Non-Cooperative Games

An Introduction to Game Theory – Part I

Lars-Åke Lindahl

bookboon
The eBook company

LARS-ÅKE LINDAHL

NON-COOPERATIVE GAMES

AN INTRODUCTION TO GAME THEORY – PART I

Non-Cooperative Games: An Introduction to Game Theory – Part I
1st edition

ISBN 978-87-403-2133-3
Peer review by Prof. Erik Ekström, Professor in mathematics, Uppsala University

CONTENTS

Preface

Game theory is a mathematical theory that can be used to model and study both conflicts and cooperations. The different models are called games with the participating parties as players. Games are classified as *non-cooperative games* or *cooperative games* depending on whether the focus is on what a player can achieve alone without cooperation with other players or what groups of players can achieve by cooperation.

A non-cooperative game is said to have strategic form if each player makes his choice of action once and for all and independently of the choices of his opponents. A typical example of a game that satisfies this is stone-paper-scissors. Extensive games model situations where players can consider and modify their plans of action as the situation develops. Chess and bridge are two typical examples of such games, but they differ in that chess players have always perfect information about the current position of the game, while a bridge player, who is in turn to bid or play a card, has only incomplete information about the current situation.

In games, there are generally no outcomes that are perceived as the best by all players. A main task in game theory is therefore to identify solutions that in some sense can be regarded as acceptable to all players, and to provide conditions for such solutions to exist. In non-cooperative games, such a solution is provided by the Nash equilibrium which is an outcome with the property that no player benefits from unilaterally changing his choice of action. In cooperative games, where the main problem is to allocate a result in such a way that no player or group of players feel aggrieved, there are a number of proposed solutions, and we will study three of these in volume II, namely the core, the nucleolus and the Shapley solution.

Game theoretical concepts and results are used in several scientific areas, perhaps primarily in economics, biology, and computer science. For example, the competition between companies in a market can be modeled as a game with equilibrium prices corresponding to Nash equilibria. In biology, game theory can be used to analyze evolutionary stable phenomena; the players in an evolutionary game are genes with biologically and hereditary encoded strategies, and evolutionarily stable phenomena, such as sex ratio, can be seen as special Nash equilibria.

Game theory is an unobjectionable mathematical theory, but the use of game theory predictions in a concrete conflict situation requires, of course, that there is a good match between model and reality. The models assume that the players are rational and that they know the rules of the game, which among other things includes that they know how the opponents value all possible outcomes. In complex situations, this is of course impossible, but game theoretical reasoning can often explain afterwards why it happened as it happened.

This book in two volumes is intended to give an introduction to Game Theory, and volume I treats, as indicated by the title, the non-cooperative theory. The presentation is elementary in the sense that no advanced knowledge of Mathematics is required. One semester of university studies of Mathematics should be enough to fully understand the text apart from a few sections whose proofs can be skipped without loss of context. You must master the sum and the product symbols, the function concept and the usual set operations (union, intersection, set membership), and know what is meant by probabilities (on a finite sample space). But above all, you must not be afraid of mathematical reasoning, which includes defining abstract concepts and introducing new symbols. The most advanced reasoning in this volume occurs in Sections 1.5, 2.3, 4.2, 7.1, 7.2 and 9.3.

Volume II on Cooperative Game Theory is essentially independent of this volume I and can therefore very well be read and studied before volume I.

A basic course on game theory could consist of Chapters 1.1–1.4, 2.1–2.2, 2.4–2.5, 3, 5, 6, 8, and 9.1–9.2 in volume I and Chapters 10.1–10.4, 11, and 12.1–12.2 in volume II.

Lars-Åke Lindahl

Notation

The mathematical symbols used in this volume are, with few exceptions, standard notation, and those who are not will be explained as they are introduced. Thus, $\mathbf{R}$ denotes the set of all real numbers, while $\mathbf{R}_+$ is the set of all non-negative real numbers. We write $f\colon X \to \mathbf{R}$ to indicate that the real-valued function f is defined on the set X.

Cartesian product is a central concept; the product $A_1 \times A_2$ of two arbitrary sets A_1 and A_2 consists of all *ordered pairs* (a_1, a_2) that can be formed by choosing the element a_1 from the set A_1 and the element a_2 from the set A_2.

The product of more than two sets is defined analogously. Given n sets $A_1, A_2, \dots, A_n$, $A_1 \times A_2 \times \cdots \times A_n$, or more briefly $\prod_{j=1}^{n} A_j$, denotes the set of all *ordered n-tuples* $(a_1, a_2, \dots, a_n)$ that can be formed by choosing $a_1 \in A_1$, $a_2 \in A_2$, …, $a_n \in A_n$. Ordered n-tuples will also be called *vectors*.

We will use the following convention for naming n-tuples. If the members of an n-tuple are given as $a_1, a_2, \dots, a_n$, we use the same letter a as the name of the n-tuple. We write, for example, $a = (a_1, a_2, \dots, a_n)$, $b = (b_1, b_2, \dots, b_n)$ and $x = (x_1, x_2, \dots, x_n)$. We use the same convention for Cartesian products so that $A = A_1 \times A_2 \times \cdots \times A_n$, etc.

Very often we need to consider the $(n-1)$-tuple that is formed when the element at location i in an n-tuple is deleted. If $a = (a_1, a_2, \dots, a_n)$ we use a_{-i} to denote the $(n-1)$-tuple $(a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$.

Thus, in the case $n = 3$, $a_{-1} = (a_2, a_3)$, $a_{-2} = (a_1, a_3)$, and $a_{-3} = (a_1, a_2)$.

We use an analogous notation for product sets; if $A = A_1 \times A_2 \times \cdots \times A_n$, then A_{-i} denotes the product set with $n-1$ factors that is formed by deleting the ith factor A_i.

Given an n-tuple $a = (a_1, a_2, \dots, a_n)$, we also need a simple way to denote the n-tuple that is obtained by replacing the element a_i at location i with an arbitrary element b. We write (a_{-i}, b) for this n-tuple, which means that

$$(a_{-i}, b) = (a_1, \dots, a_{i-1}, b, a_{i+1}, \dots, a_n).$$

Thus, in the case $n = 3$, $(a_{-1}, b) = (b, a_2, a_3)$, $(a_{-2}, b) = (a_1, b, a_3)$, and $(a_{-3}, b) = (a_1, a_2, b)$. Of course, $(a_{-i}, a_i) = a$ for all n-tuples a.

We use an analogous writing for product sets, so that

$$(A_{-i}, B) = A_1 \times \cdots \times A_{i-1} \times B \times A_{i+1} \times \cdots \times A_n.$$

Chapter 1

Utility Theory

Classical utility theory deals with the analysis of situations where individuals must make decisions and seeks to explain their choices of actions in terms of preferences, utility and expectations.

1.1 Preference relations and utility functions

We meet daily situations where we have to choose between different actions: Blazer or sweater? Tea or coffee? Car or bus to work?, and so on. Most of our choices are made by habit, unknowingly or without reflecting more closely on the matter, but in a new situation or if it is an important decision, we try to choose the best option. A prerequisite for doing so is that we are able to evaluate and rank the different alternatives. It is such decision problems that we will study in this chapter.

A decision problem consists of a set A of possible *actions* and a *preference relation* $\succeq$ on the set A, where $a \succeq b$ should be interpreted as "the action a is better than or just as good as the action b". Instead of $a \succeq b$ we also write $b \preceq a$.

A relation on A has to meet certain minimum requirements in order to work as a preference relation, and they are given in the following definition.

Definition 1.1.1 A *preference relation* $\succeq$ on a set A is a relation satisfying the following two axioms:

Axiom 1 (Completeness) *For all $a, b \in A$, $a \succeq b$ or $b \succeq a$.*

Axiom 2 (Transitivity) *If $a \succeq b$ and $b \succeq c$, then $a \succeq c$.*

Completeness means that all alternatives in A can be compared with each other. In particular, $a \succeq a$ for all $a \in A$.

We say that the alternatives a and b are *equivalent* and write $a \sim b$, if $a \succeq b$ and $b \succeq a$. We also say that the decision-maker is *indifferent* to such alternatives. The relation $\sim$ is an equivalence relation on the set A.

We write $a \succ b$ and say that "a is better than b", if $a \succeq b$ and $a \not\sim b$,

Definition 1.1.2 Let $\succeq$ be a preference relation on the set A. An element $a^* \in A$ is called *maximal* (with respect to the preference relation) if $a^* \succeq a$ for all $a \in A$ (and *minimal* if $a \succeq a^*$ for all $a \in A$).

If B is a subset of A, $a^* \in B$ and $a^* \succeq b$ for all $b \in B$, then we say that the element a^* is *maximal in* B with respect to the given preference relation.

If a^* is maximal in B, then there is in other words no better alternative in B.

It is now easy to define what should be meant by a rational decision-maker. A *rational decision-maker* is a person who, knowing his set of available actions and his preferences, chooses the maximum action.

It is easy to give examples of preference relations without maximal elements and of preference relations with more than one maximal element.

If the set A is finite, which will often be the case in the future, there is always a maximal element, because it follows from Axioms 1 and 2 that the elements of A can be arranged in a finite chain $a_1 \succeq a_2 \succeq a_3 \cdots \succeq a_n$, and a_1 is then a maximal element.

Example 1.1.1 A person's preference for the five dishes on a restaurant's menu on a given day looks like this:

$$\textit{Cod} \succ \textit{Entrecote} \sim \textit{Veal} \succ \textit{Salmon} \sim \textit{Pasta}.$$

The unique optimal choice is *Cod*. □

A common way to specify preferences is to put points on the various alternatives; the better alternative, the higher the score. This leads to the concept of utility function.

Definition 1.1.3 Let $\succeq$ be a preference relation on the set A. The function $u\colon A \to \mathbf{R}$ represents the preference relation and is called a corresponding *(ordinal) utility function* if

$$a \succeq b \iff u(a) \geq u(b).$$

Example 1.1.2 Consider the menu in Example 1.1.1. We get a corresponding utility function u by defining

$$u(\textit{Cod}) = 3, \quad u(\textit{Entrecote}) = u(\textit{Veal}) = 2,$$
$$u(\textit{Salmon}) = u(\textit{Pasta}) = 1.$$ □

Can Example 1.1.2 be generalized, i.e. is every preference relation represented by a utility function? The answer is obviously yes for preference relations on finite sets, but for infinite sets additional topological conditions, to be discussed in the next section, are needed.

Proposition 1.1.1 *Every preference relation on a finite set is represented by a utility function.*

Proof. Arrange the elements of the finite set in increasing order

$$a_1 \preceq a_2 \preceq \cdots \preceq a_n,$$

and define the function u inductively by setting $u(a_1) = 1$ and

$$u(a_k) = \begin{cases} u(a_{k-1}) & \text{if } a_k \sim a_{k-1}, \\ u(a_{k-1}) + 1 & \text{if } a_k \succ a_{k-1} \end{cases}$$

for $k = 2, 3, \ldots, n$. The function u obviously represents the preference relation. □

A utility function u is not uniquely determined by its preference relation, because it is only the order between the function values that matters, not the values themselves.

Proposition 1.1.2 *Two utility functions u and v represent the same preference relation if and only if $v = f \circ u$ for some strictly increasing function f (defined on the range of u).*

Proof. Suppose that the preference relation $\succeq$ is represented by the utility function u and that $v = f \circ u$, where the function f is strictly increasing. By monotonicity, we then have

$$u(a) \geq u(b) \;\Leftrightarrow\; v(a) = f(u(a)) \geq f(u(b)) = v(b),$$

which shows that the function v also represents the preference relation $\succeq$.

Conversely, suppose that the functions u and v both represent the preference relation $\succeq$. In particular, we then have

$$u(a) = u(b) \;\Leftrightarrow\; a \sim b \;\Leftrightarrow\; v(a) = v(b),$$

which means that we obtain a function f on the range of u by defining

$$f(u(a)) = v(a)$$

for all $a \in A$. Then $v = f \circ u$ by definition, and it follows from the definition of utility functions that

$$u(a) \geq u(b) \;\Leftrightarrow\; a \succeq b \;\Leftrightarrow\; v(a) \geq v(b) \;\Leftrightarrow\; f(u(a)) \geq f(u(b)).$$

The above equivalences imply that f is strictly increasing. □

A common way of defining a preference relation $\succeq$ on a set A is to start from a function $u\colon A \to \mathbf{R}$ and then define $\succeq$ by the equivalence

$$a \succeq b \;\Leftrightarrow\; u(a) \geq u(b).$$

The relation $\succeq$ thus defined is obviously complete and transitive, that is a preference relation. The preference relation $\succeq$ is said to be *induced* by the function u, and u is by definition a utility function that represents the preference relation.

Example 1.1.3 Most people are greedy in the sense that they prefer more money to less, which means that their preference for wealth, the difference between assets and debts measured in dollars, say, is given by the usual order $\geq$ on the set $\mathbf{R}$ of real numbers, and every strictly increasing function on $\mathbf{R}$ is a corresponding ordinal utility function.

However, the preference relation gives us no clue as to how persons value various changes of the wealth. A salary increase of 1 000 dollars per month is probably valued more by a person with a monthly salary of 10 000 dollars than by a person with 100 000 dollars per month, but in order to describe this we need to assign our utility functions additional properties. The values $u(x)$ of the utility function has to describe the benefit of x in such a way that it makes sense to compare the utility change when x is changed from a to b with the utility change when x is changed from c to d by comparing the differences $u(b) - u(a)$ and $u(d) - u(c)$. Utility functions with this feature are called *cardinal utility functions*.

Let us, as an example, compare the following two conceivable cardinal utility functions: $u(x) = x$ and $v(x) = \ln x$. Persons with the first mentioned utility function value a salary increase of 1 000 dollars equally regardless of the salary before the increase, because $u(a + 1000) - u(a) = 1\,000$ for every a.

For persons with v as utility function, the value of a salary increase is larger at the lower of the two monthly incomes, since $v(11\,000) - v(10\,000) = \ln 1.1 \approx 0.095$ and $v(101\,000) - v(100\,000) = \ln 1.01 \approx 0.010$. □

Exercises

1.1 Show that the indifference relation $\sim$ is an equivalence relation, i.e. that

(i) $a \sim a$ for all $a \in A$
(ii) $a \sim b \Rightarrow b \sim a$
(iii) $a \sim b \ \&\ b \sim c \Rightarrow a \sim c$.

1.2 The *lexicographic order* $\succeq$ on $\mathbf{R}^2$ is defined by $(x_1, x_2) \succeq (y_1, y_2)$ if and only if either $x_1 > y_1$ or $x_1 = y_1$ and $x_2 \geq y_2$. Show that the lexicographic order is a preference relation.

1.3 Define a relation $\geq$ on $\mathbf{R}^2$ by

$$(x_1, x_2) \geq (y_1, y_2) \Leftrightarrow x_1 \geq y_1 \ \&\ x_2 \geq y_2.$$

Is $\geq$ a preference relation?

1.4 Alice, Bella and Clara are football fans, and they are trying to agree on which is the best of the three London clubs Arsenal FC, Chelsea FC and West Ham United FC. This is not easy, because their individual preferences are quite different. Alice ranks the team in the order Arsenal, Chelsea, West Ham, Bella ranks them in the order West Ham, Arsenal, Chelsea, and Clara in the order Chelsea, West Ham, Arsenal. Our three football friends decide, therefore, to determine the preferences of the group by majority decisions. In the vote between Arsenal and Chelsea, Chelsea wins by two votes to one (Alice and Bella voting for Arsenal), so Arsenal $\succ$ Chelsea according to the preferences of the group.

a) Which of the teams Arsenal and West Ham and which of the teams Chelsea and West Ham is the best team according to our group of football fans?

b) Is the resulting relation a preference relation?

c) Can the relation be described by a utility function?

1.5 A cardinal utility function u with respect to wealth is called *risk averse* if it is increasing and strictly concave (which holds if $u'(x) > 0$ and $u''(x) < 0$ for all x). The utility function is called *risk neutral* if it is of the form $u(x) = ax + b$ with $a > 0$.

a) Show that a risk averse person (i.e. a person with a risk averse utility function) values the increase in wealth with 1 000 dollars higher if the original wealth is 10 000 dollars than if it is 1 million dollars.

b) How does a risk neutral person react in the same situation?

c) Generalize a) by formulating how an increase with h affects a risk averse person's utility if the original wealth is a and b, respectively.

1.6 What increase in income is required for a person with 30 000 SEK in monthly salary and $u(x) = \ln x$ as cardinal utility function for her to feel the same satisfaction as 1 000 SEK more in monthly salary gave her when her monthly salary was 10 000 SEK?

1.2 Continuous preference relations

In order to be able to say something interesting about preference relations on infinite sets, additional properties are needed in addition to the defining properties completeness and transitivity. Continuity is such a property and is defined as follows for preference relations on subsets of $\mathbf{R}^n$.

Definition 1.2.1 A preference relation $\succeq$ on a subset X of $\mathbf{R}^n$ is called *continuous* if the following condition is fulfilled:

For all convergent sequences $(x_k)_{k=1}^{\infty}$ and $(y_k)_{k=1}^{\infty}$ of points in X such that $x_k \succeq y_k$ for all k and whose limits $x = \lim_{k\to\infty} x_k$ and $y = \lim_{k\to\infty} y_k$ belong to X, we have $x \succeq y$.

A subset of $\mathbf{R}^n$ is closed if and only if it contains the limit of every convergent sequence of points from the set. The following proposition is therefore an immediate consequence of the above continuity definition.

Proposition 1.2.1 *If X is a closed subset of $\mathbf{R}^n$, $\succeq$ is a continuous preference relation on X and a is an arbitrary element of X, then $\{x \in X \mid x \succeq a\}$ is a closed set.*

Proof. If $(x_k)_{k=1}^{\infty}$ is an arbitrary convergent sequence of elements in the set $F = \{x \in X \mid x \succeq a\}$ with limit x, then x belongs to the set X since it is closed. By applying the continuity definition with $y_k = a$ for all k, we then conclude that $x \succeq a$, which means that the limit x lies in F. The set F is thus closed. □

A preference relation, which is induced by a continuous function, is necessarily continuous because of the following proposition.

Proposition 1.2.2 *A preference relation $\succeq$ on X, which is represented by a continuous utility function $u\colon X \to \mathbf{R}$, is continuous.*

Proof. Assume $x_k \succeq y_k$ for all k, where $(x_k)_{k=1}^{\infty}$ and $(y_k)_{k=1}^{\infty}$ are convergent sequences in X with limits x and y, respectively, in X. Since $u(x_k) \geq u(y_k)$ for all k, and u is continuous, we obtain the inequality $u(x) \geq u(y)$ by passing to the limit. Hence $x \succeq y$, and this proves that the preference relation is continuous. □

The converse of Proposition 1.2.2 is not true; a utility function representing a continuous preference relation is not necessarily continuous, because if the preference relation is represented by u, it is also represented by the composition $f \circ u$ for every strictly increasing function f, whether continuous or not. However, the following kind of converse holds.

Proposition 1.2.3 *Suppose X is a connected subset of $\mathbf{R}^n$. Every continuous preference relation on X is then represented by some continuous utility function.*

We will not use Proposition 1.2.3, so we omit its proof which is somewhat complicated.

An important property of continuous preference relations on compact sets is the existence of maximal elements. (Recall that a subset of $\mathbf{R}^n$ is *compact* if it is closed and bounded.)

Proposition 1.2.4 *Every continuous preference relation $\succeq$ on a compact set X has a maximal element.*

Proof. Let $F_a = \{x \in X \mid x \succeq a\}$. The set F_a is a closed subset of X for every $a \in X$, according to Proposition 1.2.1.

Now consider the intersection $\bigcap_{a \in X} F_a$ of all the sets F_a. An element $b \in X$ belongs to this intersection if and only if b belongs to each F_a, i.e. if and only if $b \succeq a$ for all $a \in X$, i.e. if and only if b is a maximal element. Therefore, all we need to prove is that the intersection $\bigcap_{a \in X} F_a$ is nonempty.

We do so by contradiction, assuming that the intersection is empty. According to a standard theorem in Topology, compact sets are characterized by the property that for every family of closed subsets with an empty intersection there is a finite subfamily of sets with empty intersection. In our case, this means that there are finitely many elements $a_1, a_2, \ldots, a_m$ such that $\bigcap_{j=1}^m F_{a_j} = \emptyset$, and we may assume, of course, that the elements are indexed so that $a_1 \succeq a_2 \succeq \cdots \succeq a_m$. But then $a_1 \in F_{a_j}$ for all j, and we conclude that the element a_1 belongs to the intersection $\bigcap_{j=1}^m F_{a_j}$, which therefore has to be nonempty. This is a contradiction and proves that there has to be a maximal element. □

Exercises

1.7 Is the lexicographic order on $\mathbf{R}^2$ (see Exercise 1.2) a continuous preference relation?

1.8 Prove that a preference relation $\succeq$ on X is continuous if and only if the following condition is met: For all convergent sequences $(x_k)_{k=1}^\infty$ in X with limit $x = \lim_{k\to\infty} x_k$ in X and all $a \in X$,

$$x_k \succeq a \text{ for all } k \Rightarrow x \succeq a \quad \text{and} \quad x_k \preceq a \text{ for all } k \Rightarrow x \preceq a.$$

1.3 Lotteries

For a rational person, the decision is easy once the alternatives are identified and ranked. For example, assume that Charlie is allowed to choose one of three objects a, b and c, and that his preference order is $a \succ b \succ c$. Then he chooses a, of course. But what is best for Charlie if the choice is between getting the object b and participating in a lottery with 50% chance of getting the object a and equal chance of getting the object c? A rational choice seems to presume that Charlie has an idea of how good b is compared to a and to c, and also to depend on his attitude towards risk. His preference relation gives no clue about this.

We conclude that it is sometimes necessary to make decisions and choose alternatives that in one way or another depend on chance. No advanced probabilistic arguments will be used in this book, however, so it is enough to be familiar with the most basic probability concepts, and we will only work with finite sample spaces. We will adhere to the usual game theory tradition by calling probability distributions "lotteries".

