best responses, can player I, in turn, be made indifferent. This gives the unique mixed equilibrium of the game, $((2/3, 1/3), (1/6, 5/6, 0))$. The resulting expected costs for the three columns l , w , r are $2/3$, $2/3$, $4/5$, where indeed player II assigns positive probability only to the smallest-cost columns l and w . The expected payoffs for both rows are $2/3$.

So far, nothing is new. Now, we consider the lower envelope of expected payoffs to player I (as a function of his own mixed strategy defined by x , which for simplicity we call "the mixed strategy x ") from that player's own perspective: It represents the worst possible expected payoff to player I, given the possible responses of player II. If player I wants to "secure" the maximum of this worst possible payoff, he should maximise over the lower envelope with a suitable choice of x . This maximum is easily found for $\hat{x} = 1/3$, shown by a dashed line in Figure 3.12.

The strategy \hat{x} is called a *max-min strategy*, and the resulting payoff is called his *max-min payoff*. In general, let A be the matrix of payoffs to player I and let X and Y be the sets of mixed strategies of player I and player II, respectively. Then a max-min strategy of player I is a mixed strategy \hat{x} with the property

$$\min_{y \in Y} \hat{x}Ay = \max_{x \in X} \min_{y \in Y} xAy, \quad (3.7)$$

which means the following: The left-hand side of (3.7) shows the smallest payoff that player I gets when he plays strategy \hat{x} , assuming that player II chooses a strategy y that

minimises player I's payoff. The right-hand side of (3.7) gives the largest such “worst-case” payoff over all possible choices of x . The equality in (3.7) states that \hat{x} is the strategy of player I that achieves this best possible worst-case payoff. It is therefore also sometimes called a *security* strategy of player I, because it defines what player I can secure to get (in terms of expected payoffs).

As a side observation, it is easy to see that the payoff $\min_{y \in Y} xAy$ which describes the “worst case” payoff to player I when he plays x can be obtained by only looking at the *pure* strategies, the columns j , of player II. In other words, $\min_{y \in Y} xAy = \min_j (xA)_j$. The reason is that with a mixed strategy y , player II cannot produce a worse expected payoff to player I than with a pure strategy j , because the expectation represents a “weighted average” of the pure-strategy payoffs $(xA)_j$ with the probabilities y_j as weights, so it is never worse than the smallest of these pure-strategy payoffs (which is proved in the same way as the best-response condition Theorem 3.1). However, we keep the minimisation over all mixed strategies y in (3.7) because we will later exchange taking the maximum and the minimum. For player I, the max-min strategy \hat{x} is better if he can use a mixed strategy, as the football penalty game shows: If he was restricted to playing pure strategies only, his (pure-strategy) max-min payoff $\max_i \min_j a(i, j)$ would be 0.5, with pure max-min strategy L , rather than the max-min payoff $2/3$ when he is allowed to use a mixed strategy.

The max-min payoff in (3.7) is well defined and does not depend on what player II does, because we always take the minimum over player II's choices y . The function $\min_{y \in Y} xAy$ is a continuous function of x , and x is taken from the compact set X , so there is some \hat{x} where the maximum $\max_{x \in X} \min_{y \in Y} xAy$ of that function is achieved. The maximum (and thus the max-min payoff) is unique, but there may be more than one choice of \hat{x} , for example if in the game in Figure 3.12, the payoff 0.7 for the strategy pair (L, w) was replaced by 0.6, in which case the lower envelope would have its maximum 0.6 along an interval of probabilities that player I chooses R , with any \hat{x} in that interval defining a max-min strategy.

Recall how the max-min strategy \hat{x} is found graphically over the lower envelope of expected payoffs. The max-min strategy \hat{x} identifies the “peak” of that lower envelope. The “hat” accent \hat{x} may be thought of as representing that peak.

⇒ Try Exercise 3.8 on page 103, which is a variant of the football penalty game.

A max-min strategy, and the max-min payoff, only depends on the player's *own* payoffs. It can therefore also be defined for a game with general payoffs which are not necessarily zero-sum. In that case, a max-min strategy represents a rather pessimistic view of the world: Player I chooses a strategy \hat{x} that maximises his potential loss against an opponent who has nothing else on her mind than harming player I. In a *zero-sum* game, however, that view is eminently rational: Player II is not trying to harm player I, but she is merely trying to maximise her own utility. This is because her payoffs are directly opposed those of player I, her payoff matrix being given by $B = -A$. So the expression $\min_{y \in Y} xAy$ in (3.7) can also be written as $-\max_{y \in Y} xBy$ which is the (negative of the) best response payoff to player II. This is exactly what player II tries to do when maximising her payoff against the strategy x of player I. As described above, the lower envelope

of payoffs to player I is just the negative of the upper envelope of payoffs to player II, which defines her best responses.

So a max-min strategy is a natural strategy to look at when studying a zero-sum game. Our goal is now to show that in a zero-sum game, a max-min strategy is the same as an equilibrium strategy. This has some striking consequences, for example the uniqueness of the equilibrium payoff in any equilibrium of a zero-sum game, which we will discuss.

We consider such a max-min strategy not only for player I but also for player II, whose payoff matrix is $B = -A$. According to (3.7), this is a mixed strategy \hat{y} of player II with the property

$$\min_{x \in X} xB\hat{y} = \max_{y \in Y} \min_{x \in X} xBy,$$

which we re-write as

$$-\max_{x \in X} xA\hat{y} = -\min_{y \in Y} \max_{x \in X} xAy$$

or, omitting the minus signs,

$$\max_{x \in X} xA\hat{y} = \min_{y \in Y} \max_{x \in X} xAy. \quad (3.8)$$

Equation (3.8) expresses everything in terms of a single matrix A of payoffs to player I, which are costs to player II that she tries to minimise. In order to simplify the discussion of these strategies, a max-min strategy of player II (in terms of her own payoffs) is therefore also called a *min-max* strategy (with the min-max understood in terms of her own costs, which are the payoffs to her opponent). So a min-max strategy of player II is a strategy \hat{y} that minimises the worst-case cost $\max_{x \in X} xA\hat{y}$ that she has to pay to player I, as stated in (3.8).

The first simple relationship between the max-min payoff to player I in a zero-sum game and the min-max cost to player II is

$$\max_{x \in X} \min_{y \in Y} xAy \leq \min_{y \in Y} \max_{x \in X} xAy. \quad (3.9)$$

The left-hand side is the “secure” payoff that player I can guarantee to himself (on average) when using a max-min strategy \hat{x} , the right-hand side the “secure” cost that player II can guarantee to herself to pay at most (on average) when using a min-max strategy \hat{y} . It is clear that this is the only direction in which the inequality can hold if that “guarantee” is to be meaningful. For example, if player I could guarantee to get a payoff of 10, and player II could guarantee to pay at most 5, something would be wrong. The proof of (3.9) is similarly simple:

$$\max_{x \in X} \min_{y \in Y} xAy = \min_{y \in Y} \hat{x}Ay \leq \hat{x}A\hat{y} \leq \max_{x \in X} xA\hat{y} = \min_{y \in Y} \max_{x \in X} xAy.$$

The main theorem on zero-sum games states the *equation* “max min = min max”, to be understood for the payoffs, that is, equality instead of an inequality in (3.9). This is also called the “minimax” theorem, and it is due to von Neumann.

The next theorem clarifies the connection between max-min strategies (which for player II we called min-max strategies) and the concept of Nash equilibrium strategies. Namely, for zero-sum games they are the same. Furthermore, the existence of a Nash equilibrium strategy pair (x^*, y^*) implies the minimax theorem, that is, equality in (3.9).

Theorem 3.8 (von Neumann’s minimax theorem) Consider a zero-sum game with pay-off matrix A to player I. Let X and Y be the sets of mixed strategies of player I and II, respectively, and let $(x^*, y^*) \in X \times Y$. Then (x^*, y^*) is a Nash equilibrium if and only if x^* is a max-min strategy and y^* is a min-max strategy, that is,

$$\min_{y \in Y} x^* A y = \max_{x \in X} \min_{y \in Y} x A y, \quad \max_{x \in X} x A y^* = \min_{y \in Y} \max_{x \in X} x A y. \quad (3.10)$$

Furthermore,

$$\max_{x \in X} \min_{y \in Y} x A y = \min_{y \in Y} \max_{x \in X} x A y. \quad (3.11)$$

Proof. Consider a max-min strategy \hat{x} in X and a min-max strategy \hat{y} in Y , which fulfil (3.10) in place of x^* and y^* , respectively. Such strategies exist because X and Y are compact sets. Let (x^*, y^*) be any Nash equilibrium of the game, which exists by Nash’s Theorem 3.3. The equilibrium property is equivalent to the “saddle point” property

$$\max_{x \in X} x A y^* = x^* A y^* = \min_{y \in Y} x^* A y, \quad (3.12)$$

where the left equation states that x^* is a best response to y^* and the right equation states that y^* is a best response to x^* (recall that $B = -A$, so player II tries to minimise her expected costs $x A y$). Then

$$\hat{x} A \hat{y} \leq \max_{x \in X} x A \hat{y} = \min_{y \in Y} \max_{x \in X} x A y \leq \max_{x \in X} x A y^* = x^* A y^* \quad (3.13)$$

and

$$\hat{x} A \hat{y} \geq \min_{y \in Y} \hat{x} A y = \max_{x \in X} \min_{y \in Y} x A y \geq \min_{y \in Y} x^* A y = x^* A y^*, \quad (3.14)$$

so all these inequalities hold as equalities. With the second inequality in both (3.13) and (3.14) as an equality we get (3.10), that is, y^* is a min-max and x^* is a max-min strategy. Having the first inequality in both (3.13) and (3.14) as an equality gives

$$\max_{x \in X} x A \hat{y} = \hat{x} A \hat{y} = \min_{y \in Y} \hat{x} A y,$$

which says that (\hat{x}, \hat{y}) is a Nash equilibrium. Furthermore, we obtain (3.11). \square

One consequence of Theorem 3.8 is that in any Nash equilibrium of a zero-sum game, a player receives the same payoff, namely the max-min payoff to player I in (3.11), which is the min-max cost to player II. This payoff is also called the *value* of the zero-sum game.

Another important consequence of the Theorem 3.8 is that the concept of “Nash equilibrium strategy” in a zero-sum game is *independent* of what the opponent does, because it is the same as a max-min strategy (when considered in terms of the player’s own pay-offs). This is far from being the case in general-payoff games, as, for example, the game of “chicken” demonstrates. In particular, if player I’s strategies x^* and \bar{x} are both part of a Nash equilibrium in a zero-sum game, they are part of *any* Nash equilibrium. So Nash equilibrium strategies are interchangeable: if (x^*, y^*) and (\bar{x}, \bar{y}) are Nash equilibria, so are (x^*, \bar{y}) and (\bar{x}, y^*) . The reason is that x^* and \bar{x} are then max-min strategies, and y^* and \bar{y} are min-max strategies, which has nothing to do with what the opponent does.

A further consequence is that if (x^*, y^*) is a Nash equilibrium and y^* is the *unique* best response (for example, this is often the case when y^* is a pure strategy), then y^* is also the unique min-max strategy. The reason is that if there was another min-max strategy \hat{y} that is different from y^* , then (x^*, \hat{y}) would be a Nash equilibrium, and consequently \hat{y} would be a best response to x^* , contradicting the uniqueness of the best response y^* . Similarly, if, say, \hat{x} is the unique max-min strategy of player I (for example, obtained by plotting the lower envelope against x as in Figure 3.12), it is also the only Nash equilibrium strategy of player I.

These properties of zero-sum games make equilibrium strategies, or equivalently max-min/min-max strategies, a particularly convincing solution concept. They are also called “optimal” strategies for the players.

3.15 Exercises for Chapter 3

Exercise 3.1 shows how to find all mixed equilibria of a small bimatrix game, which is however not a 2×2 game. It is useful to attempt this exercise as early as possible, because this will help you appreciate the methods developed later. The given solution to this exercise does not use the upper envelope method, but you should try using that method once you have learned it; it will help you solve the exercise faster. Exercise 3.2 is a (not easy) continuation of Exercise 2.7, and asks about mixed equilibria in a three-player game tree. In Exercise 3.3, you are asked to derive the “difference trick”. Exercise 3.4 demonstrates the power of the upper envelope method for $2 \times n$ games. In Exercise 3.5, you should understand the concept of dominated strategies with the help of the upper-envelope diagram, which provides a useful geometric insight. A degenerate game is treated in Exercise 3.6. Exercise 3.7 shows that in non-degenerate games, Nash equilibria are always played on square games that are part of the original game, in the sense that both players mix between the same number of pure strategies. Exercise 3.8 is a variant of the football penalty.

Exercise 3.1 Find all Nash equilibria of the following 3×2 game, including equilibria consisting of mixed strategies.

		II	
		l	r
I	T	1 0	0 6
	M	0 2	2 5
	B	3 3	4 3

Exercise 3.2 Consider the three-player game tree in Figure 2.22 on page 66, already considered in Exercise 2.7.

- (a) Is the following statement true or false? Justify your answer.

For each of players I, II, or III, the game has a Nash equilibrium in which that player plays a mixed strategy that is not a pure strategy.

- (b) Is the following statement true or false? Justify your answer.

In every Nash equilibrium of the game, at least one player I, II, or III plays a pure strategy.

Exercise 3.3 In the following 2×2 game, A, B, C, D are the payoffs to player I, which are real numbers, no two of which are equal. Similarly, a, b, c, d are the payoffs to player II, which are real numbers, also no two of which are equal.

		II	
		left	right
I	Top	a A	b B
	Bottom	c C	d D

- (a) Under which conditions does this game have a mixed equilibrium which is not a pure-strategy equilibrium?

[Hint: Think in terms of the arrows, and express this using relationships between the numbers, as, for example, $A > C$.]

- (b) Under which conditions in (a) is this the only equilibrium of the game?
- (c) Consider one of the situations in (b) and compute the probabilities $1 - p$ and p for playing “Top” and “Bottom”, respectively, and $1 - q$ and q for playing “left” and “right”, respectively, that hold in equilibrium.
- (d) For the solution in (c), give a simple formula for the quotients $(1 - p)/p$ and $(1 - q)/q$ in terms of the payoff parameters. Try to write this formula such that denominator and numerator are both positive.
- (e) Show that the mixed equilibrium strategies do not change if player I’s payoffs A and C are replaced by $A + E$ and $C + E$, respectively, for some constant E , and similarly if player II’s payoffs a and b are replaced by $a + e$ and $b + e$, respectively, for some constant e .

Exercise 3.4 Consider the following 2×5 game:

		II				
		<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>
I	<i>T</i>	2 0	4 2	3 1	5 0	0 3
	<i>B</i>	7 1	4 0	5 4	0 1	8 0

- (a) Draw the expected payoffs to player II for all her strategies a, b, c, d, e , in terms of the probability p , say, that player I plays strategy B . Indicate the best responses of player II, depending on that probability p .
- (b) Using the diagram in (a), find all pure or mixed equilibria of the game.

Exercise 3.5 Consider the quality game in Figure 2.10.

- (a) Use the “goalpost” method twice, to draw the upper envelope of best-response payoffs to player I against the mixed strategy of player II, and vice versa.
- (b) Explain which of these diagrams shows that a pure strategy dominates another pure strategy.
- (c) Explain how that diagram shows that if a strategy s dominates a strategy t , then t is never a best response, even against a *mixed* strategy of the other player.

Exercise 3.6 Find all equilibria of the following degenerate game, which is a variant of the inspection game in Figure 3.2. The difference is that player II, the inspectee, derives no gain from acting illegally (playing r) even if she is not inspected (player I choosing T).

		II	
		<i>l</i>	<i>r</i>
I	<i>T</i>	0 0	0 −10
	<i>B</i>	0 −1	−90 −6

Exercise 3.7 Prove proposition 3.6.

Exercise 3.8 Consider the duel in a football penalty kick. The following table describes this as a zero-sum game. It gives the probability of scoring a goal when the row player (the striker) adopts one of the strategies L (shoot left), R (shoot right), and the column player (the goalkeeper) uses the strategies l (jump left), w (wait then jump), r (jump right). The row player is interested in maximising and the column player in minimising the probability of scoring a goal.

		min		
		l	w	r
Max	L	0.6	0.7	1.0
	R	1.0	0.8	0.7

- (a) Find an equilibrium of this game in pure or mixed strategies, and the equilibrium payoff. Why is the payoff unique?
- (b) Now suppose that player I has an additional strategy M (shoot in the middle), so that the payoff matrix is

		min		
		l	w	r
Max	L	0.6	0.7	1.0
	R	1.0	0.8	0.7
	M	1.0	0.0	1.0

Find an equilibrium of this game, and the equilibrium payoff.
 [Hint: The result from (a) will be useful.]

Chapter 4

Game trees with imperfect information

This chapter studies extensive games, which are game trees with imperfect information. It relies on chapters 2 and 3.

4.1 Learning objectives

After studying this chapter, you should be able to:

- explain extensive games and the concept of information sets;
- describe pure and reduced pure strategies for extensive games, and the corresponding strategic form and reduced strategic form of the game;
- define the concept of perfect recall, and see if this holds for a given game;
- explain the difference between behaviour and mixed strategies, and state Kuhn's theorem;
- construct simple extensive games from descriptions of games with imperfect information, like the game in Exercise 4.3 on page 133;
- find all equilibria of simple extensive games, and represent the mixed strategies in these equilibria as behaviour strategies.

4.2 Further reading

Extensive games with imperfect information are a standard concept in game theory. They are treated here in greater detail than elsewhere, in particular as concerns information sets and the condition of perfect recall.

One alternative reference is Chapter 11 of

- Osborne, Martin J., and Ariel Rubinstein *A Course in Game Theory*. (MIT Press, 1994) [ISBN 0262650401].

A short treatment of extensive games is given in Section 2.4 of

- Gibbons, Robert *A Primer in Game Theory* [in the United States sold under the title *Game Theory for Applied Economists*]. (Prentice Hall / Harvester Wheatsheaf, 1992) [ISBN 0745011594].

Information sets were introduced by Harold Kuhn in the following article:

- Kuhn, Harold W. “Extensive games and the problem of information”, in: *Contributions to the Theory of Games II*, eds. H. W. Kuhn and A. W. Tucker, *Annals of Mathematics Studies*, Vol. 28 (1953), Princeton Univ. Press, pp. 193–216.

The “Annals of Mathematics Studies” where this article appeared is a book series. Kuhn’s article defines perfect recall in a way that is not, at first glance, the same as here, but it should be possible to understand that the definitions are the same after having studied this chapter. The definition 4.1 of perfect recall that we give is due to the following article by Reinhard Selten.

- Selten, Reinhard “Reexamination of the perfectness concept for equilibrium points in extensive games”, *International Journal of Game Theory*, Vol. 4 (1975), pp. 25–55.

4.3 Introduction

Typically, players do not always have full access to all the information which is relevant to their choices. Extensive games with *imperfect information* model exactly what information is available to the players when they make a move. Modelling and evaluating strategic information precisely is one of the strengths of game theory.

The central concept of *information sets* is introduced in Section 4.4 by means of a detailed example. We then give the general definition of extensive games, and the definition of strategies and reduced strategies, which generalise these concepts as they have been introduced for game trees with perfect information. Essentially, all one has to do is to replace “decision node” by “information set”.

The concept of *perfect recall* states a condition on the structure of the information sets of a player, which have to be in a certain sense “compatible” with the game tree. (Perfect recall is not about how to interpret the game.) The related concept of *behaviour* strategies and Kuhn’s theorem is then studied in depth.

The final Section 4.12 treats *subgames* and subgame perfect Nash equilibria (SPNE).

4.4 Information sets

Consider the situation faced by a large software company (player I) after a small startup firm (player II) has announced deployment of a key new technology. The large company has a large research and development operation. It is generally known that they have researchers working on a wide variety of innovations. However, only the large company knows for sure whether or not they have made any progress on a product similar to the startup’s new technology. The startup believes that there is a 50 percent chance that the

large company has developed the basis for a strong competing product. For brevity, when the large company has the ability to produce a strong competing product, the company will be referred to as having a “strong” position, as opposed to a “weak” one.

The large company, after the announcement by the small company, has two choices. It can counter by announcing that it too will release a competing product. Alternatively, it can choose to cede the market for this product. The large company will certainly condition its choice upon its private knowledge, and may choose to act differently when it has a strong position than when it has a weak one. If the large company has announced a product, the startup is faced with a choice: it can either negotiate a buyout and sell itself to the large company, or it can remain independent and launch its product. The startup does not have access to the large firm’s private information on the status of its research. However, it does observe whether or not the large company announces its own product, and may attempt to infer from that choice the likelihood that the large company has made progress of its own.

When the large company does not have a strong product, the startup would prefer to stay in the market over selling out. When the large company does have a strong product, the opposite is true, and the startup is better off selling out instead of staying in.




Figure 4.1 Game with imperfect information for player II, indicated by the information set that contains two decision nodes with the same moves a and b . At a leaf of the tree, the top payoff is to player I, the bottom payoff to player II (as always).

Figure 4.1 shows a game tree with imperfect information (also called “extensive game”) that models this situation. From the perspective of the startup, whether or not the large company has done research in this area is random. This is here modelled by an initial chance move, where the left branch stands for a “weak” player I (as described above), and the right branch for a “strong” player I. Both branches have here probability $1/2$.

When the large company is in a weak position, it can choose to cede the market to the startup. This is given by move T , with payoffs $(0, 16)$ to the two players. It can

also announce a competing product, modeled by move B , in the hope that the startup company, player II, will sell out, choosing move a , with payoffs 12 and 4 to players I and II. However, if player II decides instead to stay in the market (move b) it will even profit from the increased publicity and gain a payoff of 20, with a loss of -4 to player I.

In contrast, when player I is in a strong position, it can again cede the market to player II (move C , with the same payoff pair $(0, 16)$ as before), or instead announce its own product (move D). It will quickly become apparent that player I should never choose C , but we keep it as a possibility. After move D of the “strong” player I, the payoffs to the two players are $(12, 4)$ if the startup sells out (move a) and $(20, -4)$ if the startup stays in (move b).

In addition to a game tree with perfect information, the nodes of the players are enclosed by ovals which define *information sets*. These are simply sets of decision nodes, subject to certain restrictions that we will describe shortly in generality. The interpretation is that a player cannot distinguish among the nodes in an information set, given his knowledge at the time he makes the move. Because his knowledge at all nodes in an information set is the same, he makes the same choice at each node in that set. Here, the startup company, player II, must choose between move a (sell out) and move b (stay in the market). These are the two choices at player II’s information set, which has two nodes according to the different histories of play, which player II cannot distinguish.

Because player II is not informed about her position in the game, backward induction can no longer be applied. If player II knew whether player I is in a weak or strong position, then it would be better for her to choose b when facing a weak player I (left node in Figure 4.1), and to choose a when facing a strong player I (right node of the information set).

One “optimal move” that can be seen directly from the game tree is the choice between C or D of player I. Choosing C gives player I a payoff of zero, whereas D gives him either 12 or 20, depending on whether player II chooses a or b . Both payoffs are larger than zero, so player I should choose D at any rate. So the game reduces to a decision between move T and B for player I, and between a and b for player II, assuming player I always chooses D .

The game tree shows that no pair of these moves for the two players is optimal against the other. For player I, if player II chooses b , then move T is better (because 0 is a higher payoff than -4). If player I chooses T , then the left node of the information set of player II is not reached. In that case, when player II has the choice between a and b , she knows with certainty that she is at the right node of her information set. Hence, she gets the higher payoff 4 when choosing a (rather than -4 when choosing b), that is, player II will “sell out” to the big company and choose move a . If player II, in turn, chooses a , then player I’s choice between T and B is in favour of B , which gives the higher payoff 12 compared to 0. If player I chooses B , then player II will be with equal probability at either the left or the right node in her information set. Player II then has to use expected payoffs when comparing a and b . Clearly, move a will give her payoff 4, whereas move b will give her the expected payoff $1/2 \cdot 20 + 1/2 \cdot (-4) = 8$. That expected payoff exceeds the payoff for move a . That is, against B , player II is better off choosing b . We have now come

full circle: The sequence of best-response moves for the players is $b \rightarrow T \rightarrow a \rightarrow B \rightarrow b$, so there is no deterministic behaviour of the players that defines an equilibrium.

The game in Figure 4.1 can be analysed in terms of the moves alone (we will give such an analysis at the end of this section), because for each player, only two choices matter; this follows from the initial separate consideration that excludes move C as a rational choice.

However, in general it is useful to consider the *strategic form* of the game that has already been considered for game trees with perfect information. This involves systematically considering the *strategies* that players have in the game. The Nash equilibrium concept is also defined in terms of strategies, and not in terms of moves.

		II	
		a	b
I	TC	16 0	16 0
	TD	10 6	6 10
	BC	10 6	18 -2
	BD	4 12	8 8

		II	
		a	b
I	TD	10 6	6 10
	BD	4 12	8 8

Figure 4.2 Strategic form of the game in Figure 4.1 (left), and the 2×2 game (right) that results after eliminating the strictly dominated strategies TC and BC , with best response payoffs.

The left table in Figure 4.2 gives the strategic form of the game in Figure 4.1. As in the case of game trees with perfect information, a strategy of a player defines a move for each of the player's decision nodes. The difference, however, is that such a move is only specified *once* for all the decision nodes that belong to an information set, because by definition the choice of that move is the same no matter where the player is in the information set. In Figure 4.1, this means that player II has only two moves, which are the two moves a and b at her only information set. Because player II has only one information set, her strategies are the same as her moves. Player I has two information sets, so a strategy for player I is a pair of moves, one move for the first and one move for the second information set. These pairs of moves are TC, TD, BC, BD .

The four strategies of player I correspond to the rows on the left in Figure 4.2, whereas the columns are the strategies a and b of player II. Each of the eight cells in that table gives the expected payoffs as they are calculated from the payoffs at the leaves and the chance probabilities. For example, the strategy pair (BC, a) gives expected payoff $1/2 \cdot 12 + 1/2 \cdot 0 = 6$ to player I and payoff $1/2 \cdot 4 + 1/2 \cdot 16 = 10$ to player II. These expectations

come from the possible chance moves, both of which have probability $1/2$ according to Figure 4.1. So when the chance move goes to the left, the strategy pair (BC, a) means the play of the game follows move B and then move a , reaching the leaf with payoffs 12, 4 to players I, II. When the chance move goes to the right, the strategy pair (BC, a) means that move C leads immediately to the leaf with payoffs 0, 16.

The left 4×2 game in Figure 4.2 has two strictly dominated strategies, namely TC , which is strictly dominated by TD , and BC , which is strictly dominated by BD . This is in agreement with the earlier observation that for player I, move C is always worse than move D , no matter what player II does. In a certain respect one has to be careful, though, to conclude that a “worse move” always leads to a *strictly* dominated strategy (the dominated strategy, for example TC instead of TD , would be obtained by replacing a given move, like D as part of the strategy TD , by the “worse move”, here C). Namely, even if the payoffs from the “worse move” (in the example C) are strictly worse than those resulting from the given move (D), one has to make sure that the decision point with the moves in question (C and D) is *always reached with positive probability*. (Only in this case the expected payoffs from the two strategies will always differ.) This is here the case because the decision point of player I with moves C and D is reached by a chance move. If, in contrast, the chance move was replaced by a move of a third player, for example, then if that third player does not move to the right, the choice between C and D would *not affect the payoff* to player I. So in this case, we could only conclude that TD *weakly* dominates TC . In short, it is safer to look at the strategic form of the game, even though a careful look at the game tree may eliminate some strategies right away.

Elimination of the strictly dominated strategies TC and BC gives the 2×2 game on the right in Figure 4.2. There, each player has two strategies, which correspond to the two moves discussed earlier. The best response payoffs indicated by boxes show that this game has no pure-strategy equilibrium.

The equilibrium of the game involves both players randomising. The mixed strategy probabilities can be determined in the familiar way. Player I randomises with equal probability $1/2$ between TD and BD so that the expected payoff to player II is 7 for both a and b . Conversely, player II chooses a with probability $1/4$ and b with probability $3/4$. Then player I is indifferent, receiving an overall expected payoff of $36/4 = 9$ in each case.