Definition 1.3.1 Let A be a finite set. A *lottery p over A* is a probability distribution on A, i.e. a function $p\colon A \to [0,1]$ such that $\sum_{a\in A} p(a) = 1$.

The set of all lotteries over A will be denoted by $\mathcal{L}(A)$.

EXAMPLE 1.3.1 Suppose that the set A consists of three elements b, c and n, which stand for the alternatives of getting a bike, a car and nothing, respectively. We obtain a lottery over A by defining $p(b) = 0.005$, $p(c) = 0.0001$, and $p(n) = 0.9949$.

The comparison between receiving a bike for sure and participating in the lottery p, with probability of winning a bike equal to 0.005, can be seen as a comparison between the two lotteries q and p, where $q(b) = 1$ and $q(c) = q(n) = 0$. □

For lotteries with an absolutely certain outcome, like the lottery q in the example above, we use the following notation.

Definition 1.3.2 For each $a \in A$, δ_a is the lottery on A defined

$$\delta_a(x) = \begin{cases} 1 & \text{if } x = a, \\ 0 & \text{if } x \neq a. \end{cases}$$

The event a is the only possible outcome of the lottery δ_a, and for that reason we will often identify the lottery δ_a with the event a, and thus regard the set A as a subset of the lottery set $\mathcal{L}(A)$.

If p and q are two lotteries over A and α and β are two non-negative numbers with sum equal to 1, then the function $\alpha p + \beta q$ is also a lottery over A. The set $\mathcal{L}(A)$ of all lotteries is, in other words, a *convex set*. More generally, any convex sum of lotteries is a lottery, i.e. the sum $\sum_{i=1}^{m} \alpha_i p_i$ is a lottery if $p_1, p_2, \ldots, p_m$ are lotteries and $\alpha_1, \alpha_2, \ldots, \alpha_m$ are non-negative numbers with sum 1.

We can realize the lottery $p = \sum_{i=1}^{m} \alpha_i p_i$ in two steps. Let L be a lottery with m different outcomes, where outcome no. i occurs with probability α_i and is a ticket to the lottery p_i. The lottery p is then the combined effect of executing lottery L and then, depending on the outcome i, proceeding with lottery p_i.

Every lottery p over A is a convex combination of lotteries with absolutely certain outcomes since

$$p = \sum_{a \in A} p(a)\delta_a.$$

A lottery p over $A = \{a_1, a_2, \ldots, a_n\}$ is completely determined by the n-tuple $(p(a_1), p(a_2), \ldots, p(a_n))$. The map $p \mapsto (p(a_1), p(a_2), \ldots, p(a_n))$ is thus a bijection between the set $\mathcal{L}(A)$ of all lotteries over A and the compact, convex subset $\{(x_1, x_2, \ldots, x_n) \mid x_1, x_2, \ldots, x_n \geq 0,\ x_1 + x_2 + \cdots + x_n = 1\}$ of $\mathbf{R}^n$, which for $n = 2$ is the line segment between the points $(1, 0)$ and $(0, 1)$, and for $n = 3$ is the triangle with vertices at $(1, 0, 0)$, $(0, 1, 0)$ and $(0, 0, 1)$.

The vertices of the image set correspond to the certain lotteries δ_{a_i}, that is to the events a_i.

Example 1.3.2 The lottery set $\mathcal{L}(A)$ in Example 1.3.1 can be identified with the triangle

$$\{(x_1, x_2, x_3) \mid x_1 + x_2 + x_3 = 1, x_1, x_2, x_3 \geq 0\},$$

with the given lottery p corresponding to the point $(0.005, 0.0001, 0.9949)$. The vertex $(1, 0, 0)$ corresponds to the lottery δ_b with a bike as a sure win. □

Definition 1.3.3 Let $u\colon A \to \mathbf{R}$ be a function defined on A, and let p be an arbitrary lottery over A. The sum

$$\tilde{u}(p) = \sum_{a\in A} u(a)p(a)$$

is called the *expected value* of u with respect to the lottery p.

By varying p we get a function $p \mapsto \tilde{u}(p)$, which is defined on the lottery set $\mathcal{L}(A)$. The function $\tilde{u}\colon \mathcal{L}(A) \to \mathbf{R}$ is called the *expected utility function* associated with u, when u is a utility function on A.

The function $\tilde{u}\colon \mathcal{L}(A) \to \mathbf{R}$ is *affine* on its domain of definition, that is

$$\tilde{u}(\alpha p + \beta q) = \alpha\tilde{u}(p) + \beta\tilde{u}(q).$$

for all lotteries p and q and all nonnegative real numbers α, β with sum equal to 1. The function is obviously continuous; $\tilde{u}(p)$ is simply a first degree polynomial in the variables $p(a_1), p(a_2), \ldots, p(a_n)$. Writing $x_i = p(a_i)$ and $c_i = u(a_i)$, we have $\tilde{u}(p) = c_1x_1 + c_2x_2 + \cdots + c_nx_n$.

The expected value of u with respect to a lottery δ_a, which results in the outcome a with certainty, is of course equal to $u(a)$. Thus, by identifying the lottery δ_a with a, we get

$$\tilde{u}(a) = \tilde{u}(\delta_a) = u(a).$$

Suppose for example that $u(a)$ is the financial value of the outcome a of the lottery p. By the law of large numbers, the average of the values obtained when the lottery is repeated a large number of times should come close to the expected value $\tilde{u}(p)$.

Example 1.3.3 Consider again the lottery in Example 1.3.1, and assume that the bike is worth 5 100 SEK, the car is worth 200 100 SEK, and the lottery ticket costs 100 SEK. Let the function u give the value of the outcomes reduced by the ticket price so that $u(b) = 5\,000$, $u(c) = 200\,000$, and $u(n) = -100$. The expected value of u with respect to the explicit lottery p of Example 1.3.1 is then

$$\tilde{u}(p) = 5\,000 \cdot 0.005 + 200\,000 \cdot 0.0001 - 100 \cdot 0.9949 = -54.49 \text{ SEK.} \quad \square$$

1.4 Expected utility

Let us return to Charlie, who has to choose between the object b and a lottery p with 50% chance of getting the object a and equal chance of getting the object c. His preference order is $a \succ b \succ c$. With a randomly selected corresponding utility function u, i.e. function with $u(a) > u(b) > u(c)$, a possible way to evaluate the lottery option would be to use the expected value of u with respect to the lottery, i.e. $\tilde{u}(p) = \frac{1}{2}u(a) + \frac{1}{2}u(c)$, and to choose the lottery if $\tilde{u}(p) > u(b)$, and to be indifferent between the two alternatives if $\tilde{u}(p) = u(b)$.

The problem with this approach is that the outcome is not uniquely determined by his preferences but depends on the utility function used to represent them. With $u(a) = 5$, $u(b) = 2$ and $u(c) = 1$, he obtains $\tilde{u}(p) = \frac{1}{2}(5+1) = 3 > u(b)$ with the conclusion that he should choose the lottery. But with $u(a) = 5$, $u(b) = 4$ and $u(c) = 1$, he instead obtains $\tilde{u}(p) = 3 < u(b)$, and now object b is his best choice.

To evaluate lotteries, i.e. outcomes that depend on chance, using expected values of utility functions, it is not enough to use the functions' ordinal properties. We have to use their numerical values, and we say that we have selected a *cardinal utility function* when we have fixed a utility function and in addition to its ordinal properties also use its values to form expected values.

A natural problem in this context is to determine the relationship between cardinal functions on a set A that give rise to the same evaluation of the lotteries over A, i.e. induce the same preference relation on the lottery set $\mathcal{L}(A)$. The simple answer is given by our next proposition.

Proposition 1.4.1 *Let u and v be two functions on a (finite) set A. The two expected utility functions $\tilde{u}$ and $\tilde{v}$ induce the same preference relation $\succeq$ on the lottery set $\mathcal{L}(A)$ if and only if there exist two real constants $C > 0$ and D such that*

$$v(a) = Cu(a) + D \tag{1}$$

for all $a \in A$.

Remark. A function $f\colon \mathbf{R} \to \mathbf{R}$ is called *affine* if $f(x) = Cx+D$, where C and D are constants. The function is *order preserving*, i.e. $x > y \;\Leftrightarrow\; f(x) > f(y)$, if $C > 0$.

Proposition 1.4.1 means, in other words, that $\tilde{u}$ and $\tilde{v}$ induce the same preference relation on $\mathcal{L}(A)$ if and only if $v(x) = f(u(x))$ for some order preserving affine function f.

Proof. First suppose that $v(a) = Cu(a) + D$ with $C > 0$, and denote by $\succeq_u$ and $\succeq_v$ the two preference relations on the lottery set $\mathcal{L}(A)$ induced by $\tilde{u}$ and $\tilde{v}$, respectively. By the definition of expected values,

$$\begin{aligned}\tilde{v}(p) &= \sum_{a\in A} v(a)p(a) = \sum_{a\in A}(Cu(a)+D)p(a)\\ &= C\sum_{a\in A} u(a)p(a) + D\sum_{a\in A} p(a) = C\tilde{u}(p) + D\end{aligned}$$

for all lotteries p, and this implies that

$$\begin{aligned}p \succeq_v q \;&\Leftrightarrow\; \tilde{v}(p) \geq \tilde{v}(q) \;\Leftrightarrow\; C\tilde{u}(p) + D \geq C\tilde{u}(q) + D\\ &\Leftrightarrow\; \tilde{u}(p) \geq \tilde{u}(q) \;\Leftrightarrow\; p \succeq_u q,\end{aligned}$$

which proves that the two preference relations $\succeq_v$ and $\succeq_u$ are identical.

Conversely, assume that $\tilde{u}$ and $\tilde{v}$ induce the same preference relation $\succeq$ on the lottery set. In particular, we then have the equivalences

$$\tilde{u}(p) = \tilde{u}(q) \iff p \sim q \iff \tilde{v}(p) = \tilde{v}(q),$$

which describe indifference.

Let now b and c be two elements of A such that

$$u(c) \geq u(a) \geq u(b)$$

for all $a \in A$, i.e. c is a maximum point and b is a minimum point of u. We have to treat two cases.

Case 1, $u(c) = u(b)$**.**

This is a trivial case and implies that $u(a) = u(b)$ for all $a \in A$, i.e. the function u is constant. Consequently,

$$\tilde{u}(\delta_a) = u(a) = u(b) = \tilde{u}(\delta_b),$$

i.e. $\delta_a \sim \delta_b$ for all $a \in A$. Thus, all lotteries with a sure outcome are equivalent, and it follows that the function v is constant, too, since

$$v(a) = \tilde{v}(\delta_a) = \tilde{v}(\delta_b) = v(b)$$

for all a. We conclude that $v(a) - v(b) = u(a) - u(b)$ $(= 0)$ for all a, which proves the relationship (1) with $C = 1$ and $D = v(b) - u(b)$.

Case 2, $u(c) > u(b)$**.**

Fix an element $a \in A$ and let us compare the lottery δ_a with the lotteries $p = \alpha\delta_c + (1 - \alpha)\delta_b$ for different values of $\alpha \in [0, 1]$. We have equivalence $\delta_a \sim p$ for exactly one value of α, because equivalence holds if and only if

$$\tilde{u}(\delta_a) = \tilde{u}(p) = \alpha\tilde{u}(\delta_c) + (1 - \alpha)\tilde{u}(\delta_b),$$

i.e. if and only if

$$u(a) = \alpha u(c) + (1 - \alpha)u(b),$$

which gives us the unique solution

$$\alpha = \frac{u(a) - u(b)}{u(c) - u(b)}.$$

Note that the solution α really satisfies the condition $0 \leq \alpha \leq 1$ and that

$$1 - \alpha = \frac{u(c) - u(a)}{u(c) - u(b)}.$$

By assumption, $\tilde{v}$ induces the same preference relation as $\tilde{u}$, and this implies that the indifference condition $\delta_a \sim p$ can also be expressed using the function $\tilde{v}$. It follows that

$$\begin{aligned} v(a) &= \tilde{v}(\delta_a) = \tilde{v}(p) = \alpha v(c) + (1-\alpha)v(b) \\ &= \frac{u(a)-u(b)}{u(c)-u(b)}\,v(c) + \frac{u(c)-u(a)}{u(c)-u(b)}\,v(b) \\ &= \frac{v(c)-v(b)}{u(c)-u(b)}\,u(a) + \frac{u(c)v(b)-u(b)v(c)}{u(c)-u(b)}. \end{aligned}$$

This proves the relationship (1) between u and v with

$$C = \frac{v(c)-v(b)}{u(c)-u(b)} > 0 \qquad \text{and} \qquad D = \frac{u(c)v(b)-u(b)v(c)}{u(c)-u(b)},$$

and concludes the proof. □

In order to be able to express Proposition 1.4.1 in a simple way, we make the following definition.

Definition 1.4.1 Two cardinal utility functions u and v on a set are called *equivalent* if they are related by an order preserving affine function, i.e. if the relationship is given by equation (1).

Using this definition, we can formulate Proposition 1.4.1 as follows:

Two cardinal utility functions on a set induce via their expected utilities the same preference relation on the lottery set if and only if they are equivalent.

Exercise

1.9 Anne has been hired as the CEO of a company with the following remuneration: a fixed annual salary of 2 million SEK and a variable bonus. The size of the bonus will depend on the company's profit, which in any given year may be poor, mediocre, good or very good with a probability that is estimated to be 0.1, 0.4, 0.3, and 0.2, respectively. The bonus is set to be 0, 0.5, 1, and 2 million SEK, respectively.

a) Calculate Anne's expected annual remuneration.

b) Calculate Anne's expected utility if $u(x) = \ln x$ is her cardinal utility function.

c) Suppose Anne as an alternative is offered a fixed annual salary of 2 850 000 SEK with no bonus. Should she accept this offer?

d) What is the minimum annual salary without bonus that Anna is willing to accept?

1.5 von Neumann–Morgenstern preferences

We have seen that in order to prioritize in situations that depend on chance, we need to know our preferences for different lotteries, and that the concept of expected utility gives rise to preference relations on the lottery set. But to calculate the expected utility function, we need a cardinal utility function on the set A of original options, and it is doubtful, to say the least, whether we have access to such a function in real life situations. Maybe the decision-maker is instead able to rank lotteries intuitively without having to start from the utility of the different options in the set A.

A natural question in this context is then whether there are any natural preference relations on the lottery set, other than those generated by cardinal utility functions. John von Neumann and Oskar Morgenstern showed in their groundbreaking work *Theory of Games and Economic Behavior* that the answer is no, if by a natural preference relation is meant a preference relation that satisfies some natural additional requirements beyond completeness and transitivity. Decision and game theory is therefore usually based on the concept of expected utility.

In this section, we present a variant of von Neumann–Morgenstern's results.

Definition 1.5.1 Let A be a finite set. A preference relation $\succeq$ on the lottery set $\mathcal{L}(A)$ is called a *von Neumann–Morgenstern preference*, abbreviated *vNM-preference*, if it satisfies the following two axioms:

Axiom 3 (Continuity axiom) *For all lotteries p, q, $r \in \mathcal{L}(A)$ and all convergent sequences $(\alpha_n)_1^\infty$ of real numbers in the interval $[0,1]$ with $\lim \alpha_n = \alpha$,*

$$\alpha_n p + (1-\alpha_n)q \succeq r \text{ for all } n \;\Rightarrow\; \alpha p + (1-\alpha)q \succeq r$$

and

$$\alpha_n p + (1-\alpha_n)q \preceq r \text{ for all } n \;\Rightarrow\; \alpha p + (1-\alpha)q \preceq r.$$

Axiom 4 (Independence of irrelevant alternatives) *For all p, $q \in \mathcal{L}(A)$ such that $p \sim q$ and all $r \in \mathcal{L}(A)$, $\frac{1}{2}p + \frac{1}{2}r \sim \frac{1}{2}q + \frac{1}{2}r$.*

It follows in particular from the independence axiom that a person, who is indifferent between two alternatives a and a', for each alternative b also has to be indifferent between the options V and V', where V represents a 50–50 chance of getting a or b, and V' represents a 50–50 chance of getting a' or b.

The independence axiom is formulated to be as weak as possible, and there is nothing special about the weight $\frac{1}{2}$. As we will see below, it follows from Axioms 3 and 4 that $\alpha p + (1-\alpha)r \sim \alpha q + (1-\alpha)r$ for all numbers $\alpha \in [0,1]$ and all lotteries r whenever $p \sim q$.

The axiom of independence of irrelevant alternatives may appear harmless, but sometimes decision-makers act in a way that is not consistent with the axiom. Consider the following choice between two lotteries:

- Lottery p_1 promises 100 000 dollars with certainty.
- Lottery p_2 promises 125 000 dollars with probability 0.8, and nothing with probability 0.2.

Most test persons seem to prefer p_1 to p_2. Now consider the choice between the following two lotteries:

- Lottery q_1 promises 100 000 dollars with probability 0.05.
- Lottery q_2 promises 125 000 dollars with probability 0.04.

In this situation, most people tend to prefer lottery q_2 to q_1. But the lotteries q_i can be formed from the lotteries p_i by mixing them with an irrelevant alternative r in the same proportions. If r is the lottery which gives nothing with probability 1, then $q_i = 0.05p_i + 0.95r$. This psychological paradox is called Allais' paradox.

Our next proposition shows, however, that it might be reasonable to require that the properties of Axioms 3 and 4 are fulfilled.

Proposition 1.5.1 *Preference relations induced by expected utility functions are vNM-preferences.*

Proof. Suppose that the preference relation $\succeq$ is induced by the expected utility function $\tilde{u}$.

To prove that $\succeq$ satisfies the first implication of the continuity axiom, we assume that $\alpha_n p + (1-\alpha_n)q \succeq r$ for a convergent sequence $(\alpha_n)_1^\infty$ of numbers in the interval $[0,1]$ with limit α. Then

$$\alpha_n \tilde{u}(p) + (1-\alpha_n)\tilde{u}(q) = \tilde{u}(\alpha_n p + (1-\alpha_n)q) \geq \tilde{u}(r).$$

By letting n go to infinity in the inequality between the extreme ends, we obtain the following inequality

$$\tilde{u}(\alpha p + (1-\alpha)q) = \alpha\tilde{u}(p) + (1-\alpha)\tilde{u}(q) \ \geq \tilde{u}(r),$$

which shows that $\alpha p + (1-\alpha)q \succeq r$.

The second implication is proven analogously.

To prove that the independency axiom is satisfied, assume that $p \sim q$, i.e. that $\tilde{u}(p) = \tilde{u}(q)$, and let r be an arbitrary lottery. It follows that

$$\tilde{u}(\tfrac{1}{2}p + \tfrac{1}{2}r) = \tfrac{1}{2}\tilde{u}(p) + \tfrac{1}{2}\tilde{u}(r) = \tfrac{1}{2}\tilde{u}(q) + \tfrac{1}{2}\tilde{u}(r) = \tilde{u}(\tfrac{1}{2}q + \tfrac{1}{2}r),$$

that is $\frac{1}{2}p + \frac{1}{2}r \sim \frac{1}{2}q + \frac{1}{2}r$.

Hence, Axioms 3 and 4 are satisfied, which means that $\succeq$ is a vNM-preference. □

That a person should be aware of his preferences on a cardinal scale is a strong assumption, but we still need it in order to use the concept of expected utility. A cineast may prefer movie X to movie Y and movie Y to movie Z, but can he say meaningfully that his pleasure of watching movie X instead of movie Y is for example three times as great as his pleasure of watching movie Y instead of movie Z? Probably not.

One way of trying to determine our cineast's preferences on a cardinal scale would be to let him choose beween a cinema ticket to movie Y and a lottery ticket that with probability α gives him a ticket to movie X and with probability $1-\alpha$ a ticket to movie Z. For $\alpha = 1$ he will surely prefer the lottery ticket, because he will then with certainty get a ticket to the movie he prefers, i.e. X. Now let us lower the probability α. If $\alpha = 0.75$, he may still prefer the lottery ticket to the cinema ticket to Y, because the odds for him to get the X-ticket are still high enough. But for $\alpha = 0$, he will certainly prefer the movie ticket to Y, because in this case the lottery ticket means a ticket to movie Z. Somewhere in the range between 0 and 1 there should be a probability that makes him indifferent between the cinema ticket to movie Y and the lottery ticket; let us assume that this probability is $\alpha = 0.25$.

Since cardinal utility functions are equivalent under order preserving affine transformations, we are allowed to fix our cineast's utility values at X and Z so that $u(\mathrm{X}) = 1$ and $u(\mathrm{Z}) = 0$. The expected utility of the lottery ticket when $\alpha = 0.25$ is

$$0.25u(\mathrm{X}) + 0.75u(\mathrm{Z}) = 0.25,$$

and since he is indifferent between the lottery ticket and the cinema ticket to Y, his utility of this movie is now determined to be $u(\mathrm{Y}) = 0.25$.

The movie example suggests that it should always be possible to construct cardinal utility functions for decision-makers, who know how to rank lotteries in a consistent way.

This is really true for vNM-preferences and is the gist of the following proposition, which is the converse of Proposition 1.5.1. As already mentioned, the result was first proven by von Neumann och Morgenstern with slightly different assumptions.

Proposition 1.5.2 *Let $\succeq$ be a vNM-preference on the lottery set $\mathcal{L}(A)$ of a finite set A.*

(a) The preference relation $\succeq$ is then induced by an affine utility function U on $\mathcal{L}(A)$, and the utility function U is uniquely determined up to order preserving affine transformations.

(b) There exists a cardinal utility function u on A such that $U = \tilde{u}$, i.e. such that

$$U(p) = \sum_{a \in A} u(a)p(a)$$

for all lotteries p.

By combining Propositions 1.5.1 and 1.5.2, we get the following immediate corollary.

Corollary 1.5.3 *A preference relation on a lottery set $\mathcal{L}(A)$ is a vNM-preference if and only if it is induced by the expected utility function of a cardinal utility function on the set A.*

Cardinal utility functions, whose expected values are used to represent vNM-preferences for lotteries like the function u in Proposition 1.5.2, are called *Bernoulli utility functions* in honor of Daniel Bernoulli, who used utility functions with decreasing marginal utility in a treatise in 1738.