The game in Figure 4.1 is a competition between two firms that is more familiar when viewed as a simple case of the card game *Poker*. We can think of player I being dealt a weak or strong hand, which is with equal probability either worse or better than the hand of player II (this is represented by the left and right chance move, respectively). Player I knows whether his hand is weak or strong (in a real Poker game, player I knows his hand but not necessarily whether it is better or worse than that of player II). Player I can then either “fold” (moves T and C), or “raise” (moves B and D). In turn, player II can either “pass” (move a) or “meet”. The “fold” and “pass” moves entail a fixed outcome where the players receive payoffs that are independent of the hands they hold. Only the combinations of moves “raise” and “meet”, corresponding to the strategy pair (BD, b) , lead to outcomes where the players’ hands matter, as can be seen from Figure 4.1 (the reason that other strategy pairs give non-constant payoffs in Figure 4.2 is due to the different effect of the chance move in each case).

We can say that when player I has a weak hand and decides to “raise” (move B), this amounts to *bluffing*, pretending that he has a strong hand, because player II does not know this. The purpose of bluffing, however, is not only to make player II “pass”. Indeed, player II would not “pass” (move a) if player I bluffed all the time, as analysed before. One main effect of bluffing is to leave player II *uncertain* as to whether player I has a strong or weak hand. The best possible outcome for player I is if player II meets a strong hand of player I (moves D and b), because this gives player I maximum payoff. If player I raises only if he has a strong hand, then player II would rather pass, which is not in the interest of player I. Above, we have reasoned that no deterministic combination of moves leads to an equilibrium. As described, this is made very vivid in terms of the Poker game interpretation.

To conclude the discussion of this example, we can also deduce in this case directly from the game tree the probabilities for the moves T and B , and for a and b , that are necessary in order to reach an equilibrium. Even without constructing the strategic form, we have seen that no deterministic choice of these moves gives an equilibrium. Hence, both players must randomise, and for that purpose they must be *indifferent* between their choices. For player I, this indifference is very easy to determine: If he has to decide between T and B , then he gets the certain payoff 0 with T , so indifference results only when the expected payoff for B is also 0. If player II chooses a and b with probabilities $1 - q$ and q , respectively, that expected payoff to player I is $12(1 - q) - 4q$, so the equation $0 = 12(1 - q) - 4q = 12 - 16q$ determines uniquely $q = 3/4$.

In turn, player II has to be indifferent between her moves a and b in order to randomise in that way. Here the situation is somewhat more complicated, because the expected payoff to player II depends on the probabilities of being either at the left or at the right node in her information set. Assuming, as we do, that player I makes the rational choice and chooses D with certainty, the probability of being at the right node is $1/2$, whereas the probability of being at the left node is $1/2 \cdot p$, where p is the probability of player I choosing B . Because there is a non-zero probability $1 - p$ for T , the game ends with probability $1/2 \cdot (1 - p)$ at the leftmost leaf of the tree, and then the information set of player II will not be reached at all. Technically, what has to be computed is the *conditional* probability that player II will be at the left node versus being at the right node. Before going into the details, we note that from the description so far, the probability for the left node is p times the probability for the right node. Only these relative weights matter. Indifference of player II between a and b then gives the following equation, with the weighted payoffs for a and b on the left and right, respectively:

$$p \cdot 4 + 1 \cdot 4 = p \cdot 20 + 1 \cdot (-4), \quad (4.1)$$

which simplifies to $4p + 4 = 20p - 4$ or $8 = 16p$, that is, $p = 1/2$.

The conditional probabilities for player II being at the left and right node are computed as follows: The left node has absolute probability $1/2 \cdot p$ of being reached, the right node absolute probability $1/2$. The total probability of being at any of the two nodes (which are disjoint events) is the sum of these two probabilities, namely $(p + 1)/2$. The conditional probability is then the absolute probability divided by the total probability, that is, for the left node, $\frac{p/2}{(p+1)/2}$ or $\frac{p}{p+1}$, and for the right node, $\frac{1/2}{(p+1)/2}$ or $\frac{1}{p+1}$. Using

these conditional probabilities instead of the “weights” in (4.1) would mean dividing both sides of equation (4.1) by $p + 1$. The first step in solving that amended equation would be to multiply both sides by $p + 1$, giving back (4.1), so what we have done is perfectly in order (as long as the total probability of reaching the information set in the first place is nonzero).

One more time, we interpret the equilibrium probabilities $p = 1/2$ and $q = 3/4$ in terms of the game tree: Player I, by choosing B with probability $1/2$, and D with certainty, achieves that player II is half as likely to be at the left node of her information set as at her right node. Therefore, the conditional probabilities for left versus right node are $1/3$ and $2/3$, respectively ($1/4$ and $1/2$ divided by the total probability $3/4$ of reaching the information set). Exactly in that case, player II is indifferent between a and b : The *conditional expected payoffs* to player II are then (clearly) 4 for move a , and $1/3 \cdot 20 + 2/3 \cdot (-4) = 12/3 = 4$ for move b . The *overall* expected payoff to player II has to take the absolute probabilities into account: With probability $1/4$ (left chance move, player I choosing T), she gets payoff 16, with probability $3/4$, the mentioned conditional payoff 4, overall $1/4 \cdot 16 + 3/4 \cdot 4 = 4 + 3 = 7$. For player I, we have observed the indifference between T and B , so when the chance move goes to the left, player I gets payoff zero. When the chance move goes to the right, player I when choosing D will get $(1 - q) \cdot 12 + q \cdot 20 = 1/4 \cdot 12 + 3/4 \cdot 20 = 18$. The overall payoff to player I is $1/2 \cdot 0 + 1/2 \cdot 18 = 9$.

4.5 Extensive games

In this section, we give the general definition of an *extensive game*, also known as a game in extensive form, or as a game tree with imperfect information.

Similar to the game trees with perfect information considered earlier, the basic structure of an extensive game is a *tree*. As before, it has three types of nodes: Terminal nodes, called *leaves*, with a payoff for each player; non-terminal nodes which are either *chance* nodes (typically drawn as squares), where the next node is chosen randomly according to a given probability that is specified in the tree, or non-terminal nodes that are *decision* nodes, which belong to a particular player. The outgoing edges from each decision node are labelled with the names of choices or *moves* which uniquely identify the respective edges.

The decision nodes are partitioned into *information sets*. “Partitioned” means that every decision node belongs to exactly one information set, and no information set is empty. When drawing the extensive game, information sets are typically shown as ovals around the nodes that they contain; some people draw information sets by joining the nodes in the set with a single dashed line instead.

The *interpretation* of an information set is that in the course of the game, a player is only told that he is at *some* node in an information set, but not at which particular node. Consequently, the information sets have to fulfil certain conditions: In an information set, all decision nodes belong to the same player; all decision nodes in the set have the same number of outgoing edges; and the set of moves (given by the labels on the outgoing edges) is the same for each node in the information set. In Figure 4.1, for example, the

information set of player II has two moves, a and b . When the player has reached an information set, he makes a choice by picking a move, which is by definition the *same* move no matter where he is in the information set.

These conditions are trivially fulfilled if every information set is a singleton (contains only one node). If all information sets are singletons, then the game has *perfect information* and the information sets could as well be omitted. This gives a game tree with perfect information as considered earlier. Only for information sets that contain two or more nodes do these conditions matter.




Figure 4.3 Game with an information set where two nodes share a path, which is usually not allowed in an extensive game.

In addition to the above conditions, we will impose additional constraints on the structure of the information sets. The first of these is that no two nodes in an information set share a path, as in the example in Figure 4.3. The single-player game in that figure has two decision nodes of player I, which belong to a single information set that has the two moves L and R . The only other non-terminal node is a chance move. Every condition about extensive games is met. However, the two decision nodes of player I share a path, because the first node is the root, from which there is a path to the second node. Information sets with this feature have a very problematic interpretation, because we require that a *move*, for example L , at such an information set is the same no matter where the player is. However, in Figure 4.3, what would it mean for player I to make move L ? It certainly means going left from the root, but if afterwards the chance move goes to the right so that player I can move again, does he automatically have to make move L again?

The problem with the information set in Figure 4.3 is that, by the interpretation of information sets, player I must have *forgotten* if he made a move at all when called to choose between L and R , and it is not even clear whether he is allowed to choose “again” or not. An information set can represent that a player forgets what he knew or did earlier. In that case, the player is said to have *imperfect recall*. However, players with imperfect recall do not conform to our idea of “rational” players. Moreover, their forgetfulness is not very adequately represented by information sets. We will therefore assume that all players have *perfect recall*, which is defined, and discussed in detail, in Section 4.8 below. In particular, we will see that perfect recall implies that no two nodes in an information set are connected by a path in the tree (see Proposition 4.2 below), so that a situation like in Figure 4.3 cannot occur.

To summarise the definition of game trees with imperfect information: The only new feature added to the game tree is the concept of information sets, which model a player's *lack* of information about where he is in the game.

4.6 Strategies for extensive games and the strategic form

The definition of a *strategy* in an extensive game generalises the definition of a strategy for a game tree with perfect information: A (pure) strategy of a player specifies a move for each information set of the player.

A formal definition of strategies requires some notation to refer to the information sets of a player, and the moves available at each information set. Suppose player i is a player in the extensive game (like player I or player II in a two-player game). Then H_i denotes the set of information sets of player i . A particular information set of player i is typically called h . We denote the set of *choices* or moves at the information set h by C_h for $h \in H_i$. For example, if h is the left information set of player I in Figure 4.1, then $C_h = \{T, B\}$.

The set Σ_i of strategies of player i is then

$$\Sigma_i = \prod_{h \in H_i} C_h, \quad (4.2)$$

that is, the cartesian product of the choice sets C_h of player i . An element of Σ_i is a tuple (vector) of moves, one move c from C_h for each information set h of player i .

If N is the set of players, like $N = \{I, II\}$ in the notation we use for a two-player game, then the set of pure *strategy profiles* is $\prod_{i \in N} \Sigma_i$, so a strategy profile specifies one pure strategy for each player.

The *strategic form* of the extensive game is given by the set N of players, the set of strategies Σ_i for each player i , and the expected payoff that each player gets for each strategy profile. If the game has no chance moves, then the payoff for each strategy profile is given deterministically by the leaf of the game tree that is reached when the players use their strategies. If the game has chance moves, more than one leaf may be reached with a given profile, and then expected payoffs have to be computed, as in the example of Figure 4.2.

With the help of information sets, it is possible to model *simultaneous play* in an extensive game. This is not possible with a game tree with perfect information where the sequential moves mean that the player who moves second is informed about the move of the first player.

In particular, given a game in strategic form, there is an extensive game that has exactly this strategic form. For example, consider an $m \times n$ game. It does not matter which player moves first, so we can assume this is player I, who has m moves, say r_1, r_2, \dots, r_m for the m rows of the game, at the root of the game tree. Each of these moves leads to a decision node of player II. All these nodes belong to a single information set. At this information set, player II has n moves c_1, c_2, \dots, c_n corresponding to the n columns of the




Figure 4.4 A 3×2 game and an extensive game that has this game as its strategic form.

game. Each move of player II leads to a leaf of the game tree. Because m possible moves of player I are succeeded by n moves of player II, the game tree has mn leaves. If the given game has payoffs a_{ij} and b_{ij} to player I and II when they play row i and column j , respectively, then these two payoffs are given at the leaf of the game tree that is reached by move r_i followed by move c_j . Because each player has only one information set, his strategies coincide with his moves. The strategic form of the constructed extensive game is then the original $m \times n$ game. Its rows are marked with the strategies r_1, \dots, r_m and its columns with c_1, \dots, c_n because this is how we chose to name the moves of the two players in the extensive game. Figure 4.4 gives an example.

\Rightarrow Exercise 4.1 on page 133 asks you to repeat this argument, with any one of the two players moving first and the other player moving second.

4.7 Reduced strategies

Consider the extensive game in Figure 4.5. Player II has three information sets, call them h , h' , and h'' , with $C_h = \{l, r\}$, $C_{h'} = \{a, b\}$, and $C_{h''} = \{c, d\}$. Consequently, she has $2 \times 2 \times 2 = 8$ pure strategies. Each pure strategy is a triple of moves, an element of $\Sigma_{II} = C_h \times C_{h'} \times C_{h''}$, for example lac (we write the triple of moves by just writing the three moves next to each other for brevity). If player I chooses T , then this strategy lac of player II will lead to the leftmost leaf of the game tree with the payoff pair $(1, 5)$ for players I, II, because the play of the game follows the moves l , T , and then a . Clearly, if player I plays T and player II plays the strategy lad , play will lead to the same leaf. In fact, irrespective of the strategy chosen by player I, the strategies lac and lad of player II have *always* the same effect. This can be seen directly from the game tree, because when player II makes move l , she will *never* reach the information set h'' where she has to decide between c and d .

This leads us to the concept of a *reduced strategy* which is already known for a game tree with perfect information. Namely, a reduced strategy of player i specifies a move for each information set of player i , *except* for those information sets of player i that are




Figure 4.5 Extensive game with one information set for player I and three information sets h, h', h'' for player II.

unreachable because of an *own* earlier move (that is, because of a move of player i). When writing down a reduced strategy, we give the moves for each information set (in a fixed order of the information sets), where an unspecified move is denoted by an extra symbol, the star “*”, which stands for any move at that unreachable information set.

In Figure 4.5, the information set h'' of player II is unreachable after her own earlier move l , and her information set h' is unreachable after her own earlier move r . Her reduced strategies are therefore la^*, lb^*, r^*c, r^*d . Note that the star stands for any unspecified move that would normally be at that place when considering a full (unreduced) strategy. So player II’s reduced strategy la^* can be thought of representing both her pure strategies lac and lad .

The *reduced strategic form* of an extensive game is the same as its strategic form, except that for each player, only reduced strategies are considered and not pure strategies. By construction, the resulting expected payoffs for each player are well defined, because the unspecified moves never influence the leaves of the game tree that are reached during play. Figure 4.6 gives the reduced strategic form of the game in Figure 4.5.

4.8 Perfect recall

What do the information sets of player II in Figure 4.5 describe? At the information set h , player II knows for certain that she is at the root of the game. At the information set h' , she does not know the move of player I, but she knows that she made the earlier move l at the root. At the information set h'' , she again does not know the move of player I, but she knows that she made the move r earlier at the root.

For comparison, consider the game in Figure 4.7, which has the same tree and the same payoffs as the game in Figure 4.5, except that the information sets h' and h'' have been united to form a new set h''' which has four nodes. Moreover, in order to conform

I \ II		la^*	lb^*	$r*c$	$r*d$
T		5	4	3	0
		1	3	0	1
B		0	2	3	4
		0	0	2	0

Figure 4.6 Reduced strategic form of the game in Figure 4.5. In the strategies of player II, a star * indicates a unspecified move at an information set that is unreachable due to an own earlier move. The numbers in boxes are best-response payoffs.




Figure 4.7 Extensive game where player II does not have perfect recall.

with the definition of an extensive game, the moves at h''' have to be the same at all nodes; they are here called a and b which replace the earlier moves c and d at h'' . When player II is to make a move at h''' , she knows nothing about the game state except that she has made some move at the beginning of the game, followed by either move T or B of player I. In particular, the interpretation of h''' implies that player II *forgot* her earlier move l or r that she made at the root. We say that player II has *imperfect recall* in the game in Figure 4.7.

Imperfect recall creates difficulties when interpreting the game, and contradicts our idea of “rational” players. All the extensive games that we will analyse have *perfect recall*, which we now proceed to define formally. For example, Figure 4.5 shows a game with perfect recall. There, player II has two information sets h' and h'' that have more than one node, so these are information sets where player II has some lack of information. We can test whether this lack of information involves her earlier moves. Namely, each node in an information set, say in h' , is reached by a unique path from the root of the game tree. Along that path, we look at the moves that the players make, which is move l of player II, followed by move T of player I for the left node of h' , or followed by move B for the right

node of h' . (In general, this may also involve chance moves.) Among these moves, we *ignore* all moves made by the other players (which is here only player I) or by chance. Consequently, we see that each node in h' is preceded by move l of player II. That is, each node in the information set has the *same sequence of own earlier moves*. This is the definition of perfect recall.

Definition 4.1 Player i in an extensive game has *perfect recall* if for every information set h of player i , all nodes in h are preceded by the same sequence of moves of player i . The extensive game is said to have perfect recall if all its players have perfect recall.

To repeat, the crucial condition about the nodes in an information set h of player i is that they are preceded by the same *own* earlier moves. We have to disregard the moves of the other players because there must be something that player i is lacking information about if player i does not have perfect information.

⇒ Make sure you clearly understand the distinction between perfect information and perfect recall. How are the respective concepts defined?

We can verify that all players have perfect recall in Figure 4.5. We have already verified the condition of definition 4.1 for h' . The other information sets of player II are h and h'' . For h , the condition holds trivially because h has only a single node (so we can make the easy observation that in a game with perfect information, that is, singleton information sets, all players have perfect recall). For h'' , both nodes in h'' are preceded by the sequence of own moves that consists of the single move r . Player I has only one information set, and for both nodes in that set the sequence of own earlier moves is *empty*, because player I has not made a move before. It is useful to consider the empty sequence, often denoted by \emptyset like the empty set, as a sequence of moves, so that definition 4.1 can be applied.

In Figure 4.7, player II does *not* have perfect recall because, for example, the leftmost node of her information set h''' is preceded by her own earlier move l , whereas the rightmost node is preceded by the move r . So the sequences of own earlier moves are not the same for all nodes in h''' .

One way to understand definition 4.1 is via the interpretation of information sets. They represent a player's lack of knowledge about the game state, that is, at which of the nodes in the information set the player is at that moment. When the sequence of own earlier moves is the *same* for all these nodes, then the player can indeed not get additional information about where he is in the information set from knowing what moves he already played (as should normally be the case). If, on the other hand, some nodes in the information set are preceded by *different* own earlier moves, then the player can only fail to know where he is in the information set when he has *forgotten* these moves. For that reason, we say he has imperfect recall. So a player with perfect recall cannot have such information sets, because they would not accurately reflect his knowledge about the game state.

However, definition 4.1 is not about how to *interpret* the game. Perfect recall is a *structural property of the information sets*. This property is easily verified by considering




Figure 4.8 Single-player game where player I does not have perfect recall. Moves at different information sets are always considered as distinct, even if they are named the same (for example, here B).

for each information set the nodes therein and checking that they are reached by the same sequence of earlier own moves, which can be done without interpreting the game at all.

Figure 4.8 shows a game with only one player, who has an information set (with moves L and R) that contains two nodes. Both nodes are preceded by an own earlier move named “ B ”, so it seems that player I has perfect recall in this game. However, this would obviously not be the case if the two singleton information sets had moves that are named differently, for example T and B at the left information set and C and D at the right information set (as in Figure 4.1). Indeed, player I does *not* have perfect recall in the game in Figure 4.8 because moves at *different* information sets are always considered *distinct*. A formal way of saying this is that the choice sets C_h and $C_{h'}$ are disjoint whenever $h \neq h'$. This can be assumed without loss of generality, either by re-naming the moves suitably, or by assuming that in addition to the name of the move that is indicated in the game tree, part of the “identity” of a move $c \in C_h$ is the information set h where that move is made.

In the game in Figure 4.8, player I has imperfect recall, because he forgot his *earlier knowledge*, namely the outcome of the chance move. In contrast, player II in Figure 4.7 has not forgotten any earlier knowledge (all she knew earlier was that she makes the first move in the game, which she still knows when she reaches the information set h'''), but she has forgotten the move that she made. For that reason, it is sometimes said that a player has perfect recall if she never forgets what she “knew or did earlier”. Definition 4.1 is a precise way of stating this.

⇒ Exercise 4.2 on page 133 tests your understanding of perfect recall. It also gives you an opportunity to apply the methods that you learned in Chapter 3.

4.9 Perfect recall and order of moves*

This section concerns perfect recall and the “time structure” of a game. First, we note that a situation like in Figure 4.3 cannot occur.

Proposition 4.2 *If a player has perfect recall, no two nodes in an information set of that player are connected by a path.*

Proof. Suppose otherwise, that is, h is an information set of a player who has perfect recall, so that there are nodes in h where one node precedes the other on a path in the game tree. If there are several such nodes, let u be the earliest such node (that is, there is no other node in h from which there is a path to u), and let v be a node in h so that there is a path from u to v . (An example is Figure 4.3 with u as the root of the game tree and v as the second node in player I’s information set.) On the path from u to v , the player makes a move c at the decision node u . This move c precedes v . However, there is no other node in h where move c precedes u , by the choice of u , so u is not preceded by c . Consequently, the sets of the player’s own moves that precede v and u are already different, so the sequences of own earlier moves for the two nodes certainly differ as well. So the player does not have perfect recall, which is a contradiction. \square

Figure 4.9 shows a game where a chance move decides whether player I is to move first, choosing between T and B , or player II, who can choose between l and r . Move T of player I terminates the game, and so does move r of player II. However, with move B (respectively, move l), play continues, and then it is the other player’s turn to move, who is not informed whether he or she moves first or second. That is, the time order in which the players move depends on the chance move. Nevertheless, both players have perfect recall, because both nodes in a player’s information set are preceded by the same (empty) sequence of own earlier moves.




Figure 4.9 Extensive game where players do not know who moves first.

The game looks somewhat strange because it has crossing tree edges. However, these are easier to draw than crossing information sets. This game cannot be drawn with horizontal information sets whose nodes are on the same level (and tree edges going downward) because of the different time order of the nodes.




Figure 4.10 Single-player extensive game with an information set h containing two nodes that have the same set $\{B, L\}$ of own earlier moves, but different sequences LB and BL of own earlier moves.

Our next question is whether we can state definition 4.1 less restrictively by saying that a player has perfect recall if for any of his information sets h , any two nodes in h are preceded by the same earlier moves. In other words, does it suffice to say that the set of own earlier moves is the same for all nodes in an information set? (As before, we always assume that moves at different information sets are distinct from each other.) Figure 4.10 shows a game where both nodes in the last information set h of player I are preceded by move B and move L , which are his earlier moves. However, at the left node of h , he made move L first and move B second (so this is the sequence LB of own earlier moves), and at the right node of h it is the other way around. According to definition 4.1, player I already fails to have perfect recall because of the different sequences.

However, suppose we are more generous as concerns what a player I has to remember, in this case only that player I knows, when reaching h , that he made moves B and L , but does not remember in which order. In other words, suppose that only the set of own earlier moves matters for the definition of perfect recall, but not the sequence. This condition is *not* violated for h in Figure 4.10. We claim that even with this seemingly laxer definition of perfect recall, some other information set of the player violates this condition, so the player cannot have perfect recall according to definition 4.1.

Proposition 4.3 *Suppose that for every information set h of player i , any two nodes in h are preceded by the same set of moves of player i . Then player i has perfect recall.*

Proof. We have to show that for any two nodes in an information set h of player i , not only the set but also the sequence of own earlier moves is the same. Suppose, by contradiction, that there is an information set where two nodes are preceded by the same set of own earlier moves, but where the corresponding sequences (that is, the order of moves) differ. What we will prove is that in some *other* information set even the sets of own earlier moves differ. In other words, if for every information set, all nodes are preceded by the same set of earlier moves, then for every information set, all nodes are also preceded by the same sequence of earlier moves.

That is, consider two nodes in an information set so that there are two earlier moves that have their order exchanged, like B and L in Figure 4.10. (The argument holds generally.) Then for one path in the game tree (in Figure 4.10 when the chance move goes left), move B is made before move L (and L is made at some node u , say; see Figure 4.10), and for some other path, move L is made (at some node v) before move B (and move B is made at node w , say). The nodes u and v belong to the same information set because at both nodes player i makes the move L . This means that node u is preceded by B . But if node v was also preceded by move B , where that move B is made at node x , then there would be a path from x to w , where x and w belong to the same information set (where move B is made), contradicting the Proposition 4.2. So B belongs to the set of moves that precede node u , but not node v , which contradicts the assumption that all nodes in an information set of player i are preceded by the same set of earlier moves. \square

In other words, we can be “lax” when checking the definition of perfect recall: All we have to do is to look at every information set of the player, and check if “the earlier moves” are the same (irrespective whether they are considered as a set or as a sequence of earlier moves).

A second, somewhat cruder lesson can be learned from the examples in figures 4.3, 4.8, and 4.10: A game with a *single* player is likely not have perfect recall if it has information sets with two or more nodes.

\Rightarrow Construct a one-player game with perfect recall where not all information sets are singletons. Why do you need a chance move?

4.10 Behaviour strategies

Our next topic concerns ways of randomising in an extensive game. Unlike games with perfect information, which always have a Nash equilibrium in pure strategies that can be found by backward induction, games with imperfect information may require randomised strategies, as the example in Figure 4.1 has demonstrated.

We find all Nash equilibria of the game in Figure 4.5 that has the reduced strategic form shown in Figure 4.6. The only equilibrium in pure strategies is the strategy pair (T, la^*) . In order to find further mixed equilibria, if any, we use the upper envelope




Figure 4.11 Upper envelope of expected payoffs to player II for the game in Figure 4.6.

method. Figure 4.11 shows the expected payoffs to player II as a function of the probability that player I chooses B . As this figure shows, the only pairs of strategies between which player II can be indifferent are la^* and lb^* when $\text{prob}(B) = 1/3$; or lb^* and r^*c when $\text{prob}(B) = 1/2$; or r^*c and r^*d when $\text{prob}(B) = 3/4$. The first of these cannot be played in an equilibrium because against both la^* and lb^* , the best response of player I is T , so that player I cannot be made indifferent between T and B , which is necessary for player I wanting to mix between T and B in equilibrium. In the other two cases, however, the pure best responses of player I differ, and it is possible to obtain Nash equilibria. These equilibria are $((1/2, 1/2), (0, 2/5, 3/5, 0))$ and $((1/4, 3/4), (0, 0, 1/3, 2/3))$.

What do these mixed strategies mean in terms of the extensive game? For player I, the mixed strategy $(1/2, 1/2)$ means that he chooses T and B with equal probability. For player II, the mixed strategy $(0, 2/5, 3/5, 0)$ amounts to choosing lb^* with probability $2/5$, and r^*c with probability $3/5$. The interpretation of such a mixed strategy is that each player chooses *at the beginning of the game* a reduced pure strategy, which is a plan of action of what do at each reachable information set. The player then sticks to this plan of action throughout the game. In particular, if player II has chosen to play lb^* , she first makes the move l with certainty, and then move b at her second information set h' . Her reduced pure strategy r^*c is a similar plan of action.

There is a different way of playing such a mixed strategy, which is called a *behaviour strategy*. A behaviour strategy is a randomised choice of a *move* for each information set. That is, rather than using a lottery to determine a (reduced) pure strategy at the beginning of the game, the player uses a “local” lottery whenever she reaches an information set. In our example, player II’s mixed strategy $(0, 2/5, 3/5, 0)$ is then played as follows: At her first information set h , she plays moves l and r with probabilities $2/5$ and $3/5$, respectively. When she reaches her information set h' (which is only possible when she has made the earlier move l), she chooses move a with probability zero and move b with probability one; this “behaviour” at h' is deterministic. Similarly, when she reaches her information set h'' (having made move r earlier), then she chooses move c with certainty.

Definition 4.4 In an extensive game, a *behaviour strategy* β of player i is defined by a probability distribution on the set of moves C_h for each information set h of player i . It is given by a probability $\beta(c)$ for each $c \in C_h$, that is, a number $\beta(c) \geq 0$ for each $c \in C_h$ so that $\sum_{c \in C_h} \beta(c) = 1$. In a *reduced* behaviour strategy, these numbers are left unspecified if h is unreachable because all earlier own moves by player i that allow her to reach h have probability zero under β .

The concept of a reduced behaviour strategy is illustrated by the other mixed equilibrium of the game in Figure 4.5, with the mixed strategy $(0, 0, 1/3, 2/3)$ of player II. This is a randomisation between the two pure strategies $r*c$ and $r*d$. It can also be played as a (reduced) behaviour strategy. Namely, at her first information set h , player II chooses l and r with probability 0 and 1, respectively. Then the information set h' is unreachable because it is only reachable when player II makes move l earlier; this event has probability zero. Consequently, the behaviour probabilities for making moves a and b are not specified, as in a reduced behaviour strategy. (To obtain an “unreduced” behaviour strategy, one would specify arbitrary probabilities for a and b .) At her information set h'' , player II chooses c and d with probability $1/3$ and $2/3$, respectively.