Proof of Proposition 1.5.2. The proof of part (a) is elementary but long, so we start by proving that assertion (b) follows easily from assertion (a). Let U be the utility function given by (a), and define a function u on A by setting

$$u(a) = U(\delta_a)$$

for all $a \in A$. Each lottery p is a convex combination $p = \sum_{a \in A} p(a)\delta_a$ of the lotteries δ_a, and since U is an affine function, it follows that

$$U(p) = \sum_{a \in A} p(a)U(\delta_a) = \sum_{a \in A} p(a)u(a) = \tilde{u}(p),$$

which proves assertion (b).

The proof of assertion (a) rests on a series of lemmas.

Lemma 1.5.4 *Let p_1, p_2, q_1, q_2 be lotteries in $\mathcal{L}(A)$, let $(\lambda_n)_1^\infty$ and $(\mu_n)_1^\infty$ be two convergent sequences of numbers in the interval $[0, 1]$ with limits λ and μ, and suppose that*

$$\lambda_n p_1 + (1 - \lambda_n)p_2 \sim \mu_n q_1 + (1 - \mu_n)q_2$$

for all n. Then

$$\lambda p_1 + (1 - \lambda)p_2 \sim \mu q_1 + (1 - \mu)q_2.$$

Proof. Write $p = \lambda p_1 + (1-\lambda)p_2$ and $q = \mu q_1 + (1-\mu)q_2$. By the definition of $\sim$ in terms of $\succeq$, it is for symmetry reasons enough to prove the implication

$$\lambda_n p_1 + (1-\lambda_n)p_2 \succeq \mu_n q_1 + (1-\mu_n)q_2 \text{ for all } n \;\Rightarrow\; p \succeq q.$$

Suppose this implication is false, i.e. that

$$\lambda_n p_1 + (1-\lambda_n)p_2 \succeq \mu_n q_1 + (1-\mu_n)q_2$$

for all n but $p \prec q$. We will show that this assumption results in a contradiction.

First assume that there exists an element r in the lottery set $\mathcal{L}(A)$ such that $p \prec r \prec q$. We then have $\mu_n q_1 + (1-\mu_n)q_2 \succ r$ for all but finitely many indices n, because otherwise $\mu_{n_k} q_1 + (1-\mu_{n_k})q_2 \preceq r$ for some infinite subsequence (n_k), and it follows from Axiom 3 that $q \preceq r$, which contradicts our assumption $r \prec q$.

Using transitivity, we conclude that

$$\lambda_n p_1 + (1-\lambda_n)p_2 \succ r$$

for all but finitely many indices n, and using Axiom 3 we now obtain the conclusion $p \succeq r$, which is a contradiction.

So there is no element r such that $p \prec r \prec q$, and we conclude that $\mu_n q_1 + (1-\mu_n)q_2 \succeq q$ for all but finitely many indices n, because otherwise there exists an infinite subsequence (n_k) such that $\mu_{n_k} q_1 + (1-\mu_{n_k})q_2 \prec q$, which means that $\mu_{n_k} q_1 + (1-\mu_{n_k})q_2 \preceq p$, and Axiom 3 now implies that $q \preceq p$, which contradics our assumption $p \prec q$.

This completes our proof by contradiction of the lemma. □

We are now able to generalize the axiom of independence of irrelevant alternatives.

Lemma 1.5.5 *The implication*

$$p \sim q \;\Rightarrow\; \lambda p + (1-\lambda)r \sim \lambda q + (1-\lambda)r$$

is valid for all lotteries $p, q, r \in \mathcal{L}(A)$ and all numbers $\lambda \in [0,1]$.

Proof. We start by proving the implication for all scalars λ of the form

$$\lambda = \sum_{k=1}^{n} \epsilon_k 2^{-k}, \tag{2}$$

where $\epsilon_k = 0$ or $= 1$, using induction over n.

The case $n = 1$, $\epsilon_1 = 0$ is trivial, since $0p + 1r = 0q + 1r = r$, and the case $n = 1$, $\epsilon_1 = 1$ is Axiom 4.

Therefore, assume that the implication in the lemma is valid for all numbers λ of the form (2) with $n-1$ instead of n, and write the number λ occurring in (2) as $\lambda = \epsilon_1/2 + \lambda'/2$ with $\lambda' = \sum_{k=1}^{n-1} \epsilon_{k+1} 2^{-k}$. Due to the induction assumption,

$$\lambda' p + (1-\lambda')r \sim \lambda' q + (1-\lambda')r. \tag{3}$$

Using equation (3) and Axiom 4 we obtain the following equivalence in the case $\epsilon_1 = 0$:

$$\begin{aligned}\lambda p + (1-\lambda)r &= \tfrac{1}{2}\lambda' p + (1 - \tfrac{1}{2}\lambda')r \\ &= \tfrac{1}{2}\big(\lambda' p + (1-\lambda')r\big) + \tfrac{1}{2}r \sim \tfrac{1}{2}\big(\lambda' q + (1-\lambda')r\big) + \tfrac{1}{2}r \\ &= \tfrac{1}{2}\lambda' q + (1 - \tfrac{1}{2}\lambda')r = \lambda q + (1-\lambda)r;\end{aligned}$$

and the following equivalence in the case $\epsilon_1 = 1$:

$$\begin{aligned}\lambda p + (1-\lambda)r &= (\tfrac{1}{2} + \tfrac{1}{2}\lambda')p + (\tfrac{1}{2} - \tfrac{1}{2}\lambda')r \\ &= \tfrac{1}{2}p + \tfrac{1}{2}\big(\lambda' p + (1-\lambda')r\big) \sim \tfrac{1}{2}p + \tfrac{1}{2}\big(\lambda' q + (1-\lambda')r\big) \\ &\sim \tfrac{1}{2}q + \tfrac{1}{2}(\lambda' q + (1-\lambda')r) = (\tfrac{1}{2} + \tfrac{1}{2}\lambda')q + (\tfrac{1}{2} - \tfrac{1}{2}\lambda')r \\ &= \lambda q + (1-\lambda)r.\end{aligned}$$

This completes the induction step, and the lemma is proven for all numbers λ of the special form (2).

If λ is an arbitrary number in the interval $[0,1]$, we choose a sequence $(\lambda_n)_1^\infty$ of numbers of the special form (2) and which converges to λ as $n \to \infty$. Then $\lambda_n p + (1-\lambda_n)r \sim \lambda_n q + (1-\lambda_n)r$ for all n, and it now follows from Lemma 1.5.4 that $\lambda p + (1-\lambda)r \sim \lambda q + (1-\lambda)r$. □

Lemma 1.5.6 *Suppose that $q \prec p$ and let x be an element of the lottery set such that $q \preceq x \preceq p$. Then there exists a number $\mu \in [0,1]$ such that $x \sim \mu p + (1-\mu)q$.*

Proof. The set $\Lambda = \{\lambda \in [0,1] \mid \lambda p + (1-\lambda)q \succeq x\}$ is bounded below and nonempty, since it contains 1. Let $\mu = \inf \Lambda$, and choose a sequence $(\mu_n)_1^\infty$ of numbers in Λ that converges to μ. Since $\mu_n p + (1-\mu_n)q \succeq x$ for all n, it follows from the continuity axiom that $\mu p + (1-\mu)q \succeq x$.

We will prove that the converse relation $\mu p + (1-\mu)q \preceq x$ also holds. This is obvious in the case $\mu = 0$, because then $\mu p + (1-\mu)q = q \preceq x$. So assume that $\mu > 0$, and choose a sequence $(\mu_n)_1^\infty$ of numbers from the interval $[0,\mu[$ that converges to μ. It follows from the infimum definition that $\mu_n \notin \Lambda$, i.e. that $\mu_n p + (1-\mu_n)q \prec x$, for all n, and Axiom 3 now yields the conclusion $\mu p + (1-\mu)q \preceq x$. □

Lemma 1.5.7 *Suppose that $q \prec p$ and $0 < \lambda < 1$. Then*

$$q \prec \lambda p + (1-\lambda)q \prec p.$$

Proof. We prove the lemma by contradiction assuming that

$$p \preceq \lambda p + (1-\lambda)q. \tag{4}$$

It follows from Lemma 1.5.6 and the assumption (4) that there is a number $\mu \in [0,1]$ such that

$$p \sim \mu(\lambda p + (1-\lambda)q) + (1-\mu)q = \mu\lambda p + (1-\mu\lambda)q,$$

and we will prove using induction that

$$p \sim \mu^n \lambda^n p + (1-\mu^n\lambda^n)q \tag{5}$$

for all positive integers n. The case $n = 1$ is already done, so assume that (5) holds for a certain positive integer n. By combining this induction hypothesis with Lemma 1.5.5, we conclude that

$$\begin{aligned} p &\sim \mu\lambda p + (1-\mu\lambda)q \sim \mu\lambda\big(\mu^n\lambda^n p + (1-\mu^n\lambda^n)q\big) + (1-\mu\lambda)q \\ &= \mu^{n+1}\lambda^{n+1}p + (1-\mu^{n+1}\lambda^{n+1}). \end{aligned}$$

This proves (5) with n replaced by $n+1$ and completes the induction step.

Since $\mu^n\lambda^n \to 0$ as $n \to \infty$, it now follows from (5) and Lemma 1.5.4 that

$$p \sim 0p + 1q = q,$$

which is a contradiction and proves that $\lambda p + (1-\lambda)q \prec p$.

The other half $q \prec \lambda p + (1-\lambda)q$ of the lemma is proven using an analogous proof by contradiction. □

Lemma 1.5.8 *Suppose that $q \prec p$ and $0 \le \mu < \lambda \le 1$. Then*

$$\mu p + (1-\mu)q \prec \lambda p + (1-\lambda)q.$$

Proof. The assertion is a special case of Lemma 1.5.7 if $\mu = 0$ or $\lambda = 1$, so assume that $0 < \mu < \lambda < 1$. Then $0 < \mu/\lambda < 1$, and using Lemma 1.5.7 we first obtain the relation

$$\lambda p + (1-\lambda)q \succ q,$$

and then the asserted relation

$$\begin{aligned}\lambda p + (1-\lambda)q &\succ \frac{\mu}{\lambda}\big(\lambda p + (1-\lambda)q\big) + (1-\frac{\mu}{\lambda})q \\ &= \mu p + (1-\mu)q.\end{aligned}$$ □

Lemma 1.5.9 *Let p and q be lotteries on a finite set satisfying $q \prec p$.*

(i) For each lottery x such that $q \preceq x \preceq p$, there exists a unique real number $\mu = \mu_{q,p}(x)$ in the interval $[0,1]$ such that $x \sim \mu p + (1-\mu)q$.

(ii) For all lotteries x, y satisfying $q \preceq x \preceq p$ and $q \preceq y \preceq p$, and for all scalars $\lambda \in [0,1]$,

$$x \succeq y \;\Leftrightarrow\; \mu_{q,p}(x) \ge \mu_{q,p}(y) \tag{6}$$

and

$$\mu_{q,p}(\lambda x + (1-\lambda)y) = \lambda \mu_{q,p}(x) + (1-\lambda)\mu_{q,p}(y). \tag{7}$$

Proof. By Lemma 1.5.6, there is a number $\mu \in [0,1]$ such that

$$x \sim \mu p + (1-\mu)q,$$

and Lemma 1.5.8 implies that the number μ is unique and satisfies equivalence (6).

If $x \sim \mu_1 p + (1-\mu_1)q$ and $y \sim \mu_2 p + (1-\mu_2)q$, then, by applying Lemma 1.5.5 twice, we obtain

$$\begin{aligned}\lambda x + (1-\lambda)y &\sim \lambda(\mu_1 p + (1-\mu_1)q) + (1-\lambda)y \\ &\sim \lambda(\mu_1 p + (1-\mu_1)q) + (1-\lambda)(\mu_2 p + (1-\mu_2)q) \\ &= \big(\lambda\mu_1 + (1-\lambda)\mu_2\big)p + \big(1 - (\lambda\mu_1 + (1-\lambda)\mu_2)\big)q,\end{aligned}$$

and this proves equation (7). □

From now on we assume that there exist at least two non-equivalent lotteries in $\mathcal{L}(A)$, because the preference relation is obviously represented by every constant function on $\mathcal{L}(A)$ in the trivial case when $p \sim q$ for all $p, q \in \mathcal{L}(A)$, and every representing utility function has to be constant.

If the lottery set has a least element q and a greatest element p (and $q \prec p$), then Lemma 1.5.9 provides us with an affine utility function U that represents the preference relation, namely the function $U = \mu_{q,p}$. The uniqueness part of Proposition 1.5.2 follows from our next lemma, which will also be needed to extend our existence result to the general case.

Lemma 1.5.10 *Suppose that U and V are two affine utility functions on $\mathcal{L}(A)$, both representing the preference relation $\succeq$. Then there exist two real numbers $C > 0$ and D such that $V = CU + D$.*

Proof. Fix two elements p and q in $\mathcal{L}(A)$ such that $q \prec p$. Since $U(p) \neq U(q)$, there exist uniquely determined real numbers C and D such that $CU(p)+D = V(p)$ and $CU(q) + D = V(q)$, and C is a positive number since $U(p) > U(q)$ and $V(p) > V(q)$.

Let $x \in \mathcal{L}(A)$ be arbitrary. Then, either $q \preceq x \preceq p$, $q \prec p \prec x$ or $x \prec q \prec p$.

In the first mentioned case, there exists a number $\mu \in [0,1]$ such that $x \sim \mu p + (1-\mu)q$, and it follows that

$$\begin{aligned} V(x) &= V(\mu p + (1-\mu)q) = \mu V(p) + (1-\mu)V(q) \\ &= \mu(CU(p)+D) + (1-\mu)(CU(q)+D) \\ &= C(\mu U(p) + (1-\mu)U(q)) + D = C\,U(\mu p + (1-\mu)q) + D \\ &= CU(x) + D. \end{aligned}$$

In the second case, there is instead a number μ such that $0 < \mu < 1$ and $p \sim \mu x + (1-\mu)q$, and it now follows that $U(p) = \mu U(x) + (1-\mu)U(q)$ and $V(p) = \mu V(x) + (1-\mu)V(q)$. Thus,

$$\begin{aligned} V(x) &= \frac{1}{\mu}(V(p) - V(q)) + V(q) = \frac{1}{\mu}(CU(p) - CU(q)) + CU(q) + D \\ &= CU(x) + D. \end{aligned}$$

Analogously, $V(x) = CU(x) + D$ also in the third case. □

Lemma 1.5.9 shows how to define an affine utility function $\mu_{q,p}$ that represents $\succeq$ on the subinterval $\{x \in \mathcal{L}(A) \mid q \preceq x \preceq p\}$ of the lottery set. What remains is to show how to join these functions to an affine utility function U on all of $\mathcal{L}(A)$.

To this end, fix two elements p_0, $p_1 \in \mathcal{L}(A)$ with $p_0 \prec p_1$, and let p, q be two arbitrary elements in $\mathcal{L}(A)$ such that $q \preceq p_0 \prec p_1 \preceq p$. Define the function $U_{q,p}$ on the interval $\{x \in \mathcal{L}(A) \mid q \preceq x \preceq p\}$ by

$$U_{q,p}(x) = \frac{\mu_{q,p}(x) - \mu_{q,p}(p_0)}{\mu_{q,p}(p_1) - \mu_{q,p}(p_0)}.$$

Then $U_{q,p}$ is an affine function which represents $\succeq$ on the given interval, and $U_{q,p}(p_0) = 0$ and $U_{q,p}(p_1) = 1$.

If p', q' are two other elements of $\mathcal{L}(A)$ with $q' \preceq p_0 \prec p_1 \preceq p'$, then of course $U_{q',p'}(p_0) = 0$ and $U_{q',p'}(p_1) = 1$, whence it follows from Lemma 1.5.10 that $U_{q,p}(x) = U_{q',p'}(x)$ for all x belonging to the intersection of the two intervals $\{x \in \mathcal{L}(A) \mid q \preceq x \preceq p\}$ and $\{x \in \mathcal{L}(A) \mid q' \preceq x \preceq p'\}$.

Therefore, we obtain a uniquely defined function U by for each given $x \in \mathcal{L}(A)$ first choosing $p \succ q$ such that the interval $\{x \in \mathcal{L}(A) \mid q \preceq x \preceq p\}$ contains x as well as p_0 och p_1, and then defining

$$U(x) = U_{q,p}(x).$$

The function U is order preserving and affine, since it has those properties on each subinterval of the lottery set, and U is therefore a utility function that represents $\succeq$ on the whole of the lottery set $\mathcal{L}(A)$.

This completes the proof of Proposition 1.5.2. □

Exercises

1.10 Prove that the following property, known as *Archimedean property*, follows from the continuity axiom: For all lotteries p, q, $r \in \mathcal{L}(A)$ such that $p \succ q \succ r$, there exist numbers α, β in the interval $]0,1[$ such that $\alpha p + (1-\alpha)r \succ q$ and $q \succ \beta p + (1-\beta)r$.

1.11 Give an example of a preference relation on $\mathcal{L}(A)$ that is not a vNM-preference. [Hint: Two elements in A are sufficient.]

Chapter 2

Strategic Games

Strategic games serve as models of interactive decision making in which each decision-maker chooses his action once and for all, independently and without knowlege of the other participants' choices. The model consists of a finite set of players and for each player a set of actions. The players' combined choices of actions is called an outcome, and the preferences of each player is defined on the set of all possible outcomes. This distinguishes strategic games from pure decision problems since each player has to take into account the possible actions of the other players when making his own choice. There is therefore in general no best choice of action for a given player. So game theory seldom has an answer to the question how to play in a specific game when played once. Instead, the focus is on describing different equilibrium solutions, i.e. outcomes with the property that no player profits by unilaterally deviating from his choice of action.

2.1 Definition and examples

We now give the formal definition of a strategic game.

Definition 2.1.1 A *strategic game* $\langle N, (A_i), (\succeq_i)\rangle$ consists of

- a finite set N of *players* − the number of players is assumed to be at least two, and we will number them 1,2, . . . , n, which means that we identify N with the set $\{1, 2, 3, \dots, n\}$;
- for each player $i \in N$ a set A_i of possible *actions*;
- for each player i a *preference relation* $\succeq_i$ on the product set

$$A = A_1 \times A_2 \times \cdots \times A_n.$$

The game is called *finite* if A is a finite set (i.e. if all action sets A_i are finite).

The preference relations $\succeq_i$ i are often defined by utility functions u_i, usually called *payoff functions*, and we then use the notation $\langle N, (A_i), (u_i)\rangle$ for the game $\langle N, (A_i), (\succeq_i)\rangle$.

The product set $A = A_1 \times A_2 \times \cdots \times A_n$ consists by definition of all n-tiples $a = (a_1, a_2, \ldots, a_n)$ of the different players' actions $a_i \in A_i$. Such n-tiples of actions are called *outcomes* of the game. The product set A thus consists of all possible outcomes. If the sets A_i consist of ν_i possible actions, then there are of course $\nu_1\nu_2\cdots\nu_n$ possible outcomes .

The third property in the definition of a strategic game means that a player's preference relation (or payoff function) should be defined on the set of outcomes. Since a player can only control his own actions, it is in general impossible for him to achieve the outcome which is optimal for him. What should be meant by an optimal solution is therefore different for strategic games than for pure decision problems where the decision-maker has full control over the situation. So one of the main tasks in game theory is to identify solutions that in some sense can be considered as optimal or stable.

We will mainly study finite strategic games and mostly exemplify with two-person games. To define a finite game $\langle\{1,2\}, (A_i), (u_i)\rangle$ with two players we have to specify the two action sets $A_1 = \{r_1, r_2, \ldots, r_m\}$ and $A_2 = \{c_1, c_2, \ldots, c_n\}$, and the payoffs $\alpha_{ij} = u_1(r_i, c_j)$ and $\beta_{ij} = u_2(r_i, c_j)$ for all outcomes (r_i, c_j). This is easily done using payoff tables as follows:

	c_1	c_2	$\cdots\cdots$	c_n
r_1	$(\alpha_{11}, \beta_{11})$	$(\alpha_{12}, \beta_{12})$	$\cdots\cdots$	$(\alpha_{1n}, \beta_{1n})$
r_2	$(\alpha_{21}, \beta_{21})$	$(\alpha_{22}, \beta_{22})$	$\cdots\cdots$	$(\alpha_{2n}, \beta_{2n})$
$\vdots$				
r_m	$(\alpha_{m1}, \beta_{m1})$	$(\alpha_{m2}, \beta_{m2})$	$\cdots\cdots$	$(\alpha_{mn}, \beta_{mn})$

Player 1 is thus always the player who chooses a row and we will often call him the *row player*, player 2 being the *column player*. The specific names of the elements in the sets A_1 and A_2 are irrelevant, of couse, so the game is entirely determined by the two $m\times n$-matrices $[\alpha_{ij}]$ and $[\beta_{ij}]$, which contain all information about the two players' payoff functions.

The matrix $[\alpha_{ij}]$ is the row player's *payoff matrix*, and $[\beta_{ij}]$ is the column player's payoff matrix.

We now illustrate the strategic game concept with some classical examples.

EXAMPLE 2.1.1 (*Prisoner's dilemma*) Our first example is intended to illustrate a situation where cooperation would be profitable but where lack of trust may induce a party to abandon the most profitable action.

Two members of a criminal gang, both suspected of a serious crime, are arrested and imprisoned without possibility to communicate with each other. There are sufficient evidence to convict each of them for a lesser crime but not enough evidence to convict them for the principal charge, unless at least one of the prisoners confesses and can be used as a witness against the other. So the prosecutor offers each prisoner a bargain. If both deny, then they will both receive a sentence of one year in prison for the minor crime. If only one of them confesses, he will be freed and used as a witness against the other, who will receive a sentence of five years. And if both prisoners confess, each will be sentenced to three years in prison.