A behaviour strategy can be played by “delaying” the player’s random choice until she has reached an information set, when she makes the next move. However, nothing prevents the player making these random choices in advance for each information set, which will then give her a pure strategy that she plays throughout the game. Then the behaviour strategy is used like a mixed strategy. In other words, *behaviour strategies can be considered as special mixed strategies*.

A behaviour strategy β of player i defines a mixed strategy μ in the following way: What we need is the probability $\mu(\pi)$ for a particular pure strategy π in Σ_i ; recall from (4.2) that Σ_i is the set of pure strategies of player i . This pure strategy π defines a move $\pi(h)$ for each information set $h \in H_i$. When player i uses the behaviour strategy β , that move has a certain probability $\beta(\pi(h))$. Because these random moves at the information sets h of player i are made independently, the probability that *all* the moves specified by π are made is the product of these probabilities. That is, the mixed strategy probability is

$$\mu(\pi) = \prod_{h \in H_i} \beta(\pi(h)). \quad (4.3)$$

We have been using equation (4.3) when considering a mixed strategy, for example $\mu = (0, 0, 1/3, 2/3)$ in Figure 4.5, as a behaviour strategy. That behaviour strategy β was given by $\beta(l) = 0$, $\beta(r) = 1$, $\beta(c) = 1/3$, and $\beta(d) = 2/3$, so that $2/3 = \mu(r*d) = \beta(r) \cdot \beta(*) \cdot \beta(d)$. This is expressed as a product of *three* behaviour probabilities, one for each information set, because the pure strategies are given by a move for each information set; note that in (4.3), the product has to be taken over all information sets h of player i . There is one complication here, because we are already considering reduced behaviour strategies and reduced pure strategies: What is meant by $\beta(*)$? This means the behaviour probability for making *any* move at h' (in this case, as part of the pure strategy $r*d$), which is unspecified. Clearly, that probability is equal to 1. So we take $\beta(*) = 1$ whenever applying (4.3) to reduced strategies. Equivalently, we may just omit the unspecified behaviour probabilities from the product, writing simply $\mu(r*d) = \beta(r) \cdot \beta(d)$.

The convention $\beta(*) = 1$ is easily seen to agree with the use of (4.3) for unreduced strategies. For example, consider the unreduced strategic form where $r*d$ stands for the two pure strategies rad and rbd . Each of these has a certain probability under both μ and β (when extended to the unreduced strategic form), where μ is the representation of β as a mixed strategy. Whatever these probabilities are (obtained by splitting the probability for $r*d$ into probabilities for rad and rbd), we certainly have

$$\mu(rad) + \mu(rbd) = \beta(r)\beta(a)\beta(d) + \beta(r)\beta(b)\beta(d) = \beta(r) \cdot (\beta(a) + \beta(b)) \cdot \beta(d).$$

The left-hand side is the proper interpretation of the mixed strategy probability $\mu(r*d)$, the right-hand side of $\beta(r) \cdot \beta(*) \cdot \beta(d)$, in agreement with $\beta(*) = \beta(a) + \beta(b) = 1$.




Figure 4.12 Game with four pure strategies of player II. Here, not all mixed strategies of player II are representable as behaviour strategies.

While every behaviour strategy can be considered as a mixed strategy, the converse does not hold. Figure 4.12 shows a game where player II has the four pure strategies ac , ad , bc , bd , which are also her reduced pure strategies. Consider the mixed strategy $\mu = (1/2, 0, 0, 1/2)$ of player II, where she chooses the pure strategy ac , that is, moves a and c , with probability $1/2$, otherwise the pure strategy bd . Clearly, every move is made with positive probability overall (in fact, probability $1/2$ for each single move). Then if (4.3) was true, then every move combination would have to have positive probability as well! However, this is not the case, because the strategy ad , for example, has probability zero under μ . The reason is that when player II uses μ , the random choice between the moves a and b is *correlated* with the random choice between c and d . That is, knowing that move a has been made gives some information about whether move c or d has been made (here, the correlation is maximal: move a implies move c). If these random choices were made *independently*, as required in a behaviour strategy, then the random choice between a and b would not affect the random choice between c and d .

In this example, the “behaviour” of player II at the left information set in Figure 4.12 described by μ is to choose a with probability $1/2$ and b with probability $1/2$, and ditto moves c and d . This defines a behaviour strategy β of player II. This behaviour strategy defines the mixed strategy $(1/4, 1/4, 1/4, 1/4)$ because it chooses each pure strategy with equal probability. This mixed strategy is clearly different from μ .

A second example using Figure 4.12 is the mixed strategy $\mu' = (1/6, 1/3, 1/4, 1/4)$. The corresponding behaviour strategy β' has move probabilities $\beta'(a) = \beta'(b) = 1/2$,

because $\beta'(a) = \mu'(ac) + \mu'(ad) = 1/6 + 1/3$, and $\beta'(c) = \mu'(ac) + \mu'(bc) = 1/6 + 1/4 = 5/12$ and $\beta'(d) = \mu'(ad) + \mu'(bd) = 1/3 + 1/4 = 7/12$. Again, making these moves independently with β' results in different probabilities for the pure strategies, for example in $\beta'(a) \cdot \beta'(c) = 5/24$ for ac , which is different from $\mu'(ac)$.

What we have done in this example is to take a mixed strategy μ of a player and “observe” at each information set of that player the “behaviour” resulting from μ . This behaviour defines a certain behaviour strategy β . In a similar way, β' is obtained from μ' . The behaviour strategy may differ from the original mixed strategy because with the behaviour strategy, random choices at different information sets are no longer correlated. However, what does *not* change are the probabilities of reaching the nodes of game tree, given some (pure or mixed) strategy of the other players. We say that the strategies μ and β are *equivalent* (similarly, μ' and β' are equivalent). This equivalence holds in any extensive game with perfect recall. This is the topic of the next section.

⇒ You should now attempt Exercise 4.3 on page 133.

4.11 Kuhn’s theorem: behaviour strategies suffice

Information sets were introduced by Harold Kuhn in 1953. His article also defined perfect recall and behaviour strategies, and proved an important connection between the two concepts. This connection is the central property of extensive games and is often called “Kuhn’s theorem”. It says that a *player who has perfect recall can always replace a mixed strategy by an equivalent behaviour strategy*.

Why should we even care about whether a player can use a behaviour instead of a mixed strategy? The reason is that a behaviour strategy is much *simpler to describe* than a mixed strategy, significantly so for larger games. As an example, consider a simplified Poker game where each player is dealt only a single card, with 13 different possibilities for the rank of that card. A game tree for that game may start with a chance move with 13 possibilities for the card dealt to player I, who learns the rank of his card and can then either “fold” or “raise”. Irrespective of the rest of the game, this defines already $2^{13} = 8192$ move combinations for player I. Assuming these are all the pure strategies of player I, a mixed strategy for player I would have to specify 8191 probabilities (which determine the remaining probability because these probabilities sum to one). In contrast, a behaviour strategy would specify one probability (of folding, say) for each possible card, which are 13 different numbers. So a behaviour strategy is much less complex to describe than a mixed strategy.

This example also gives an intuition *why* a behaviour strategy suffices. The moves at the different decision nodes of player I can be randomised independently because they concern disjoint parts of the game. There is no need to correlate the moves for different cards because in any one game, only one card is drawn.

In general, a player may make moves *in sequence* and a priori there may be some gain in correlating an earlier move with a subsequent move. This is where the condition of perfect recall comes in. Namely, the knowledge about any earlier move is already

captured by the information sets because perfect recall says that the player knows all his earlier moves. So there is no need to “condition” a later move on an earlier move because all earlier moves are known.

For illustration, consider the game in Figure 4.7 which does not have perfect recall. It is easily seen that this game has the same strategic form, shown in Figure 4.6, as the game in Figure 4.5, except that the pure strategies of player II are la , lb , ra , rb instead of the reduced strategies la^* , lb^* , r^*c , r^*d . Here, the mixed strategy $(0, 2/5, 3/5, 0)$ of player II, which selects the pure strategy lb with probability $2/5$, and ra with probability $3/5$, *cannot* be expressed as a behaviour strategy. Namely, this behaviour strategy would have to choose moves l and r with probability $2/5$ and $3/5$, respectively, but the probabilities for moves a and b , which would be $3/5$ and $2/5$, are *not* independent of the earlier move (because otherwise the pure strategy la , for example, would not have probability zero). So this mixed strategy, which is part of a Nash equilibrium, cannot be expressed as a behaviour strategy. Using a mixed strategy gives player II more possibilities of playing in the game.

In contrast, the mixed strategy $(0, 2/5, 3/5, 0)$ of player II in the perfect-recall game in Figure 4.5 chooses lb^* and r^*c with probability $2/5$ and $3/5$, respectively. As shown earlier, this is in effect a behaviour strategy, because correlating move l with move b , and move r with move c , is done via the information sets h' and h'' , which are separate in Figure 4.5 but merged in Figure 4.7. That is, player II *knows* that she made move l when she has to make her move b at h' , and she knows that she made move r when she has to make her move c at h'' .

Before stating Kuhn's theorem, we have to define what it means to say that two randomised strategies μ and ρ of a player are *equivalent*. This is the case if for any fixed strategies of the other players, every node of the game tree is reached with the same probability when the player uses μ as when he uses ρ . (Sometimes μ and ρ are also called “realisation equivalent” in this case because every node in the game tree has the same “realisation probability” when the player uses μ as when he uses ρ .) Two equivalent strategies μ and ρ always give every player the same expected payoff because the leaves of the game tree are reached with the same probability no matter whether μ or ρ is used.

When are two strategies μ and ρ equivalent? Suppose these are strategies of player i . Consider a particular node, say u , of the game tree. Equivalence of μ and ρ means that, given that the other players play in a certain way, the probability that u is reached is the same when player i uses μ as when he uses ρ . In order to reach u , the play of the game has to make a certain sequence of moves, including chance moves, that lead to u . This move sequence is unique because we have a tree. In that sequence, some moves are made by chance or by players other than player i . These moves have certain given probabilities because the chance move probabilities and the strategies adopted by the other players are fixed. We can assume that the overall probability for these moves is not zero, by considering fixed strategies of the other players which allow to reach u (otherwise u is trivially reached with the same zero probability with μ as with ρ). So what matters in order to compare μ and ρ are the moves by player i on the path to u . If the probability for playing the *sequence of moves of player i* on the path to u is the *same* under μ and ρ , and

if this holds for every node u , then μ and ρ are equivalent. This is the key ingredient of the proof of Kuhn's theorem.

Theorem 4.5 (Kuhn [1953]) *If player i in an extensive game has perfect recall, then for any mixed strategy μ of player i there is an equivalent behaviour strategy β of player i .*

Proof. We illustrate the proof with the extensive game in Figure 4.13, with player II as player i . The payoffs are omitted; they are not relevant because only the probabilities of reaching the nodes of the game tree matter.




Figure 4.13 Game illustrating sequences of moves of player II. Payoffs are omitted.

Given μ , we want to construct an equivalent behaviour strategy β . Essentially, the probability $\beta(c)$ for any move c at an information set h of player i will be the “observed” probability that player i makes move c when she uses μ . In order to obtain this probability, we consider *sequences* of moves of player i . Such a sequence is obtained by considering any node x of the game tree and taking the moves of player i on the path from the root to x ; the resulting sequence shall be denoted by $\sigma^i(x)$. For example, if x is the leaf w of the game tree in Figure 4.13 and $i = \text{II}$, then that sequence is $\sigma^{\text{II}}(w) = lbc$, that is, move l followed by b followed by c . If all players are considered, then node w is reached by the sequence of moves $lLbc$ which includes move L of player I, but that move is ignored in the definition of $\sigma^{\text{II}}(w)$; for player I, the sequence of moves defined by w is just the move L , that is, $\sigma^{\text{I}}(w) = L$.

According to definition 4.1, the condition of perfect recall states that the sequence of own earlier moves is the same for all nodes in an information set. That is,

$$h \in H_i \implies \forall u, v \in h \quad \sigma^i(u) = \sigma^i(v).$$

An example is given in Figure 4.13 by the information set h where player II has the two moves c and d , which has two nodes u and v with $\sigma^{\text{II}}(u) = \sigma^{\text{II}}(v) = lb$. Consequently,

the sequence of moves *leading to* h can be defined uniquely as $\sigma_h = \sigma^i(u)$ for any $u \in h$, where i is identified as the player who moves at h , that is, $h \in H_i$. Furthermore, we can extend the sequence σ_h that leads to h by any move c at h , which we write as the longer sequence $\sigma_h c$, to be read as “ σ_h followed by c ”. In the example, we have $\sigma_h = lb$, and it can be extended either by move c or by move d , giving the longer sequences lbc and lbd . Any sequence of moves of player i is either the empty sequence \emptyset (for example $\sigma^i(r) = \emptyset$ if r is the root of the tree), or it has a last move c made at some information set h of player i , and then it can be written as $\sigma_h c$. So these sequences $\sigma_h c$ for $h \in H_i$, $c \in C_h$ and \emptyset are all possible sequences of moves of player i .

We now consider the *probability* $\mu[\sigma]$ that such a sequence σ of player i is played when player i uses μ . This probability $\mu[\sigma]$ is simply the combined probability under μ for all the pure strategies that *agree* with σ in the sense that they prescribe all the moves in σ . For example, if $\sigma = lb$ in Figure 4.13, then the pure strategies of player i that agree with σ are the strategies that prescribe the moves l and b , which are the pure strategies lbc and lbd . So

$$\mu[lb] = \mu(lbc) + \mu(lbd), \quad (4.4)$$

where $\mu(lbc)$ and $\mu(lbd)$ are the mixed strategy probabilities for the pure strategies lbc and lbd .

The longer the sequence σ is, the smaller is the set of pure strategies that agree with σ . Conversely, shorter sequences have more strategies that agree with them. The extreme case is the empty sequence \emptyset , because every pure strategy agrees with that sequence. Consequently, we have

$$\mu[\emptyset] = 1. \quad (4.5)$$

We now look at a different interpretation of equation (4.4). In this special example, lbc and lbd are also sequences of moves, namely $lbc = \sigma_h c$ and $lbd = \sigma_h d$ with $\sigma_h = lb$. So the combined probability $\mu[lbc]$ of all pure strategies that agree with lbc is simply the mixed strategy probability $\mu(lbc)$, and similarly $\mu[lbd] = \mu(lbd)$. Clearly,

$$\mu[lb] = \mu[lbc] + \mu[lbd].$$

In general, considering any information set h of player i , we have the following important equation:

$$\mu[\sigma_h] = \sum_{c \in C_h} \mu[\sigma_h c]. \quad (4.6)$$

This equation holds because because the sequence σ_h that leads to h is extended by some move c at h , giving the longer sequence $\sigma_h c$. When considering the pure strategies that agree with these extended sequences, we obtain exactly the pure strategies that agree with σ_h .

We are now in a position to define the behaviour strategy β that is equivalent to μ . Namely, given an information set h of player i so that $\mu[\sigma_h] > 0$, let

$$\beta(c) = \frac{\mu[\sigma_h c]}{\mu[\sigma_h]}. \quad (4.7)$$

Because of (4.6), this defines a probability distribution on the set C_h of moves at h , as required by a behaviour strategy. In (4.7), the condition of perfect recall is used because otherwise the sequence σ_h that leads to the information set h would not be uniquely defined.

What about the case $\mu[\sigma_h] = 0$? This means that no pure strategy that agrees with σ_h has positive probability under μ . In other words, h is unreachable when player i uses μ . We will show soon that $\mu[\sigma] = \beta[\sigma]$ for all sequences σ of player i . Hence, h is also unreachable when the player uses β , so that $\beta(c)$ for $c \in C_h$ does not need to be defined, that is, β is a *reduced* behaviour strategy according to definition 4.4. Alternatively, we can define some behaviour at h (although it will never matter because player i will not move so as to reach h) in an arbitrary manner, for example by giving each move c at h equal probability, setting $\beta(c) = 1/|C_h|$.

We claim that β , as defined by (4.7), is equivalent to μ . We can define $\beta[\sigma]$ for a sequence σ of player i , which is the probability that all the moves in σ are made when the player uses β . Clearly, this is simply the product of the probabilities for the moves in σ . That is, if $\sigma = c_1 c_2 \cdots c_n$, then $\beta[\sigma] = \beta(c_1)\beta(c_2)\cdots\beta(c_n)$. When $n = 0$, then σ is the empty sequence \emptyset , in which case we set $\beta[\emptyset] = 1$. So we get by (4.7)

$$\beta[c_1 c_2 \cdots c_n] = \beta(c_1)\beta(c_2)\cdots\beta(c_n) = \frac{\mu[c_1]}{\mu[\emptyset]} \frac{\mu[c_1 c_2]}{\mu[c_1]} \cdots \frac{\mu[c_1 c_2 \cdots c_{n-1} c_n]}{\mu[c_1 c_2 \cdots c_{n-1}]} = \mu[c_1 c_2 \cdots c_n]$$

because all the intermediate terms cancel out, and by (4.5). (This is best understood by considering an example like $\sigma = lbc = c_1 c_2 c_3$ in Figure 4.13.) Because $\beta[\sigma] = \mu[\sigma]$ for all sequences σ of player i , the strategies β and μ are equivalent, as explained before Theorem 4.5. The reason is that any node u of the game tree is reached by a unique sequence of moves of *all* players and of chance. The moves of the other players have fixed probabilities, so all that matters is the sequence $\sigma^i(u)$ of moves of player i , which is some sequence σ , which has the same probability no matter whether player i uses β or μ . \square

4.12 Subgames and subgame perfect equilibria

In an extensive game with perfect information, a subgame is any subtree of the game, given by a node of the tree as the root of the subtree and all its descendants. This definition has to be refined when considering games with imperfect information, where every player has to *know* that they have entered the subtree.

Definition 4.6 In an extensive game, a *subgame* is any subtree of the game tree so that every information set is either disjoint to or a subset of the nodes in the subtree.

In other words, if any information set h contains some node of the subtree, then h must be completely contained in the subtree. This means that the player to move at h knows that he or she is in the subgame. Information sets that have no nodes in the subtree do not matter.




Figure 4.14 Extensive game and its reduced strategic form. Only one Nash equilibrium of the game is subgame perfect.

In Figure 4.14, the subtree with player I's decision node as root defines a subgame of the game. On the other hand, any of the two nodes in player II's second information set are not roots of subgames, because this information set is neither contained in nor disjoint to the respective subtree.

A trivial subgame is either the game itself, or any leaf of the game tree. Some games have no other subgames, for example the games in figure 4.1 or 4.5.

⇒ Test your understanding of subgames with Exercise 4.4 on page 133.

If every player has perfect recall, then any mixed strategy can be replaced by an equivalent behaviour strategy by Kuhn's Theorem 4.5. The strategies in any Nash equilibrium can therefore be assumed to be behaviour strategies. A behaviour strategy can easily be restricted to a subgame, by ignoring the behaviour probabilities for the information sets that are disjoint to the subgame.

Definition 4.7 In an extensive game with perfect recall, a *subgame perfect Nash equilibrium* (SPNE) is a profile of behaviour strategies that defines a Nash equilibrium for every subgame of the game.

In Figure 4.14, the reduced strategic form of the extensive game is shown on the right. If one ignores the third strategy r^* of player II in this strategic form, the resulting 2×2 game, with columns called a and b instead of la and lb , defines the strategic form of the subgame that starts with the move of player I. This subgame has a unique mixed equilibrium, where player I chooses T, B with probabilities $2/3, 1/3$, and player II chooses a, b with probabilities $1/2, 1/2$. The resulting payoffs to player I and II are 2 and 4, which can be substituted for the subgame to obtain a smaller game where player II chooses between l and r . Then it is optimal for player II to choose l with certainty. This defines a behaviour strategy for each player, which is an SPNE.

The step of computing an equilibrium of a subgame and substituting the resulting equilibrium payoffs for the subgame generalises the backward induction step in games with perfect information. In games with perfect information, this backward induction step amounts to finding an optimal move of a single player in a very simple subgame where each move leads directly to a payoff. The optimal move defines an equilibrium in that subgame. In a game with imperfect information, the subgame may involve more than one player, and the equilibrium is usually more difficult to find.

The described SPNE of the extensive game in Figure 4.12 is the mixed-strategy equilibrium $((2/3, 1/3), (1/2, 1/2, 0))$ of the reduced strategic form, which also defines a pair of behaviour strategies. Another Nash equilibrium is the pair (B, r^*) of reduced pure strategies. The strategy r^* is only a reduced behaviour strategy for player II, because it does not specify the probabilities for the moves a and b at her second information set (which is unreachable when she makes move r). However, no matter how the behaviour probabilities for player II's moves a and b are specified, they are not part of a Nash equilibrium of the subgame when player I makes move B . Consequently, we can safely say that (B, r^*) does not correspond to an SPNE of the game.

In general, just as for games with perfect information, an equilibrium of an extensive game can only be called subgame perfect if each player's behaviour (or pure) strategy is not reduced, but is fully specified and defines a probability distribution (or deterministic move) for each information set of the player.

We described for the example in Figure 4.14 how to find an SPNE, which generalises the backward induction process. Starting with each leaf of the game tree (which represents a trivial subgame), one considers any subgame that does not have any non-trivial subgames on its own. This subgame has an equilibrium by Nash's theorem. The mixed strategies in this equilibrium can be represented as behaviour strategies by Kuhn's theorem. The resulting expected payoffs to the players in the equilibrium are then assigned to a new leaf that substitutes the considered subgame, and then this is iterated. (Unlike in games with perfect information, however, the step of finding an equilibrium of a subgame may involve a large game, for example if the game has no non-trivial subgames at all, where the next subgame considered after each leaf is the entire game.) This construction gives an SPNE; we omit the proof that we obtain indeed an equilibrium of the entire game, which is analogous to the theorem that asserts that backward induction defines an equilibrium. We summarise the existence of SPNE in the last theorem of this chapter.

Theorem 4.8 *Any extensive game with perfect recall has an SPNE in behaviour strategies.*

4.13 Exercises for Chapter 4

Exercise 4.1 shows how, with the help of information sets, any strategic-form game (where the players move simultaneously) can be represented as an extensive game. Exercise 4.2 asks to look at a number of extensive games, to check whether they have perfect recall, and to find all their equilibria using the methods from earlier chapters. Exercise 4.4

shows that information sets can intersect a path in the game tree only once, and tests your understanding of the concept of subgames. In Exercise 4.3, you have to construct an extensive game with imperfect information, which tests your understanding of information sets; for the analysis of this game, the theory of zero-sum games is useful.

Exercise 4.1 Explain two ways, with either player I or player II moving first, of representing an $m \times n$ bimatrix game as an extensive game that has the given game as its strategic form. Giving an example may be helpful. In each case, what are the number of decision nodes of player I and of player II, respectively? How many terminal nodes does the game tree have? Why do we need information sets here?

Exercise 4.2 Which of the six pictures in Figure 4.15 define extensive games where each player has perfect recall, and which do not? The games (a) and (b) are zero-sum games, with payoffs to player I. For each of the “legal” extensive games with perfect recall, find *all* Nash equilibria in pure or mixed strategies. Justify your answers.

Exercise 4.3 Consider the following two-person zero-sum game: A deck has three cards (of different ranks: High, Middle, Low), and each player is dealt a card. All deals are equally likely, and of course the players get different cards. A player does not know the card dealt to the other player. After seeing his hand, player I has the option to Raise (R) or Fold (F). When he folds, he loses one unit to the other player. When he raises, player II has the option to meet (m) or pass (p). When player II chooses “pass”, he will have to pay player I one unit. When player II chooses “meet”, the higher card wins, and the player with the lower card has to pay two units to the winning player.

- (a) Draw a game in extensive form that models this game, with information sets, and payoffs to player I at the leaves.
- (b) Simplify this game by assuming that at an information set where a player’s move is always at least as good as the other move, no matter what the other player’s move or the chance move were or will be, then the player will choose that move. Draw the simplified extensive form game.
- (c) What does the simplification in (b) have to do with weakly dominated strategies? Why is the simplification legitimate here?
- (d) Give the strategic form of the simplified game from (b) and find an equilibrium of that game. What are the behaviour probabilities of the players in that equilibrium? What is the unique payoff to the players in equilibrium?

Exercise 4.4 For an extensive game with perfect recall, show that the the information set containing the root (starting node) of any *subgame* is always a singleton (that is, contains only that node).

[Hint: Use Proposition 4.2 on page 120.]

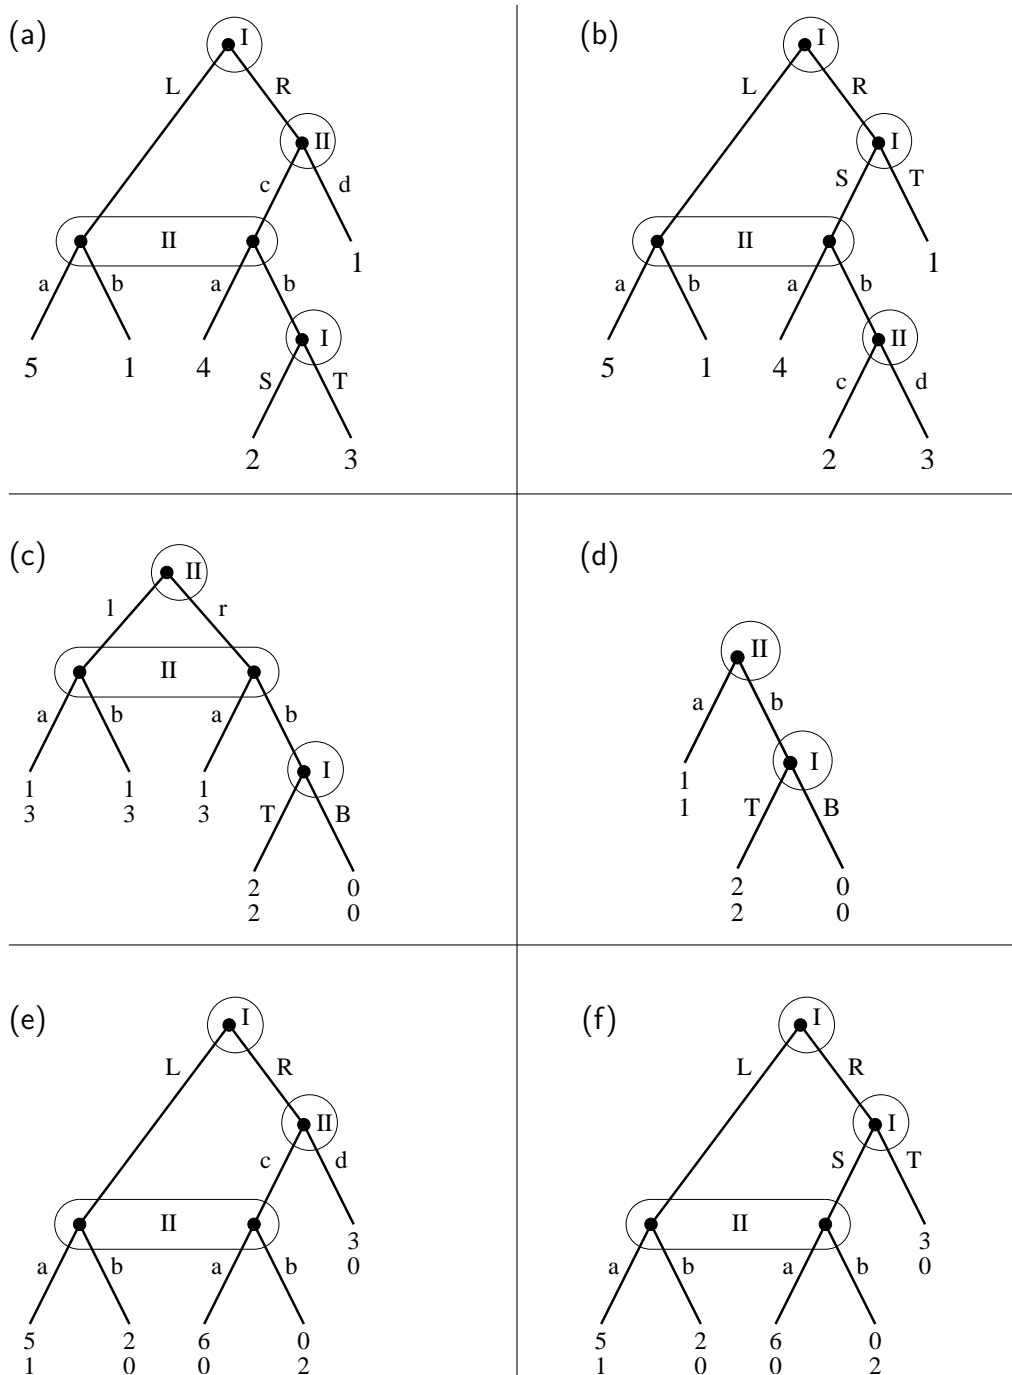


Figure 4.15 Extensive games for Exercise 4.2.

Chapter 5

Bargaining

This chapter presents models of bargaining. Chapter 2 has to be understood first, as well as the concept of max-min strategies from Section 3.14.

5.1 Learning objectives

After studying this chapter, you should be able to:

- explain the concepts of bargaining sets, bargaining axioms, Nash bargaining solution, threat point, and Nash product;
- draw bargaining sets derived from bimatrix games, and for bargaining over a “unit pie” when players have concave utility functions;
- explain, and solve specific examples of, the model of bargaining over multiple rounds of demands and counter-demands, and of stationary SPNE for infinitely many rounds.

5.2 Further reading

Chapters 15 and 7 of

- Osborne, Martin J., and Ariel Rubinstein *A Course in Game Theory*. (MIT Press, 1994) [ISBN 0262650401].

The original article on the iterated-offers bargaining model is

- Rubinstein, Ariel “Perfect equilibrium in a bargaining model.” *Econometrica*, Vol. 50 (1982), pp. 97–109.

5.3 Introduction

This chapter treats bargaining situations. This is a very practical game-theoretic problem that, for example, buyers and sellers encounter in real life. The two players have something to gain by reaching an agreement. This approach belongs to “cooperative game theory”, which tries to guide players that have decided to cooperate.

A basic bargaining model is that of “splitting a pie” in a way that is acceptable to both players, but where within that range a player can only gain at the expense of the other. The question is how to split the pie fairly. In his undergraduate thesis, John Nash proposed a system of “axioms”, or natural conditions, that should allow an impartial judge to arrive at a fair “bargaining solution”. Nash showed that these axioms always entail a unique solution, which is also easy to calculate. Nash’s development is presented in sections 5.4 to 5.6.

The later sections of this chapter go back to the game-theoretic models considered earlier which, unlike the Nash bargaining solution, are “non-cooperative” in nature, meaning that players have no recourse to a judge but have to follow their own incentives as modelled with actions and payoffs in the game. Here, bargaining is modelled by players alternately making offers which the other player can accept or reject. The analysis of this game uses the study of commitment games learned at the end of Chapter 2. The most basic form is the “ultimatum game” that consists of a single take-it-or-leave-it offer which can only be accepted or rejected.

With alternating offers over several rounds, there is a chance at each stage that the game does not proceed to the next round so that players wind up with nothing because they could not agree. The effect of this is a “discount factor” which quantifies to what extent an earlier agreement is better than a later one. By thinking ahead, the optimal strategies of the players define sequences of alternating offers, with agreement in the first round. This can be done with a fixed number of rounds, or with “stationary strategies” for an infinite number of bargaining rounds where, in essence, the game repeats every two rounds when it comes back to the same player to make an offer.

The nice conclusion of this alternating-offers bargaining model is that when the discount factor represents very patient players, it gives the same outcome as Nash’s axiomatic bargaining solution. Nash’s “cooperative” solution concept is thereby given a good foundation in terms of self-interested behaviour.

5.4 Bargaining sets

Consider the variant of the prisoner’s dilemma game on the left in Figure 5.1. Player I’s strategy B dominates T , and r dominates l , so that the unique Nash equilibrium of the game is (B, r) with payoff 1 to both players. The outcome (T, l) would give payoffs 2 and 3 to players I and II, which would clearly be better.

Suppose now that the players can talk to each other and have some means of enforcing an agreement, for example by entering some legal contract that would be very costly to break. We do not model how the players can talk, agree, or enforce contracts, but instead try to give guidelines about the agreement they can reach.

The first step is to model what players can achieve by an agreement, by describing a “bargaining set” of possible payoff pairs that they can bargain over. Suppose that the players can agree to play any of the cells of the payoff table. The resulting utility pairs (u, v) are dots in a two-dimensional diagram with horizontal coordinate u (utility to player I)

and vertical coordinate v (utility to player II). These dots are shown on the right in Figure 5.1, marked with the respective strategy pairs, for example T, l marking the utility pair $(2, 3)$.




Figure 5.1 Left: 2×2 game similar to the prisoner's dilemma. Right: corresponding bargaining set S (dark grey), and its bargaining solution (U, V) .

The *convex hull* (set of convex combinations) of these utility pairs is shown in light and dark grey in Figure 5.1. Any such convex combination is obtained by a joint lottery that the players have decided to play as a possible result of their agreement. For example, the utility pair $(1.5, 2)$ would be the expected utility if the players have decided to toss a fair coin, and, depending on the outcome, to either play (T, l) or else (B, r) ; we assume that this play is part of the players' contract that they have to adhere to.

Making the bargaining set convex with the help of lotteries is not as strange as it sounds. For example, if two people find something that both would like to have but that cannot be divided, tossing a coin to decide who gets it is an accepted fair solution.

Not every point in the convex hull is "reasonable" because we always assume that the players have the option *not to agree*. For example, player I would not agree that both players choose only the strategy pair (T, r) , where player I gets payoff zero. The reason is that without any agreement, player I can use his *max-min* strategy B that guarantees him a payoff of 1 irrespective of what the other player does. As shown in the earlier section on zero-sum games, a max-min strategy may be a mixed strategy; in the prisoner's dilemma game it is a pure strategy, in fact the dominating strategy. We ask that every agreed outcome of the bargaining situation should give each player at least his or her max-min payoff (which is, as usual, defined in terms of that player's own payoffs).

The *bargaining set* S derived from a two-player game in strategic form is thus defined as the convex hull of utility pairs (as specified in the game), with the additional constraint that $u \geq u_0$ and $v \geq v_0$ for all (u, v) in S , where u_0 is the max-min payoff to player I and v_0 is the max-min payoff to player II. With $u_0 = v_0 = 1$, the resulting set S is shown in darker grey in Figure 5.1.

In this example, the "Nash bargaining solution" (which is not to be confused with "Nash equilibrium", a very different concept) will give the utility pair $(U, V) = (2.5, 2.5)$,

as we will explain (the fact that $U = V$ is a coincidence). This expected payoff pair is a convex combination of the utility pairs $(2, 3)$ and $(5, 0)$. It is achieved by a joint lottery of playing T, l with probability $5/6$ and B, l with probability $1/6$. In effect, player II always plays l and player I randomises between T and B with probabilities $5/6$ and $1/6$. However, this is not a “mixed strategy” but rather a lottery performed by a third party, with an outcome that both players (in particular player II) have decided to accept.

5.5 Bargaining axioms

We consider bargaining situations modelled as bargaining sets, and try to obtain a “solution” for any such situation. The conditions about this model are called *axioms* and concern the bargaining set itself, as well as the solution one wants to obtain for any such set.

Definition 5.1 (Axioms for bargaining set) A bargaining set S with threat point (u_0, v_0) is a non-empty subset of \mathbb{R}^2 so that

- (a) $u \geq u_0, v \geq v_0$ for all $(u, v) \in S$;
- (b) S is compact;
- (c) S is convex.

Each pair (u, v) in a bargaining set S represents a utility u for player I and a utility v for player II which the players can achieve by reaching a specific agreement. At least one agreement is possible because S is non-empty. The threat point (u_0, v_0) is the pair of utilities that each player can (in expectation) guarantee for himself or herself. Axiom (b) in definition 5.1 states that S is bounded and closed, so it is possible to define a “bargaining solution” by some sort of maximisation process. The convexity axiom (c) means that players have access to joint lotteries as part of their agreement, if necessary.

In Section 5.4, we have constructed a bargaining set S from a bimatrix game, and now argue that it fulfils the axioms in definition 5.1. This is clear for (b) and (c). By construction, (a) holds, but we have to show that S is not empty. The threat point (u_0, v_0) is obtained with u_0 as the max-min payoff to player I (in terms of his payoffs) and v_0 as the max-min payoff to player II (in terms of her payoffs). Thus, if player I plays his max-min strategy \hat{x} and player II plays her max-min strategy \hat{y} , then the payoff pair $(u, v) = (a(\hat{x}, \hat{y}), b(\hat{x}, \hat{y}))$ that results in the game is an element of S because $u \geq u_0$ and $v \geq v_0$, so S is not empty.

\Rightarrow Try Exercise 5.1 on page 162.

However, in the construction of a bargaining set from a bimatrix game, the threat point (u_0, v_0) need not itself be an element of S , as shown in Exercise 5.1. For this reason, the threat point (u_0, v_0) is explicitly stated together with the bargaining set. If the threat point belongs to S and (a) holds, then there can be no other threat point. The axioms in definition 5.2 below would be easier to state if the threat point was included in S , but then we would have to take additional steps to obtain such a bargaining set from a bimatrix

game. In our examples and drawings we will typically consider sets S that include the threat point.

From now on, we simply assume that the bargaining set S fulfils these axioms. We will later consider bargaining sets for other situations, which are specified directly and not from a bimatrix game.

In the following definition, the axioms (d)–(h) continue (a)–(c) from definition 5.1. We discuss them afterwards.

Definition 5.2 (Axioms for bargaining solution) For a given bargaining set S with threat point (u_0, v_0) , a *bargaining solution* $N(S)$ is a pair (U, V) so that

- (d) $(U, V) \in S$;
- (e) the solution (U, V) is *Pareto-optimal*, that is, for all $(u, v) \in S$, if $u \geq U$ and $v \geq V$, then $(u, v) = (U, V)$;
- (f) it is *invariant under utility scaling*, that is, if $a, c > 0$, $b, d \in \mathbb{R}$ and S' is the bargaining set $\{(au + b, cv + d) \mid (u, v) \in S\}$ with threat point $(au_0 + b, cv_0 + d)$, then $N(S') = (aU + b, cV + d)$;
- (g) it *preserves symmetry*, that is, if $u_0 = v_0$ and if any $(u, v) \in S$ fulfils $(v, u) \in S$, then $U = V$;
- (h) it is *independent of irrelevant alternatives*: If S, T are bargaining sets with the same threat point and $S \subset T$, then either $N(T) \notin S$ or $N(T) = N(S)$.

The notation $N(S)$ is short for “Nash bargaining solution”. The solution is a pair (U, V) of utilities for the two players that belongs to the bargaining set S by axiom (d).

The Pareto-optimality stated in (e) means that it should not be possible to improve the solution for one player without hurting the other player. Graphically, the solution is therefore on the upper-right (or “north-east”) border of the bargaining set, which is also called the *Pareto-frontier* of the set.

Invariance under utility scaling as stated in (f) means the following. The utility function of a player can be changed by changing its “scale” and “base point” without changing its meaning, just like a temperature scale. Such a change for player I means that instead of u we consider $au + b$ where a and b are fixed real numbers and $a > 0$. Similarly, for player II we can replace v by $cv + d$ for fixed c, d and $c > 0$. When drawing the bargaining set S , this means that (by adding b) the first coordinate of any point in that set is moved by the amount b (to the right or left depending on whether b is positive or negative), and stretched by the factor a (if $a > 1$, and shrunk if $a < 1$). In (f), the set S' and the new threat point are obtained from S by applying such a “change of scale” for both players.

Because this scaling does not change the meaning of the utility values (they still represent the same preference), the bargaining solution should not be affected either, as stated in axiom (f). In a sense, this means the bargaining solution should not depend on a direct comparison of the players’ utilities. For example, if you want to split a bar of chocolate with a friend you cannot claim more because you like chocolate “ten times more than she does” (which you could represent by changing your utility from u to $10u$),

nor can she claim that you should get less because you already have 5 bars of chocolate in your drawer (represented by a change from u to $u + 5$).

Preservation of symmetry as stated in (g) means that a symmetric bargaining set S , where for all (u, v) we have $(u, v) \in S \iff (v, u) \in S$, and with a threat point of the form (u_0, u_0) , should also have a bargaining solution of the form $N(S) = (U, U)$.




Figure 5.2 Bargaining sets S (dark grey) and T (dark and light grey). Independence of irrelevant alternatives states that if S is a subset of T , then $N(T)$ must not be a point in S different from $N(S)$, like P .

The independence of irrelevant alternatives stated in (h) is the most complicated axiom. As illustrated in Figure 5.2, it means that if a bargaining set S is enlarged to a larger set T that has additional points (u, v) , but without changing the threat point, then either T has a bargaining solution $N(T)$ that is one of these new points, or else it is the same solution as the solution $N(S)$. It is not allowed that the solution for T moves to a different point, like P in Figure 5.2, that still belongs to S .

As a coarse analogy, suppose you are offered the choice between chicken and beef as an airplane meal, decide to take chicken, and when the flight attendant says: “oh, we also have fish” you say “then I take beef”. This does not comply with axiom (h): “fish” is the irrelevant alternative (because you do not take it) and the choice between chicken and beef should be independent of that.

The stated axioms (a)–(h) for a bargaining set and a possible bargaining solution are reasonable requirements that one can defend as general principles.

5.6 The Nash bargaining solution

This section shows that the bargaining axioms (a)–(h) in definitions 5.1 and 5.2 imply that a bargaining solution exists, which is, moreover, unique.

Theorem 5.3 *Under the Nash bargaining axioms, every bargaining set S containing a point (u, v) with $u > u_0$ and $v > v_0$ has a unique solution $N(S) = (U, V)$, which is obtained*

as the point (u, v) that maximises the product (also called the “Nash product”) $(u - u_0)(v - v_0)$ for $(u, v) \in S$.

Proof. First, consider the set $S' = \{(u - u_0, v - v_0) \mid (u, v) \in S\}$, which is the set S translated so that the threat point (u_0, v_0) is simply the origin $(0, 0)$. Maximising the Nash product now amounts to maximising $u \cdot v$ for $(u, v) \in S$. Let (U, V) be the utility pair (u, v) where this happens, where by assumption $UV > 0$. We re-scale the utilities so that $(U, V) = (1, 1)$, by replacing S with the set $S'' = \{(u/U, v/V) \mid (u, v) \in S\}$.

We now consider the set $T = \{(u, v) \mid u \geq 0, v \geq 0, u + v \leq 2\}$. For the bargaining set T , the solution is $N(T) = (1, 1)$, because T is a symmetric set, and $(1, 1)$ is the only symmetric point on the Pareto-frontier of T . We claim that $S \subseteq T$. Then, by the independence of irrelevant alternatives, $N(S)$ must be equal to $N(T)$ because this point $(1, 1)$ is an element of S , so because it is not the case that $N(T) \notin S$, we must have $N(S) = N(T)$.

Proving $S \subseteq T$ uses the convexity of S . Namely, suppose there was a point (\bar{u}, \bar{v}) in S which is not in T . Because $\bar{u} \geq 0, \bar{v} \geq 0$, this means $\bar{u} + \bar{v} > 2$. The idea is now that even if the Nash product $\bar{u} \cdot \bar{v}$ is not larger than 1 (which is the Nash product UV), then the Nash product uv of a suitable convex combination $(u, v) = (1 - \varepsilon)(U, V) + \varepsilon(\bar{u}, \bar{v})$ is larger than 1. Here, ε will be a small positive real number.




Figure 5.3 Proof that the point (\bar{u}, \bar{v}) , which is outside the triangle T (dark and light grey), cannot belong to S (dark grey), which shows that $S \subseteq T$.

This is illustrated in Figure 5.3: The hyperbola $\{(u, v) \mid uv = 1\}$ touches the set T at the point $(1, 1)$. The tangent to the hyperbola is the line through points $(2, 0)$ and $(0, 2)$, which is the Pareto-frontier of T . Because the point (\bar{u}, \bar{v}) is to the right of that tangent, the line segment that connects $(1, 1)$ and (\bar{u}, \bar{v}) must intersect the hyperbola and therefore contain points (u, v) to the right of the hyperbola, where $uv > 1$. This is a contradiction, because $(u, v) \in S$ since S is convex, but the maximum Nash product of a point in S is 1.

So consider the Nash product for the convex combination (u, v) of $(1, 1)$ and (\bar{u}, \bar{v}) :

$$\begin{aligned}
uv &= [(1 - \varepsilon)U + \varepsilon\bar{u}][(1 - \varepsilon)V + \varepsilon\bar{v}] \\
&= [U + \varepsilon(\bar{u} - U)][V + \varepsilon(\bar{v} - V)] \\
&= UV + \varepsilon[\bar{u} + \bar{v} - (U + V) + \varepsilon(\bar{u} - U)(\bar{v} - V)].
\end{aligned}$$

Now, because $\bar{u} + \bar{v} > 2$, but $U + V = 2$, the term $\bar{u} + \bar{v} - (U + V)$ is positive, and $\varepsilon(\bar{u} - U)(\bar{v} - V)$ can be made smaller than that term in absolute value, for suitably small $\varepsilon > 0$, even if $(\bar{u} - U)(\bar{v} - V)$ is any negative number. Consequently, $uv > 1 = UV$ for sufficiently small positive ε , which shows that UV was not the maximum Nash product in S , which contradicts the construction of (U, V) . This proves that, indeed, $S \subseteq T$, and hence $N(S) = N(T)$ as claimed.

In order to show that the point $(U, V) = (1, 1)$ is unique, we can again look at the hyperbola, and argue similarly as before. Namely, suppose that (\bar{u}, \bar{v}) is a point in S different from $(1, 1)$ so that $\bar{u} \cdot \bar{v} = 1$. Then this point and $(1, 1)$ are two points on the hyperbola, and points on the line segment joining these two points belong to S because S is convex. However, they have a larger Nash product. For example, the mid-point $1/2 \cdot (1, 1) + 1/2 \cdot (\bar{u}, \bar{v})$ has the Nash product $(1/2 + \bar{u}/2)(1/2 + \bar{v}/2) = 1/4 + \bar{u}/4 + \bar{v}/4 + \bar{u} \cdot \bar{v}$. Because $\bar{u} \cdot \bar{v} = 1$, this is equal to $1/2 + (\bar{u} + \bar{v})/4$, which is larger than 1 if and only if $(\bar{u} + \bar{v})/4 > 1/2$ or $\bar{u} + \bar{v} > 2$. The latter condition is equivalent to $\bar{u} + 1/\bar{u} > 2$ or, because $\bar{u} > 0$, to $\bar{u}^2 + 1 > 2\bar{u}$ or $\bar{u}^2 - 2\bar{u} + 1 > 0$, that is, $(\bar{u} - 1)^2 > 0$, which is true because $\bar{u} \neq 1$. So, again because S is convex, the Nash product has a unique maximum in S . \square

The assumption that S has a point (u, v) with $u > u_0$ and $v > v_0$ has the following purpose: Otherwise, for $(u, v) \in S$, all utilities u of player I or v of player II (or both) fulfil $u = u_0$ or $v = v_0$, respectively. Then the Nash product $(U - u_0)(V - v_0)$ for $(U, V) \in S$ would always be zero, so it has no unique maximum. In this rather trivial case, the bargaining set S is either just the singleton $\{(u_0, v_0)\}$ or a line segment $\{u_0\} \times [v_0, V]$ or $[u_0, U] \times \{v_0\}$. In both cases, the Pareto-frontier consists of a single point (U, V) , which defines the unique Nash solution $N(S)$.

\Rightarrow If the bargaining set is a line segment as discussed in the preceding paragraph, why is its Pareto-frontier a single point? Is this always so if the line segment is not parallel to one of the coordinate axes?

It can be shown relatively easily that the maximisation of the Nash product as described in Theorem 5.3 gives a solution (U, V) that fulfils the axioms in definitions 5.1 and 5.2.

\Rightarrow In exercise 5.2, you are asked to verify this for axiom (f).

We describe how to find the bargaining solution (U, V) for the bargaining set S in Figure 5.1. The threat point (u_0, v_0) is $(1, 1)$. The Pareto-frontier consists of two line segments, a first line segment joining $(1, 3.5)$ to $(2, 3)$, and a second line segment joining $(2, 3)$ to $(4, 1)$, where the latter is part of a longer line segment that extends from $(2, 3)$ to $(5, 0)$ (the first line segment is also part of the longer segment that connects $(4, 0)$ to $(2, 3)$, but this will not be needed). The Nash product $(u - u_0)(v - v_0)$ for a point (u, v) on the Pareto-frontier is the size of the area of the rectangle between the threat point (u_0, v_0) and the point (u, v) . In this case, the diagram shows that the second line segment

creates the larger rectangle in that way. A point on the second line segment is of the form $(1-p)(2,3) + p(5,0)$ for $0 \leq p \leq 1$, that is, $(2+3p, 3-3p)$. Because $(u_0, v_0) = (1, 1)$, the resulting Nash product is $(1+3p)(2-3p)$ or $2+3p-9p^2$. The unique maximum over p of this expression is obtained for $p = 1/6$ because that is the zero of its derivative $3-18p$, with negative second derivative so this is a maximum. The point (U, V) is therefore given as $(2.5, 2.5)$. (In general, such a maximum can result for $p < 0$ or $p > 1$, which implies that the left or right endpoint of the line segment with $p = 0$ or $p = 1$, respectively, gives the largest value of the Nash product. Then one has to check the adjoining line segment for a possibly larger Nash product, which may occur on the corner point.) This value of p allows us to represent the Nash bargaining solution as a convex combination of payoff pairs that correspond to cells in the bimatrix game, namely via a lottery that chooses T, l with payoff pair $(2, 3)$ with probability $1-p = 5/6$ and B, r with payoff pair $(5, 0)$ with probability $p = 1/6$. Both players have agreed in advance to accept the outcome of this lottery (even though with probability $1/6$ the pair $(5, 0)$ gives payoff zero to player II, even less than at the threat point where she gets $v_0 = 1$; however, her expected payoff 2.5 before the lottery is higher than v_0).

The following proposition gives a useful condition to find the bargaining solution geometrically. It is of particular importance later, in Section 5.12.

Proposition 5.4 *Suppose that the Pareto-frontier of a bargaining set S with threat point $(0,0)$ is given by $\{(u, f(u)) \mid 0 \leq u \leq 1\}$ for a decreasing and continuous function f with $f(0) = 1$ and $f(1) = 0$. Then the bargaining solution (U, V) is the unique point $(U, f(U))$ where the bargaining set has a tangent with slope $-f(U)/U$. If f is differentiable, then this slope is the derivative $f'(U)$ of f at U , that is, $f'(U) = -f(U)/U$.*




Figure 5.4 Characterisation of the Nash bargaining solution for differentiable (left) and non-differentiable bargaining frontier (right).

Proof. We first consider the case that the function f that defines the Pareto-frontier is differentiable. The bargaining solution maximises the Nash product $u \cdot f(u)$ on the Pareto-frontier. This requires the derivative with respect to u to be zero, that is, $f(u) + u \cdot f'(u) = 0$ or $f'(u) = -f(u)/u$, so U has to solve this equation as claimed. The solution is also

unique: $f'(u)$ is negative and decreasing with u because the bargaining set is convex (so f is a concave function, see the definition in (5.1) below); furthermore, the function $f(u)$ is decreasing and $1/u$ (where we can assume $u > 0$) is strictly decreasing, so that $-f(u)/u$ is strictly increasing, and hence there is at most one u with $f'(u) = -f(u)/u$.

The left picture in Figure 5.4 shows the line connecting the origin $(0,0)$ with (U,V) , which has the slope V/U , that is, $f(U)/U$, as indicated by the angle α near the origin. The tangent of the Pareto-frontier at the point (U,V) slopes downwards, with negative slope $f'(U)$. The equation $f'(U) = -f(U)/U$ states that this tangent intersects the horizontal axis at the same angle α . In other words, the two lines form a “roof” with apex (U,V) with equal slopes α on both sides.

The same “roof” construction is possible when f is not differentiable. Every point $(U, f(U))$ defines a unique slope with angle α for the line that reaches this point from the origin. As shown in the right picture of Figure 5.4, there may be more than one tangent at this point if the Pareto-frontier has a “kink” at that point. The possible tangents have slopes from an interval, and it is not difficult to show that for different points these intervals cannot have more than one slope in common. (This holds because the bargaining set is convex; for example, all tangents to the set at a point below (U,V) in the figure have the same slope, which is the slope of the line segment that connects (U,V) to $(1,0)$ which forms that part of the Pareto-frontier.) Consequently, there is at most one point $(U, f(U))$ on the Pareto-frontier with a tangent of slope $-f(U)/U$. This point defines the Nash bargaining solution (U,V) . To see this, observe that the function $v = UV/u$ of u (with the constant UV) goes through the point (U,V) . At this point, the derivative of $-UV/u^2$ of this function is $-V/U$, so it has the same tangent that we consider for the bargaining set, as shown in the figure. Because the hyperbola $v = c/u$ with $c = UV$ touches the bargaining set at the point (U,V) , a higher constant c is not possible, which shows that the product uv is maximised for $(u,v) = (U,V)$, so this is the bargaining solution. \square

\Rightarrow Before Proposition 5.4, we have found the Nash bargaining solution (U,V) for the game in Figure 5.1. Do this now with the help of Proposition 5.4.
[Hint: The slope in question has an angle of 45 degrees.]

5.7 Splitting a unit pie

In the remainder of this chapter, we consider a bargaining situation where player I and player II have to agree to split a “pie” into an amount x for player I and y for player II. The total amount to be split is normalised to be one unit, so this is called “a unit pie”. The possible splits (x,y) have to fulfil $x \geq 0$, $y \geq 0$, and $x + y \leq 1$. If the two players cannot agree, they both receive zero, which defines the threat point $(0,0)$. Splits (x,y) so that $x + y < 1$ are not optimal, but we admit such splits in order to obtain a convex set.

If player I gets utility x and player II gets utility y , then the bargaining set is the triangle of all (x,y) with $x \geq 0$, $y \geq 0$ and $x + y \leq 1$ shown in Figure 5.5. The split (x,y) is Pareto-optimal if $x + y = 1$, so the Pareto-frontier is given by the splits $(x, 1 - x)$, shown as the bold line in the figure. Because the threat point is $(0,0)$, the Nash product is $x(1 - x)$




Figure 5.5 Bargaining over a unit pie with $u(x) = x$, $v(y) = y$. The bargaining solution is $(U, V) = (1/2, 1/2)$.

which is maximised for $x = 1/2$. The Nash bargaining solution $(U, V) = (1/2, 1/2)$ results also from the symmetry axiom in definition 5.2(g).




Figure 5.6 Left: utility function $u(x) = \sqrt{x}$, and why it is concave. Right: bargaining set S in (5.2) when $u(x) = \sqrt{x}$ and $v(y) = y$.

More generally, we consider the split of a unit pie where the players have utility functions u and v that are not the identity function. We assume that the utility functions are $u : [0, 1] \rightarrow [0, 1]$ for player I and $v : [0, 1] \rightarrow [0, 1]$ for player II, which are normalised so that $u(0) = v(0) = 0$ and $u(1) = v(1) = 1$. These utility functions are assumed to be increasing, concave and continuous.

The function u (and similarly v) is said to be concave if any two points on the graph of the function are connected by a line segment that is *nowhere above* the graph of the function. That is,

$$(1-p)u(x) + pu(x') \leq u((1-p)x + px') \quad \text{for all } x, x' \text{ and } p \in [0, 1]. \quad (5.1)$$

Figure 5.6 shows on the left the function $u(x) = \sqrt{x}$, with three white dots that are the points $(x, u(x))$, $((1-p)x + px', u((1-p)x + px'))$, and $(x', u(x'))$ on the graph of the

function. The black dot is the convex combination $(1-p) \cdot (x, u(x)) + p \cdot (x', u(x'))$, which is below the middle white dot, as asserted in (5.1). This holds generally and shows that u is concave. A function that is twice differentiable is concave if its second derivative is never positive. Furthermore, it can be shown that concavity implies continuity, except at the endpoints (here 0 and 1) of the domain of the function.