We can model the situation as a strategic game as follows. The players are the two prisoners, and both have the same action set, viz. $\{Deny, Confess\}$. This means that there are four outcomes, ranked as follows by prisoner 1:

$$(Confess, Deny) \succ_1 (Deny, Deny) \succ_1 (Confess, Confess) \succ_1 (Deny, Confess)$$

since he prefers shortest possible imprisonment. Analogously, prisoner 2 has the following preferences

$$(Deny, Confess) \succ_2 (Deny, Deny) \succ_2 (Confess, Confess) \succ_2 (Confess, Deny).$$

A utility function u_1, which represents the preferences of prisoner 1, has to satisfy

$$\begin{aligned} u_1(Confess, Deny) > u_1(Deny, Deny) > u_1(Confess, Confess) \\ > u_1(Deny, Confess), \end{aligned}$$

and such a function is given by for instance

$$\begin{aligned} &u_1(Confess, Deny) = 0, &&u_1(Deny, Deny) = -1, \\ &u_1(Confess, Confess) = -3, &&u_1(Deny, Confess) = -5. \end{aligned}$$

Similarly, we define for prisoner 2

$$\begin{aligned} &u_2(Deny, Confess) = 0, &&u_2(Deny, Deny) = -1, \\ &u_2(Confess, Confess) = -3, &&u_2(Confess, Deny) = -5. \end{aligned}$$

The game can now be represented in a compact way by the following payoff table:

	Deny	*Confess*
Deny	$(-1,-1)$	$(-5,0)$
Confess	$(0,-5)$	$(-3,-3)$

How should prisoner 1 reason in this situation? Best for him if he confesses and prisoner 2 denies, because he will then be released. But what if prisoner 2 also confesses; then he must stay in prison for three years. Maybe, it is better for him to deny, hoping that prisoner 2 will deny, too, because that means just one year in prison. But can he really trust prisoner 2 who faces the risk of five years in prison by denying. So prisoner 2 may confess and then he, instead, will have to serve five years in prison by denying instead of confessing. Both players face the same dilemma, of course. □

Our next example models situations in which persons with conflicting interests would benefit from coordinating their behavior, but where due to lack of communication they run the risk of ending up in a highly unfavorable outcome.

Example 2.1.2 (*Battle of Sexes*) A couple has agreed to meet after work but both of them have forgotten whether they decided to go to the concert hall or to the football stadium. The wife (player 1) would prefer attending the concert whereas her husband (player 2) would rather watch the football match. Both would prefer to be together at the same event rather than to go to different places, but if they end up in different places, they are both indifferent between attending the concert or watching football. If they cannot communicate, where should they go?

Their dilemma can be modeled as a strategic game, where both players have $\{Concert, Football\}$ as their sets of action and with payoffs given by the following table:

	Concert	*Football*
Concert	$(2, 1)$	$(0, 0)$
Football	$(0, 0)$	$(1, 2)$

If the parties were allowed to communicate and negotiate, they would certainly agree to attend the same event, but without such possibilities there is no guarantee that they do not end up in an outcome that both like least. □

Example 2.1.3 (*Matching Pennies*) Each of two persons chooses simultaneously one side of a coin, i.e. *Head* or *Tail*. If their choices coincide, person 1 gets a dollar from person 2, and if they differ, person 1 pays a dollar to person 2. Both persons only care about the amount of money they receive, which means that their preferences are given by the following utility functions:

$$u_1(Head, Head) = u_1(Tail, Tail) = 1, \;\; u_1(Head, Tail) = u_1(Tail, Head) = -1$$
$$u_2(Head, Head) = u_2(Tail, Tail) = -1, \;\; u_2(Head, Tail) = u_2(Tail, Head) = 1.$$

The game is now given by the following table

	Head	*Tail*
Head	$(1, -1)$	$(-1, 1)$
Tail	$(-1, 1)$	$(1, -1)$

The interests of the players are diametrically opposed in this game. Player 1 wants to choose the same alternative as player 2, while player 2 wants to choose the opposite alternative. □

Example 2.1.4 (*Hawk–Dove*) Two animals are fighting over a prey. Both animals can behave aggressively (hawkish) or passively (dovish). Each animal prefers to act hawkish if its opponent is dowish and to act dovish if its opponent is hawkish, and each animal, whether being hawkish or not, prefers that its opponent behaves like a dove.

The situation can be modeled as a strategic game with the two combatants as players and with payoffs satisfying

$$u_1(H,D) > u_1(D,D) > u_1(D,H) > u_1(H,H)$$
$$u_2(D,H) > u_1(D,D) > u_1(H,D) > u_1(H,H),$$

and an example of payoffs satisfying this is given by the following table:

	Hawk	*Dove*
Hawk	$(0,0)$	$(3,1)$
Dove	$(1,3)$	$(2,2)$

□

The philosopher Jean-Jaques Rousseau has described the conflict between individual safety and social cooperation using the following example.

Example 2.1.5 (*Stag Hunt*) A hunting party goes hunting in a forest with plenty of hares and a stately stag. Each hunter can individually choose to hunt the stag or to hunt a hare, and each hunter must make his choice without knowing the choices of the others. If everyone in the team is focused on stag hunting, then they will manage to shoot the stag that is shared equally between the hunters, but the stag escapes if any of the hunters is engaged in hare hunting. Each hunter engaged in hare hunting also manages to shoot a hare which he may keep for himself. Each hunter prefers, however, a portion of the deer to their own hare.

The hunt can be modeled as a strategic game with the n hunters as players and with $\{Stag, Hare\}$ as set of actions for each player. Every hunter prefers that all hunters take part in the stag hunt in front of any other option and is indifferent between what the others do if he chooses to shoot a hare. Every hunter also prefers shooting a hare to participating in the stag hunt if somebody else should choose to shoot a hare. Finally, every hunter is indifferent between all alternatives where he partipates in the stag hunt and somebody else does not (because he gets no prey in these cases). The preferences of player 1 are thus given as follows:

$$(Stag, Stag, \ldots, Stag) \succ_1 (Hare, X_2, \ldots, X_n) \sim_1 (Hare, Y_2, \ldots, Y_n)$$
$$\succ_1 (Stag, Z_2, \ldots, Z_n) \sim_1 (Stag, W_2, \ldots, W_n).$$

where the letters X_i, Y_i, Z_i and W_i stand for *Stag* or *Hare* with not all of $Z_2, \ldots, Z_n$ and not all of $W_2, \ldots, W_n$ being equal to *Stag*.

The payoff table for the stag hunt with just two players has the following form:

	Stag	*Hare*
Stag	$(2, 2)$	$(0, 1)$
Hare	$(1, 0)$	$(1, 1)$

The Stag Hunt is, like the Prisoner's Dilemma, an example of a situation where cooperation would be profitable, but where lack of trust may force a party to secure what he can achieve on his own and in this way abandon a better alternative. □

2.2 Nash equilibrium

What action should a player of a specific strategic game choose under the assumption that he wants to maximize his utility? The problem, of course, is that what is best for a player depends on the other players' choices. So a player who wants to maximize his utility must have an idea about how the other players are going to play. And that is not included in the rules of the game, which only assume that the action sets and the preferences of all players are common knowledge.

So game theory provides no general answers to the question how to play a strategic game when played a single time. However, what game theory does is to single out outcomes with special features, for instance outcomes that are stable in the sense that no player has anything to gain by unilaterally deviating from the action contained in the outcome. This stable outcome is the celebrated Nash equilibrium, that we are now going to define formally.

Definition 2.2.1 Let $\langle N, (A_i), (\succeq_i)\rangle$ be a strategic game with n players. An outcome $a^* = (a_1^*, \dots, a_n^*)$ is called a *Nash equilibrium* or *Nash solution* if no player benifits from unilaterally replacing his action in the outcome a^* with any other action, i.e. if for every player i and every action $a_i \in A_i$ we have

$$a^* = (a_{-i}^*, a_i^*) \succeq_i (a_{-i}^*, a_i).$$

As shown in the examples below there are strategic games without any Nash equilibrium and games with multiple Nash equilibria.

EXAMPLE 2.2.1 (*Prisoner's Dilemma*) The outcome (*Confess*, *Confess*) is a Nash solution of the game

	Deny	*Confess*
Deny	$(-1,-1))$	$(-5,0)$
Confess	$(0,-5)$	$(-3,-3)$

because

$$u_1(\mathit{Confess}, \mathit{Confess}) = -3 > -5 = u_1(\mathit{Deny}, \mathit{Confess}) \qquad \text{and}$$
$$u_2(\mathit{Confess}, \mathit{Confess}) = -3 > -5 = u_2(\mathit{Confess}, \mathit{Deny}),$$

i.e. no player benefits from unilaterally changing from confessing to denying.

The Nash solution is unique, because an outcome where at least one of the players denies is worse for that player then the outcome where he confesses and the other player's action remains unchanged.

The Prisoner's Dilemma game is an illustration of the fact that a Nash solution is not necessarily the collectively best outcome. Here, the alternative (*Deny*, *Deny*), giving both of them one year in prison, is obviously better than the Nash solution, which results in three years inprisonment each. If the prisoners were allowed to cooperate, they would choose the first mentioned alternative, of course. □

EXAMPLE 2.2.2 (*Battle of Sexes*) The game

	Concert	*Football*
Concert	$(2,1)$	$(0,0)$
Football	$(0,0)$	$(1,2)$

has two Nash solutions, namely (*Concert*, *Concert*) and (*Football*, *Football*). The first mentioned Nash solution is the row player's best solution and the second mentioned Nash solution is the column player's best solution. Moreover, both Nash solutions are better for both players than the remaining two outcomes. Both players are of course well aware of this fact but without additional information it is still impossible for them to coordinate their actions so that they end up in one of the two Nash solutions. Similar problems often arise in games with more than one Nash solution. □

EXAMPLE 2.2.3 (*Matching Pennies*) The game

	Head	*Tail*
Head	$(1, -1)$	$(-1, 1)$
Tail	$(-1, 1)$	$(1, -1)$

has no Nash equilibrium. The outcome (*Head*, *Head*) is not a Nash equilibrium since player 2 profits from exchanging *Head* for *Tail*, the outcome (*Head*, *Tail*) is not a Nash equilibrium since player 1 now profits from exchanging *Head* for *Tail*, the outcome (*Tail*, *Tail*) is not a Nash equilibrium, since player 2 profits from exchanging *Tail* for *Head*, and nor is the outcome (*Tail*, *Head*) a Nash solution since player 1 profits from exchanging *Tail* for *Head*. □

EXAMPLE 2.2.4 (*Stag Hunt*) The stag hunt game in Example 2.1.5 (with n hunters) has two Nash solutions, viz. $(Stag, Stag, \dots, Stag)$, when all hunters are engaged in the stag hunting, and $(Hare, Hare, \dots, Hare)$, when all hunters shoot hares. Obviously, no hunter profits by unilaterally deviating from the first mentioned solution since he will then lose a part of the deer, nor does anyone profit by unilaterally deviating from the second solution, since this means getting nothing instead of a hare.

It is also easy to see that there are no other Nash solutions. An outcome of the form $(X_1, \dots, X_{i-1}, Stag, X_{i+1}, \dots, X_n)$, where hunter no. i is a stag hunter and at least one of the remaining persons is hunting hares (i.e. some $X_k = Hare$) can not be a Nash solution, because hunter i gets nothing and profits by exchanging stag hunting for hare hunting. □

We will now give a useful alternative characterization of Nash equilibria. If all players except player i announce their choices of action in advance, player i faces a pure decision problem. His best response, given that the other players have chosen a_{-i}, is to choose an action in A_i that is maximal with respect to the restriction of his preference relation $\succsim_i$ to the set $\{a_{-i}\} \times A_i$. This motivates the following concept.

Definition 2.2.2 Let a_{-i} be a vector of actions for all players except player i in a strategic game $\langle N, (A_i), (\succeq_i)\rangle$. The set

$$B_i(a_{-i}) = \{x_i \in A_i \mid (a_{-i}, x_i) \succeq_i (a_{-i}, y_i) \text{ for all } y_i \in A_i\}$$

is called the *best-responce set* of player i given the other players' choices a_{-i}. The set valued function B_i, with A_{-i} as domain of definition, is called the *best-response function* of player i.

In other words, player i's best-response set $B_i(a_{-i})$ consists of the actions $x_i \in A_i$ that make the outcome (a_{-i}, x_i) to a maximal element in the set $\{a_{-i}\} \times A_i$ with respect to the the player's preference relation $\succeq_i$.

All best-response sets are obviously nonempty if the game is finite, and it follows from Proposition 1.2.4 that all best-response sets are nonempty if the action sets A_i are compact and the preference relations $\succeq_i$ are continuous.

Using best-response functions we can reformulate the Nash equilibrium definition as follows.

Proposition 2.2.1 *An outcome a^* in a strategic game is a Nash equilibrium if and only if $a^*_i \in B_i(a^*_{-i})$ for every player i.*

Proof. The action a_i^* in a Nash equilibrium a^* is optimal for player i given that the other players have chosen a_{-i}^*. □

EXAMPLE 2.2.5 (*Hawk–Dove*) The game with payoff table

	Hawk	*Dove*
Hawk	$(0,0)$	$(3,1)$
Dove	$(1,3)$	$(2,2)$

has the following best-response sets:

$$B_1(\textit{Hawk}) = \{\textit{Dove}\},\ \ B_1(\textit{Dove}) = \{\textit{Hawk}\}$$
$$B_2(\textit{Hawk}) = \{\textit{Dove}\},\ \ B_2(\textit{Dove}) = \{\textit{Hawk}\}.$$

The outcome $(\textit{Dove}, \textit{Hawk})$ is a Nash solution since

$$\textit{Dove} \in B_1(\textit{Hawk}) \qquad \text{and} \qquad \textit{Hawk} \in B_2(\textit{Dove}).$$

Analogously, $(\textit{Hawk}, \textit{Dove})$ is a Nash solution. The outcome $(\textit{Hawk}, \textit{Hawk})$ is not a Nash solution, because $\textit{Hawk} \notin B_1(\textit{Hawk})$.

So in a Nash equilibrium one of the combatants should be hawkish and the other should be dowish and withdraw from the battle. Just like in the game Battle of Sexes, it seems unclear how this balance will be achieved. In nature, this tends to be resolved automatically, however. After some fights, the pecking order is established, and the combatants know who should behave hawkish and who should behave dovish. □

The characterization in Proposition 2.2.1 gives rise to a simple algorithm for calculating Nash equilibria in finite two-person games. An example clarifies the situation.

EXAMPLE 2.2.6 Consider the game

	k_1	k_2	k_3	k_4	k_5	k_6
r_1	$(2,1)$	$(4,3)$	$(7,2)$	$(7,4)$	$(0,5)$	$(3,2)$
r_2	$(4,0)$	$(5,4)$	$(1,6)$	$(0,4)$	$(0,3)$	$(5,1)$
r_3	$(1,3)$	$(5,3)$	$(3,2)$	$(4,1)$	$(1,0)$	$(4,3)$
r_4	$(4,3)$	$(2,5)$	$(4,0)$	$(1,0)$	$(1,5)$	$(2,1)$

The row player's best response to the column player's action k_1 is r_2 or r_4. We mark this in the table by putting an asterisk after $u_1(r_2, k_1)$ and $u_1(r_4, k_1)$, i.e. after the maximum fours. The row player's best response if the column player chooses action k_2 is r_2 or r_3, so we put asterisks after the corresponding payoffs $u_1(r_2, k_2)$ and $u_1(r_3, k_2)$ in column number 2, etc.

Similarly, k_5 is the best response of the column player if the row player chooses action r_1, because k_5 maximizes the function $y \mapsto u_2(r_1, y)$. We therefore put an asterisk after $u_2(r_1, k_5)$, i.e. after the five in row 1 and column 5, etc. Our algorithm results in the following table:

	k_1	k_2	k_3	k_4	k_5	k_6
r_1	$(2, 1)$	$(4, 3)$	$(7^*, 2)$	$(7^*, 4)$	$(0, 5^*)$	$(3, 2)$
r_2	$(4^*, 0)$	$(5^*, 4)$	$(1, 6^*)$	$(0, 4)$	$(0, 3)$	$(5^*, 1)$
r_3	$(1, 3^*)$	$(5^*, 3^*)$	$(3, 2)$	$(4, 1)$	$(1^*, 0)$	$(4, 3^*)$
r_4	$(4^*, 3)$	$(2, 5^*)$	$(4, 0)$	$(1, 0)$	$(1^*, 5^*)$	$(2, 1)$

When we are done, we note that there are two pairs of payoffs with double asterisks, namely $(5, 3)$ and $(1, 5)$. The double asterisks at $(5, 3)$ means that r_3 is the best response of the row player if the column player chooses action k_2, and k_2 is the best response of the column player if the row player chooses action r_3. Hence, (r_3, k_2) is a Nash solution. Analogously, (r_4, k_5) is a Nash solution, too, and there are no other Nash solutions. □

Example 2.2.7 (*Splitting money*) Two persons are bargaining how to split \$10 and agree to use the following procedure. Each of them announces the amount he wants to receive by naming an integer in the interval $[0, 10]$. If the sum of these integers does not exceed 10, then they receive the amounts that they named and the remainder is destroyed. If the sum is greater than 10 and the numbers are different, then the person who named the smaller amount gets it and the other person gets the remainder. If the sum exceeds 10 and the numbers are equal, than both get \$5.

The best-response functions B_1 and B_2 of the two players are identical and given by

$$B_i(0) = \{10\}, B_i(1) = \{9, 10\}, B_i(2) = \{8, 9, 10\}, B_i(3) = \{7, 8, 9, 10\},$$
$$B_i(4) = \{6, 7, 8, 9, 10\}, B_i(5) = \{5, 6, 7, 8, 9, 10\},$$
$$B_i(6) = \{5, 6\}, B_i(7) = \{6\}, B_i(8) = \{7\}, B_i(9) = \{8\}, B_i(10) = \{9\}.$$

The graphs of the two best-response functions B_2 and B_1 are shown in the left and the middle panel of Figure 2.1. The right panel contains boths graphs and shows that the intersection of the two consists of the four points $(5, 5)$, $(5, 6)$, $(6, 5)$ and $(6, 6)$. Since a point (x_1, x_2) belongs to the intersection of the two graphs if and only if $x_1 \in B_1(x_2)$ and $x_2 \in B_2(x_1)$, this means that $(5, 5)$, $(5, 6)$, $(6, 5)$ and $(6, 6)$ are the four Nash equilibria of the game.

□

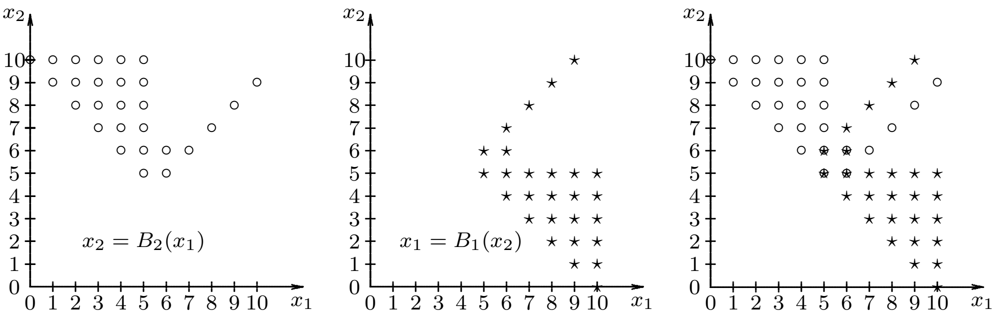

Figure 2.1. The left panel shows the graph of the best-response function of person 2, the middle panel shows the graph of the best-response function of person 1, and the right panel shows both graphs.

Exercises

2.1 Two children play the game *Rock–paper–scissors*. (Paper beats rock, rock beats scissors and scissors beat paper.) Formulate the game as a strategic game and determine if there are any Nash equilibria.

2.2 Determine the best-response functions of the players in the games Prisoner's Dilemma, Matching Pennies and Stag Hunt, and verify that the Nash solutions of the games are the ones computed in the examples above.

2.3 Find the Nash equilibria of the following game:

	k_1	k_2	k_3	k_4
r_1	$(-3,-4)$	$(2,-1)$	$(0,6)$	$(1,1)$
r_2	$(2,0)$	$(2,2)$	$(-3,0)$	$(1,-2)$
r_3	$(2,-3)$	$(-5,1)$	$(-1,-1)$	$(1,-3)$
r_4	$(-4,3)$	$(2,-5)$	$(1,2)$	$(-3,1)$

2.4 An outcome $a^* = (a_1^*, \dots, a_n^*)$ is called a *strict Nash equilibrium* of the game $\langle N, (A_i), (\succeq_i)\rangle$ if

$$(a_{-i}^*, a_i^*) \succ_i (a_{-i}^*, a_i).$$

for each player i and each action $a_i \in A_i$ such that $a_i \neq a_i^*$. Calculate the Nash equilibria of the game

	L	M	R
T	$(1,1)$	$(1,0)$	$(0,0)$
B	$(1,0)$	$(0,1)$	$(1,2)$

Is any of them a strict Nash equilibrium?

2.5 *Guess 2/3 of the average* is a game where each of n players selects a real number from the interval $[0, 100]$, and a cash prize is split evenly among the players whose numbers are closest to 2/3 of the average of the n numbers chosen. Show that the game's unique Nash solution consists of all players choosing the number 0.

2.6 Calculate the Nash equilibria in the Guess 2/3 of the average game if the players are only allowed to choose integers in the interval $[0, 100]$.

2.3 Existence of Nash equilibria

Matching Pennies is an example of a finite strategic game with no Nash equilibrium, and finite games without Nash equilibria are not uncommon. There are, however, important classes of finite games that always have Nash equilibria and an example of such a class is given by the so called finite extensive games with perfect information that we will study in Chapter 8.

In this section, we present an existence theorem of Nash. The theorem is not directly applicable to finite games because the action sets in Nash's theorem are assumed to be convex, and finite sets are of course not convex unless they are singleton sets. Nash's theorem may, however, be applied to the so-called mixed extension of a finite game. The action sets in the mixed extension of a finite game are lottery sets, and these sets are convex and compact. Mixed extensions of games will be the theme of the next chapter.