The bargaining set of resulting utility pairs to the two players is given by

$$S = \{(u(x), v(y)) \mid x \geq 0, y \geq 0, x + y \leq 1\}. \quad (5.2)$$

It can be shown that S is convex and compact because u and v are increasing, concave and continuous as assumed. Intuitively, this holds because S consists of utility pairs (u, v) where the possible values of $u(x)$ give the horizontal coordinate u , and the values of $v(y)$ the vertical coordinate v . In comparison to x and y , concavity means that these coordinates are “diluted” for small values of x and y and are more “compressed” for larger values of x and y . Hence, one can view S as a deformation of the triangle in Figure 5.5 with the Pareto-frontier bulging outward.

For $u(x) = \sqrt{x}$ and $v(y) = y$, the set S is shown on the right in Figure 5.6. The Pareto-frontier is given by points $(u, v) = (u(x), v(1-x))$ for $0 \leq x \leq 1$. Because $u = \sqrt{x}$ and $v = 1-x = 1-u^2$, these are pairs $(u, 1-u^2)$, so v is easy to draw as the function $1-u^2$ of u . In general, one may think of traversing the points $(u(x), v(1-x))$ on the Pareto-frontier when changing the parameter x , from the upper endpoint $(0, 1)$ of the frontier when $x = 0$, to its right endpoint $(1, 0)$ when $x = 1$.

The Nash bargaining solution (U, V) is found on the Pareto-frontier $(u(x), v(1-x))$ by maximising the Nash product $u(x)v(1-x)$. In Figure 5.6, the Nash product is $\sqrt{x}(1-x)$ or $x^{1/2} - x^{3/2}$ with derivative $\frac{1}{2}x^{-1/2} - \frac{3}{2}x^{1/2}$ or $\frac{1}{2}x^{-1/2}(1-3x)$, which is zero when $x = 1/3$, so the pie is split into $1/3$ for player I and $2/3$ for player II. The corresponding pair (U, V) of utilities is $(1/\sqrt{3}, 2/3)$ or about $(0.577, 0.667)$.

\Rightarrow Do Exercise 5.4 on page 163.

5.8 The ultimatum game

Assume that two players have to split a unit pie, as described in the previous section. In the following, we will develop the “alternating offers” bargaining game applied to this situation. This is a non-cooperative game which gives a specific model of how players may interact in a bargaining situation. The outcome of this non-cooperative game, in its most elaborate version of “stationary strategies” (see Section 5.11), will approximate the Nash bargaining solution. The model justifies the Nash solution via a different approach, which is not axiomatic, but instead uses the concept of subgame perfect equilibria in game trees.

The model is that the pie is split according to the suggestion of one of the players. Suppose player I suggests splitting the pie into x for player I and $1-x$ for player II, so that the players receive utility $u(x)$ and $v(1-x)$, respectively. The amount x proposed by




Figure 5.7 Discrete version of the ultimatum game. The arrows denote the SPNE.

player I is also called his *demand* because that is the share of the pie that he receives in the proposed split.

In the most basic form, this model is called the *ultimatum game*, where player II can either accept the split or reject it. If player II accepts, the payoff pair is $(u(x), v(1-x))$, and if she rejects, it is $(0, 0)$, so neither player receives anything. Figure 5.7 shows a discretised version of this game where the possible demand x is a multiple of 0.01. Hence, player I has 101 possible actions for $x \in \{i/100 \mid i = 0, 1, \dots, 100\}$, and for each proposed x player II can either choose A (accept) or R (reject). Player II therefore has 2^{101} pure strategies, given by the different combinations of A and R for each possible x .

The game is shown in Figure 5.7, and we look for a subgame perfect Nash equilibrium (SPNE) of the game, that is, we apply backward induction. Whenever $1-x > 0$, the utility $v(1-x)$ to player II is positive; this holds because v is concave: otherwise, $0 = v(1-x) = v(x \cdot 0 + (1-x) \cdot 1) < 1-x = x \cdot v(0) + (1-x) \cdot v(1)$, which contradicts (5.1). So the SPNE condition implies that player II accepts any positive amount $1-x$ offered by player I. These choices, determined by backward induction, are shown as arrows in Figure 5.7, and they are unique except when $x = 1$ where player I demands the whole pie. Then player II is indifferent between accepting and rejecting because she will get nothing in either case. Consequently, in terms of pure strategies, both $AA \cdots AA$ and $AA \cdots AR$ are SPNE strategies of player II. Here, the first 100 A's denote acceptance for each $x = 0, 0.01, \dots, 0.99$, and the last A or R is the choice when $x = 1$, shown by a white or black arrow, respectively, in Figure 5.7. Given this strategy of player II, player I's best response is to maximise his payoff, which he does as shown by the corresponding white or black arrow for player I. If player II accepts the demand $x = 1$, then player I demands $x = 1$, and the pair $(1, AA \cdots AA)$ is an SPNE, and if player II rejects the demand $x = 1$, then player I demands $x = 0.99$ (because for $x = 1$ he would get nothing), and the resulting SPNE is $(0.99, AA \cdots AR)$. These are the two SPNE in pure strategies.

In essence, player II receives virtually nothing in the ultimatum game because of the backward induction assumption. This does not reflect people's behaviour in real bargaining situations of this kind, as confirmed in laboratory experiments where people play the ultimatum game. Typically, such an experiment is conducted where player I states,

unknown to player II, his demand x , and player II declares simultaneously, unknown to player I, her own “reserve demand” z ; the two players may even be matched anonymously from a group of many subjects that take the respective roles of player I and II. Whenever $1 - x \geq z$, that is, the offer to player II is at least as large as her reserve demand, the two players get x and $1 - x$, otherwise nothing. (So player II may declare a small reserve demand like $z = 0.01$ but still receive 0.4 because player I’s demand is $x = 0.6$; both can be regarded as reasonable choices, but they are not mutual best responses.) In such experiments, player II often makes positive reserve demands like $z = 0.4$, and player I makes cautious demands like $x = 0.5$, which contradict the SPNE assumption. A better description of these experiments may therefore be Nash equilibria that are not subgame perfect, with “threats” by player II to reject offers $1 - x$ that are too low. Nevertheless, we assume SPNE because they allow definite conclusions, and because the model will be refined, and made more realistic, by allowing rejections to be followed by “counter-demands”.

The described two SPNE of the discrete ultimatum game are found whenever the unit pie can be split into $M + 1$ possible splits $(x, 1 - x)$ for $x = i/M$, for $i = 0, 1, \dots, M - 1, M$, like $M = 100$ above. Of her 2^{M+1} pure strategies, player II can play only two in an SPNE, namely always choosing A whenever $x < 1$, and either A or R in response to $x = 1$. The respective demand by player I is $x = 1$ or $x = 1 - 1/M$. With a very fine discretisation where M becomes very large, these two SPNE are nearly identical, with player I demanding all (or nearly all) of the pie, and player II receiving almost nothing.

⇒ Exercise 5.3 on page 162 asks you to find all Nash equilibria of the ultimatum game.




Figure 5.8 Continuous version of the ultimatum game.

In the “continuous” version of the ultimatum game, the possible demand x by player I is any number in the interval $[0, 1]$, which defines an infinite number of strategies for player I. The strategies for player II are even more complicated, namely given by a function $f: [0, 1] \rightarrow \{A, R\}$, where $f(x) = A$ means that player II accepts the demand x , and $f(x) = R$ means that player II rejects it. (Equivalently, a strategy of player II is an arbitrary subset of $[0, 1]$ which consists exactly of those demands x that she accepts.) This infinite game is indicated in Figure 5.8. The infinitely many choices x are indicated by the

triangle, whose baseline represents the interval $[0, 1]$. Every point of that baseline has a separate decision point for player II where she can decide between A and R . Only a single decision point is drawn in the picture.

We want to show that the continuous ultimatum game has a *unique* SPNE. This requires an additional assumption, namely that the utility function u of player I is *strictly* increasing.

We first claim that concavity of u implies that u is not only increasing, but strictly increasing on an interval $[0, x']$, where $u(x) = 1$ for all x with $x' \leq x \leq 1$, so that u is constant on the interval $[x', 1]$. (In particular, if $x' = 1$, then u is strictly increasing throughout.) Namely, the set $I = \{x \in [0, 1] \mid u(x) = 1\}$ is not empty (it contains 1) and closed because u is continuous, so let x' be the smallest element of I , with $u(x') = u(1) = 1$. Then I is the interval $[x', 1]$ because for $x \in [x', 1]$ the value $u(x)$ must be at least as high as the horizontal line that connects $u(x')$ and $u(1)$, but $u(x) > u(1)$ is not allowed. We have $x' > 0$ because $u(0) = 0 < 1$. To see that u is strictly increasing on the interval $[0, x']$, suppose there was x'' and \hat{x} in that interval with $x'' < \hat{x}$ but $u(x'') = u(\hat{x}) < u(x') = 1$. Then it is easy to see (a drawing will help) that $u(\hat{x})$ is below the line that connects $u(x'')$ and $u(x')$ in contradiction to concavity. This proves the claim.

In the case where $x' < 1$, player I is perfectly satisfied by receiving only x' , which is less than a full share of the pie. Then any demand x with $x' \leq x < 1$ is accepted by player II because she receives $v(1 - x) > 0$, and x is a possible strategy of player I in an SPNE. In that case, we do not get a unique SPNE. We therefore require that u is strictly increasing, or equivalently (as explained, because u is concave) that $u(x) = 1$ implies $x = 1$. So we make this additional assumption, and similarly assume that player II's utility function v is strictly increasing.

We claim that the *only* SPNE of the continuous ultimatum game is where player I demands the whole pie, with $x = 1$, and where player II accepts any demand, including the demand $x = 1$ on the equilibrium path where she is indifferent between accepting and rejecting. Clearly, this is one SPNE of the game, but why is there no other? One argument is that this is the only “limit” of the two SPNE found in any discretised version of the game. A direct reasoning is as follows. Clearly, player II has to accept whenever $x < 1$ because then she receives a positive payoff $v(1 - x)$. Suppose there is an SPNE where player I's demand is x so that $x < 1$. Because u is strictly increasing, player I would get more by demanding a little more, but still less than 1, for example $(x + 1)/2$, which player II still accepts. So $x = 1$ is the only strategy by player I in an SPNE. Then on the equilibrium path, player II is indifferent, and can choose either A or R in response to $x = 1$, and in fact she could randomise between A and R . Suppose that R is played with any positive probability $\varepsilon > 0$, giving player I expected payoff $(1 - \varepsilon)u(1) + \varepsilon u(0)$, which is $1 - \varepsilon$. But then player I's best response to that strategy of player II would not be $x = 1$. Instead, player I could improve his payoff, for example by demanding $1 - \varepsilon/2$ because then player II would accept for sure and player I receives $u(1 - \varepsilon/2) \geq 1 - \varepsilon/2 > 1 - \varepsilon$. However, no demand less than 1 is possible in an SPNE, as we saw earlier.

\Rightarrow Show that concavity of u implies the inequality $u(1 - \varepsilon/2) \geq 1 - \varepsilon/2$ stated in the preceding paragraph.

In summary, with demands that can be chosen continuously, the demand of a player in an SPNE makes the other player *indifferent between accepting and rejecting*, which the other player nevertheless *accepts* with certainty. This will also be the guiding principle in SPNE of more complicated games, which we study next.

5.9 Alternating offers over two rounds

Next, we extend the bargaining model so that once a player has rejected the demand of the other player, she can make a counter-demand that, in turn, the other player can accept or reject. In order to find an SPNE of this game by backward induction, these alternating offers have to end after some time. The simplest case is bargaining over two rounds, discussed in this section. In the first round, player I makes a demand x , which player II can accept or reject. In the latter case, she can make a counter-demand y that represents the share of the pie she claims for herself, so the pie is split into $1 - y$ for player I and y for player II. If player I accepts that demand, he receives utility $u(1 - y)$ and player II utility $v(y)$. If player I rejects the final demand, both players receive zero.




Figure 5.9 Two-round bargaining with a chance move with probability δ that a second round takes place.

In this game, player II can make the demand in the final round, which is the ultimatum game with the roles of the players exchanged. In that game, player II can demand the

whole pie, so there would be no reason for her to accept anything offered to her in the first round. For that reason, the model has an additional feature shown in Figure 5.9. Namely, when player II rejects the first demand, there is a positive probability $1 - \delta$ that bargaining terminates with no agreement, where we assume $0 < \delta < 1$. This is the unsuccessful outcome, with payoff zero to both, given as the result of a chance move. The game proceeds to player II's counter-demand only with probability δ .




Figure 5.10 The game of figure 5.9 with δ as discount factor applied to second-round utility values.

By computing expected values, the chance move can be eliminated from the game, which gives the game in Figure 5.10. In that game, the pair $(u(1 - y), v(y))$ of utilities when agreement is reached in the second round is multiplied by δ , which is a reduction of the payoffs because $\delta < 1$. The usual interpretation of δ is that of a *discount factor*, which expresses that later payoffs are worth less than earlier payoffs.

A discount factor applied to future payoffs is realistic. In monetary terms, it can express lost interest payments (for money today as opposed to money in the future). More importantly, future payoffs are less secure because, for example, the buyer or seller in a deal may change his or her mind, or because of other unforeseen circumstances. This is modelled with the chance move in Figure 5.9. In addition, this chance move provides a good reason why both players have the same discount factor. Alternatively, one may use different discount factors δ_I and δ_{II} for the two players, which give second-round payoffs $\delta_I u(1 - y)$ and $\delta_{II} v(y)$. If these discount factors are different, a higher discount factor represents a more patient player because his or her future payoff would differ less from the present payoff. However, the resulting game would be more complicated to analyse,

and it would be more difficult to establish a connection to the Nash bargaining solution. For that reason, we use only one discount factor, in agreement with the model that δ represents the probability that bargaining continues into the next round.

An SPNE of the game is found as follows. By backward induction, we have to analyse the last round first. When player II makes her demand y to player I, he can only accept or reject. This is the ultimatum game analysed earlier, so player II demands the whole pie $y = 1$, and player I accepts any offer made to him, in particular this demand $y = 1$ (recall that player I's strategy is a function of the demand y that he faces), and player I will get nothing. Here, it does not matter if the last stage represents the subgame starting with player II's demand y in Figure 5.9, or the same game with discounted payoffs in Figure 5.10, because the latter payoffs are merely scaled with the positive scalar δ , so these are the same games in strategic terms, with the same equilibria. In either game, backward induction shows that if player II rejects player I's demand x in the first round and chooses R , the payoffs for both players are zero for player I, and $\delta v(1)$, which is δ , for player II.

According to the next backward induction step, what should player I do in the first round? Player II's reaction is unique except when she is indifferent. This means that she will accept if she is offered a utility $v(1 - x)$ (now undiscounted, because this is the first round) that exceeds the amount δ that she can expect to get by rejection; she is indifferent if $v(1 - x) = \delta$; and she will reject if $v(1 - x) < \delta$. By the same reasoning as that used in the continuous ultimatum game, the demand x by player I in the SPNE is chosen so that $v(1 - x) = \delta$, and player II accepts on the equilibrium path. (If she did not accept, player I would offer her a little bit more and make her accept, but player I must make her indifferent in order to obtain an SPNE.)

The resulting SPNE is unique. Player I demands x so that $v(1 - x) = \delta$, which is accepted immediately by player II, so that the game terminates with player I receiving $u(x)$ and player II receiving $v(1 - x)$, which is equal to δ . The full equilibrium strategies are (with this x): Player II accepts any demand x' so that $0 \leq x' \leq x$, and rejects any $x' > x$. Any counter-demand by player II in the second round is $y = 1$, and player I accepts any amount in the second round. The strategies have to specify the actions of the players in the second stage in order to perform the backward induction analysis.

Player I's demand x in the first round is easily found graphically as shown in Figure 5.11. This shows, as the outer curve, the Pareto-frontier of the bargaining set, here for the utility functions $u(x) = \sqrt{x}$ and $v(y) = y$ used earlier. Recall that for $0 \leq x \leq 1$, one obtains any point $(u(x), v(1 - x))$ on that Pareto-frontier. Multiplying any such pair of real numbers with the factor δ gives a "shrunk" version of the Pareto-frontier, shown as the curve with its lower endpoint $(\delta, 0)$ and upper endpoint $(0, \delta)$. This second curve is traversed when looking at the points $(\delta u(1 - y), \delta v(y))$ when y is changed from 0 to 1, which are the discounted payoffs after agreement in the second round. In the subgame given by the second round, player I can only accept or reject, so this ultimatum game has the upper endpoint $(0, \delta)$ of the curve, denoted by A in Figure 5.11, as its payoff pair. Player I's consideration in round one is to find a point B on the curve of first-round payoffs so that player II is indifferent between accepting and rejecting (where she will accept in the SPNE). So this point B has the same vertical coordinate δ as A , that is, B is of the




Figure 5.11 Graphical solution of the two-round bargaining game in Figure 5.10. The inner curve is the set of discounted payoff pairs for the second round.

form $(u(x), \delta)$. As shown by the arrow, B is found by moving horizontally from point A until hitting the outer curve. In the example, x is the solution to the equation $v(1 - x) = \delta$, which is simply $x = 1 - \delta$ because v is the identity function. The resulting utility $u(x)$ to player I is $\sqrt{1 - \delta}$ (which is larger than $1 - \delta$).

5.10 Alternating offers over several rounds

The bargaining game with three rounds of alternating offers is shown in Figure 5.12. Player I makes a demand x in the first round, which can be accepted or rejected by player II. If she rejects, she can make a counter-demand y , which in turn player I can accept or reject. If player I accepts, both players' payoffs are discounted, giving the expected payoff pair $(\delta u(1 - y), \delta v(y))$. If player I rejects, then he can make in the third and final round a counter-counter-demand s , which if accepted gives the pair of twice-discounted payoffs $(\delta^2 u(s), \delta^2 v(1 - s))$. If player II rejects the final demand s , both players get zero. These expected payoffs result because any rejection independently incurs a risk that, with probability $1 - \delta$, the game terminates with no agreement and payoff pair $(0, 0)$, and continues with probability δ into the next round. Figure 5.12 shows directly the discounted payoffs and not the chance moves between the bargaining rounds, which would have to be included in a detailed model like Figure 5.9.

The SPNE of the three-round game is found by backward induction. The third and final round is an ultimatum game where player I demands $s = 1$, which is accepted by player II. The set of twice discounted payoffs $(\delta^2 u(s), \delta^2 v(1 - s))$ is the inner curve on the left in Figure 5.13. Player I's demand $s = 1$ gives the payoff pair $(\delta^2, 0)$, shown as point A in the figure. In the previous round, player II maximises her payoff $\delta v(y)$ by making player I indifferent between accepting and rejecting, so y is the solution to the equation $\delta u(1 - y) = \delta^2$. This is point B in the figure, which is the point on the curve of second-round payoff pairs $(\delta u(1 - y), \delta v(y))$ with the same horizontal coordinate as A .




Figure 5.12 Bargaining over three rounds.

With y determined like that, player I's demand x in the first round is the solution to the equation $v(1-x) = \delta v(y)$ because this makes player II indifferent between accepting and rejecting in the first round. The corresponding point $C = (u(x), v(1-x))$ on the set of first-round payoff pairs has the same vertical coordinate as B . On the equilibrium path, player I demands x in the first round, which is immediately accepted by player II. The complete equilibrium strategies are as before, where a player accepts any demand up to and including the demand in the SPNE, and otherwise rejects and makes the described counter- or counter-counter-demand in the next round.

The game with four bargaining rounds has the same structure as the game with three rounds, except that if player II rejects player I's demand s in round three, she can make a final demand t of her own in the fourth and final round, which player I can only accept or reject. If player I accepts, the players receive the three-times discounted payoffs $(\delta^3 u(1-t), \delta^3 v(t))$. The corresponding SPNE is solved graphically as shown on the right in Figure 5.13. The last round is an ultimatum game where player II can demand




Figure 5.13 Graphical solution of the bargaining game over three rounds (left) and four rounds (right).

the whole pie with $t = 1$, which determines the payoff pair $(0, \delta^3)$, which is point A in the figure. In the previous round, the third round of the game, player I demands s so that player II is indifferent between accepting and rejecting, which gives point B , which is the pair of discounted payoffs $(\delta^2 u(s), \delta^2 v(1 - s))$ so that $\delta^2 v(1 - s) = \delta^3$. In the round before that, player II makes player I indifferent between accepting and rejecting with the demand y so that $\delta u(1 - y) = \delta^2 u(s)$, shown as point C . In the first round, player I demands x so as to make player II indifferent, according to $v(1 - x) = \delta v(y)$, which defines point $D = (u(x), v(1 - x))$. This is also the equilibrium outcome because that demand x is accepted by player II.

In our example with utility functions $u(x) = \sqrt{x}$ and $v(y) = y$, the SPNE of the game with three rounds is found as follows (according to the left picture in Figure 5.13). The vertical arrow from point A to point B gives the equation $\delta u(1 - y) = \delta^2$ or $u(1 - y) = \sqrt{1 - y} = \delta$, that is, $y = 1 - \delta^2$. The horizontal arrow from B to C represents the equation $v(1 - x) = \delta v(y)$, that is, $1 - x = \delta(1 - \delta^2) = \delta - \delta^3$, which has the solution $x = 1 - \delta + \delta^3$. This is player I's demand in the first round in the SPNE; his payoff is $u(x) = \sqrt{1 - \delta + \delta^3}$.

In the game with four rounds, the right picture in Figure 5.13 gives the following SPNE. The arrow from A to B gives the equation $\delta^2 v(1 - s) = \delta^3$ or $v(1 - s) = \delta$, which determines player I's demand $s = 1 - \delta$ in round three.

The arrow from B to C gives the equation $\delta u(1 - y) = \delta^2 u(s)$ or $u(1 - y) = \delta u(s)$, that is, $\sqrt{1 - y} = \delta \sqrt{1 - \delta}$ or equivalently $1 - y = \delta^2(1 - \delta) = \delta^2 - \delta^3$, which has the solution $y = 1 - \delta^2 + \delta^3$ that determines player II's demand y in the second round. Finally, the arrow from C to D gives the equation $v(1 - x) = \delta v(y)$ or $1 - x = \delta(1 - \delta^2 + \delta^3)$, which gives player I's demand x in the first round as $x = 1 - \delta + \delta^3 - \delta^4$.

Note that in the game of bargaining over four rounds, the arrow from A to B defines the same equation as in the two-round game in Figure 5.11; essentially, the fourth round in the four-round game describes the same situation, and gives the same picture, as the second round in the two-round game, except that all payoffs are shrunk by the factor δ^2 .

The bargaining game with alternating offers can be defined for any number of rounds. In each odd-numbered round (the first, third, fifth round, etc.), player I makes a demand, which can be accepted or rejected by player II. After rejection, the game continues with probability δ into the next round. Player II can make a demand in each even-numbered round. If the total number of rounds is even, player II can make the last demand, which is the whole pie, which is advantageous for player II. If the total number of rounds is odd, player I can make the last demand, and in turn demand the whole pie for himself.

The graphical solutions in figures 5.11 and 5.13 can be extended to any number of rounds, with payoffs in the k th round given by $(\delta^{k-1}u(z), \delta^{k-1}v(1-z))$ for some $z \in [0, 1]$. These payoff pairs define a curve that is the original Pareto-frontier of the bargaining set shrunk by the factor δ^{k-1} . The examples for a total number of two, three, or four rounds show that the resulting SPNE defines a point $(u(x), v(1-x))$ which the players receive because they agree in round one.

In general, if the total number of rounds is odd, player I is always better off than if the total number of rounds is even, which favours player II who can make the last demand. However, with an increasing total number of rounds, the resulting points $(u(x), v(1-x))$ move towards each other. We should expect that with a very large number of rounds, it matters less and less who can make the last demand (because the future discounted payoffs from which the backward induction starts are very small), and expect some convergence. However, we will not study the “limit” of these finite games, but instead consider a new concept, a game with an infinite number of rounds.

5.11 Stationary strategies

We now analyse the bargaining game with alternating offers with an *infinite* number of rounds, that is, the game goes on forever. Formally, this is a game Γ with perfect information represented by an infinite tree. It can be defined recursively as shown in Figure 5.14. In the first and second round, player I and player II each make a demand x and y , respectively, from the interval $[0, 1]$, which the other player can accept or reject. After rejection, a chance move terminates the game with probability $1 - \delta$ where both players receive payoff zero. With probability δ , the game continues into the next round where the player who has just rejected the offer can make her counter-demand. With the beginning of the third round, this is again player I, so that this defines the same game Γ that we started with in round one. Correspondingly, Γ is simply appended as the beginning of the subgame that starts after the second chance move that follows player I's rejection of player II's counter-demand y . (Note that infinitely many copies of Γ are substituted at each node that results from any x followed by any y .)

An SPNE of this infinite game cannot be defined by backward induction because the game goes on forever, so that there is no final move from which one could start the induction. However, the game is defined recursively and is the same game after two rounds. Hence, it is possible to define strategies of the players that repeat themselves any two rounds. Such strategies are called *stationary* and mean that player I always demands x whenever he can make a demand (in each odd-numbered round, starting with




Figure 5.14 The bargaining game Γ with an infinite number of rounds, which repeats itself after any demand x of player I in round one and any demand y of player II in round two after each player has rejected the other's demand; each time, the next round is reached with probability δ .

the first round), and player II always demands y (in each even-numbered round, starting with the second round). The SPNE condition states that these strategies should define an equilibrium in any subgame of the game. Any such subgame starts either with a chance move, or with a demand by one of the players.

In addition to the players's demands in each round, the SPNE should also specify whether the other player accepts or rejects the offer. After rejection, the resulting expected payoffs automatically shrink by a factor of δ . Consequently, a player's demand should not be so high as to result in rejection, but should instead make it optimal for the other player to accept. The demand itself is maximal subject to this condition, that is, it is, as before, chosen so as to make the other player indifferent between accepting and rejecting, and the other player accepts in the SPNE.

We now show that an SPNE exists, and how to find it, using a graph similar to the left picture in Figure 5.13. Like there, we draw the three curves of payoff pairs $(u(x), v(1-x))$

resulting from player I's demand x in round one if this demand is accepted by player II, the discounted payoffs $\delta(u(1-y), \delta v(y))$ if player I accepts player II's demand y in round two, and the third-round twice-discounted payoffs $\delta^2(u(s), \delta^2 v(1-s))$ if player II accepts player I's demand s in round three. The demands x and y are chosen so as to make the other player indifferent between acceptance and rejection. For the demand s in round three, we want to fulfil the requirement of stationary strategies, that is, $s = x$.




Figure 5.15 Finding stationary strategies in an SPNE of the infinite game, by following the arrows from A to B to C to D with the requirement that $A = D$.

Figure 5.15 illustrates how to fulfil these requirements. Suppose that player I's demand s in round three defines the point $A = (\delta^2 u(s), \delta^2 v(1-s))$, relatively high up on the inner curve as shown in the first picture. In the previous round two, point $B = (\delta u(1-y), \delta v(y))$ has the same horizontal coordinate as point A because player II tries to make player I indifferent with her demand y , according to the equation

$$\delta u(1-y) = \delta^2 u(s). \quad (5.3)$$

In the first round, point C is found similarly as the point $(u(x), v(1-x))$ that has the same vertical coordinate as B , because player I demands x so as to make player II indifferent between accepting and rejecting, according to

$$v(1-x) = \delta v(y). \quad (5.4)$$

The demand x is the same as demand s , as required in stationary strategies, if point C defines the same relative position on the outer curve as A on the inner curve. In Figure 5.15, we have shown point D as $\delta^2 C$, that is, point C shrunk by a factor of δ^2 , which is the expected payoff pair two rounds later. In the left picture, A and D do not coincide, because $s < x$, that is, the third-round demand s was too small. The second picture in Figure 5.15 shows a point A on the inner curve corresponding to a demand s that is too favourable for player I, where the arrows from A to B to C according to equations (5.3) and (5.4) give a first-round demand x with $x < s$.