In order to formulate Nash's theorem we will need some convexity notions. First recall that a subset X of $\mathbf{R}^n$ is called *convex* if, for each pair of points in the set, the set contains the entire segment between the points, i.e. if

$$x, y \in X,\, 0 < \alpha < 1 \;\Rightarrow\; \alpha x + (1-\alpha)y \in X.$$

Definition 2.3.1 Let X be a convex subset of $\mathbf{R}^n$. A function $f\colon X \to \mathbf{R}$ is called *quasiconcave* if the sets

$$\{x \in X \mid f(x) \geq t\}$$

are convex for all $t \in \mathbf{R}$.

A preference relation $\succeq$ on X is called *quasiconcave* if the sets

$$\{x \in X \mid x \succeq a\}$$

are convex for all $a \in X$.

Proposition 2.3.1 *Preference relations $\succeq$ with quasiconcave utility functions u are quasiconcave.*

Proof. The asssertion is a direct consequence of the quasiconcavity definitions and the relationship

$$\{x \in X \mid x \succeq a\} = \{x \in X \mid u(x) \geq u(a)\}. \qquad \square$$

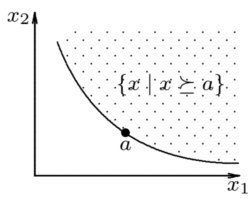

Figure 2.2. A preference relation $\succeq$ is quasiconcave if the sets $\{x \mid x \succeq a\}$ are convex for all a.

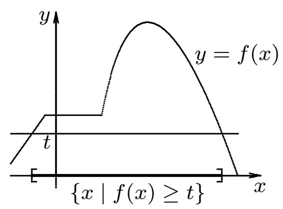

Figure 2.3. The graph of a quasiconcave one-variable function f. The sets $\{x \mid f(x) \geq t\}$ are intervals or empty for all t.

Proposition 2.3.2 *Suppose that X is a convex and compact subset of $\mathbf{R}^n$ and that $\succeq$ is a continuous and quasiconcave preference relation on X. The set of maximal elements of the preference relation is then nonempty, convex and compact.*

Proof. The set M of maximal elements is nonempty according to Proposition 1.2.4, and if m is a maximal element, then $M = \{x \in X \mid x \succeq m\}$. Hence, M is convex by Definition 2.3.1 and closed by Proposition 1.2.1. A closed subset of a compact set i compact. □

Proposition 2.3.3 (Nash's theorem) *A strategic game $\langle N, (A_i), (\succeq_i)\rangle$ has a Nash equilibrium if the following conditions are fulfilled for every player $i \in N$:*

- *The action set A_i is a convex compact subset of some $\mathbf{R}^{n_i}$.*
- *The preference relation $\succeq_i$ is continuous.*
- *The restriction of the preference relation $\succeq_i$ to the set $\{a_{-i}\} \times A_i$ is quasiconcave for each outcome $a \in A$.*

Remark. The last condition of the theorem means that the sets

$$\{(a_{-i}, x_i \in A) \mid (a_{-i}, x_i) \succeq_i a\}$$

are convex subsets of $\{a_{-i}\} \times A_i$ for every outcome a, and this holds if and only if the sets $\{x_i \in A_i \mid (a_{-i}, x_i) \succeq_i a\}$ are convex subsets of A_i for each $a \in A$.

Proof. It follows from Proposition 2.3.2 that the set M of all maximal elements (a_{-i}, m_i) in $\{a_{-i}\} \times A_i$ of the preference relation $\succeq_i$ is a nonempty, convex and compact subset of $\{a_{-i}\} \times A_i$. The best-response set $B_i(a_{-i})$ of player i, which equals the projection of M on the ith factorA_i, is therefore nonempty, convex and compact, too.

Let $A = A_1 \times A_2 \times \cdots \times A_n$, and define a map ϕ from A to the set $\mathcal{P}(A)$ of all subsets of A by putting

$$\phi(a) = B_1(a_{-1}) \times B_2(a_{-2}) \times \cdots \times B_n(a_{-n}).$$

A Nash equilibrium is by Proposition 2.2.1 a point $a^* \in A$ with the property $a^* \in \phi(a^*)$, and such a point is called a *fixed-point* of the map ϕ. To prove the existence of a fixed-point of ϕ we will use Kakutani's fixed-point theorem which is stated below.

To this end, we note that

1. A is a nonempty, compact and convex subset of $\mathbf{R}^m$ for some m.
2. For each $a \in A$, $\phi(a)$ is a nonempty, convex subset of A.
3. The graph
$$G = \{(a, b) \mid a \in A \text{ and } b \in \phi(a)\}$$
of the map ϕ is a closed set.

Assertion 1 follows from the assumption that each factor A_i is compact and convex, and assertion 2 follows from the fact that each factor $B_i(a_{-i})$ in the definition of $\phi(a)$ is nonempty and convex.

A set is closed if and only if the limit point of every convergent sequence of points in the set belongs to the set. So to prove assertion 3, we assume that $\big((a^k, b^k)\big)_1^\infty$ is a convergent sequence of points in the graph G with limit point (a, b). This means that $a^k \in A$ and $b_k \in \phi(a^k)$ for all k, and that $a^k \to a$ and $b^k \to b$ as $k \to \infty$. Since A is a closed set, the limit a belongs to A. Moreover, by definition b_i^k belongs to the best-response set $B_i(a_{-i}^k)$ of player i for each i, which means that

$$(a_{-i}^k, b_i^k) \succeq_i (a_{-i}^k, y_i)$$

for all $y_i \in A_i$. Now let $k \to \infty$ in the above relation; by continuity

$$(a_{-i}, b_i) \succeq_i (a_{-i}, y_i)$$

for all $y_i \in A_i$, which shows that $b_i \in B_i(a_{-i})$ for all i, and that consequently $b \in \phi(a)$. Hence, the limit point (a, b) lies in G, and this concludes the proof of assertion 3.

The existence of a fixed-point of the map ϕ is now an immediate consequence of Kakutani's fixed-point theorem, which reads as follows and is proven in Appendix 2 of Part II on cooperative games. □

Theorem 2.3.4 (Kakutani's fixed-point theorem) *Let ϕ be a set-valued map that is defined on a nonempty subset X of $\mathbf{R}^n$, and suppose that the following conditions are fulfilled:*

- *X is a compact and convex set;*
- *$\phi(x)$ is a nonempty, convex subset of X for each $x \in X$;*
- *ϕ has a closed graph.*

Then ϕ has a fixed-point, i.e. there is a point $x^* \in X$ such that $x^* \in \phi(x^*)$.

As a further application of Kakutani's fixed-point theorem, we now prove a result for symmetric games.

Definition 2.3.2 A strategic two-person game $\langle\{1, 2\}, (A_i), (\succeq_i)\rangle$ is *symmetric* if $A_1 = A_2$ and

$$(a_1, a_2) \succeq_1 (b_1, b_2) \;\Leftrightarrow\; (a_2, a_1) \succeq_2 (b_2, b_1)$$

for all (a_1, a_2), $(b_1, b_2) \in A_1 \times A_2$.

In a symmetric game both players have the same action set, and player 1 values the outcome (a_1, a_2) in the same way as player 2 values the outcome (a_2, a_1). If the preferences are given by the utility function u_1 and u_2, respectively, we may therefore assume that $u_1(a_1, a_2) = u_2(a_2, a_1)$ for all outcomes (a_1, a_2). The best-response sets of the two players satisfy $B_1(x) = B_2(x)$ for all $x \in A_1$.

Proposition 2.3.5 *Consider a symmetric strategic two-person game, and suppose that the common action set A of the two players is convex and compact, that the preference relation $\succeq_1$ of player 1 is continuous and that the restriction of $\succeq_1$ to the set $A \times \{a\}$ is quasiconcave for each $a \in A$. Then the game has a* symmetric Nash equilibrium*, i.e. a Nash equilibrium of the form (a^*, a^*).*

Proof. Player 1's best-response function $B_1 \colon A \to \mathcal{P}(A)$ meets the assumptions of Kakutani's fixed-point theorem. Hence, B_1 has a fixed-point $a^* \in A$, and since the game is symmetric, $a^* \in B_1(a^*)$ implies that the outcome (a^*, a^*) is a Nash solution. □

Exercises

2.7 Is the lexicographic order on $\mathbf{R}^2$ quasiconcave?

2.8 Decide whether the following set-valued maps ϕ, defined on the interval $[0,1]$, have closed graphs.

a) $\phi(x) = [0, x]$

b) $\phi(x) = [0, x[$

c) $\phi(x) = \begin{cases} [0, x] & \text{if } 0 \le x < 1, \\ \{0\} & \text{if } x = 1 \end{cases}$

d) $\phi(x) = \begin{cases} [\frac{3}{4}, 1] & \text{if } 0 \le x < \frac{1}{2}, \\ [x, 1] & \text{if } \frac{1}{2} \le x \le 1 \end{cases}$

e) $\phi(x) = \begin{cases} [0, 1] & \text{if } x = 0, \\ \{\sin \frac{1}{x}\} & \text{if } 0 < x \le 1 \end{cases}$

f) $\phi(x) = \begin{cases} [-1, 1] & \text{if } x = 0, \\ \{\sin \frac{1}{x}\} & \text{if } 0 < x \le 1 \end{cases}$

2.4 Maxminimization

In every finite strategic game with preferences given by utility functions, there is for each player a utility amount, the player's safety level, that the player can guarantee himself regardless of how the other players behave.

Example 2.4.1 Consider the game in Example 2.2.6 with the following payoff table:

	k_1	k_2	k_3	k_4	k_5	k_6
r_1	$(2,1)$	$(4,3)$	$(7,2)$	$(7,4)$	$(0,5)$	$(3,2)$
r_2	$(4,0)$	$(5,4)$	$(1,6)$	$(0,4)$	$(0,3)$	$(5,1)$
r_3	$(1,3)$	$(5,3)$	$(3,2)$	$(4,1)$	$(1,0)$	$(4,3)$
r_4	$(4,3)$	$(2,5)$	$(4,0)$	$(1,0)$	$(1,5)$	$(2,1)$

Let us study how player 2 should play if he is extremely pessimistic and assumes that whatever action he takes, the opponent always happens to choose an action that is most unfavorable for him.

If player 2 chooses k_1, then the worst payoff 0 is obtained when player 1 chooses r_2. If he chooses k_2, the worst payoff is 3, obtained if player 1 chooses r_1 or r_3. The worst payoff the actions k_3, k_4 and k_5 yield is 0, and the worst payoff that action k_6 yields is 1. So by choosing action k_2, player 2 can at least guarantee himself a payoff of 3 utility units, independently of his opponent's choice of action, and this is also the largest payoff he can guarantee himself. Thus, 3 is his safety level, and the action k_2, which guarantees him at least this amount, is called his maxminimizing action. The action k_2 maximizes the function

$$f_2(y) = \min_{x \in A_1} u_2(x, y).$$

Analogously, we find player 1's safety level and maxminimizing action by maximizing the function

$$f_1(x) = \min_{y \in A_2} u_1(x, y)$$

over all $x \in A_1$. Player 1 has two maxminimizing actions, namely r_3 and r_4, which both guarantee him utility 1, the safety level of player 1.

The outcomes (r_3, k_2) and (r_4, k_2), which occur if both players choose maxminimizing actions, result in the payoffs $(5, 3)$ and $(2, 5)$, respectively. This happens to be better for player 1 than his safety level in both cases, whereas player 2 exceeds his safety level only in the second case. □

The reasoning in Example 2.4.1 can be generalized and leads to the following general definition.

Definition 2.4.1 Let $\langle N, (A_i), (u_i)\rangle$ be a strategic game and define for each player i a function f_i on A_i by

$$f_i(a_i) = \inf\{u_i(x_{-i}, a_i) \mid x_{-i} \in A_{-i}\}.$$

The quantity

$$\ell_i = \sup\{f_i(a_i) \mid a_i \in A_i\}$$

is called the *safety level* of player i. If there is an action $\hat{a}_i \in A_i$ such that the supremum is attained, i.e. such that $f_i(\hat{a}_i) = \ell_i$, then we call $\hat{a}_i$ a *maxminimizing action* for player i.

If player i has a maxminimizing action, then we are of course allowed to write $\ell_i = \max\{f_i(a_i) \mid a_i \in A_i\}$ instead of $\ell_i = \sup\{f_i(a_i) \mid a_i \in A_i\}$.

A player's safety level is not necessarily a finite number but could be equal to $+\infty$ or $-\infty$, and finiteness does of course not guarantee that the supremum is attained, i.e. that there exists a maxminimizing action. In any case, given any number $s < \ell_i$, player i has an action which guarantees him a payoff greater than s, and if he has a maxminimizing action $\hat{a}_i$, then this action will guarantee him a payoff that is greater than or equal to ℓ_i.

The games that interest us the most are either finite games or games with compact action sets A_i and continuous utility functions u_i. In these cases, the infima and the suprema in the definitions of the functions f_i and the safety levels ℓ_i are of course attained, and we can replace inf with min and sup with max everywhere, and each player certainly has a maxminimizing action.

We have defined a player's maxminimizing action $\hat{a}_i$ using the player's utility function u_i, but since the relation between two ordinal utility functions that represent the same preference relation is given by a strictly increasing transformation, the maxminimizing action $\hat{a}_i$ will in fact only depend on the underlying preference relation $\succeq_i$. The safety level will of course depend on the choice of utility function − if $u'_i = F \circ u_i$ and a player's safetly level is equal to ℓ'_i with respect to the utility function u'_i and equal to ℓ_i with respect to the utility function u_i, then $\ell'_i = F(\ell_i)$.

Proposition 2.4.1 *Assume that the strategic game $\langle N, (A_i), (u_i)\rangle$ has a Nash equilibrium a^*, and let ℓ_i be the safety level of player i. Then*

$$u_i(a^*) \geq \ell_i.$$

A player's payoff at a Nash equilibrium is in other words at least as large as his safety level.

Proof. We have the inequality

$$f_i(a_i) = \inf\{u_i(x_{-i}, a_i) \mid x_{-i} \in A_{-i}\} \leq u_i(a^*_{-i}, a_i)$$

for each $a_i \in A_i$, and this implies that

$$\ell_i = \sup\{f_i(a_i) \mid a_i \in A_i\} \leq \sup\{u_i(a^*_{-i}, a_i) \mid a_i \in A_i\} = u_i(a^*_{-i}, a^*_i) = u_i(a^*).$$

□

EXAMPLE 2.4.2 The game in Exempel 2.4.1 has two Nash solutions (r_3, k_2) and (r_4, k_5), at which the payoffs to the players are given by the pairs $(5, 3)$ and $(1, 5)$, respectively. Both solutions give the two players payoffs that are at least as large as their safety levels, which are equal to 1 and 3, respectively. □

Exercises

2.9 Compute Nash solutions, maxminimizing actions and safety levels for the following game:

	k_1	k_2	k_3	k_4
r_1	$(3, 4)$	$(2, 1)$	$(0, 6)$	$(1, 4)$
r_2	$(2, 0)$	$(2, 2)$	$(3, 0)$	$(1, 2)$
r_3	$(2, 3)$	$(5, 1)$	$(1, 1)$	$(3, 3)$
r_4	$(4, 3)$	$(2, 5)$	$(2, 2)$	$(3, 2)$

2.10 Calculate maxminimization actions and safety levels for the game Splitting money in Example 2.2.7.

2.11 Prove that there is no Nash equilibrium in a game with a player whose safety level is $+\infty$.

2.5 Strictly competitive games

The strategic game concept is too general to allow any interesting far-reaching results about Nash solutions. To obtain such results we have to confine ourselves to some smaller class of games, and such an important class consists of games for two persons whose interests are diametrically opposed.

Definition 2.5.1 A strategic game $\langle\{1,2\},(A_i),(\succeq_i)\rangle$ with two players is called *strictly competitive* if

$$(1) \qquad (a_1,a_2) \succeq_1 (b_1,b_2) \;\Leftrightarrow\; (b_1,b_2) \succeq_2 (a_1,a_2)$$

for all outcomes (a_1,a_2) and (b_1,b_2).

EXAMPLE 2.5.1 The game Matching Pennies in Example 2.1.3 is strictly competitive, because the preferences of the two players fulfill the following conditions

$$(Kr, Kr) \sim_1 (Kl, Kl) \succ_1 (Kr, Kl) \sim_1 (Kl, Kr)$$
$$(Kl, Kr) \sim_2 (Kr, Kl) \succ_2 (Kl, Kl) \sim_2 (Kr, Kr)$$

which means that the condition (1) is satisfied. □

Suppose that the preference relation $\succeq_1$ is represented by the utility function u_1. It then follows from condition (1) that

$$(b_1,b_2) \succeq_2 (a_1,a_2) \;\Leftrightarrow\; u_1(a_1,a_2) \geq u_1(b_1,b_2\ ) \Leftrightarrow\; -u_1(b_1,b_2) \geq -u_1(a_1,a_2),$$

which means that $-u_1$ is a utility function representing $\succeq_2$. In a strictly competitive game it is thus possible to choose the players' ordinal utility functions u_1 and u_2 so that $u_1 + u_2 = 0$.

Definition 2.5.2 A two-person strategic game $\langle\{1,2\},(A_i),(u_i)\rangle$ with cardinal utility functions is called a *zero-sum game* if $u_1 + u_2 = 0$.

Zero-sum games are strictly competitive. Every outcome gives one player the same amount of utility units as the other player loses. Thus, we can regard the outcome as a transfer of utility from one player to the other; no utility is supplied from the outside or disappears,

In zero-sum games it is obviously sufficient to specify the utility function for one of the players. For finite zero-sum games we will do so by giving the payoff matrix of player 1 (the row player), and we call this matrix the *game's payoff matrix.*

The Nash equilibria in a strictly competitive game are completely determined by the preferences of one of the players. It follows immediately from the definition of Nash equilibrium and the equivalence (1) that the outcome (a_1^*, a_2^*) is a Nash equilibrium if and only if

$$(a_1^*, a_2) \succeq_1 (a_1^*, a_2^*) \succeq_1 (a_1, a_2^*)$$

for all $(a_1, a_2) \in A_1 \times A_2$.

This condition can of course also be expressed in terms of utility functions; if u_1 is a utility function for player 1, then (a_1^*, a_2^*) is a Nash equilibrium if and only if the *saddle point inequality*

$$u_1(a_1^*, a_2) \geq u_1(a_1^*, a_2^*) \geq u_1(a_1, a_2^*) \tag{2}$$

holds for all $(a_1, a_2) \in A_1 \times A_2$.

In a strictly competitive game $\langle\{1,2\}, (A_i), (u_i)\rangle$, every Nash equilibrium is in other words a saddle point of player 1's utility function u_1. Conversely, every saddle point (a_1^*, a_2^*) of u_1 with the property that a_1^* maximizes the function $u_1(\cdot, a_2^*)$ and a_2^* minimizes the function $u_1(a_1^*, \cdot)$, is a Nash equilibrium.

To verify the saddle point inequality it is of course enough to check that

$$u_1(a_1^*, a_2) \geq u_1(a_1, a_2^*)$$

for all outcomes (a_1, a_2), because by taking $a_1 = a_1^*$ we get the left part and by taking $a_2 = a_2^*$ we get the right part of inequality (2).

In a strictly competitive game $\langle\{1,2\}, (A_i), (u_i)\rangle$, the maxminimizing actions of both players are also completely determined by one player's utility function. A maxminimizing action for player 1 is by definition an action $\hat{a}_1 \in A_1$ such that

$$\inf_{a_2 \in A_2} u_1(\hat{a}_1, a_2) = \sup_{a_1 \in A_1} \inf_{a_2 \in A_2} u_1(a_1, a_2),$$

and a maxminimizing action for player 2 is similarly an action $\hat{a}_2 \in A_2$ such that

$$\inf_{a_1 \in A_1} u_2(a_1, \hat{a}_2) = \sup_{a_2 \in A_2} \inf_{a_1 \in A_1} u_2(a_1, a_2),$$

but we can replace the function u_2 with the function $-u_1$ in the last equality, because $-u_1$ is, as noted above, a utility function for player 2. Using the general relation $\inf -f = -\sup f$ between the supremum of a function f and the infimum of its negative, we conclude that the maxminimizing action $\hat{a}_2$ is also obtained as the solution to the problem

$$\sup_{a_1 \in A_1} u_1(a_1, \hat{a}_2) = \inf_{a_2 \in A_2} \sup_{a_1 \in A_1} u_1(a_1, a_2).$$

This is the equality to be used in the sequel in order to characterize the maxminimizing actions of player 2 in strictly competitive games.

We now turn to the relation between Nash equilibria and maxminimizing actions in strictly competitive games.

Proposition 2.5.1 *Let $\hat{a}_1$ and $\hat{a}_2$ be maxminimizing actions of players 1 and 2, respectively, in a strictly competitive game $\langle \{1, 2\}, (A_i), (u_i) \rangle$ and assume that*

$$\inf_{a_2 \in A_2} u_1(\hat{a}_1, a_2) = \sup_{a_1 \in A_1} u_1(a_1, \hat{a}_2).$$

The outcome $(\hat{a}_1, \hat{a}_2)$ is then a Nash equilibrium.

Proof. Under the given assumptions we have

$$u_1(\hat{a}_1, a_2) \geq \inf_{a_2 \in A_2} u_1(\hat{a}_1, a_2) = \sup_{a_1 \in A_1} u_1(a_1, \hat{a}_2) \geq u_1(a_1, \hat{a}_2)$$

for all outcomes (a_1, a_2), and this implies that the saddle point inequality is satisfied by the outcome $(\hat{a}_1, \hat{a}_2)$. □

Proposition 2.5.2 *Suppose that* (a_1^*, a_2^*) *is a Nash equilibrium in the strictly competitive strategic game* $\langle\{1,2\},(A_i),(u_i)\rangle$. *Then*

(i) a_1^* *and* a_2^* *are maxminimizing actions of players 1 and 2, respectively;*

(ii) $\sup_{a_1\in A_1}\inf_{a_2\in A_2} u_1(a_1,a_2) = \inf_{a_2\in A_2}\sup_{a_1\in A_1} u_1(a_1,a_2)$;

(iii) $u_i(a_1^*, a_2^*)$ *is the safety level of player* i *for* $i = 1, 2$.

Proof. We start by proving the inequality

$$\sup_{a_1\in A_1}\inf_{a_2\in A_2} u_1(a_1,a_2) \le \inf_{a_2\in A_2}\sup_{a_1\in A_1} u_1(a_1,a_2), \tag{3}$$

for arbitrary functions $u_1\colon A_1\times A_2 \to \mathbf{R}$. To this end, let $a_1 \in A_1$ be arbitrary and take the greatest lower bound of the function $u_1(a_1,\cdot)$. This results in the inequality

$$\inf_{a_2\in A_2} u_1(a_1,a_2) \le u_1(a_1,a_2),$$

which holds for all $a_2 \in A_2$ and all $a_1 \in A_1$. By considering both sides of this inequality as functions of a_1 and comparing their least upper bounds, we conclude that

$$\sup_{a_1\in A_1}\inf_{a_2\in A_2} u_1(a_1,a_2) \le \sup_{a_1\in A_1} u_1(a_1,a_2)$$

for all $a_2 \in A_2$. The right side of this inequality is a function of a_2 and the greatest lower bound of this function is by definition greater than or equal to the the left side, which is exactly what inequaility (3) says.

Using the saddle point inequality (2) for the Nash equilibrium (a_1^*, a_2^*), we obtain the following chain of inequalities

$$\begin{aligned}\sup_{a_1\in A_1}\inf_{a_2\in A_2} u_1(a_1,a_2) &\ge \inf_{a_2\in A_2} u_1(a_1^*,a_2) = u_1(a_1^*,a_2^*)\\ &= \sup_{a_1\in A_1} u_1(a_1,a_2^*) \ge \inf_{a_2\in A_2}\sup_{a_1\in A_1} u_1(a_1,a_2).\end{aligned}$$

Combining this inequality with the converse inequality (3) we conclude that the extreme ends in the above inequality are equal. Hence, equality prevails everywhere. Thus,

$$\inf_{a_2\in A_2} u_1(a_1^*,a_2) = \sup_{a_1\in A_1}\inf_{a_2\in A_2} u_1(a_1,a_2) = u_1(a_1^*,a_2^*)$$

and

$$\sup_{a_1\in A_1} u_1(a_1,a_2^*) = \inf_{a_2\in A_2}\sup_{a_1\in A_1} u_1(a_1,a_2) = u_1(a_1^*,a_2^*),$$

which proves statements (i) and (ii) and that $u_1(a_1^*, a_2^*)$ is the safety level of player 1. For symmetry reasons, the safety level for player 2 is equal to $u_2(a_1^*, a_2^*)$. □

Our next proposition is a corollary to the previous two propositions.

Proposition 2.5.3 *The following holds for strictly competitive games with at least one Nash equilibrium.*

(i) An outcome (a_1^*, a_2^*) is a Nash equilibrium if and only if a_1^* and a_2^* are maxminimizing actions of player 1 and player 2, respectively.
(ii) The payoff to a player is the same at all Nash equilibria of the game.
(iii) If (a_1^*, a_2^*) and $(\hat{a}_1, \hat{a}_2)$ are two Nash equilibria, then the two outcomes $(a_1^*, \hat{a}_2)$ and $(\hat{a}_1, a_2^*)$ are also Nash equilibria.

Proof. (i) The actions in a Nash equilibrium are maxminimizing actions of the respective players according to Proposition 2.5.2 (i). Moreover, since the game is assumed to have at least one Nash equilibrium, we know from statement (ii) of the same proposition that

$$\sup_{a_1 \in A_1} \inf_{a_2 \in A_2} u_1(a_1, a_2) = \inf_{a_2 \in A_2} \sup_{a_1 \in A_1} u_1(a_1, a_2),$$

and this means that

$$\inf_{a_2 \in A_2} u_1(\hat{a}_1, a_2) = \sup_{a_1 \in A_1} u_1(a_1, \hat{a}_2),$$

whenever $\hat{a}_1$ is a maximizing action for player 1 and $\hat{a}_2$ is a maximizing action for player 2. It now follows from Proposition 2.5.1 that each par $(\hat{a}_1, \hat{a}_2)$ of maxminimizing actions of the two players is a Nash equilibrium.

(ii) By Proposition 2.5.2, a player's payoff at a Nash equilibrium is equal to his safety level.

(iii) Let (a_1^*, a_2^*) and $(\hat{a}_1, \hat{a}_2)$ be two Nash equilibria. Statement (i) tells us that a_1^* and $\hat{a}_1$ are maxminimizing actions of player 1, and a_2^* and $\hat{a}_2$ are maxminimizing actions of player 2, and that consequently the two outcomes $(a_1^*, \hat{a}_2)$ and $(\hat{a}_1, a_2^*)$ are Nash equilibria. □

EXAMPLE 2.5.2 The two-person game G, given by the following payoff table

	k_1	k_2	k_3	k_4	
r_1	$(2,-2)$	$(-5,5)$	$(7,-7)$	$(-1,1)$	-5
r_2	$(3,-3)$	$(4,-4)$	$(3,-3)$	$(5,-5)$	3
r_3	$(1,-1)$	$(6,-6)$	$(-3,3)$	$(2,-2)$	-3
	3	6	7	5	

is strictly competitive. The numbers in the rightmost column are the numbers $\min_j u_1(r_i, k_j)$ and the last row contains the numbers $\max_i u_1(r_i, k_j)$.

The largest of the numbers $\min_j u_1(r_i, k_j)$, i.e. 3, is equal to the smallest of the numbers $\max_i u_1(r_i, k_j)$. The row player's maxminimizing action r_2 and column player's maxminimizing action k_1 thus satisfy the condition in Proposition 2.5.1, and we conclude that (r_2, k_1) is the unique Nash equilibrium of the game.

Assuming cardinal utilities, G is a zero-sum game with payoff matrix

$$U = \begin{bmatrix} 2 & -5 & 7 & -1 \\ 3 & 4 & 3 & 5 \\ 1 & 6 & -3 & 2 \end{bmatrix}.$$

(Remember our convention for finite zero-sum games that the row player's payoff matrix is the game's payoff matrix.) It is of course easy to compute possible Nash equilibria by just looking at the payoff matrix. An outcome (r_i, k_j) satisfying the saddle point inequality (2) corresponds to a location (i, j) in the payoff matrix U where the element u_{ij} is largest in its column and smallest in its row. Here, we see that the number 3 at location $(2, 1)$ satisfies this requirement. Hence, (r_2, k_1) is a Nash equilibrium and there are no others. □

Exercises

2.12 Compute the Nash equilibria of the zero-sum games with the following payoff matrices:

a) $\begin{bmatrix} 1 & 2 & 4 \\ 3 & 1 & 3 \\ 5 & 3 & 7 \end{bmatrix}$ b) $\begin{bmatrix} 1 & 2 & 4 \\ 3 & 3 & 1 \\ 2 & 2 & 3 \end{bmatrix}$.

2.13 Construct a competitive game with four Nash equilibria where both players have three options each.

Chapter 3

Two Models of Oligopoly

The competition between firms in a market can be modeled as a strategic game. In this section, we study two models of oligopoly from the mid-1800s. The first is due to the economist Augustin Cournot.

3.1 Cournot's model of oligopoly

Oligopoly is a market form where a number of firms are in competition. Normally, the number of firms is assumed to be relatively small, but this limitation does not matter in our discussion. So consider a situation where a homogeneous product is offered to the market by n firms. The cost of firm i to manufacture, market and sell q_i units of the product is $C_i(q_i)$, where C_i is an increasing function. All products are sold for a single price, and the price per unit $P(Q)$ depends on the total output $Q = q_1 + q_2 + \cdots + q_n$. An increased supply pushes the price down, so we assume that the function P, commonly called the *inverse demand function*, is decreasing.

The firms are assumed to be profit maximizing, but the catch is, of course, that a firm's profit depends not only on its own output but also on the output of the other firms. The revenue for firm i, which produces and sells q_i units, is $q_iP(q_1 + q_2 + \cdots + q_n)$, and its profit is thus

$$V_i(q_1, q_2, \ldots, q_n) = q_iP(q_1 + q_2 + \cdots + q_n) - C_i(q_i).$$

The monopoly case $n = 1$ with just one firm in the market is special. A monopolistic firm has full control of the situation, and its problem is simply to maximize the one-variable function $V(x) = xP(x) - C(x)$, where C is the firm's cost function.

An oligopoly with $n \geq 2$ firms involved may be perceived as a strategic game in which the players are the firms, the action set A_i of each firm i is the set of its possible outputs, and the payoff of each firm is its profit function V_i. For simplicity, we assume that $A_i = \mathbf{R}_+$, the set of all non-negative real numbers, for each firm i.

This strategic game has, under appropriate assumptions, a Nash solution, and the Nash solution was identified by Cournot as a reasonable solution to the firms' problems of determining their optimal output quantities. Of course, Cournot could not make use of any game theory terminology in his analysis of the situtation. Nash seems, on the other hand, not to have known Cournot's economical works when he developed his theory around 1950.

We can determine possible Nash equilibria $q^* = (q_1^*, q_2^*, \dots, q_n^*)$ in the oligopoly game by utilizing the fact that q_i^* is the best response of firm i, given that the other firms have chosen the output quantities q_{-i}^*. The outcome q^* is, in other words, a Nash equilibrium if and only if, for each company i, q_i^* maximizes the profit function

$$x_i \mapsto V_i(q_{-i}^*, x_i) = x_i P(q_1^* + \dots + x_i + \dots + q_n^*) - C_i(x_i).$$

Let us assume that $q_i^* > 0$ and that the functions P and C_i are differentiable at the points Q^* and q_i^*, respectively, where $Q^* = q_1^* + \dots + q_n^*$. The derivative of firm i's profit function is then equal to zero at the maximum point q_i^*, so it follows after differentiation that the Nash equilibrium q^* is a solution to the following system of equations:

$$\text{(1)} \qquad \begin{cases} P(Q) + q_1 P'(Q) = C_1'(q_1) \\ P(Q) + q_2 P'(Q) = C_2'(q_2) \\ \qquad \vdots \\ P(Q) + q_n P'(Q) = C_n'(q_n) \\ q_1 + q_2 + \dots + q_n = Q \end{cases}$$

From now on we assume that the firms have identical linear cost functions, i.e. that $C_i(q_i) = cq_i$ for some positive constant c, and that $P'(Q^*) < 0$. The inverse demand function P is in particular strictly decreasing at the point Q^*. Under these conditions, the Nash equilibrium q^* has to be symmetric, i.e. $q_1^* = q_2^* = \dots = q_n^*$.

To prove this, we need only subtract the first equation of the system (1) from equation number i, which results in the following equation for the Nash equilibrium

$$(q_i^* - q_1^*)P'(Q^*) = c - c = 0,$$

with the conclusion that $q_i^* - q_1^* = 0$ for all firms i.

The firms' total output at the Nash equilibrium is therefore $Q^* = nq_1^*$, and we can now compute q_1^* by solving the equation

$$\text{(2)} \qquad P(nx) + xP'(nx) = c,$$

an equation in just one variable x.

EXAMPLE 3.1.1 Let us determine the Nash equilibrium in Cournot's oligopoly game when the firms have a common linear cost function $C(q_i) = cq_i$ and the inverse demand function is of the form

$$P(Q) = \begin{cases} a - Q & \text{if } 0 \le Q \le a, \\ 0 & \text{if } Q > a. \end{cases}$$

We assume that $a > c$, for otherwise, all outputs are unprofitable.

All output vectors $q = (q_1, q_2, \dots, q_n)$ with total output $Q > a - c$ are nonprofitable for each firm i with a positive output q_i, and such a firm can improve the situation unilaterally by reducing its output to zero. Hence, the total output Q^* is less than $a - c$ at a Nash equilibrium q^*, which means that the inverse demand function has the form $P(Q) = a - Q$ in a neighborhood of the equilibrium point.

Let us first treat the monopoly case $n = 1$. The single firm's profit when $Q = q_1 \le a - c$ is now given by the expression

$$V_1(q_1) = q_1(a - q_1) - cq_1 = q_1(a - c - q_1),$$

and the graph of the profit function is shown in Figure 3.1. $V_1(q_1)$ is a quadratic polynomial with maximum at the point $q_{\text{mon}} = \frac{1}{2}(a - c)$. This is

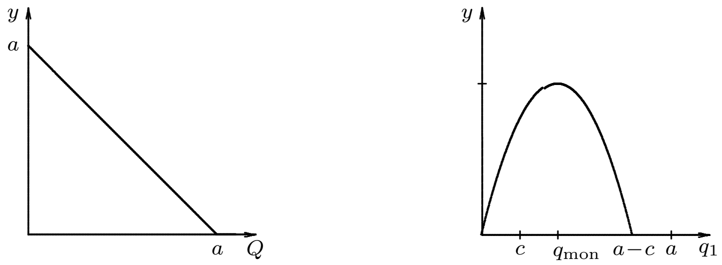

Figure 3.1. The left panel shows the graph of the inverse demand function $P(Q)$ in Example 3.1.1. The righ panel shows the profit curve $y = V_1(q_1)$ for a monopolistic firm.

exactly the solution we get from equation (2), which for $n = 1$, $P(Q) = a - Q$ and $C(q_1) = cq_1$ has the form $a - x - x = c$. The monopolistic firm's optimal price $P(q_{\text{mon}})$ is equal to $\frac{1}{2}(a + c)$ and the firm's profit is $\frac{1}{4}(a - c)^2$.

We now turn to the oligopoly case with $n > 1$ firms. (The duopoly case $n = 2$ is illustrated by Figure 3.2.) Our initial analysis is valid because the firms have identical cost functions. All coordinates q_j^* are therefore equal at a Nash equilibrium q^*, and they are obtained as solutions to equation (2), which now reads as $a - nx - x = c$. Each firm's output is therefore equal to $\frac{1}{n+1}(a - c)$ at the Nash equilibrium, the total output is $Q^* = \frac{n}{n+1}(a - c)$, the output price is $P(Q^*) = \frac{1}{n+1}(a + nc)$, and the total profit for all firms is $\frac{n}{(n+1)^2}(a - c)^2$.

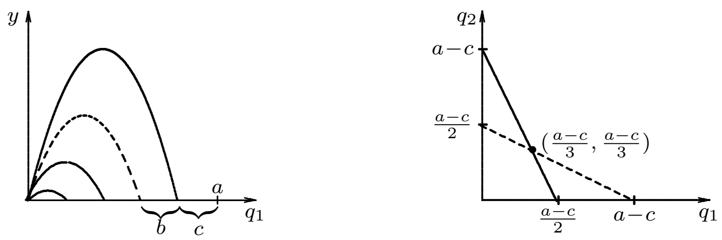

Figure 3.2. The left panel shows profit curves $y = V_1(q_1, q_2)$ for one of the duopoly firms for different values of the competing firm's output q_2; the outermost curve corresponds to $q_2 = 0$ and the dashed curve to $q_2 = b$. The righ panel shows the best response functions of the two firms; the solid line shows firm 1's best response as a function of q_2, the dashed line shows firm 2's best response as a function of q_1. The two lines intersection is the Nash equilibrium.

Note that the unit price $P(Q^*) = \frac{1}{n+1}(a+nc)$ decreases when the number n of firms increases, and that the price converges to c when the number of firms tends to infinity. The total output converges to $a-c$, and the total profit tends to 0. When the number of oligopoly firms becomes large, we have a market situation that is close to perfect competition.

By comparing the monopolistic firm's optimal price, output and profit with the oligopolistic firms' equilibrium price, total output and profit, we see that oligopoly implies a lower unit price, a greater total output and a lower total profit than what monopoly does. Oligopolists would, in other words, profit from forming a cartel, which would allow them to charge the monopoly price. That is why there are laws against cartels.

The conclusions in the preceding paragraph applies quite generally. We have the following result.

Proposition 3.1.1 *Consider a market, where the inverse demand function P has a negative derivative and the revenue function $x \mapsto xP(x)$ is concave. Oligopolistic firms with identical linear cost functions have a total output at the Nash equilibrium that is greater than the optimal output of a monopolistic firm with the same cost function.*

Proof. It follows from equation (2) that the total output Q^* $(= nq_1^*)$ at the Nash equilibrium of n oligopolistic firms is obtained as solution to the equation

$$P(x) + \frac{1}{n}xP'(x) = c,$$

and that the monopolistic firm's optimal output q_{mon} is a solution to the same equation in the case $n = 1$, i.e. to the equation

$$P(x) + xP'(x) = c.$$

The curve $y = P(x) + xP'(x)$ is decreasing, because the derivative of a concave function is decreasing, and $P(x) + xP'(x)$ is the derivative of the revenue function $xP(x)$. Since $P'(x) < 0$, we have $xP'(x) < \frac{1}{n}xP'(x)$ for $x > 0$, which implies that the decreasing curve $y = P(x) + xP'(x)$ lies below the curve $y = P(x) + \frac{1}{n}xP'(x)$. We conclude that the intersection between the curve $y = P(x) + xP'(x)$ and the line $y = c$ lies to the left of the intersection between the curve $y = P(x) + \frac{1}{n}xP'(x)$ and the same line. See figure 3.3. This means that the output q_{mon} from the monopolistic firm is less than the total output Q^* from the oligopolistic firms. The monopoly price is consequently higher than the oligopoly price. □

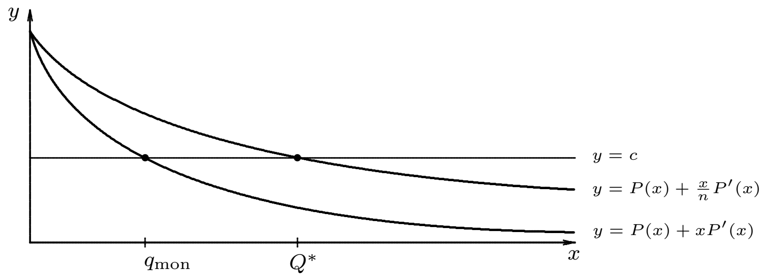

Figure 3.3. The output q_{mon} from a monopolistic firm is less than the total equilibrium output Q^* from the firms in an oligopoly.

The outputs given by the Nash equilibrium q^* are not necessarily the most profitable ones for oligopolistic firms. Under quite general conditions, all firms can in fact do better by slightly lowering their outputs.

Proposition 3.1.2 *Consider an oligopolistic market with arbitrary differentiable cost functions C_i, and suppose that q^* is a Nash equilibrium with positive coordinates q_i^* and that the derivative P' of the inverse demand function is negative at the point $Q^* = q_1^* + \cdots + q_n^*$. Then there exists a positive number ϵ such that every firm i has a bigger profit at outputs q_i satisfying $q_i^* - \epsilon < q_i < q_i^*$ than at the equilibrium output q_i^*.*

Proof. Consider the profit function for firm 1:

$$V_1(q_1, q_2, \ldots, q_n) = q_1 P(q_1 + q_2 + \cdots + q_n) - C_1(q_1).$$

Since q^* is a Nash equilibrium, the function $q_1 \mapsto V_1(q_1, q_2^*, \ldots, q_n^*)$ has a maximum when $q_1 = q_1^*$, and this implies that $\frac{\partial V_1}{\partial q_1}(q^*) = 0$. For $i \geq 2$ $\frac{\partial V_1}{\partial q_i}(q^*) = q_1^* P'(Q^*) < 0$, because of our assumptions concerning the sign of $P'(Q^*)$. Hence, all the coordinates of the gradient $\nabla V_1 = (\frac{\partial V_1}{\partial q_1}, \frac{\partial V_1}{\partial q_2}, \ldots, \frac{\partial V_1}{\partial q_n})$ of the profit function V_1 are strictly negative at the Nash equilibrium q^*, except for the first coordinate which is equal to 0.

Now, let $w = (w_1, w_2, \ldots, w_n)$ be a vector in $\mathbf{R}^n$, all of whose coordinates are *negative*, and consider the function

$$g(t) = V_1(q^* + tw) = V_1(q_1^* + w_1 t, \ldots, q_n^* + w_n t)$$

for t in a neighborhood of 0. (Compare with Figure 3.4.) By the chain rule,

$$g'(0) = \frac{\partial V_1}{\partial q_1}(q^*)w_1 + \frac{\partial V_1}{\partial q_2}(q^*)w_2 + \cdots + \frac{\partial V_1}{\partial q_n}(q^*)w_n = \nabla V_1(q^*) \cdot w.$$

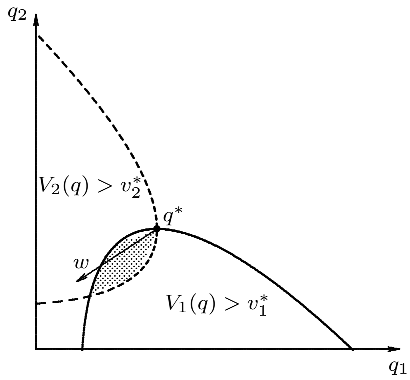

Figure 3.4. An illustration of how duopolistic firms profit from lowering their outputs from the Nash equilibrium output q^*. The profit $V_1(q)$ for firm 1 is equal to the firm's Nash equilibrium profit $v_1^* = V_1(q^*)$ along the solid curve, and it is bigger below the curve. The profit $V_2(q)$ for firm 2 is equal to its equilibrium profit $v_2^* = V_2(q*)$ along the dashed curve, and it is bigger to the left of the curve. Outputs in the shaded area give both firms bigger profits than the equilibrium profits.