In general, the pictures demonstrate that when starting from $s = 0$, these equations give a first-round demand x where $x > s$, and when starting from $s = 1$, a first-round demand x where $x < s$. (Note that these extreme starting points $s = 0$ and $s = 1$ give exactly the solutions of bargaining over two and three rounds, respectively.) Moreover, when starting from any s in $[0, 1]$, the resulting demand x is a *continuous* function of s . Because the continuous function $x - s$ is positive for $s = 0$ and negative for $s = 1$, it is zero for some intermediate value s , where $s = x$, which defines a stationary SPNE.

In short, the stationary SPNE of the infinite bargaining game is a solution to the equations (5.3), (5.4), and $s = x$. These are equivalent to the two, nicely symmetric equations

$$u(1 - y) = \delta u(x), \quad v(1 - x) = \delta v(y). \quad (5.5)$$

The first of these equations expresses that player II makes player I indifferent in every even-numbered round by her demand y . The second equation in (5.5) states that player I makes player II indifferent by his demand x in every odd-numbered round. In this stationary SPNE, player I's demand x in round one is immediately accepted by player II. As before, the full strategies specify acceptance of any lower and rejection of any higher demand. The actual demand of the other player is always accepted, so the later rounds are never reached.

In the example with utility functions $u(x) = \sqrt{x}$ and $v(y) = y$, the stationary strategies in an SPNE are found as follows. They are given by the demands x and y that solve (5.5), that is,

$$\sqrt{1 - y} = \delta \sqrt{x}, \quad 1 - x = \delta y.$$

The first equation is equivalent to $y = 1 - \delta^2 x$, which substituted into the second equation gives $1 - x = \delta - \delta^3 x$, or $1 - \delta = x(1 - \delta^3)$. Because $1 - \delta^3 = (1 - \delta)(1 + \delta + \delta^2)$, this gives the solutions for x and y as

$$x = \frac{1 - \delta}{1 - \delta^3} = \frac{1}{1 + \delta + \delta^2}, \quad y = \frac{1 + \delta}{1 + \delta + \delta^2}. \quad (5.6)$$

The pie is split in the first round into $x = 1/(1 + \delta + \delta^2)$ for player I and $1 - x = (\delta + \delta^2)/(1 + \delta + \delta^2)$ for player II. The share $1 - x$ of the pie for player II is lower than her demand y in the second round, as illustrated by the last picture in Figure 5.15, which shows that the relative position of point A on the outer curve is higher than that of point B on the middle curve. This holds because the second-round payoffs are discounted.

Does it matter for our analysis that player I makes the first demand? Of course, if player II made the first demand, one could swap the players and analyse the game with

the described method. However, the game where player II makes the first demand, and the corresponding stationary solution, is already analysed as part of the game Γ where player I moves first. The reason is that the game with player II moving first is the subgame of Γ in Figure 5.14 after the first chance move, which repeats itself as in Γ . All we would have to do is to ignore player I's first demand, and the first chance move and the corresponding discount factor δ . So we would draw the same picture of curves with discounted payoff pairs as in Figure 5.15, except that we would omit the outer curve and scale the picture by multiplication with $1/\delta$.

For the alternating-offers game with player II moving first, the resulting equations are given as before by (5.5). They have the same stationary equilibrium solution. This solution starts with player II making the first demand y , followed by player I's counter-demand x , and so on in subsequent rounds. Because in this SPNE the first demand is accepted, player II gets payoff $v(y)$ and player I gets payoff $u(1 - y)$, which is $\delta u(x)$ by (5.5). Hence, all that changes when player II makes the first demand is that player II's payoff is no longer discounted by multiplication with δ , and player I's payoff is discounted instead.

5.12 The Nash bargaining solution via alternating offers

In this final section, we show that for very patient players, when the discount factor δ approaches 1, the stationary SPNE converges to the Nash bargaining solution. This is true for the example that we have used throughout, with utility functions $u(x) = \sqrt{x}$ and $v(y) = y$, where the split in the Nash bargaining solution is $(x, 1 - x) = (1/3, 2/3)$. This is also the limit of the stationary solution in (5.6) when $\delta \rightarrow 1$ (note the convenient cancellation by $1 - \delta$ in the term $(1 - \delta)/(1 - \delta^3)$ that represents x ; this term has an undefined limit $0/0$ when $\delta = 1$, which without the cancellation requires using l'Hôpital's rule).




Figure 5.16 Why stationary strategies converge to the Nash bargaining solution.

The stationary solution to the infinite bargaining game of alternating offers is shown as the last picture in Figure 5.15, with points $B = (\delta u(1 - y), \delta v(y))$ and $C = (u(x), v(1 - x))$. Let the point F be given by $F = (u(1 - y), v(y))$, which is the “undiscounted” version of B , that is, it has the same relative position on the outer curve as B on the middle curve. As $\delta \rightarrow 1$, the points F and C get closer and closer to each other, and in the limit converge to the same point. As shown in Figure 5.16, the points F and C define a line where the line segment connecting the two points is a “cord” of the Pareto-frontier, and in the limit the whole line becomes a *tangent* to the Pareto-frontier, when the two points C and F converge (if the Pareto-frontier is differentiable, the tangent is unique, otherwise we obtain in the limit some tangent to the frontier).

We are now very close to showing that the limit of point C when $\delta \rightarrow 1$ defines the Nash bargaining solution. We use Proposition 5.4 that characterises the bargaining solution geometrically. Namely, we consider the points F and C as the upper-left and lower-right corner of a rectangle. As shown in Figure 5.16, the upper-right corner of that rectangle is the point $E = (u(x), v(y))$, and its lower-left corner is the point $G = (u(1 - y), v(1 - x))$. The equations (5.5) state that $G = \delta E$, that is, the points G and E are co-linear, and the diagonal of the rectangle through E and G defines a line through the origin $(0, 0)$. This line has slope $v(y)/u(x)$, shown by the angle α near the origin. The other diagonal of the rectangle through the points F and C has the same (negative) slope, shown by the angle α when that line intersects the horizontal u -axis.

In the limit, when $\delta \rightarrow 1$, we have that $y \rightarrow 1 - x$, and we obtain the same picture as in Figure 5.4 of a “roof” with two sides that have, in absolute value, the same slope. The left side of the roof is a line through the origin, and the right side of the roof is a tangent (of correct slope) to the Pareto-frontier of the bargaining set. This is exactly the condition of Proposition 5.4 that characterises the apex of the roof as the Nash bargaining solution (U, V) .

We summarise the main result of this chapter. The model of alternating offers with a discount factor is an interesting non-cooperative game in its own right. When the number of bargaining rounds is infinite, the game has an SPNE in stationary strategies. When the discount factor δ gets close to one, this stationary equilibrium gives a payoff pair that approaches the Nash bargaining solution.

The Nash bargaining solution is originally derived from general “axioms” that a general solution to bargaining problems should satisfy. The “implementation” of this solution is not spelled out in this axiomatic approach, where it is merely assumed that the players reach a binding agreement in some way.

The alternating-offers game justifies this solution via a completely different route, namely a model of players that interact by making demands and possible rejections and counter-demands. In the resulting SPNE with stationary strategies, each player anticipates the future behaviour of the other player. The players immediately reach an agreement, which is the Nash bargaining solution.

\Rightarrow Exercise 5.5 on page 163 gives an excellent example that concerns all aspects of the alternating-offers model.

5.13 Exercises for Chapter 5

Exercise 5.1 describes a bargaining set derived from a bimatrix game in analogy to Figure 5.1, with a somewhat different shape. Exercise 5.2 proves that the Nash bargaining solution obtained by maximising the Nash product actually fulfils one of the required axioms. Exercise 5.3 considers a discrete ultimatum game similar to Figure 5.7, where you are asked in (d) to find all its Nash equilibria; not all of these are SPNE, so they may involve the “threat” of rejecting an offer that would be accepted in an SPNE. Exercises 5.4 and 5.5 consider the problem of splitting a unit pie with a specific pair of utility functions, and the iterated-offers bargaining game.

Exercise 5.1 Draw the convex hull of payoff pairs for the following 2×2 game. What is the threat point? Show the bargaining set derived from this game, as explained in Section 5.4. Indicate the Pareto-frontier. Find the Nash bargaining solution (U, V) .

		II	
		<i>l</i>	<i>r</i>
I	<i>T</i>	4 3	1 2
	<i>B</i>	3 4	1 1

Exercise 5.2 Consider a bargaining set S with threat point $(0, 0)$. Let $a > 0$, $b > 0$, and let $S' = \{(au, bv) \mid (u, v) \in S\}$. The set S' is the set of utility pairs in S rescaled by a for player I and by b for player II. Recall that one axiomatic property of the bargaining solution is independence from these scale factors.

Show that the Nash product fulfils this property, that is, the Nash bargaining solution from S obtained by maximising the Nash product re-scales to become the solution for S' .

Note: Don't panic! This is extremely easy once you state in the right way what it means that something maximises something. You may find it useful to consider first the simpler case $b = 1$.

Exercise 5.3 We repeat the definition of the discretised ultimatum game. Let M be a positive integer. Player I's possible actions are to ask for a number x , called his demand, which is one of the (integer) numbers $0, 1, 2, \dots, M$. In response to this demand, player II can either accept (A) or reject (R). When player II accepts, then player I will receive a payoff of x and player II a payoff of $M - x$. When player II rejects, both players receive payoff zero.

- Draw the game tree for this game, as an extensive game with perfect information, for $M = 3$.
- What is the number of pure strategies in this game, for general M , for player I and for player II? What is the number of reduced pure strategies in this game, for player I and for player II?

- (c) Determine all subgame perfect equilibria of the game in pure strategies.
- (d) Determine all Nash equilibria (not necessarily subgame perfect) of this game in pure strategies.
- (e) Determine all subgame perfect equilibria of this game where both players may use behaviour strategies, in addition to those in (c).

Exercise 5.4 Consider the following bargaining problem. In the usual way, a “unit pie” is split into non-negative amounts x and y with $x + y \leq 1$. The utility function of player I is $u(x) = x$, the utility function of player II is $v(y) = 1 - (1 - y)^2$. The threat point is $(0, 0)$.

- (a) Draw the bargaining set (as set of possible utilities (u, v)) and indicate the Pareto-frontier in the picture.
- (b) Find the Nash bargaining solution. How will the pie be split, and what are the utilities of the players?

Exercise 5.5 Consider the bargaining problem of splitting a unit pie with utility functions $u, v: [0, 1] \rightarrow [0, 1]$ as in Exercise 5.4, that is, $u(x) = x$ and $v(y) = 1 - (1 - y)^2$. Assume now that the bargaining outcome is determined by the subgame perfect equilibrium of the standard alternating-offers bargaining game with k rounds. In round 1 player I makes a certain demand. If player II rejects this, she can make a certain counter-demand in round 2, and so on. A rejection in the last round k means that both players get nothing. The last player to make a demand is player I if k is odd, and player II if k is even. When agreement takes place in round i for some i with $1 \leq i \leq k$, then the bargaining set has shrunk by a factor of δ^{i-1} . That is, the utility of what the players can get after each round of rejection is reduced by multiplication with the discount factor δ .

- (a) Suppose that the number of rounds is $k = 2$. Draw the picture of shrinking bargaining sets (for some choice of δ) in these two rounds. Find the subgame perfect equilibrium, in terms of player I's demand x , for general δ . What is the demand x when $\delta = 3/4$? What are the strategies of the two players in this subgame perfect equilibrium?
- (b) Suppose that the number of rounds is $k = 3$. Draw the picture of shrinking bargaining sets in these three rounds. Find the subgame perfect equilibrium, for general δ , computing the demand of player I in round 1 (you do not need to specify the players' strategies). What is player I's demand x when $\delta = 1/4$? (Then x happens to be a rational number.)
- (c) Consider now an infinite number of rounds, and look for stationary equilibrium strategies, where player I demands the same amount x each time he makes a demand, and player II demands the same amount y every time she makes a demand. What is this demand x when $\delta = 3/5$? What is the demand x when δ goes to 1? Compare this with the result of Exercise 5.4(b).

Chapter 6

Geometric representation of equilibria

This chapter is an extension of the lecture notes “Game Theory Basics” on topics that have been first taught in autumn 2010. It comprises about one week of lectures that partially replace the old chapter 5 on bargaining, and follows on from Section 3.13.

In this chapter we consider a geometric method to identify the Nash equilibria of a two-player game with up to three strategies per player, and provide an existence proof of Nash equilibria for nondegenerate two-player games known as the “Lemke–Howson algorithm”.

6.1 Learning objectives

After studying this chapter, you should be able to:

- draw best response diagrams for bimatrix games with up to three strategies per player;
- identify the Nash equilibria of the game with the help of the labels used in these diagrams, and explain the connection of labels and the best response condition;
- show the computation of one Nash equilibrium with the Lemke–Howson algorithm using the diagram.

6.2 Further reading

The best response diagrams are to my knowledge not found in undergraduate textbooks. A general description of the Lemke–Howson algorithm is found in my chapter of a postgraduate textbook,

- B. von Stengel (2007), Equilibrium computation for two-player games in strategic and extensive form. Chapter 3 of *Algorithmic Game Theory*, eds. N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, Cambridge Univ. Press, Cambridge, pages 53–78.

6.3 Introduction

The following sections present methods for finding Nash equilibria in mixed strategies for two-player games in strategic form.

We recall the notation used in Chapter 3. Player I has m and player II has n strategies. For any pair (i, j) of strategies, each player receives a payoff, which is independent of the payoff to the other player. In particular, we do *not* require the zero-sum property. As usual, the m strategies of player I can be considered as rows and the n pure strategies of player II as columns in a table representing the strategic form. The game is then completely specified by two $m \times n$ matrices A and B containing the payoffs to player I and player II, respectively. The game is therefore also called a *bimatrix* game, given by the matrix pair (A, B) .

In the following, we give a geometric illustration of equilibria for such bimatrix games. This is particularly helpful for games with up to three pure strategies per player, in particular 3×3 games, and can even be pictured (although less easily drawn) if a player has four strategies. A number of considerations can be demonstrated with 3×3 games which would be too simple or pointless with 2×2 games, so that this geometric approach is very useful.

Given this geometric approach, we will then present an *algorithm* – a systematic method which can be performed by a computer – for finding one equilibrium of a bimatrix game, known as the *Lemke–Howson* algorithm. This method provides also a constructive and elementary proof that every bimatrix has at least one equilibrium. Moreover, the method shows that a nondegenerate bimatrix game has always an *odd* number of Nash equilibria.

6.4 Best response regions

Consider an $m \times n$ bimatrix game (A, B) , with payoff matrix A for the row player (player I) and B for the column player (player II). Throughout, we will use as an example the 3×3 game

$$A = \begin{bmatrix} 0 & 3 & 0 \\ 1 & 0 & 1 \\ -3 & 4 & 5 \end{bmatrix}, \quad B = \begin{bmatrix} 0 & 1 & -2 \\ 2 & 0 & 3 \\ 2 & 1 & 0 \end{bmatrix}. \quad (6.1)$$

Note, however, that the presented method does not require that both players have the same number of pure strategies. That is, $m \neq n$ is quite possible. In that respect, a 3×2 game, for example, would be more typical, but this would be less suitable to illustrate the power of the construction.

The **first** step of the construction is a unique numbering of the pure strategies of the players to identify them more easily. Let the set of pure strategies of player I be given by

$$M = \{1, \dots, m\}$$

and the set of pure strategies of player II by

$$N = \{m+1, \dots, m+n\}.$$

That is, the numbering of the strategies is continued for player II to produce two disjoint sets M and N . The pure strategies will be used to identify equilibria by considering them as *labels* of certain geometrically defined sets. That is, a *label* is defined as an element of $M \cup N$. We will circle these numbers to identify them as such labels. In our example (6.1) we therefore have

$$M = \{\textcircled{1}, \textcircled{2}, \textcircled{3}\}, \quad N = \{\textcircled{4}, \textcircled{5}, \textcircled{6}\}.$$

It is helpful to mark the rows of A with the pure strategies $\textcircled{1}$, $\textcircled{2}$, $\textcircled{3}$ of player I and the columns of B with the pure strategies $\textcircled{4}$, $\textcircled{5}$, $\textcircled{6}$ of player II (that is, writing $\textcircled{1}$, $\textcircled{2}$, $\textcircled{3}$ to the left of A and $\textcircled{4}$, $\textcircled{5}$, $\textcircled{6}$ at the top of B).

The **second** step is the geometric representation of the sets of *mixed strategies* of player I and player II which we recall from Section 3.8. A mixed strategy of player I is a probability distribution on the set M of his pure strategies. This is an m -dimensional vector x with nonnegative components x_i for i in M that sum up to one. In the same manner, y is a mixed strategy of player II if it has n nonnegative components y_j for j in N with $\sum_{j \in N} y_j = 1$. The sets X and Y of mixed strategies of player I and II are therefore defined as

$$X = \{x \in \mathbb{R}^M \mid \sum_{i \in M} x_i = 1, x_i \geq 0 \text{ for } i \in M\},$$

$$Y = \{y \in \mathbb{R}^N \mid \sum_{j \in N} y_j = 1, y_j \geq 0 \text{ for } j \in N\}.$$

We use here the notation \mathbb{R}^N instead of \mathbb{R}^n to indicate that the components y_j of a vector y in \mathbb{R}^N have the elements j in N as subscripts, which are here not the first n natural numbers, but, according to the above definition of N , the numbers $m+1, \dots, m+n$. Instead of \mathbb{R}^M one could also write \mathbb{R}^m .

For $n = 3$ as in the above example (6.1) the set Y is the set of all convex combinations of the unit vectors $\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$, $\begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$, $\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$ (we use x and y as column vectors). That is, Y is a triangle, as Figure 6.1 illustrates, which is similar to Figure 3.4. Note the numbering of the components y_4, y_5, y_6 of y .

This triangle is part of the plane that goes through the three unit vectors and is therefore a two-dimensional geometric figure. For $n = 4$ one would obtain a tetrahedron, for $n = 2$ merely a line segment. In general, the “convex hull” of the unit vectors is called a *simplex*. This geometric object has one dimension less than the space containing the unit vectors under consideration, because the components of any vector contained in that simplex add up to one, that is, they fulfill one linear equation.

The **third** and most important step in the construction is the subdivision of the strategy sets into *best response regions*. Such a best response region is that set of mixed strategies of a player where a pure strategy of the *other* player is a best response. We first illustrate this with our example (6.1) with $n = 3$ and consider a mixed strategy y




Figure 6.1 Mixed strategy simplex of player II for a 3×3 game.

of player II. Then $y = (y_4, y_5, y_6)^\top$, where the superscript “ \top ” means transposition, to write the column vector y as a transposed row vector. For given y , player I has the pure strategies ①, ② and ③ as possible reponses and then gets the expected payoffs for

$$\text{①} : 3y_5, \quad \text{②} : y_4 + y_6, \quad \text{③} : -3y_4 + 4y_5 + 5y_6.$$

When is ① a best response to y ? Apparently when the expected payoff for ① is at least as high as that for ② and ③, that is, if the inequalities

$$3y_5 \geq y_4 + y_6, \quad 3y_5 \geq -3y_4 + 4y_5 + 5y_6 \quad (6.2)$$

hold, or equivalently

$$3y_5 \geq y_4 + y_6, \quad 3y_4 \geq y_5 + 5y_6. \quad (6.3)$$

The inequalities in (6.3) have been obtained by sorting the linear coefficients of the variables y_4, y_5, y_6 so that they appear only once and with positive sign on one side of each inequality; for example, rather than writing $3y_5 - y_4 \geq y_6$ we write $3y_5 \geq y_4 + y_6$ by adding y_4 on both sides.

The best response region for the strategy ① is therefore that subset of Y where the inequalities (6.2) are fulfilled. These are in general $m - 1$ linear inequalities that describe the comparison with the other pure strategies of player I. The general definition of a best response region for a pure strategy i in M of player I is

$$Y(i) = \{y \in Y \mid i \text{ is a best response to } y\}$$

and correspondingly for a pure strategy j in N of player II

$$X(j) = \{x \in X \mid j \text{ is a best response to } x\},$$

where $X(j)$ is defined by $n - 1$ inequalities.

In addition to these inequalities we have the linear constraints that define Y and X , respectively. These are the n respectively m inequalities $y_j \geq 0$ for $j \in N$ and $x_i \geq 0$ for $i \in M$, and the equation that these components sum to one because they are probabilities. Each best response region is therefore a so-called “polytope”, a set of vectors that is bounded and described by linear inequalities and equalities.




Figure 6.2 Pairwise comparison of ①, ②, ③ to construct the best response regions for these three strategies of player I in the game (6.1).

In our example we obtain $Y(\textcircled{1})$, $Y(\textcircled{2})$, $Y(\textcircled{3})$ as follows. Recall that we drew Y as a triangle. Then the set where *both* ① and ② are best responses is a line segment in Y defined by the equation

$$3y_5 = y_4 + y_6. \quad (6.4)$$

As mentioned, it is useful to sort this linear equation in y_4 , y_5 , and y_6 so that all coefficients are nonnegative (which is here already the case), in a similar manner as (6.3) is an equivalent way of writing (6.2). The reason is that then one can easily identify two points

on the mentioned line segment that represents the intersection of $Y(\textcircled{1})$ and $Y(\textcircled{2})$. Each of these points is on one side of the triangle Y , namely for $y_6 = 0$ and for $y_4 = 0$. The resulting solutions to (6.4) are $(y_4, y_5, y_6) = (3/4, 1/4, 0)$ and $(y_4, y_5, y_6) = (0, 1/4, 3/4)$. Both are solutions (y_4, y_5, y_6) to (6.4), and hence any convex combination of these is also a solution: In general, consider a vector a and scalar a_0 and let the vectors y and y' fulfill the same linear equation $a^\top y = a_0$ and $a^\top y' = a_0$, and let $p \in [0, 1]$. Then $a^\top (y(1-p) + y'p) = (a^\top y)(1-p) + (a^\top y')p = a_0(1-p) + a_0p = a_0$.

Then, a simple sign test in the corners of the triangle shows on which side of said line segment we can find the respective best response region: For example, for $(y_4, y_5, y_6) = (1, 0, 0)$ the pure strategy $\textcircled{2}$ is obviously better for player I than $\textcircled{1}$. The three pictures in Figure 6.2 show geometrically the best response regions that would result if player I had only two pure strategies, disregarding the respective third strategy. The first picture shows this for the two strategies $\textcircled{1}, \textcircled{2}$, the second one for $\textcircled{1}, \textcircled{3}$, and the last one for $\textcircled{2}, \textcircled{3}$.

\Rightarrow Construct the second and third picture in Figure 6.2 in analogy to the above construction of the first picture that uses the “indifference line” between $\textcircled{1}$ and $\textcircled{2}$ in equation (6.4). See also Exercise 6.1(a).

The intersections of these sets, which were obtained by pairwise comparison, for example in (6.3) for $Y(\textcircled{1})$, then yield the following picture, where the circled numbers $\textcircled{1}, \textcircled{2}, \textcircled{3}$ label the sets $Y(\textcircled{1}), Y(\textcircled{2}), Y(\textcircled{3})$.




Figure 6.3 Best response regions as obtained from the intersections of the pairwise comparisons shown in Figure 6.2.

So far we have subdivided the mixed strategy set of one player, like here Y for player II, into regions where a pure strategy of the other player is a best response. We call such a pure strategy the *label* of a mixed strategy y . That is, y has label i in M if and only if $y \in Y(i)$. A mixed strategy y can have more than one label (as the points on the line segments in the last picture do), namely if y has more than one best response. For example, $(y_4, y_5, y_6) = (3/4, 1/4, 0)$ has the two labels $\textcircled{1}$ and $\textcircled{2}$.

The **fourth** step of the construction is now the introduction of *own* pure strategies as additional labels, that is, labels of player II for mixed strategies in Y . Namely, y gets the

label j in N if $y_j = 0$. That is, a mixed strategy y is also labelled with those own strategies that it does *not* choose with positive probability. Formally, these labels are defined for $i \in M$ and $j \in N$ according to

$$X(i) = \{x \in X \mid x_i = 0\},$$

$$Y(j) = \{y \in Y \mid y_j = 0\}$$

and we say x in X has the label k in $M \cup N$ if $x \in X(k)$. Correspondingly, y in Y is said to have label k if $y \in Y(k)$, for any k in $M \cup N$. In our example (6.1) these labels are then represented as shown in Figure 6.4.




Figure 6.4 Mixed strategy sets X and Y for the game (6.1) subdivided into best response regions. The dots identify points x^1, \dots, x^7 and y^1, \dots, y^7 that have three labels.

The picture on the right is Y with the already constructed labels ①, ②, and ③, where the additional labels ④, ⑤, ⑥ represent the sides of the triangle. In the interior of the triangle, all pure strategies of player II have positive probability, so they have only the best responses of player I as labels (in any dimension, the “own” labels define the sides of the mixed strategy simplex). This also means that the corner points no longer have to be annotated with the unit vectors, because they can be identified by the only own strategy that does not have probability zero. For example, the left corner of Y in the above picture where $(y_4, y_5, y_6) = (1, 0, 0)$ is identified by the labels ⑤ and ⑥ because there $y_5 = 0$ and $y_6 = 0$.

The left picture shows the analogous labelling of X , namely with the “own” labels ①, ②, ③ at the sides, and the labels ④, ⑤, ⑥ of the other player for the best response regions $X(④)$, $X(⑤)$, $X(⑥)$, which indicate when a pure strategy j of player II is a best response to x in X . We leave it as an exercise to construct these best response regions by drawing suitable line segments that define the intersections of any two best response regions in X .

The two strategy sets X and Y are now labelled as described. A pair of mixed strategies (x, y) carries certain labels k from $M \cup N$, which denote pure strategies of the players. For example, the pair (x, y) with $(x_1, x_2, x_3) = (0, 0, 1)$ and $(y_4, y_5, y_6) = (1, 0, 0)$ – in X

the top and in Y the leftmost point of the triangle in Figure 6.4 – has the labels ② (labels of both x and y), ① and ④ (label of x), and ⑤ and ⑥ (label of y). This is an example of a pair (x, y) of strategies which is not an equilibrium. Why? Because y is a best response to x (it uses only the strategy ④ with positive probability, and that is a best response to x because x is labelled with ④), but, conversely, the strategy ③ that x chooses with positive probability is not a best response to y because otherwise y would be labelled with ③. In other words, ③ is a label that is *missing* from the labels of x and y . We will now see easily that as long as a label is missing in that way, we cannot have an equilibrium.

6.5 Identifying equilibria

Step **five** of our construction is the identification of equilibria in the pair of labelled mixed strategy sets as for our example in Figure 6.4. Namely, an equilibrium is a **completely labelled** pair x, y of strategies. That is, a pair x, y is an equilibrium of the game if and only if every label in $M \cup N$ appears either as a label of x or of y (or both). The reason for this is the best response condition, Theorem 3.1.

An equilibrium is a pair of mixed strategies that are best responses to each other. According to Theorem 3.1, this means that all pure strategies that they choose with positive probability must be best responses. In other words, any pure strategy must be either a best response or have probability zero. This is exactly the purpose of the labelling: A label marks a pure strategy as a best response, or as pure strategy that has probability zero. Hence we obtain the earlier claim: In order to have an equilibrium, every pure strategy must appear as a label of a mixed strategy of a player.

We are thus left with a combinatorial condition that means checking for all pairs (x, y) in $X \times Y$ whether they have all labels. In our example, each of the six labels ①, ②, ③, ④, ⑤, ⑥ must appear either as a label of x or of y . Luckily, there are not many candidates of such mixed strategies where one could fulfill that condition. Figure 6.4 shows that there are only few points in either triangle that have three labels, and because six labels are needed in total, only those points have to be examined as to whether they have together all six labels. These points x^1, \dots, x^7 and y^1, \dots, y^7 are marked by dots in the figure.