We conclude that $g'(0) > 0$, because all terms $\frac{\partial V_1}{\partial q_j} w_j$ in the sum except the first are positive and the first is zero. The function g is therefore strictly increasing at $t = 0$, and this implies that there exists a positive number t_1 such that

$$V_1(q^* + tw) > V_1(q^*)$$

for all t in the interval $0 < t < t_1$. The same, of course, applies to the other profit functions V_i, and we conclude that all firms have bigger profits by the output $q^* + tw$ than by the equilibrium output q^*, provided $t > 0$ is a sufficiently small number. □

Exercises

3.1 Find the Nash equilibrium of Cournot's oligopoly model when the inverse demand function is $P(Q) = 2c(1 + Q)^{-1}$ and all firms have the same linear cost function $C(q_i) = cq_i$

3.2 Find the Nash equilibrium of Cournot's duopoly model if the inverse demand function is

$$P(Q) = \begin{cases} \frac{1}{4}Q^2 - 5Q + 26 & \text{when } 0 \le Q \le 10, \\ 1 & \text{when } Q \ge 10. \end{cases}$$

The cost of producing one unit is 1 for both firms.

3.3 Find the Nash equilibrium of Cournot's duopoly model when the inverse demand function is

$$P(Q) = \begin{cases} a - Q & \text{if } 0 \le Q \le a, \\ 0 & \text{if } Q > a \end{cases}$$

and the cost functions are:

a) $C_i(q_i) = c_i q_i$, where $0 < c_1 < c_2 < a$.

b) $C_i(q_i) = \begin{cases} 0 & \text{if } q_i = 0, \\ b + cq_i & \text{if } q_i > 0, \end{cases}$ where $b > 0$ and $0 < c < a$.

3.2 Bertrand's model of oligopoly

The output is the strategic variable in Cournot's modell; each firm chooses an output and the total output determines the price of the product. Bertrand misunderstood Cournot on that point in a review in 1883 of Cournot's book by perceiving the price of the product as the competitive element. For this reason, Bertrand's name is usually associated with oligopoly models using pricing as the firms' decision variable.

In the Bertrand model we assume, as in the Cournot model, that n firms produce an identical product, but the firms can now choose to price the product differently. The consumer demand for the product is determined entirely by the price of the product and the total demand at price p is equal to $D(p)$. The function D is called the *demand function*

Let p_i denote the price set by firm i, and let $p_{\min}$ denote the lowest price. If the firms set different prices, then all consumers buy the product from the firm with the lowest price, which produces just enough output to meet the demand. If more than one firm set the same lowest price, then all of those share the demand at that price equally. A firm whose price is higher than the lowest price sells nothing and produces no output.

We assume for simplicity that the cost of producing and selling a unit of the product is the same and equal to c for all firms.

Firms with price higher than the lowest price have no sales nor any costs, and firms with price equal to $p_{\min}$ share the total demand $D(p_{\min})$ equally without producing any surplus. The profit per unit is equal to $p_{\min} - c$, which of course is a loss if $p_{\min} < c$. The profit V_i for firm i is thus given by the formula

$$V_i(p_1, p_2, \dots, p_n) = \begin{cases} 0 & \text{if } p_i > p_{\min}, \\ \frac{1}{m}(p_i - c)D(p_i) & \text{if } p_i = p_{\min}, \end{cases}$$

where m is the number of firms setting the same lowest price $p_{\min}$.

We can view Bertrand's model as a strategic game with the n firms as players, the possible pricings as their action sets and the profit functions as their utility functions. We assume that all positive real numbers can be used as prices.

We can give a simple analysis of the Bertrand model under very general assumptions on the demand function D; the only assumptions needed are that it is positive and that the demand does not suddenly decrease drastically when the price drops. More particularly, we assume that for every $x > c$ there exist a constant $k > \frac{1}{2}$ and a number $y < x$ arbitrarily close to x such that $D(y) > kD(x)$. This certainly holds if the demand function is decreasing, which is a reasonable economic assumption, or if it is continuous. (We can choose $k = \frac{3}{4}$ in both cases.)

In Cournot's model the Nash equilibrium depends in a rather complicated way on the inverse demand function and the cost functions. In Bertrand's model, however, the picture is very simple:

The pricing $p^* = (p_1^*, p_2^*, \dots, p_n^*)$ is a Nash equilibrium in Bertrand's oligopoly model if and only if $p_{min}^* = c$ and $p_i^* = c$ for at least two firms. In a duopoly the pricing (c, c) is therefore the unique Nash equilibrium.

The result is not very encouraging for the firms; none of them makes a profit if they choose the pricing advocated by the Nash solution.

The proof of our claim about the Nash equilibrium p^* is simple. Firstly, the lowest price $p^*_{\min}$ can not be less than c, because a firm charging a price lower than c is losing money and profits from raising the price to c, which under all circumstances transforms the loss to break even.

Secondly, the lowest price can not be greater than c. If $p^*_{\min} > c$ and there is a firm i with a price p^*_i higher than $p^*_{\min}$, then this firm would profit from lowering its price to a price p_i in the interval $]c, p^*_{\min}[$, because this action would give the firm a monopoly on all sales and transform a zero result to a positive result.

If there is no such firm, i.e. if all firms set the same lowest price $p^*_{\min} > c$, any firm would profit from unilaterally lowering its price slightly, since it would in this way obtain the total profit instead of having to share it with all the other firms. (Here we need our assumption that the demand does not decrease drastically when the price drops!)

Hence, $p^*_{\min} = c$. Moreover, at least two firms must have the same lowest price, because if for example $c = p^*_1 < p^*_i$ for $i = 2, 3, \ldots, n$, then it is possible for firm 1 to increase its profit from 0 to a positive number by raising the price, but not more than that the new price p_1 is still the lowest price.

All pricings p^* with $p^*_{\min} = c$ and where at least two firms have the same lowest price are indeed Nash equilibria, because the profit is equal to zero for all firms, and no firm can obtain a positive profit by unilaterally changing its price.

Exercise

3.4 Consider Bertrand's duopoly model with a demand function D that is constant on the interval $[c, p_0[$ and makes a jump at the point p_0 so that $D(p_0) = 2D(c)$. (The assumption is not entirely unrealistic – if the price is below a certain point, the consumers might think that the product is not good, but if the price goes beyond this critical point, the demand rises sharply.) An example of such a function could be

$$D(x) = \begin{cases} 1 & \text{if } 0 < x < p_0, \\ 2p_0/x & \text{if } x \geq p_0. \end{cases}$$

Show that (p_0, p_0) is now a Nash equilibrium in addition to (c, c).

Chapter 4

Congestion Games and Potential Games

4.1 Congestion games

Congestion games is a class of games having Nash equilibra that was introduced and studied by Robert Rosenthal in 1973. We begin with an illustrative example.

Example 4.1.1 Three persons, Alice, Bella and Clara, will move from point P to point Q using the road network in Figure 4.1.

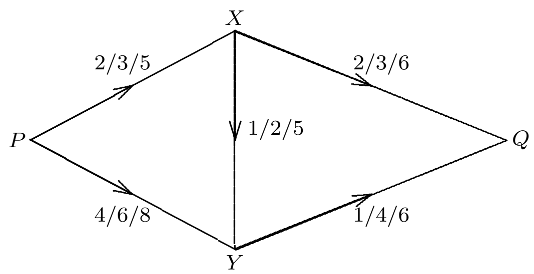

Figure 4.1. A road network that is used by three persons.

The cost (in time, for example) of using a road depends on congestion, i.e. the number of people using the road. The cost of using the road from P to X is 2 units if one person uses it, 3 units per person if two persons use it, and 5 units per person if three persons use it. We have marked this in the figure by writing 2/3/5 next to the road. The costs of the other roads PY, XY, XQ and YQ are specified in a similar way in the figure. The total travel cost for a person is obtained by adding the person's costs of the roads she uses during her trip from P to Q.

Alice and Bella are allowed to use any road from P to Q (in the direction of the arrows), but Clara has a vehicle that makes it impossible to use the road PY. All of them are interested in obtaining the lowest possible cost.

We can regard their problem as a strategic game in which Alice has three possibilities for action: the option x of using the roads PX and XQ, the option y of using the roads PY and YQ, and the option z of using the roads PX, XY and YQ. Bella has the same options x, y and z as Alice, but Clara can only use the options x and z.

If for example Alice and Bella both use the option z and Clara uses the option x, road PX will be used by three persons at a cost of 5 each, roads XY and YQ will be used by two persons at a cost of 2 and 4, respectively, for each user, and road XQ will be used by just one person (Clara) at a cost of 2 for her. Alice's total cost is 5+2+4= 11, Bella has the same cost, while Clara's cost is 5+2=7. The cost vector corresponding to the outcome (z, z, x) is thus equal to $(11, 11, 7)$.

Suppose that Alice unilaterally changes her route to alternative y. The cost vector will change to $(8, 8, 5)$, which is better for Alice (and also for the other two persons). Hence, the alternative (z, z, x) is not a Nash equilibrium.

Nor is (y, z, x), because Bella profits from unilaterally changing from z to x, as the outcome (y, x, x) results in the cost vector $(5, 6, 6)$. It is easy to verify that no one can now unilaterally improve the costs. Hence, (y, x, x) is a Nash equilibrium, and for symmetry reasons (x, y, x) is also a Nash equilibrium with corresponding cost vector $(6, 5, 6)$. □

We now give a general definition of congestion models like the one in Example 4.1.1.

Definition 4.1.1 A *congestion model* $M = \langle N, R, (A_i)_{i\in N}, (\kappa_j)_{j\in R}\rangle$ consists of

- a set $N = \{1, 2, \dots, n\}$ of players;
- a finite set R of *resources*;
- for each player i a nonempty set A_i of subsets of R;
- for each resource $r \in R$ a vector $\kappa_r = (\kappa_r(1), \kappa_r(2), \dots, \kappa_r(n))$ consisting of real numbers.

Example 4.1.2 The model in Example 4.1.1 is a congestion model with

- $N = \{1, 2, 3\} = \{$Alice, Bella, Clara$\}$;
- $R = \{PX, XQ, XY, PY, YQ\}$;
- $A_1 = A_2 = \{\{PX, XQ\}, \{PY, YQ\}, \{PX, XY, YQ\}\}$ and $A_3 = \{\{PX, XQ\}, \{PX, XY, YQ\}\}$
- $\kappa_{PX} = (2, 3, 5)$, $\kappa_{XQ} = (2, 3, 6)$, $\kappa_{XY} = (1, 2, 5)$, $\kappa_{PY} = (4, 6, 8)$ and $\kappa_{YQ} = (1, 4, 6)$. □

In the congestion model $M = \langle N, R, (A_i)_{i\in N}, (\kappa_j)_{j\in R}\rangle$, every element of the set A_i is thus a set of resources, and we interpret a player's choice of alternative $a_i \in A_i$ as a choice to use the resources contained in a_i

The number $\kappa_r(k)$ should be interpreted as each player's cost of using the resource r when exactly k players use the resource, i.e. when $r \in A_i$ for exactly k players i.

For each outcome $a = (a_1, a_2, \dots, a_n) \in A = A_1 \times A_2 \times \cdots \times A_n$ and each resource $r \in R$, we define

$$n_r(a) = \text{the number of } i \in N \text{ such that } r \in a_i,$$

i.e. $n_r(a)$ is the number of players that use the resource r.

The outcome a comes with a total cost $c_i(a)$ for player i, which is obtained by adding the costs of all the resources that player i uses so that

$$c_i(a) = \sum_{r\in a_i} \kappa_r(n_r(a)).$$

We assume that each player i tries to minimize his total cost $c_i(a)$, i.e. maximize $-c_i(a)$, which he will try to do as a participant in the strategic game $\langle N, (A_i)_{i\in N}, (-c_i)_{i\in N}\rangle$, which is called the *congestion game* associated with the congestion model M.

The congestion game in Example 4.1.1 has a Nash equilibrium, and this is not a coincidence, because we have the following general result.

Proposition 4.1.1 *Every congestion game has a Nash equilibrium.*

Proof. We use the notation in and after Definition 4.1.1, and we will prove that there is an outcome a^* such that

$$c_i(a^*_{-i}, a_i) \geq c_i(a^*)$$

for all players i and all actions $a_i \in A_i$. We will do so by constructing a function $\Phi\colon A \to \mathbf{R}$ such that

$$c_i(a_{-i}, x_i) - c_i(a) = \Phi(a_{-i}, x_i) - \Phi(a) \tag{1}$$

for all players i, all $a \in A$ and all $x_i \in A_i$.

Assume for a moment that we have such a function Φ. Since its domain of definition is finite, there exists a (not necessarily unique) minimum point $a^* \in A$. The outcome a^* is a Nash equilibrium, because

$$c_i(a^*_{-i}, a_i) - c_i(a^*) = \Phi(a^*_{-i}, a_i) - \Phi(a^*) \geq 0$$

for all players i and all actions $a_i \in A_i$ according to equation (1).

It remains to define the function Φ, which we do for $a \in A$ by letting

$$\Phi(a) = \sum_{r\in R} \sum_{k=1}^{n_r(a)} \kappa_r(k).$$

Thus, to compute $\Phi(a)$, for each resource r we first determine the number $n_r(a)$ of players that use the resource r and then calculate the sum S_r of the costs of using this resource $1, 2, \ldots, n_r(a)$ times. Finally, we have to add all the sums S_r.

Let us now compare $n_r(a) = n_r(a_{-i}, a_i)$ with $n_r(a_{-i}, x_i)$ when x_i is an arbitrary action in A_i. We recall that x_i and a_i are subsets of the resource set R, and that $n_r(a_{-i}, x_i)$ is the number of players $k \in N \setminus \{i\}$ that use the resource r (i.e. the number of players $k \in N \setminus \{i\}$ such that $r \in a_k$), plus 1 if player i also uses the resource r. Therefore,

$$n_r(a_{-i}, x_i) = \begin{cases} n_r(a_{-i}, a_i) & \text{if } r \in x_i \cap a_i, \\ n_r(a_{-i}, a_i) & \text{if } r \notin x_i \cup a_i, \\ n_r(a_{-i}, a_i) + 1 & \text{if } r \in x_i \setminus a_i, \\ n_r(a_{-i}, a_i) - 1 & \text{if } r \in a_i \setminus x_i. \end{cases}$$

It follows that

$$\begin{aligned}
c_i(a_{-i},x_i) - c_i(a_{-i},a_i) &= \sum_{r\in x_i} \kappa_r(n_r(a_{-i},x_i)) - \sum_{r\in a_i} \kappa_r(n_r(a_{-i},a_i)) \\
&= \sum_{r\in x_i\setminus a_i} \kappa_r(n_r(a_{-i},x_i)) + \sum_{r\in x_i\cap a_i} \kappa_r(n_r(a_{-i},x_i)) \\
&\quad - \sum_{r\in a_i\setminus x_i} \kappa_r(n_r(a_{-i},a_i)) - \sum_{r\in x_i\cap a_i} \kappa_r(n_r(a_{-i},a_i)) \\
&= \sum_{r\in x_i\setminus a_i} \kappa_r(n_r(a_{-i},a_i)+1) - \sum_{r\in a_i\setminus x_i} \kappa_r(n_r(a_{-i},a_i))
\end{aligned}$$

and that

$$\begin{aligned}
\Phi(a_{-i},x_i) - \Phi(a_{-i},a_i) &= \sum_{r\in R}\Big(\sum_{k=1}^{n_r(a_{-i},x_i)} \kappa_r(k) - \sum_{k=1}^{n_r(a_{-i},a_i)} \kappa_r(k)\Big) \\
= \Big(\sum_{r\in R\setminus(x_i\cup a_i)} + \sum_{r\in x_i\setminus a_i} + \sum_{r\in a_i\setminus x_i} + \sum_{r\in x_i\cap a_i}\Big)&\Big(\sum_{k=1}^{n_r(a_{-i},x_i)} \kappa_r(k) - \sum_{k=1}^{n_r(a_{-i},a_i)} \kappa_r(k)\Big) \\
= \sum_{r\in x_i\setminus a_i} \kappa_r(n_r(a_{-i},a_i)+1) - \sum_{r\in a_i\setminus x_i} \kappa_r(n_r(a_{-i},a_i)).
\end{aligned}$$

Thus we have shown that

$$c_i(a_{-i},x_i) - c_i(a_{-i},a_i) = \Phi(a_{-i},x_i) - \Phi(a_{-i},a_i),$$

which is equation (1). □

EXAMPLE 4.1.3 The function Φ in the above proof is an example of a so called potential function, functions to be studied in detail in the next section, but let us end this part by computing the potential function Φ for the congestion game in Example 4.1.1. The outcome $a = (x,x,x)$, the case when all players choose the roads PX and XQ, uses the resources $r = PX$ and $r = XQ$ three times each, while the remaining resources are not used at all. Hence,

$$\Phi(x,x,x) = \sum_{k=1}^{3} \kappa_{PX}(k) + \sum_{k=1}^{3} \kappa_{XQ}(k) = (2+3+5) + (2+3+6) = 21.$$

The remaining 17 values of the potential function are computed in a similar way, but we should of course use symmetry which implies that $\Phi(a') = \Phi(a'')$

if a' and a'' are permutations of each other. We get the following table of the function values.

a	(x,x,x)	(x,x,z)	(x,y,x)	(x,y,z)	(x,z,x)	(x,z,z)
$\Phi(a)$	21	17	15	17	17	20

a	(y,x,x)	(y,x,z)	(y,y,x)	(y,y,z)	(y,z,x)	(y,z,z)
$\Phi(a)$	15	17	19	24	17	23

a	(z,x,x)	(z,x,z)	(z,y,x)	(z,y,z)	(z,z,x)	(z,z,z)
$\Phi(a)$	17	20	17	23	20	29

The potential function assumes it minimum value 15 at (x,y,x) and at (y,x,x), the Nash equilibria of the game. □

Exercise

4.1 Consider the congestion game in Example 4.1.1, and assume that the third player Clara can also use the road PY and thus has access to the same alternatives x, y and z as the other two players. Determine, under these conditions, the minimum points of the potential function and thus the game's Nash equilibria.

4.2 Potential games

The fact that every congestion game has a Nash equilibrium is, as shown in the proof above, a direct consequence of the existence of a function Φ that satisfies equation (1). This makes it possible to generalize the result to all games with a similar potential function, which motivates the following definition.

Definition 4.2.1 A function $\Phi\colon A \to \mathbf{R}$ on the set $A = A_1 \times A_2 \times \cdots \times A_n$ of all possible outcomes of a strategic game $\langle N, (A_i), (u_i)\rangle$ is called an

- *(exact) potential function* if

$$u_i(a_{-i}, x_i) - u_i(a_{-i}, a_i) = \Phi(a_{-i}, x_i) - \Phi(a_{-i}, a_i)$$

  for all players $i \in N$, all outcomes $a \in A$ and all actions $x_i \in A_i$;
- *ordinal potential function* if

$$u_i(a_{-i}, x_i) - u_i(a_{-i}, a_i) > 0 \;\Leftrightarrow\; \Phi(a_{-i}, x_i) - \Phi(a_{-i}, a_i) > 0$$

  for all players $i \in N$, all outcomes $a \in A$ and all actions $x_i \in A_i$.

A strategic game that has an exact potential function is called a *potential game*, and a strategic game that has an ordinal potential function is called an *ordinal potential game*.

Potential games were introduced and studied by Dov Monderer and Lloyd Shapley in 1996. Every potential game is obviously an ordinal potential game, and the proof of Proposition 4.1.1 shows that every congestion game is a potential game. Conversely, Monderer and Shapley have proved that every potential game is equivalent to a congestion game.

Proposition 4.2.1 *Every maximum point of the potential function of an ordinal potential game is a Nash equilibrium of the game. Particularly, each finite ordinal potential game has at least one Nash equilibrium.*

Proof. Let $\langle N, (A_i), (u_i)\rangle$ be a strategic game with an ordinal potential function Φ, let a^* be a maximum point of Φ, and assume that a^* is not a Nash equilibrium. Then there is a player i with an action $x_i \in A_i$ such that $u_i(a^*_{-i}, x_i) - u_i(a^*_{-i}, a^*_i) > 0$. But this implies that $\Phi(a^*_{-i}, x_i) - \Phi(a^*_{-i}, a^*_i) > 0$, which contradicts our assumption that $a^* = (a^*_{-i}, a^*_i)$ is a maximum point of Φ. Hence, a^* must be a Nash equilibrium.

The ordinal potential function Φ of a finite game certainly has a maximum point since its domain of definition is finite. Hence, finite ordinal games have Nash equilibria. □

Example 4.2.1 Prisoner's dilemma, i.e. the game with payoff functions given by the table

	Deny	*Confess*
Deny	$(-1,-1)$	$(-5,0)$
Confess	$(0,-5)$	$(-3,-3)$

is a potential game. It is easy to verify that we get an exact potential function Φ by defining $\Phi(\textit{Deny,Deny}) = 0$, $\Phi(\textit{Deny,Confess}) = \Phi(\textit{Confess,Deny}) = 1$ and $\Phi(\textit{Confess,Confess}) = 3$. The maximum point (*Confess,Confess*) is a Nash equilibrium. □

Example 4.2.2 Cournot's model, which was studied in Section 3.1, consists of n profit maximizing firms that compete with a product on a common market. If each firm i produces q_i units of the product, then it is possible to sell all the products for the price of $P(Q)$ per unit, where P is the inverse demand function and $Q = q_1 + q_2 + \cdots + q_n$. The cost of producing one unit of the product is assumed to be the same for all firms and equal to c, which means that the profit of firm i is

$$V_i(q_1, q_2, \ldots, q_n) = q_iP(Q) - cq_i.$$

We assume that each firm i can produce any positive amount of the product, i.e. its set A_i of possible actions is equal to $]0, +\infty[$.