We use these points to find the equilibria of the game, according to the following table.

x	labels	labels	y
x^1	① ② ④		
x^3	② ③ ⑤		
x^6	③ ④ ⑥		
x^7	① ③ ⑥		
x^2	② ④ ⑤	① ③ ⑥	y^6
x^4	③ ④ ⑤	① ② ⑥	y^4
x^5	① ④ ⑥	② ③ ⑤	y^2
		③ ④ ⑤	y^1
		② ⑤ ⑥	y^3
		① ② ③	y^5
		③ ④ ⑥	y^7

(6.5)

The three middle rows are equilibria because they describe those pairs of points that together have all labels. The first equilibrium is the strategy pair (x^2, y^6) with labels ② ④ ⑤ and ① ③ ⑥ and $x^2 = (1/2, 0, 1/2)$ and $y^6 = (1/4, 3/4, 0)$. The second equilibrium is the strategy pair (x^4, y^4) with labels ③ ④ ⑤ and ① ② ⑥ and $x^4 = (2/3, 1/3, 0)$ and $y^4 = (3/4, 1/4, 0)$. The third equilibrium is the strategy pair (x^5, y^2) with labels ① ④ ⑥ and ② ③ ⑤ and $x^5 = (0, 2/3, 1/3)$ and $y^2 = (1/2, 0, 1/2)$.

It is not mere luck that only finitely many points in X have m labels and, similarly, finitely many points in Y of n labels, but this is generally so unless there are certain special dependencies between the payoffs. Namely, consider how one obtains a label of a mixed strategy: Clearly, every mixed strategy – say y of player II – has at least one label because there is at least one pure best response to it. But every extra label means that y has to fulfill an extra linear equation: If that label is a strategy j of player II, it means that $y_j = 0$. If that label is a strategy i of player I, then the expected payoff for that strategy i must equal the expected payoff for the label that y already has, which is another linear equation. In our example, such an equation determines a line as used in the above construction of best response regions. That is, every label uses up one freedom for choosing y , and after $n - 1$ labels in addition to the first label, there is at most a point y left that fulfills the resulting $n - 1$ linear equalities in addition to the equation $\sum_{j \in N} y_j = 1$. The only case where this does not happen is when these linear equations are not independent, which is not normally the case. So, in general, one would expect that only finitely many points in X have m labels and finitely many points in Y have n labels. Only these can be parts of pairs that together have the required $m + n$ labels, namely all the labels in $M \cup N$ that we are looking for. In other words, a game usually only has a finite number of mixed strategy equilibria.

The geometric-combinatorial method can be sped up somewhat by the observation that at most one point y in Y has “matching” labels for a point in x (namely all those labels in $M \cup N$ that x does not have). For example, going through the above list of labels of points in X we note that the point x^3 with labels ② ③ ⑤ would have to have matching labels ① ④ ⑥ of corresponding point in Y , but no such point in Y exists because the region labelled with ① does not have a common point with the side of the triangle labelled ④, for example. In that way, it suffices to go through the points in X that have three labels and look for their possible “partners”, without having to check all points in Y .

6.6 The Lemke–Howson algorithm

The construction of mixed strategy sets and their labels provides a quick way of finding all equilibria of a bimatrix game. However, we have not yet shown that there is even one equilibrium. We continue to use the labelling technique to prove that there is an equilibrium, which can in fact be turned into an algorithm which is known as the Lemke–Howson algorithm, named after the people who discovered it.

For that purpose – step **six** in our construction – we add another point to each strategy set which is not a point in that set. This is the zero vector in m -space for the set X , call it $\mathbf{0}_x$, and the zero vector $\mathbf{0}_y$ in n -space for the set Y . We then connect $\mathbf{0}_x$ by line segments

to all the corners x^1, x^3 and x^7 of X (the unit vectors), and similarly $\mathbf{0}_y$ with line segments to the corners y^1, y^3 and y^7 of Y . The resulting pair of diagrams is shown in Figure 6.5.




Figure 6.5 Subdivision of the mixed strategy sets X and Y into labelled best regions, together with connecting their vertices with the respective zero vectors $\mathbf{0}_x$ and $\mathbf{0}_y$.

The pair $(\mathbf{0}_x, \mathbf{0}_y)$ does not belong to the strategy sets, but it is a so-called **artificial equilibrium** because it shares with an equilibrium the property that is completely labelled. Namely, every pure strategy has “probability” zero, so every i in M is a label of $\mathbf{0}_x$, and every j in N is a label of $\mathbf{0}_y$.

The pair $(\mathbf{0}_x, \mathbf{0}_y)$ serves as a starting pair for our construction. The main trick is now to relax the condition that the pair of points under consideration must be completely labelled. That, we allow one label to be *missing*. Any label can serve that purpose, but it must be fixed once it is chosen. For example, assume we decide that ① is the missing label. Then there are additional points in the diagram besides $(\mathbf{0}_x, \mathbf{0}_y)$ that have all labels except ①. This is the line segment connecting $(0, 0, 0)$ to x^3 which is the unit vector $(1, 0, 0)$ (because ① is missing, the first component of that vector is allowed to be positive). So we move to x^3 in X which has the three labels ②, ③, and, as a new label, ⑤ (which is the best response to x^3). In the other set, we have stayed at the artificial point $\mathbf{0}_y$. However, now

observe that we have in X a point with three labels, and in the set Y extended by $\mathbf{0}_y$ a point with the three labels ④, ⑤, ⑥, and label ① is still missing, so there must be a duplicate label – obviously the label ⑤ that was just “picked up” in X . So we now can safely drop label ⑤ by leaving $\mathbf{0}_y$ (and still have only ① as a missing label) and move to y^7 in Y which has labels ④, ⑥, and, as new label, ③.

We have now arrived at (x^3, y^7) which is a pair of points in $X \times Y$, and will not leave these strategy sets again. Observe that we are still missing label ① and therefore have a duplicate label which is the most recently encountered label ③ in Y . We can therefore drop this label in X , where we move from x^3 which has labels ②, ③, ⑤ by dropping label ③, that is, by moving along the line segment defined by the points in X labelled with both ② and ⑤ to the new point x^2 with labels ②, ④, ⑤. The new label of x^2 is ④, which is now duplicate. We therefore drop label ④ from y^7 in Y and move on to the point y^6 with labels ①, ③, ⑥. Here, the new label is ①, and this terminates our path because that was the only label that has been missing all along, and which is the only missing label, so that we have found an equilibrium, namely (x^2, y^6) , which is the first equilibrium in the table (6.5).

To repeat, the entire path of strategy pairs defined by the missing label ① is $(\mathbf{0}_x, \mathbf{0}_y) \rightarrow (x^3, \mathbf{0}_y) \rightarrow (x^3, y^7) \rightarrow (x^2, y^7) \rightarrow (x^2, y^6)$. The sequence of pairs starts and ends at a pair that is completely labelled, where the starting pair is the artificial equilibrium $(\mathbf{0}_x, \mathbf{0}_y)$. The first step is determined by the missing label, here ①. The steps alternate between changing either x or y in the pair (x, y) , and keeping the other strategy fixed. Each step from the second step onwards is determined by the duplicate label because one label is missing, and all other labels are present, and x has m and y has n labels, so one of the $m + n$ labels of (x, y) is duplicate. That duplicate label has just been picked up in one strategy (say x) and is then dropped in the other strategy (then y), always preserving the property that all $m + n - 1$ labels, except the missing label, are present throughout the path of points (this path includes the edge of changing a strategy, e.g. the edge that joins x^3 to x^2 defines the edge in $X \times Y$ that joins (x^3, y^7) to (x^2, y^7) above).

This is the general argument: First, we start from the artificial equilibrium $(\mathbf{0}_x, \mathbf{0}_y)$. Then we decide on a missing label, which gives us the freedom to move along a line segment in one strategy set or the other (both also connected to the artificial equilibrium). The result will be a path where we move alternately along a line segment in one set and stay fixed in the other (these are indeed line segments, because one linear equality is dropped, so the fact that there is one label less defines a line for the resulting linear equations, rather than a point). At the end of such a line segment, there will be one new label that we encounter. Either that is the missing label: In that case we have found an equilibrium. Or it is another label, which is not the missing label, and hence a duplicate label because it was not in the strategy set that we currently moved in (because we just found this in that strategy set as a new label), so, because all other labels have been present, it must be also a label belonging to the point that is currently kept fixed in the other strategy set. We then drop that label in the other set, continuing along a **unique** line segment there. Eventually, we must reach an equilibrium.

Finding an equilibrium is crucially based on the observation that the method produces a unique path with the artificial equilibrium and another equilibrium (the one that we find)

as endpoints. The reason has just been indicated: any m labels in X define a unique point, and dropping a particular of these m labels a unique line segment having that point as one of its endpoints. Similarly, n labels in Y define a unique point there, which again has a unique line segment resulting from dropping one of these labels. In that way, the path progresses in a unique manner (alternating between the two sets). Also, this means that the path cannot **fork** because that would mean that it comes back to itself. Namely, this would mean there is a Y-shaped intersection, given by line segments defined by the $m + n - 1$ labels (all labels but the missing one), and at such an intersection one could progress non-uniquely, which is not the case. Finally, there are only finitely many points where we switch on the path because there are only finitely many combinations of m labels in X and of n labels in Y , so the path must terminate.

The preceding argument has often be cast in the story of the *haunted castle*: In the castle there is princess who sits in a room with only one door, the castle has only one door leading to the outside, and every other room in the castle has exactly two doors. A knight comes to rescue the princess, and easily finds her as follows. The knight enters the castle from the outside door and enters a room that either has the princess in it, or it is a room with two doors, in which case he goes through the other door of that room, and continues in that manner. In that way the knight passes through a number of rooms in the castle, but cannot enter a room again, because that room would have needed a third door to enter it. The knight can also not get out of the castle because it has only one door to the outside, which is the one he entered. So the sequence of rooms (the castle has only a finite number of rooms) must end somewhere, which can only be the room with only one door that contains the princess. In this analogy, the room with the princess is the equilibrium, and the outside door, the starting point, is the artificial equilibrium.

6.7 Odd number of Nash equilibria

The analogy in the last section of the castle with the princess as the Nash equilibrium is not quite right because a game may have more than one Nash equilibrium. A better analogy is the following: The castle has only one room with an outside door, and any room has either one or two doors. (This requires a princess for every room that has only one door, see Figure 6.6.) Then by entering from the outside and leaving each encountered room that has two doors through the other door, the knight will eventually find a room with only one door. When starting from a room with a princess in it, following the rooms may either lead back to the outside, or to a room with yet another princess. So the princesses are the endpoints of paths that are defined by sequences of rooms with two doors, which either end with a single-door room with a princess in it, or the outside. Clearly, the number of princesses is odd. In addition, there may be “circuits” of rooms that all have two doors, like at the bottom left in Figure 6.6.

Applied to the Lemke–Howson algorithm, the construction does not depend on starting at the artificial equilibrium which is the completely labelled pair. One could start at any equilibrium of the game itself. Decide on a missing label, for example ① as in the preceding example. This will lead away from that equilibrium along a unique line segment, and then continue in the manner described. Starting from the equilibrium (x^2, y^6)




Figure 6.6 Floorplan of a castle with one outside door and each room having only one or two doors. Each room with only door has a princess in it. There is an odd number of them.

just found given by the pair with labels ② ④ ⑤ and ① ③ ⑥ in X and Y , this just leads backwards along the path that was just computed to the artificial equilibrium. (Each end-points of that path has a unique adjacent line segment, and the intermediate points have two such line segments, one in the forward and one in the backward direction, in different strategy sets.)

Starting from any other equilibrium, however, must, when using the same missing label, lead to yet another equilibrium. In our example, we know by the earlier inspection in (6.5) that there is also the equilibrium (x^4, y^4) with labels ③ ④ ⑤ and ① ② ⑥ in X and Y . With missing label ①, this time dropped from y^4 in Y , leads to the pair (x^4, y^3) with labels ③ ④ ⑤ and ② ⑤ ⑥ with duplicate label ⑤. Moving away from x^4 in X leads to the pair (x^6, y^3) with labels ③ ④ ⑥ and ② ⑤ ⑥ where the newly picked up label ⑥ is duplicate. Dropping ⑥ from y^3 in Y gives the next pair (x^6, y^2) with labels ③ ④ ⑥ and

② ③ ⑤ where ③ is picked up in Y and therefore duplicate. Finally, dropping ③ in X leads to the pair (x^6, y^2) with labels ① ④ ⑥ and ② ③ ⑤ where the missing label ① is found, which is the third equilibrium in (6.5).

That is, any missing label (① in our example) defines pairs of points in the set $X \times Y$ extended by the artificial equilibrium. These are collections of pairs of points and line segments, which together define certain paths. These paths have as endpoints the missing equilibria and the artificial equilibrium, which are all different. Since these endpoints come in pairs, there is an even number of them. Apart from the artificial equilibrium, these are all equilibria of the game, so the game has an odd number of equilibria. This is the claim we announced.

The uniqueness of the Lemke–Howson path (that is, the fact that any room in the castle has only one or two doors) is due to the assumption of **nondegeneracy** of the game as defined in Definition 3.5. According to this definition, a game is nondegenerate if the number of pure best responses to any mixed strategy is never larger than the number of pure strategies that the strategy itself uses with positive probability. In other words, a pure strategy must have at most one pure best response, a strategy that uses only two pure strategies must have at most two pure best responses, and so on. We leave it as an exercise to show that this equivalent to the statement: any mixed strategy in X has at most m labels, and any mixed strategy in Y has at most n labels. Nondegeneracy is a “safe” assumption if a strategic form game is given “generically”, i.e. with all payoffs being independent and allowing for some randomness as to their exact values (within a small range around their assumed value), so that a degeneracy is very unlikely in the first place, or can be removed by perturbing that payoff a little to achieve the necessary independence if this should not be so. Even for degenerate games, however, the geometric description in terms of best response regions is useful and provides insight into the structure of the set of equilibria of the game.

The description we have given applies to 3×3 games, where the strategy sets are triangles and best response regions that subdivide these triangles are easily drawn. However, all these considerations apply to general $m \times n$ games as well, but then cannot easily be visualised any more.

6.8 Exercises for Chapter 6

Exercise 6.1 Consider the following 3×3 game, which is the game (A, B) in (6.1).

		II		
		④	⑤	⑥
I	①	0 0	1 3	-2 0
	②	2 1	0 0	3 1
	③	2 -3	1 4	0 5

- (a) Show how to obtain the top right and bottom picture in Figure 6.2 by comparing rows ① and ③, and rows ② and ③, in analogy to the comparison between rows ① and ② using the inequalities (6.2) and (6.3).
- (b) The endpoints of the lines in the pictures in Figure 6.2 can also be identified quickly by looking at the three 2×3 games that result when omitting one of the three strategies of player I and using the “difference trick”. Explain how.
- (c) Show how to divide X into best response regions, as in (a), to obtain the left picture in Figure 6.4.
- (d) Apply the Lemke–Howson algorithm for **each** initially missing label ①, ②, ③, ④, ⑤, ⑥, by giving the sequence of computed points, starting with $(\mathbf{0}_x, \mathbf{0}_y)$. What are the equilibria that are found for each missing label?
- (e) How are the Nash equilibria and the artificial equilibrium $(\mathbf{0}_x, \mathbf{0}_y)$ paired via the Lemke–Howson paths depending on the missing label? Note: this requires to run the Lemke–Howson algorithm starting from a Nash equilibrium, not only from the artificial equilibrium.

Exercise 6.2 Consider the 3×3 game (A, A^\top) with A as in (6.1).

- (a) Draw the subdivision of X and Y into labelled best response regions. Note: You do not have to do any computation but you can use one of the pictures in Figure 6.4, only changing some labels. Explain.
- (b) In this game, identify all Nash equilibria with the help of the best response regions in (a).
- (c) Apply the Lemke–Howson algorithm for **each** initially missing label ①, ②, ③, ④, ⑤, ⑥, by giving the sequence of computed points, starting with $(\mathbf{0}_x, \mathbf{0}_y)$. What are the equilibria that are found for each missing label? Note: you can save half of the work by using the symmetry of the game. Explain how.

Exercise 6.3 Consider the degenerate 2×3 game in Figure 3.10, given by

		II		
		l	c	r
I	T	2 0	1 3	0 2
	B	0 1	1 2	2 0

Construct the best response regions for both mixed strategy sets (note that X is only a line segment) and identify the sets of all Nash equilibria with the help of the labels.

Chapter 7

Linear programming and zero-sum games

This chapter is another extension of the lecture notes “Game Theory Basics” on topics that have been first taught in autumn 2010. It comprises a second week of lectures that partially replace the old chapter 5 on bargaining, and follows on from Section 3.14.

The topic of this chapter is the strong connection between linear programming and finding equilibria of zero-sum games. Linear programming is the science of solving linear optimization problems, also called linear programs. It can be applied to zero-sum games, as we explain in this chapter. The central theorem of linear programming is the duality theorem (Theorem 7.3). In Section 7.3, we first define linear programs in the standard inequality form. Finding bounds for the objective function motivates the construction of the dual linear program, as we show with an example. We then state the weak and strong duality theorems.

In Section 7.4, we show how the strong duality theorem implies the minimax theorem for zero-sum games.

Section 7.5 is concerned with linear programs in general form where variables may be nonnegative or unconstrained, and the linear constraints may be inequalities or equations. The corresponding duality theorem is a very useful tool, and it is worth acquiring practice in using it. We will give several exercises that show its use.

As non-examinable material, we also give a proof of the strong duality theorem, with the help of Farkas’s Lemma, also known as the theorem of the separating hyperplane. This is done in Section 7.6.

7.1 Learning objectives

After studying this chapter, you should be able to:

- Write down a linear program (LP) in standard inequality form (7.1) and its corresponding Tucker diagram (7.4) and understand how to read this diagram to obtain the dual LP (7.3).

- State the weak and strong duality theorems.
- Do the same for an LP in general form.
- Apply LP duality as in the exercises.

7.2 Further reading

A good introductory text for linear programming is:

- V. Chvátal (1983), *Linear Programming*. Freeman, New York.

Shorter, classical texts that give mathematical definitions more directly are:

- G. B. Dantzig (1963), *Linear Programming and Extensions*. Princeton University Press, Princeton.
- D. Gale (1960), *The Theory of Linear Economic Models*. McGraw-Hill, New York.

7.3 Linear programs and duality

We use the following notation. For positive integers m, n , the set of $m \times n$ matrices is denoted by $\mathbb{R}^{m \times n}$. An m -vector is an element of \mathbb{R}^m . Unless stated otherwise, all vectors are column vectors, so a vector x in \mathbb{R}^m is considered as an $m \times 1$ matrix. Its transpose x^\top is the corresponding row vector in $\mathbb{R}^{1 \times m}$. The vectors $\mathbf{0}$ and $\mathbf{1}$ have all components equal to zero and one, respectively, and have suitable dimension, which may vary with each use of $\mathbf{0}$ or $\mathbf{1}$. An inequality like $x \geq \mathbf{0}$ holds for all components. The identity matrix, of any dimension, is denoted by I .

A linear optimization problem or *linear program* (LP) says: optimize (maximize or minimize) a linear objective function subject to linear constraints (inequalities or equalities).

The standard *inequality* form of an LP is given by an $m \times n$ matrix A , an m -vector b and an n -vector c and says:

$$\begin{aligned} & \text{maximize } c^\top x \\ & \text{subject to } Ax \leq b, \\ & \quad x \geq \mathbf{0}. \end{aligned} \tag{7.1}$$

Thereby, x is the n -vector $(x_1, \dots, x_n)^\top$ of variables. The vector c of coefficients determines the *objective function*. The matrix A and the vector b determine the m linear *constraints*, which are here only inequalities. Furthermore, all variables x_1, \dots, x_n are constrained to be nonnegative; this is stated as $x \geq \mathbf{0}$ separately from $Ax \leq b$ because nonnegative variables are a standard case.

The LP (7.1) states maximization subject to “upper bounds” (inequalities “ \leq ”). A way to remember this is to assume that the components of b and A are positive (they do not have to be), so that these impose actual constraints on x . Clearly, if A has only positive entries, then $Ax \leq b$ can only be fulfilled for nonnegative x if $b \geq \mathbf{0}$.

In general, the LP (7.1) is called *feasible* if $Ax \leq b$ and $x \geq \mathbf{0}$ hold for some x , otherwise *infeasible*. If $c^\top x$ can be arbitrarily large for suitable x subject to these constraints, then the LP is called *unbounded*, otherwise *bounded*. The LP has an optimal solution only if it is feasible and bounded.

Example 7.1 Consider (7.1) with

$$A = \begin{bmatrix} 3 & 4 & 2 \\ 1 & 1 & 1 \end{bmatrix}, \quad b = \begin{bmatrix} 7 \\ 2 \end{bmatrix}, \quad c^\top = [8 \quad 10 \quad 5],$$

which can be stated explicitly as: for $x_1, x_2, x_3 \geq 0$ subject to

$$\begin{array}{rcl} 3x_1 + 4x_2 + 2x_3 & \leq & 7 \\ x_1 + x_2 + x_3 & \leq & 2 \\ \hline \text{maximize} & 8x_1 + 10x_2 + 5x_3 & \end{array} \quad (7.2)$$

The bar is commonly written to separate the objective function from the constraints.

One feasible solution to (7.2) is $x_1 = 0, x_2 = 1, x_3 = 1$ with objective function value 15. Another is $x_1 = 1, x_2 = 1, x_3 = 0$ with objective function value 18, which is better. (We often choose integers in our examples, but coefficients and variables are allowed to assume any real values.) How do we know when we have an optimal value?

The *dual* of an LP can be motivated by finding an upper bound to the objective function of the given LP (which is called the *primal* LP). The dual LP results by reading the constraint matrix vertically rather than horizontally, exchanging the roles of objective function and right hand side, as follows.

In the example (7.2), we multiply each of the two inequalities by some nonnegative number, for example the first inequality by $y_1 = 1$ and the second by $y_2 = 6$, and add the inequalities up, which yields

$$(3 + 6)x_1 + (4 + 6)x_2 + (2 + 6)x_3 \leq 7 + 6 \cdot 2$$

or

$$9x_1 + 10x_2 + 8x_3 \leq 19.$$

In this inequality, which holds for any feasible solution, all coefficients of the nonnegative variables x_j are at least as large as in the primal objective function, so the right hand side 19 is certainly an upper bound for this objective function. In fact, we can obtain an even better bound by multiplying the two primal inequalities by $y_1 = 2$ and $y_2 = 2$, getting

$$(3 \cdot 2 + 2)x_1 + (4 \cdot 2 + 2)x_2 + (2 \cdot 2 + 2)x_3 \leq 2 \cdot 7 + 2 \cdot 2$$

or

$$8x_1 + 10x_2 + 6x_3 \leq 18.$$

Again, all coefficients are at least as large as in the primal objective function. Thus, it cannot be larger than 18, which was achieved by the above solution $x_1 = 1, x_2 = 1, x_3 = 0$, which is therefore optimal.

In general, the dual LP for the primal LP (7.1) is obtained as follows:

- Multiply each primal inequality by some nonnegative number y_i (such as not to reverse the inequality).
- Add each of the n columns and require that the resulting coefficient of x_j for $j = 1, \dots, n$ is at least as large as the coefficient c_j of the objective function. (Because $x_j \geq 0$, this will at most increase the objective function.)
- *Minimize* the resulting right hand side $y_1 b_1 + \dots + y_m b_m$ (because it is an upper bound for the primal objective function).

So the dual of (7.1) says:

$$\begin{aligned} & \text{minimize } y^\top b \\ & \text{subject to } y^\top A \geq c^\top, \quad y \geq \mathbf{0}. \end{aligned} \tag{7.3}$$

Clearly, (7.3) is also an LP in standard inequality form because it can be written as: maximize $-b^\top y$ subject to $-A^\top y \leq -c$, $y \geq \mathbf{0}$. In that way, it is easy to see that the dual LP of the dual LP (7.3) is again the primal LP (7.1).

Both primal and dual LP are defined by the same data A, b, c . A good way to picture this is the following “Tucker diagram”:

$$\begin{array}{ccc} x \geq \mathbf{0} & & \\ & \boxed{A} & \leq \boxed{b} \\ y \geq \mathbf{0} & & \\ & \vee & \hookrightarrow \min \\ & \boxed{c^\top} & \rightarrow \max \end{array} \tag{7.4}$$

The diagram (7.4) shows the $m \times n$ matrix A with the m -vector b on the right and the row vector c^\top at the bottom. The top shows the primal variables x with their constraints $x \geq \mathbf{0}$, and the left-hand side the dual variables y with their constraints $y \geq \mathbf{0}$. The primal LP is to be read horizontally, with constraints $Ax \leq b$, and the objective function $c^\top x$ that is to be maximized. The dual LP is to be read vertically, with constraints $y^\top A \geq c^\top$ (where in the diagram (7.4) \geq is written vertically as \vee), and the objective function $y^\top b$ that is to be minimized. A way to remember the direction of the inequalities is to see that one inequality $Ax \leq b$ points “towards” A and the other, $y^\top A \geq c^\top$, “away from” A , where maximization is subject to upper bounds and minimization subject to lower bounds, apart from the nonnegativity constraints for x and y .

The fact that the primal and dual objective functions are mutual bounds is known as the “weak duality” theorem, which is very easy to prove – essentially in the way we have motivated the dual LP above.

Theorem 7.2 (Weak LP duality) *For a pair x, y of feasible solutions of the primal LP (7.1) and its dual LP (7.3), the objective functions are mutual bounds:*

$$c^\top x \leq y^\top b.$$

If thereby $c^\top x = y^\top b$ (equality holds), then these two solutions are optimal for both LPs.

Proof. In general, if u, v, w are vectors of the same dimension, then

$$u \geq \mathbf{0}, \quad v \leq w \quad \Rightarrow \quad u^\top v \leq u^\top w, \quad (7.5)$$

because $v \leq w$ is equivalent to $(w - v) \geq \mathbf{0}$ which with $u \geq \mathbf{0}$ implies $u^\top (w - v) \geq 0$ and hence $u^\top v \leq u^\top w$; note that this is an inequality between scalars which can also be written as $v^\top u \leq w^\top u$.

Feasibility of x for (7.1) and of y for (7.3) means $Ax \leq b, x \geq \mathbf{0}, y^\top A \geq c^\top, y \geq \mathbf{0}$. Using (7.5), this implies

$$c^\top x \leq (y^\top A)x = y^\top (Ax) \leq y^\top b$$

as claimed.

If $c^\top x^* = (y^*)^\top b$ for some primal feasible x^* and dual feasible y^* , then $c^\top x \leq (y^*)^\top b = c^\top x^*$ for any primal feasible x , and $y^\top b \geq c^\top x^* = (y^*)^\top b$ for any dual feasible y , so equality of the objective functions implies optimality. \square

The following “strong duality” theorem is the central theorem of linear programming.

Theorem 7.3 (Strong LP duality) *Whenever both the primal LP (7.1) and its dual LP (7.3) are feasible, they have optimal solutions with equal value of their objective functions.*

We will prove this theorem in Section 7.6. Its proof is not trivial. In fact, many theorems in economics have a hidden LP duality so that they can be proved by writing down a suitable LP and interpreting its dual LP. For that reason, Theorem 7.3 is extremely useful. A first, classical application is the proof of the minimax theorem.

7.4 The minimax theorem via LP duality

In this section, we prove the minimax theorem, that is, the existence of an equilibrium for a finite zero-sum game, with the help of Theorem 7.3.

Consider an $m \times n$ zero-sum game given by the game matrix A with payoffs a_{ij} for $1 \leq i \leq m$ and $1 \leq j \leq n$ to the row player (the maximizer). The problem of finding a min-max strategy for the column player (the minimizer) can be stated as follows: Find probabilities q_1, \dots, q_n and a real number u so that

$$\sum_{j=1}^n a_{ij} q_j \leq u \quad (1 \leq i \leq m) \quad (7.6)$$

and so that u is minimal. Then the column player does not have to pay more than u for any pure strategy i of the row player, so u is clearly the min-max cost. The problem of finding q_1, \dots, q_n and u defines already an LP, except that it is not in the standard inequality form because of the equality constraint $q_1 + \dots + q_n = 1$ for the probabilities.

We do obtain an LP with constraints $Ax \leq b$ and $x \geq \mathbf{0}$ with the game matrix A as follows:

- Assume that all payoffs a_{ij} are positive, if necessary by adding a constant to all payoffs.
- Then (7.6) implies $u > 0$ whenever q_1, \dots, q_n are probabilities, so that (7.6) is equivalent to

$$\sum_{j=1}^n a_{ij}(q_j/u) \leq 1 \quad (1 \leq i \leq m). \quad (7.7)$$

- Consider the LP

$$\begin{aligned} & \text{maximize} && \sum_{j=1}^n x_j \\ & \text{subject to} && \sum_{j=1}^n a_{ij}x_j \leq 1 \quad (1 \leq i \leq m) \\ & && x_j \geq 0 \quad (1 \leq j \leq n). \end{aligned} \quad (7.8)$$

Then the minimax theorem can be stated as follows.

Theorem 7.4 *Consider an $m \times n$ zero-sum game with payoffs $a_{ij} > 0$ to the row player. An optimal solution x^* to the LP (7.8) with u^* defined as $u^* = 1/\sum_{j=1}^n x_j^*$ defines a min-max strategy q of the column player by $q_j = x_j^* \cdot u^*$ for $1 \leq j \leq n$. The dual LP of (7.8) is*

$$\begin{aligned} & \text{minimize} && \sum_{i=1}^m y_i \\ & \text{subject to} && \sum_{i=1}^m y_i a_{ij} \geq 1 \quad (1 \leq j \leq n) \\ & && y_i \geq 0 \quad (1 \leq i \leq m). \end{aligned} \quad (7.9)$$

Its optimal solution y^ fulfills $u^* = 1/\sum_{i=1}^m y_i^*$ and defines a max-min strategy p of the row player by $p_i = y_i^* \cdot u^*$ for $1 \leq i \leq m$, so (p, q) is a Nash equilibrium of the game. The value of the game is u^* .*

Proof. The LP (7.8) is feasible with $x = \mathbf{0}$. Its Tucker diagram is

$$\begin{array}{ccc}
 & x \geq \mathbf{0} & \\
 & \boxed{} & \\
 y \geq \mathbf{0} & \boxed{A} & \leq \boxed{\mathbf{1}} \\
 & \vee & \hookrightarrow \min \\
 & \boxed{\mathbf{1}^\top} & \rightarrow \max
 \end{array} \tag{7.10}$$

which shows that (7.9) is indeed its dual LP. This dual LP is also feasible, for example with $y_1 = \max\{1/a_{1j} \mid 1 \leq j \leq n\}$ and $y_2 = \cdots = y_m = 0$. By Theorem 7.3, both LPs have optimal solutions x^* and y^* with equal objective function value $v^* = \sum_{j=1}^n x_j^* = \sum_{i=1}^m y_i^*$ where $v^* > 0$ because $y^* \neq \mathbf{0}$ since $y = \mathbf{0}$ is not feasible for (7.9). Hence, $u^* = 1/v^*$ exists as claimed. With $q = x \cdot u^*$, (7.6) holds with $u = u^*$. Hence, q is a mixed strategy of the column player who has to pay no more than u^* against any pure strategy i (and therefore also against any mixed strategy) of the row player. Similarly, $p = y \cdot u^*$ is a mixed strategy of the row player where (7.9) implies

$$\sum_{i=1}^m p_i a_{ij} \geq u^* \quad (1 \leq j \leq n),$$

that is, against any column j he gets at least payoff u^* . Hence the max-min-value and min-max-value of the game are both equal to u^* , which is the value of the game, and (p, q) is a Nash equilibrium. \square

The Tucker diagram (7.10) shows how the zero-sum game matrix A defines an LP where the column player chooses the primal variables x and the row player chooses the dual variables y . Optimal primal and dual solutions x^*, y^* are nonnegative and nonzero and can therefore be normalized, by multiplication with a constant u^* (the inverse of the optimal value of both objective functions), so that x^*, y^* becomes a pair of mixed strategies q, p , where q is a min-max strategy and p is a max-min strategy. In (7.10), the vectors $\mathbf{1}$ and $\mathbf{1}^\top$ define the primal and dual “right-hand side”. After normalizing the mixed strategies, each component 1 of these vectors becomes the min-max cost u^* in each row and max-min payoff u^* in each column. Hence, these 1’s in (7.10) represent this cost and payoff before normalization. It is important that A has only positive entries in order to guarantee that the game has a positive value, that is, $u^* > 0$, so that the normalization is possible.

7.5 General LP duality

An LP in standard inequality form is not always the most convenient way to use the duality theorem, which holds in a more general setting. In this section, we consider the general form of an LP that allows for linear equalities as constraints (in addition to linear inequalities), and for variables that are unrestricted in sign (in addition to the standard

case of nonnegative variables). These cases are closely related with respect to the duality property. As we will see, a primal equality constraint corresponds to a dual variable that is unrestricted in sign, and a primal variable that is unrestricted in sign gives rise to a dual constraint that is an equality. The other case, which we have already seen, is a primal inequality that corresponds to a dual variable that is nonnegative, or a primal nonnegative variable where the corresponding dual constraint is an inequality.

In the standard LP in inequality form (7.1), the constraint $x \geq \mathbf{0}$ looks somewhat redundant since the nonnegativity of x can easily be made part of $Ax \leq b$ by including in those inequalities n additional inequalities $-x_j \leq 0$ for $1 \leq j \leq n$. By omitting nonnegativity of x , one obtains an LP with unconstrained variables x , given by

$$\begin{aligned} & \text{maximize } c^\top x \\ & \text{subject to } Ax \leq b. \end{aligned} \tag{7.11}$$

With (7.11) as the primal LP, we can again multiply any of the inequalities in $Ax \leq b$ with a variable y_i , with the aim of finding an upper bound to the primal objective function $c^\top x$. The inequality is preserved when y_i is nonnegative, but in order to obtain an upper bound we have to require that $y^\top A = c^\top$ because the sign of any component of x is not known. That, is the dual to (7.11) is

$$\begin{aligned} & \text{minimize } y^\top b \\ & \text{subject to } y^\top A = c^\top, \quad y \geq \mathbf{0}. \end{aligned} \tag{7.12}$$

Weak duality applied to the pair of primal and dual LPs (7.11) and (7.12) states that for feasible solutions x and y the corresponding objective functions are mutual bounds, that is (including proof)

$$c^\top x = y^\top Ax \leq y^\top b. \tag{7.13}$$

Another case is an LP with nonnegative variables and equality constraints, which is often called an LP in *equality* form:

$$\begin{aligned} & \text{maximize } c^\top x \\ & \text{subject to } Ax = b, \\ & \quad x \geq \mathbf{0}. \end{aligned} \tag{7.14}$$

In the corresponding dual LP, each constraint in $Ax = b$ is again multiplied with a dual variable y_i . However, because the constraint is an equality, the variable y_i can be unrestricted in sign. The dual constraints are inequalities because the primal objective function has nonnegative variables x_j . That is, the dual LP to (7.14) is

$$\begin{aligned} & \text{minimize } y^\top b \\ & \text{subject to } y^\top A \geq c^\top. \end{aligned} \tag{7.15}$$

The weak duality theorem uses that for feasible solutions x and y to (7.14) and (7.15) we have $x \geq \mathbf{0}$ and $y^\top A - c^\top \geq \mathbf{0}$ and thus

$$c^\top x \leq y^\top Ax = y^\top b. \tag{7.16}$$

These two cases, of a primal LP (7.11) with unrestricted variables and a corresponding dual LP (7.12) with equality constraints, and of a primal LP (7.14) with nonnegative variables and (7.15) with equality constraints, and of a corresponding dual LP with unrestricted variables, show that as before the dual of the dual is the primal.

Any LP in inequality form (7.1) can be converted to equality form by introducing a *slack variable* z_i for each constraint:

$$\begin{aligned} & \text{maximize } c^\top x \\ & \text{subject to } Ax + z = b, \\ & \quad x \geq 0 \\ & \quad z \geq 0. \end{aligned} \tag{7.17}$$

This amounts to extending the constraint matrix A to the right by an identity matrix and by adding coefficients 0 in the objective function for the slack variables.

Note: Converting the inequality form (7.1) to equality form (7.17) defines a new dual LP with unrestricted variables y_1, \dots, y_m , but the former inequalities $y_i \geq 0$ reappear now explicitly via the identity matrix and objective function zeros introduced with the slack variables, as shown in the following Tucker diagram.

$$\begin{array}{c} \begin{array}{cc} x \geq 0 & z \geq 0 \end{array} \\ \begin{array}{|c|c|} \hline \begin{array}{c} y \in \mathbb{R} \\ \\ \\ \\ \end{array} & \begin{array}{c} 1 \\ 1 \\ \ddots \\ 1 \end{array} \\ \hline \end{array} \begin{array}{c} A \\ \\ \\ \\ \end{array} = \begin{array}{|c|} \hline b \\ \hline \end{array} \\ \begin{array}{cc} \vee & \vee \end{array} \\ \begin{array}{|c|c|} \hline c^\top & 0 \ 0 \ \dots \ 0 \\ \hline \end{array} \end{array}$$

An LP in *general form* has inequalities and equalities as constraints, as well as nonnegative and *unrestricted* variables. In the dual LP, the inequalities correspond to nonnegative variables and the equalities correspond to unrestricted variables and *vice versa*. The full definition of an LP in general form is as follows. Let M and N be finite sets (whose elements denote rows and columns, respectively), $I \subseteq M$, $J \subseteq N$, $A \in \mathbb{R}^{M \times N}$ (that is, A is a matrix with entries a_{ij} for $i \in M$ and $j \in N$), $b \in \mathbb{R}^M$, $c \in \mathbb{R}^N$. We first draw the Tucker diagram, shown in Figure 7.1. The big boxes contain the respective parts of the constraint matrix A , the vertical boxes on the right the parts of the right hand side b , and the horizontal box at the bottom the parts of the primal objective function c^\top .

Here the indices in I denote primal inequalities and corresponding nonnegative dual variables, whereas those in $M - I$ denote primal equality constraints and corresponding unconstrained dual variables. Analogously, the indices in J denote primal variables that are nonnegative and corresponding dual inequalities, whereas the indices in $N - J$ denote unconstrained primal variables and corresponding unconstrained dual equality constraints. Figure 7.1 is drawn as if the elements of I and J are listed first in the sets M and N , respectively, but the order of the indices in M and N does not matter.




Figure 7.1 Tucker diagram for an LP in general form.

The feasible sets for primal and dual LP are defined as follows. Consider the following set defined by linear inequalities and equalities

$$\begin{aligned}
 P = \{x \in \mathbb{R}^N \mid & \sum_{j \in N} a_{ij}x_j \leq b_i, \quad i \in I, \\
 & \sum_{j \in N} a_{ij}x_j = b_i, \quad i \in M - I, \\
 & x_j \geq 0, \quad j \in J\}.
 \end{aligned} \tag{7.18}$$

Any x belonging to P is called *primal feasible*. The *primal LP* is the problem

$$\text{maximize } c^\top x \text{ subject to } x \in P. \tag{7.19}$$

(This results when reading the Tucker diagram in Figure 7.1 horizontally.) The corresponding *dual LP* has the feasible set

$$\begin{aligned}
 D = \{y \in \mathbb{R}^M \mid & \sum_{i \in M} y_i a_{ij} \geq c_j, \quad j \in J, \\
 & \sum_{i \in M} y_i a_{ij} = c_j, \quad j \in N - J, \\
 & y_i \geq 0, \quad i \in I\}
 \end{aligned} \tag{7.20}$$

and is the problem

$$\text{minimize } y^\top b \text{ subject to } y \in D. \tag{7.21}$$

(This results when reading the Tucker diagram in Figure 7.1 vertically.) By reversing signs, one can verify that the dual of the dual LP is again the primal.

Then the duality theorem of linear programming states (a) that for any primal and dual feasible solutions, the corresponding objective functions are mutual bounds, and (b) if the primal and the dual LP both have feasible solutions, then they have optimal solutions with the same value of their objective functions.

Theorem 7.5 (General LP duality) Consider the primal-dual pair of LPs (7.19), (7.21). Then

- (a) (Weak duality.) $c^\top x \leq y^\top b$ for all $x \in P$ and $y \in D$.
- (b) (Strong duality.) If $P \neq \emptyset$ and $D \neq \emptyset$ then $c^\top x = y^\top b$ for some $x \in P$ and $y \in D$.

Theorem 7.5 can be proved with the help of Theorem 7.3 by representing an equation as two inequalities in opposite directions, and an unrestricted variable as the difference of two nonnegative variables. The details are omitted for the moment.

7.6 The Lemma of Farkas and proof of strong LP duality*

In this section, we state and prove Farkas's Lemma, also known as the theorem of the separating hyperplane, and use it to prove Theorem 7.3.

The Lemma of Farkas is concerned with the question of finding a nonnegative solution x to a system $Ax = b$ of linear equations. If $A = [A_1 \cdots A_n]$, this means that b is a nonnegative linear combination $A_1x_1 + \cdots + A_nx_n$ of the columns of A . We first observe that such linear combinations can always be obtained by taking suitable linearly independent columns of A .

Lemma 7.6 Let $A = [A_1 \cdots A_n] \in \mathbb{R}^{m \times n}$, let $v \in \mathbb{R}^m$, and let

$$C = \{Ax \mid x \in \mathbb{R}^n, x \geq 0\}. \quad (7.22)$$

Then if $v \in C$, there is a set $J \subseteq \{1, \dots, n\}$ so that the vectors A_j for $j \in J$ are linearly independent, and unique positive reals u_j for $j \in J$ so that

$$\sum_{j \in J} A_j u_j = v. \quad (7.23)$$

Proof. Because $v \in C$, we have $Au = v$ for some $u \geq 0$, so that (7.23) holds with $J = \{j \mid u_j > 0\}$. If the vectors A_j for $j \in J$ are linearly independent (in which case we simply call J independent), we are done. Otherwise, we change the coefficients u_j by keeping them nonnegative but so that at least one of them becomes zero, which gives a smaller set J .

Suppose J is not independent, that is, there are scalars x_j for $j \in J$, not all zero, so that

$$\sum_{j \in J} A_j x_j = 0,$$

where we can assume that the set $S = \{j \in J \mid x_j > 0\}$ is not empty (otherwise replace x by $-x$). Then

$$\sum_{j \in J} A_j (u_j - \lambda x_j) = v$$

for any λ . We choose the largest λ so that $u_j - \lambda x_j \geq 0$ for all $j \in J$. If $x_j \leq 0$ this imposes no constraint on λ , but for $x_j > 0$ (that is, $j \in S$) this means $u_j/x_j \geq \lambda$, so the largest λ fulfilling all these constraints is given by

$$\lambda = \min \left\{ \frac{u_j}{x_j} \mid j \in S \right\} =: \frac{u_{j^*}}{x_{j^*}} \quad (7.24)$$

which implies $u_{j^*} - \lambda x_{j^*} = 0$ so that we can remove any j^* that achieves the minimum in (7.24) from J . By replacing u_j with $u_j - \lambda x_j$, we thus obtain a smaller set J to represent v as in (7.23). By continuing in this manner, we eventually obtain an independent set J as claimed (if $v = \mathbf{0}$, then J is the empty set). Because the vectors A_j for $j \in J$ are linearly independent, the scalars u_j for $j \in J$ are unique. \square

The set C in (7.22) of nonnegative linear combinations of the column vectors of A is also called the *cone* generated by these vectors. Figure 7.2(a) gives an example with $A \in \mathbb{R}^{2 \times 4}$ and a vector $b \in \mathbb{R}^2$ that is *not* in the cone C generated by A_1, A_2, A_3, A_4 .




Figure 7.2 (a) vectors A_1, A_2, A_3, A_4 , the cone C generated by them, and a vector b not in C ; (b) a separating hyperplane H for b with normal vector $y = c - b$.

For the same vectors, Figure 7.2(b) shows a vector y with the following properties:

- $y^\top A_j \geq 0$ for all j , for $1 \leq j \leq n$, and
- $y^\top b < 0$.

The set $H = \{z \in \mathbb{R}^m \mid y^\top z = 0\}$ is called a *separating hyperplane* with normal vector y because all vectors A_j are on one side of H (they fulfill $y^\top A_j \geq 0$, which includes the case $y^\top A_j = 0$ where A_j belongs to H , like A_2 in Figure 7.2(b)), whereas b is strictly on the other side of H because $y^\top b < 0$. Farkas's Lemma asserts that such a separating hyperplane exists for any b that does not belong to C .

Lemma 7.7 (Farkas) *Let $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$. Then either*

- (a) $\exists x \in \mathbb{R}^n : x \geq 0, Ax = b$, or
- (b) $\exists y \in \mathbb{R}^m : y^\top A \geq 0, y^\top b < 0$.

In Lemma 7.7, it is clear that (a) and (b) cannot both hold because if (a) holds, then $y^\top A \geq 0$ implies $y^\top b = y^\top (Ax) = (y^\top A)x \geq 0$.

If (a) is false, that is, b does not belong to the cone C in (7.22), then y can be constructed by the following intuitive geometric argument: Take the vector c in C that is closest to b (see Figure 7.2(b)), and let $y = c - b$. We will show that y fulfills the conditions in (b).

Apart from this geometric argument, an important part of the proof is to show that the cone C is topologically *closed*, that is, it contains any point nearby. Otherwise, b could be a point near C but not in C which would mean that the distance $\|c - b\|$ for any c in C can be arbitrarily small, where $\|z\|$ denotes the Euclidean norm, $\|z\| = \sqrt{z^\top z}$. In that case, one could not define y as described. We first show that C is closed as a separate property.

Lemma 7.8 *For an $m \times n$ matrix $A = [A_1 \cdots A_n]$, the cone C in (7.22) is a closed set.*

Proof. Let z be a point in \mathbb{R}^m near C , that is, for all $\varepsilon > 0$ there is a v in C so that $\|v - z\| < \varepsilon$. Consider a sequence $v^{(k)}$ (for $k = 1, 2, \dots$) of elements of C that converges to z . By Lemma 7.6, for each k there exist a subset $J^{(k)}$ of $\{1, \dots, n\}$ and unique positive real numbers $u_j^{(k)}$ for $j \in J^{(k)}$ so that the columns A_j for $j \in J^{(k)}$ are linearly independent and

$$v^{(k)} = \sum_{j \in J^{(k)}} A_j u_j^{(k)}.$$

There are only finitely many different sets $J^{(k)}$, so there is a set J that appears infinitely often among them; we consider the subsequence of the vectors $v^{(k)}$ that use this set, that is,

$$v^{(k)} = \sum_{j \in J} A_j u_j^{(k)} = A_J u_J^{(k)} \quad (7.25)$$

where A_J is the matrix with columns A_j for $j \in J$ and $u_J^{(k)}$ is the vector with components $u_j^{(k)}$ for $j \in J$. Now, $u_J^{(k)}$ in (7.25) is a continuous function of $v^{(k)}$: In order to see this, consider a set I of $|J|$ linearly independent rows of A_J , let A_{IJ} be the square submatrix of A_J with these rows and let $v_I^{(k)}$ be the subvector of $v^{(k)}$ with these rows, so that $u_J^{(k)} = A_{IJ}^{-1} v_I^{(k)}$ in (7.25). Hence, as $v^{(k)}$ converges to z , the $|J|$ -vector $u_J^{(k)}$ converges to some u_J^* with $z = A_J u_J^*$, where $u_J^{(k)} > 0$ implies $u_J^* \geq 0$, which shows that $z \in C$. \square

Remark 7.9 In Lemma 7.8, it is important that C is the cone generated by a finite set A_1, \dots, A_n of vectors. The cone generated from an infinite set may not be closed. For example, let C be the set of nonnegative linear combinations of the (row) vectors $(n, 1)$ in \mathbb{R}^2 , for $n = 0, 1, 2, \dots$. Then $(1, 0)$ is a vector near C that does not belong to C .

\Rightarrow Prove Remark 7.9, by giving an exact description of C .

\Rightarrow Show that Farkas's Lemma 7.10 below implies that the cone C in (7.22) is an intersection of half-spaces, that is, of sets of the form $\{v \in \mathbb{R}^m \mid a^\top v \leq a_0\}$ for some $a \in \mathbb{R}^m$ and $a_0 \in \mathbb{R}$, which implies that C is closed. However, note that the proof of the lemma uses that C is closed, so this has to be shown first.

Proof of Lemma 7.7. Let $A = [A_1 \cdots A_n]$ and C as in (7.22). Assume that (a) is false, that is, $b \notin C$; we have to prove that (b) holds. Choose $c \in C$ so that $\|c - b\|$ is minimal. This minimum exists because C is closed by Lemma 7.8. Then $\|c - b\| > 0$ because $b \notin C$, and thus $\|c - b\|^2 = (c - b)^\top (c - b) > 0$, that is,

$$(c - b)^\top c > (c - b)^\top b. \quad (7.26)$$

Let $y = c - b$. We show that for all d in C

$$(c - b)^\top d \geq (c - b)^\top c. \quad (7.27)$$

If (7.27) was false, then

$$(c - b)^\top (d - c) < 0 \quad (7.28)$$

for some d in C . Consider the convex combination of d and c given by $d_\varepsilon = c + \varepsilon(d - c) = (1 - \varepsilon)c + \varepsilon d$ which belongs to C , which is clearly a convex set. Then

$$\|d_\varepsilon - b\|^2 = \|c - b\|^2 + \varepsilon \left(2(c - b)^\top (d - c) + \varepsilon \|d - c\|^2 \right)$$

which by (7.28) is less than $\|c - b\|^2$ for sufficiently small positive ε , which contradicts the minimality of $\|c - b\|^2$ for $c \in C$. So (7.27) holds.

In particular, for $d = A_j + c$ we have $y^\top A_j \geq 0$, for any $1 \leq j \leq n$, that is, $y^\top A \geq \mathbf{0}$.

Equations (7.27) and (7.26) imply $(c - b)^\top d > (c - b)^\top b$ for all $d \in C$. For $d = \mathbf{0}$ this shows $0 > y^\top b$. \square

Lemma 7.7 is concerned with finding nonnegative solutions x to a system of equations $Ax = b$. We will use a closely related “inequality form” of this lemma that concerns nonnegative solutions x to a system of inequalities $Ax \leq b$.

Lemma 7.10 (Farkas with inequalities) *Let $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$. Then either*

- (a) $\exists x \in \mathbb{R}^n : x \geq \mathbf{0}, Ax \leq b$, or
- (b) $\exists y \in \mathbb{R}^m : y \geq \mathbf{0}, y^\top A \geq \mathbf{0}, y^\top b < 0$.

Proof. Clearly, there is a vector x so that $Ax \leq b$ and $x \geq \mathbf{0}$ if and only if there are $x \in \mathbb{R}^n$ and $s \in \mathbb{R}^m$ with

$$Ax + s = b, \quad x \geq \mathbf{0}, \quad s \geq \mathbf{0}. \quad (7.29)$$

The system (7.29) is a system of equations as in Lemma 7.7 with the matrix $[A \ I]$, where I is the $m \times m$ identity matrix, instead of A , and vector $(x, s)^\top$ instead of x . The condition $y^\top [A \ I] \geq \mathbf{0}$ in Lemma 7.7(b) is then simply $y \geq \mathbf{0}, y^\top A \geq \mathbf{0}$ as stated here in (b). \square

The conversion of inequalities $Ax \leq b$ to a system of equations by means of nonnegative “slack variables” as in (7.29) is a standard trick in linear programming that we have seen before.

We can now prove the strong duality theorem.

Proof of Theorem 7.3. We assume that (7.1) and (7.3) are feasible, and want to show that there are feasible x and y so that $c^\top x \geq y^\top b$, which by Theorem 7.2 implies $c^\top x = y^\top b$. Suppose, to the contrary, that there are x, y so that

$$x \geq \mathbf{0}, \quad Ax \leq b, \quad y \geq \mathbf{0}, \quad y^\top A \geq c^\top \quad (7.30)$$

but no solution x, y to the system of inequalities

$$\begin{aligned} -A^\top y &\leq -c \\ Ax &\leq b \\ -c^\top x + b^\top y &\leq 0 \end{aligned} \tag{7.31}$$

and $x, y \geq \mathbf{0}$. By Lemma 7.10, this means there are $u \in \mathbb{R}^m$, $v \in \mathbb{R}^n$, and $t \in \mathbb{R}$ so that

$$u \geq \mathbf{0}, \quad v \geq \mathbf{0}, \quad v^\top A - tc^\top \geq \mathbf{0}, \quad t \geq 0, \quad u^\top (-A^\top) + tb^\top \geq \mathbf{0}, \tag{7.32}$$

but

$$-u^\top c + v^\top b < \mathbf{0}. \tag{7.33}$$

We derive a contradiction as follows: If $t = 0$, this means $v^\top A \geq \mathbf{0}$, $u^\top A^\top \leq \mathbf{0}$, which with (7.31), using (7.30), implies

$$v^\top b \geq v^\top Ax \geq 0 \geq u^\top A^\top y \geq u^\top c$$

in contradiction to (7.33).

If $t > 0$, then u and v are essentially primal and dual feasible solutions that violate weak LP duality, because then by (7.32) $bt \geq Au$ and $v^\top A \geq tc^\top$ and therefore

$$v^\top bt \geq v^\top Au \geq tc^\top u$$

which after division by t gives $v^\top b \geq c^\top u$, again contradicting (7.33).

So if the first two sets of inequalities in (7.31) have a solution $x, y \geq \mathbf{0}$, then there is also a solution that fulfills the last inequality, as claimed by strong LP duality. \square

7.7 Exercises for Chapter 7

Exercise 7.1 As in Section 7.4, consider a zero-sum game with an $m \times n$ matrix A of payoffs to player 1 (the maximizing row player). The minimax theorem was proved with the help of LP duality by converting the constraints in (7.6) to (7.7) and arriving at an LP (7.8). The assumption was that A has positive entries. Instead, use the duality theorem for general LPs by considering directly the inequalities in (7.6) with an unconstrained dual variable u , and the additional equality constraint that q_1, \dots, q_n are probabilities, to prove the minimax theorem. Why do you no longer need any sign restrictions about the entries of the matrix A ?

Exercise 7.2 Consider an $m \times n$ game with payoffs a_{ij} to player I. A mixed strategy p of player I is said to *dominate* a pure strategy i if the expected payoff to player I resulting from p is always larger than the payoff resulting from i no matter which pure strategy j is chosen by player II. (Note that p may be a pure strategy in the sense that this pure strategy is chosen with probability 1 by p .) That is, for all $j = 1, \dots, n$ we have

$$\sum_{k=1}^m p_k a_{kj} > a_{ij}.$$

Moreover, here we can assume $p_i = 0$ because $p_i = 1$ is clearly not possible, and then p_k for $k \neq i$ can be replaced by $p_k / \sum_{l=1, l \neq i}^m p_l$ and p_i set to zero.

- (a) The following is the matrix of payoffs to player I in a 3×2 game:

$$A = \begin{bmatrix} 6 & 0 \\ 0 & 6 \\ 2 & 2 \end{bmatrix}.$$

Show that the bottom row of player I is dominated by a mixed strategy.

- (b) Consider the following statements and argue that (1) implies (2):

- (1) Strategy i of player I is a best response to a suitable mixed strategy q of player II.
- (2) Strategy i of player I is not dominated by a mixed strategy.

- (c) The following is the matrix of payoffs to player I in a 3×2 game:

$$A = \begin{bmatrix} 6 & 0 \\ 0 & 6 \\ 4 & 4 \end{bmatrix}.$$

Show that the bottom row of player I is *not* dominated by a mixed strategy.

- (d) Show, using linear programming duality, that in (b) also statement (2) implies statement (1), so that the two statements are in fact equivalent. Since it takes some practice to do this, here are fairly detailed instructions: Write down condition (1) as a primal LP, with the following components: a linear equation that states that row i of A has expected payoff u against a mixed strategy q of player II; next, inequalities stating that all other rows k have smaller payoff than u , by introducing a variable that denotes the difference – that variable is then maximized in the linear program; and finally, that the variables that describe the mixed strategy q of player II are indeed a probability distribution. Once you have done this and have written these equalities and inequalities in a Tucker diagram, read that diagram vertically to obtain the dual linear program. The dual variables should then give you the answer (a statement about mixed strategies p for player I in this case). All you have to do is to suitably interpret the fact that the dual program has the same optimal objective function value as the primal program.