The Cournot model is an ordinal potential game under these assumptions with

$$\Phi(q_1, q_2, \ldots, q_n) = q_1q_2\cdots q_n(P(Q) - c)$$

as an ordinal potential function, because

$$\Phi(q_{-i}, x_i) - \Phi(q_{-i}, q_i) = \frac{q_1q_2\cdots q_n}{q_i}\big(V_i(q_{-i}, x_i) - V_i(q_{-i}, q_i)\big),$$

which implies that

$$V_i(q_{-i}, x_i) - V_i(q_{-i}, q_i) > 0 \;\Leftrightarrow\; \Phi(q_{-i}, x_i) - \Phi(q_{-i}, q_i) > 0.$$

Now assume that the inverse demand function P is continuous, that $P(Q) > c$ for some output Q, and that $P(Q) < c$ for all sufficiently large outputs Q. The ordinal potential function Φ certainly has a maximum under these very realistic assumptions, because Φ is negative outside a sufficiently big hypercube $[0, K]^n$, positive at some point within the hypercube, and continuous. All the coordinates of the maximum point $\bar{q} = (\bar{q}_1, \bar{q}_2, \dots, \bar{q}_n)$ must be equal, i.e. $\bar{q}_1 = \bar{q}_2 = \cdots = \bar{q}_n$, because if follows from the inequality of arithmetic and geometric means that

$$\Phi(\hat{q}, \hat{q}, \dots, \hat{q}) \geq \Phi(\bar{q}_1, \bar{q}_2, \dots, \bar{q}_n)$$

if $\hat{q} = (\bar{q}_1 + \bar{q}_2 + \cdots + \bar{q}_n)/n$, with strict inequality if not all $\bar{q}_i$ are equal.

We conclude that Cournot's modell under our assumptions has a symmetric Nash equilibrium $\bar{q}$, and if the inverse demand function is differentiable, we obtain $\bar{q}_1$ as a solution of the equation $xP'(nx) + P(nx) = c$. □

Definition 4.2.2 Let $\langle N, (A_i), (u_i)\rangle$ be an arbitrary strategic game and write as usual $A = A_1 \times A_2 \times \cdots \times A_n$. A *path* γ *from* $a \in A$ *to* $b \in A$ is a sequence $(a^{(0)}, a^{(1)}, \dots, a^{(m)})$ of outcomes in A with $a^{(0)} = a$ and $a^{(m)} = b$ and with the property that the coordinates $a_i^{(k-1)}$ and $a_i^{(k)}$ of two consecutive outcomes $a^{(k-1)}$ och $a^{(k)}$ are equal for all indices i but one, which we denote by i_k. This means that player i_k is the only player who changes his action in the step from outcome $a^{(k-1)}$ to outcome $a^{(k)}$.

The *length* of γ is m, i.e. equal to the number of outcomes in the path minus 1.

The path is a *cycle* if it starts and ends in the same outcome, i.e. if $a = b$.

A path from a to b is *simple* if all outcomes in the path are different, except possibly the starting point a and the endpoint b.

We allow paths of lenght 0. Such a path consists of just one outcome a and is also a simple cycle. If a and b are two outcomes that differ in one coordinate, then (a, b, a) is a simple cycle.

Given a path $\gamma_1 = (a^{(0)}, a^{(1)}, \dots, a^{(m)})$ from a to b and a path $\gamma_2 = (b^{(0)}, b^{(1)}, \dots, b^{(k)})$ from b to c, where in particular $a^{(m)} = b = b^{(0)}$, we denote by $\gamma_1 + \gamma_2$ the path from a to c obtained by starting with the path γ_1 and continuing along the path γ_2, i.e. the path $(a^{(0)}, a^{(1)}, \dots, a^{(m)}, b^{(1)}, \dots, b^{(k)})$ from a to c.

The path $-\gamma_1$ is of course the path γ_1 traversed backwards from b to a, i.e. the path $(a^{(m)}, a^{(m-1)}, \dots, a^{(1)}, a^{(0)})$.

Definition 4.2.3 The path $\gamma = (a^{(0)}, a^{(1)}, \dots, a^{(m)})$ is called an *improvement path* if $u_{i_k}(a^{(k)}) > u_{i_k}(a^{(k-1)})$ for each k with $1 \le k \le m$.

The utility increases along an improvement path for the player i_k who changes his action at the transition from outcome $a^{(k-1)}$ to outcome $a^{(k)}$.

It is obiously impossible to extend an improvement path from a to b with one outcome to a longer improvement path if and only if the endpoint b is a Nash equilibrium. Every not extendible improvement path must, in other words, end in a Nash equilibrium.

Definition 4.2.4 A game is said to have the *finite improvement path property (FIP-property)* if there does not exist any improvement path that can be extended to improvement paths of arbitrary length.

Games with the FIP-property have Nash equilibria, because improvement paths of maximal length must end in Nash equilibria.

Improvement paths give the players of a game with the FIP-property an opportunity of finding a Nash equilibrium if the game is repeated a number of times, and each time one of the players is allowed to improve his utility by changing his action. The game will eventually end in a Nash equilibrium.

Proposition 4.2.2 *Every finite ordinal potential game has the FIP-property.*

Proof. Let $\langle N,(A_i),(u_i)\rangle$ be a game with an ordinal potential function Φ. The sequence $\big(\Phi(a^{(k)})\big)_{k=0}^{m}$ is strictly increasing for each improvement path $(a^{(0)},a^{(1)},\dots,a^{(m)})$, since $u_{i_k}(a^{(k)}) > u_{i_k}(a^{(k-1)})$ implies $\Phi(a^{(k)}) > \Phi(a^{(k-1)})$.

The length of the improvement paths in a finite ordinal game is therefore bounded by the number of elements in A, the domain of definition of Φ. □

Example 4.2.3 The game Matching Pennies with payoff table

	Head	*Tail*
Head	$(1,-1)$	$(-1,1)$
Tail	$(-1,1)$	$(1,-1)$

does not have the FIP-property, because the improvement cycle

$$\gamma = \big((Head, Head), (Head, Tail), (Tail, Tail), (Tail, Head), (Head, Head)\big)$$

can be extended to an improvement cycle $\gamma+\gamma+\cdots+\gamma$ of arbitrary length. Also, the game has no Nash equilibrium. □

Definition 4.2.5 Let $\gamma = (a^{(0)},a^{(1)},\dots,a^{(m)})$ be a path in a strategic game $\langle N,(A_i),(u_i)\rangle$. The number

$$I(\gamma) = \sum_{k=1}^{m}\big(u_{i_k}(a^{(k)}) - u_{i_k}(a^{(k-1)})\big),$$

where as before i_k is the unique coordinate where the vectors $a^{(k-1)}$ and $a^{(k)}$ are different, is called the *index* of the path.

The above sum is empty if the path is trivial, i.e. consists of just one outcome $a^{(0)}$. An empty sum is interpreted as 0. The index of a trivial path is therefore by definition equal to 0.

Obviously, $I(-\gamma) = -I(\gamma)$, and $I(\gamma_1+\gamma_2) = I(\gamma_1)+I(\gamma_2)$ if the path γ_1 ends where the path γ_2 starts.

In games $\langle N,(A_i),(u_i)\rangle$ with an exact potential function Φ,

$$u_{i_k}(a^{(k)}) - u_{i_k}(a^{(k-1)}) = \Phi(a^{(k)}) - \Phi(a^{(k-1)}),$$

which implies that the index of an arbitrary path $\gamma = (a^{(0)},a^{(1)},\dots,a^{(m)})$ from $a=a^{(0)}$ to $b=a^{(m)}$ becomes

$$I(\gamma) = \sum_{k=1}^{m}\big(\Phi(a^{(k)}) - \Phi(a^{(k-1)})\big) = \Phi(b)-\Phi(a). \tag{2}$$

If a game has a potential function Φ, then $\Phi + C$ is also a potential function for each constant C, and this is also the only way to create new exact potential functions.

Proposition 4.2.3 *Let Φ and Ψ be two exact potential functions in a strategic game. Then there exists a constant C such that $\Psi = \Phi + C$.*

Proof. Fix an outcome $a \in A$ and choose given $b \in A$ a path γ_b from a to b. Then

$$\Psi(b) - \Psi(a) = I(\gamma_b) = \Phi(b) - \Phi(a),$$

due to equation (2), and hence $\Psi(b) = \Phi(b) + \Psi(a) - \Phi(a)$, which means that our claim is true with $C = \Psi(a) - \Phi(a)$. □

Corollary 4.2.4 *In a potential game there exists, for each outcome a, a unique potential function Ψ_a such that $\Psi_a(a) = 0$, and $\Psi_a(b) = I(\gamma_b)$ for each path γ_b from a to b.*

Proof. Let Φ be an arbitrary potential function in the game, and define $\Psi_a = \Phi - \Phi(a)$. Then, Ψ_a is also a potential function, $\Psi_a(a) = 0$, and $I(\gamma_b) = \Psi_a(b) - \Psi_a(a) = \Psi_a(b)$ if γ_b is a path from a to b according to equation (2). □

Those who have studied physics should note the following analogy between path index and the physical concept of work. Let the set of possible outcomes of the game correspond to the physical space, the utility vectors $(u_1(a), u_2(a), \dots, u_n(a))$ correspond to forces, and the paths γ correspond to space curves. The path index $I(\gamma)$ will then correspond to the work done by the force field along the corresponding space curve and, mathematically, this work is given by a curve integral. If the force field is conservative, i.e. if no work is done along closed curves, there exists a potential function by which the work along any curve is equal to the difference of the values of the potential function at the endpoint and the starting point. The next proposition describes the corresponding game theoretical result, and the proof is completely analogous to the proof of the physical result.

Proposition 4.2.5 *The following four conditions are equivalent for strategic games $\langle N, (A_i), (u_i)\rangle$:*

(i) The game is a potential game.
(ii) The index of each cycle is equal to 0.
(iii) The index of each simple cycle is equal to 0.
(iv) The index of each simple cycle of length 4 is equal to 0.

Proof. (i) $\Rightarrow$ (ii): It follows immediately from equation (2) that $I(\gamma) = 0$ for every cycle γ, if the game has a potential function Φ.

(ii) $\Rightarrow$ (iii) and (iii) $\Rightarrow$ (iv): Trivial implications.

(iv) $\Rightarrow$ (i): Assume that every simple cycle of length 4 has index 0. Fix an outcome $a = (a_1, a_2, \dots, a_n) \in A$, and let $b = (b_1, b_2, \dots, b_n)$ be an arbitrary outcome.

Corollary 4.2.4 suggests that we should obtain a potential function Φ by defining $\Phi(b)$ as the index of an arbitrary path from a to b. But there is a problem − we have to convince ourselves that this defines $\Phi(b)$ uniquely, i.e. that different paths γ_b and γ'_b from a to b always have the same index. If, instead, we had started from assumption (ii), then our problem would have been resolved, because $\gamma'_b - \gamma_b$ is a cycle, and if all cycles have index 0, then $I(\gamma'_b) - I(\gamma_b) = I(\gamma'_b - \gamma_b) = 0$. We can get around this difficulty by defining, for each outcome b, a unique path from a to b, and then defining $\Phi(b)$ as the index of this special path. Doing so, we will then be able to prove that Φ is a potential function under the sole assumption that each simple cycle of lenght 4 has index equal to 0.

To achieve this, we define recursively the outcomes $b^{(k)}$ for $k = 0, 1, \dots, n$

by

$$b^{(k)} = \begin{cases} a & \text{for } k = 0, \\ (b^{(k-1)}_{-k}, b_k) & \text{for } k \geq 1. \end{cases}$$

This means that $b^{(1)} = (b_1, a_2, \dots, a_n)$, $b^{(2)} = (b_1, b_2, a_3, \dots, a_n)$, …, $b^{(n)} = (b_1, b_2, \dots, b_n) = b$. The player who, if $b_k \neq a_k$, changes his action in the step from outcome $b^{(k-1)}$ to outcome $b^{(k)}$ is player k. ($b^{(k)} = b^{(k-1)}$ if $b_k = a_k$.)

We now define the function $\Phi\colon A \to \mathbf{R}$ by putting

$$\Phi(b) = \sum_{k=1}^{n} \big(u_k(b^{(k)}) - u_k(b^{(k-1)})\big).$$

The function value $\Phi(b)$ is equal to the index $I(\Gamma_b)$ of the uniquely defined path Γ_b from a to b which is obtained from the sequence $b^{(0)}, b^{(1)}, b^{(2)}, \dots, b^{(n)}$ by deleting any duplicates. Such duplicates occur if $b_k = a_k$ for some k.

We now prove that Φ is a potential function. To this end, let x_i be an arbitrary action in A_i, write $c = (b_{-i}, x_i)$ and note that

$$c^{(k)} = \begin{cases} b^{(k)} & \text{for } k = 0, 1, \dots, i-1, \\ (b^{(k)}_{-i}, x_i) & \text{for } k = i, i+1, \dots, n. \end{cases}$$

We have to show that

$$\Phi(c) - \Phi(b) = u_i(c) - u_i(b).$$

This is true, of course, if $x_i = b_i$, because then $c = b$. Suppose therefore that $x_i \neq b_i$. Then, for all indices $k \geq i+1$, $b^{(k-1)} \neq c^{(k-1)}$, $b^{(k)} \neq c^{(k)}$ and $b^{(k-1)} \neq c^{(k)}$, because $b^{(k-1)}_i = b^{(k)}_i = b_i$ and $c^{(k-1)}_i = c^{(k)}_i = x_i$. Since $b^{(k-1)}_k = c^{(k-1)}_k = a_k$ and $b^{(k)}_k = c^{(k)}_k = b_k$, it also follows that $b^{(k-1)} \neq b^{(k)}$ and that $c^{(k-1)} \neq c^{(k)}$, provided that $b_k \neq a_k$.

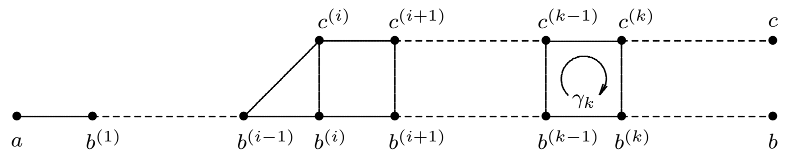

Figure 4.2. Illustration to the proof of Proposition 4.2.5. $\Phi(c) = I(\Gamma_c)$ and $\Phi(b) = I(\Gamma_b)$, where Γ_c is the upper path from a to c, and Γ_b is the lower path from a to b.

The path $\gamma_k = (b^{(k-1)}, c^{(k-1)}, c^{(k)}, b^{(k)}, b^{(k-1)}))$ is therefore a simple cycle of length 4 for all $k \geq i+1$ with $b_k \neq a_k$. See Figure 4.2. The index of the cycle is

$$I(\gamma_k) = \big(u_i(c^{(k-1)}) - u_i(b^{(k-1)})\big) + \big(u_k(c^{(k)}) - u_k(c^{(k-1)})\big) \\ + \big(u_i(b^{(k)}) - u_i(c^{(k)})\big) + \big(u_k(b^{(k-1)}) - u_k(b^{(k)})\big),$$

and since all simple cycles of lenght 4 have index 0, we conclude that $I(\gamma_k) = 0$ which after rewriting becomes

$$(3) \qquad u_k(c^{(k)}) - u_k(c^{(k-1)}) = \big(u_k(b^{(k)}) - u_k(b^{(k-1)})\big) \\ + \big(u_i(c^{(k)}) - u_i(b^{(k)})\big) - \big(u_i(c^{(k-1)}) - u_i(b^{(k-1)})\big).$$

This holds for $k \geq i+1$ also in the case $b_k = a_k$, because then $b^{(k-1)} = b^{(k)}$ and $c^{(k-1)} = c^{(k)}$.

By adding the equations (3) for $k = i+1, i+2, \dots n$, noting that $c^{(n)} = c$ and $b^{(n)} = b$, we obtain the following result:

$$\sum_{k=i+1}^{n} \big(u_k(c^{(k)}) - u_k(c^{(k-1)})\big) = \sum_{k=i+1}^{n} \big(u_k(b^{(k)}) - u_k(b^{(k-1)})\big) + \big(u_i(c) - u_i(b)\big) \\ - \big(u_i(c^{(i)}) - u_i(b^{(i)})\big).$$

Since $c^{(k)} = b^{(k)}$ when $k < i$, we also have the following equality

$$\sum_{k=1}^{i-1}\bigl(u_k(c^{(k)}) - u_k(c^{(k-1)})\bigr) = \sum_{k=1}^{i-1}\bigl(u_k(b^{(k)}) - u_k(b^{(k-1)})\bigr).$$

Now add the two sums above, from 1 to $i-1$ and from $i+1$ to n, respectively, and also add the missing term $\bigl(u_i(c^{(i)}) - u_i(c^{(i-1)})\bigr)$, and note that $c^{(i-1)} = b^{(i-1)}$. This gives us the following result

$$\begin{aligned}\Phi(c) &= \sum_{k=1}^{n}\bigl(u_k(c^{(k)}) - u_k(c^{(k-1)})\bigr)\\ &= \sum_{k=1}^{n}\bigl(u_k(b^{(k)}) - u_k(b^{(k-1)})\bigr) + \bigl(u_i(c) - u_i(b)\bigr) - \bigl(u_i(c^{(i)}) - u_i(b^{(i)})\bigr)\\ &\quad + \bigl(u_i(c^{(i)}) - u_i(c^{(i-1)})\bigr) - \bigl(u_i(b^{(i)}) - u_i(b^{(i-1)})\bigr)\\ &= \Phi(b) + u_i(c) - u_i(b),\end{aligned}$$

which implies that $\Phi(c) - \Phi(b) = u_i(c) - u_i(b)$ and proves that Φ is indeed a potential function. □

EXAMPLE 4.2.4 In Section 2.1, we discussed the game Stag Hunt. It is a strategic game with n players, where each player has two options, *Hare*: to hunt a hare, and *Stag*: to hunt the stag. Each player prefers that all players hunt the stag, but it is better for every player to hunt a hare than to hunt the stag if not everyone takes part in the stag hunting. The preferences of player i are described by the following utility function u_i:

$$u_i(a_1, a_2, \dots, a_n) = \begin{cases} 2 & \text{if } a_j = \textit{Stag} \text{ for all players } j,\\ 1 & \text{if } a_i = \textit{Hare},\\ 0 & \text{if } a_i = \textit{Stag} \text{ and } a_j = \textit{Hare} \text{ for at least one player } j.\end{cases}$$

In a simple cycle of lenght 4, only two players are actively involved and they alternate between the alternatives *Stag* and *Hare*. We may, without loss of generality, assume that players 1 and 2 are the players involved, that the cycle starts with both of them hunting stag, and that it is player 1 who begins to change. This means that the cycle γ has the following form:

$$\bigl((\textit{Stag}, \textit{Stag}, b), (\textit{Hare}, \textit{Stag}, b), (\textit{Hare}, \textit{Hare}, b), (\textit{Stag}, \textit{Hare}, b), (\textit{Stag}, \textit{Stag}, b)\bigr),$$

where b is an arbitrary action vector for the remaining $n-2$ players. Depending on whether the b vector only contains the *Stag*-option or contains at least one *Hare*, the index of the cycle is $I(\gamma) = -1+1-1+1 = 0$ or $I(\gamma) = 1+1-1-1 = 0$.

Hence, all simple cycles of length 4 have index 0, which implies that Stag Hunt is a potential game, and we get a potential function Φ by defining $\Phi(a) = I(\gamma_a)$, where γ_a is an arbitrary path from the outcome $a^{Stag} = (Stag, Stag, \dots, Stag)$ to $a = (a_1, a_2, \dots, a_n)$.

If $a_i = Hare$ for exactly k players i, where $1 \le k \le n$, then it is possible to reach a from a^{Stag} using a path γ_a of lenght k, where each step consists of a new player changing action from *Stag* to *Hare*, which increases the utility by 1 at all steps except the first one, when the utility instead decreases by 1. The path index $I(\gamma_a)$ is therefore equal to $-1+(k-1)$, i.e. $I(\gamma_a) = k-2$. This gives us the following potential function

$$\Phi(a_1, a_2, \dots, a_n) = \begin{cases} 0 & \text{if } a_i = \textit{Stag} \text{ for all players } i, \\ k-2 & \text{if } a_i = \textit{Hare} \text{ for } k \ge 1 \text{ players } i. \end{cases} \qquad \square$$

Exercises

4.2 Are the following games in Section 2.1 potential games? If appropriate, determine a potential function.

a) The game Battle of Sexes, Example 2.1.2.

b) The game Hawk or Dove, Example 2.1.4.

4.3 Prove that the game with the following payoff table

	L	R
T	$(1,0)$	$(2,0)$
B	$(2,0)$	$(0,1)$

is not an ordinal potential game but that it has the FIP-property.

Chapter 5

Mixed Strategies

Consider the game *Rock-paper-scissors* with payoff table

	rock	*scissors*	*paper*
rock	$(0,0)$	$(1,-1)$	$(-1,1)$
scissors	$(-1,1)$	$(0,0)$	$(1,-1)$
paper	$(1,-1)$	$(-1,1)$	$(0,0)$

It goes without saying that if you play this game several times, you should change alternatives from time to time, because otherwise your opponent immediately learns how to win every time. Both players should thus choose their alternatives randomly.

Suppose that the row player chooses the alternatives *rock, scissors, paper* randomly with probabilities α_1, α_2, α_3, and that the column player chooses the same alternatives randomly and independently of the row player with probabilities β_1, β_2, β_3. The probability that they will choose $(\textit{rock}, \textit{rock})$ will then become $\alpha_1\beta_1$, the probability of $(\textit{rock}, \textit{scissors})$ becomes $\alpha_1\beta_2$, etc. The expected payout (average payout if the game is repeated many times) to the row player will be

$$\tilde{u}_1(\alpha,\beta) = \alpha_1\beta_2 - \alpha_1\beta_3 - \alpha_2\beta_1 + \alpha_2\beta_3 + \alpha_3\beta_1 - \alpha_3\beta_2.$$

The game is a strictly competitive game; the column player's expected payoff is $\tilde{u}_2(\alpha,\beta) = -\tilde{u}_1(\alpha,\beta)$.

If, for example, $\alpha = (\frac{1}{2}, 0, \frac{1}{2})$ and $\beta = (\frac{1}{2}, \frac{1}{2}, 0)$ then

$$\tilde{u}_1(\alpha,\beta) = \tfrac{1}{4} + \tfrac{1}{4} - \tfrac{1}{4} = \tfrac{1}{4}$$

and $\tilde{u}_2(\alpha,\beta) = -\frac{1}{4}$, which is good for the row player, who can expect to win an average of $\frac{1}{4}$ if the game is played several times, and consequently not so good for the column player. Of course, there are better strategies for the column player.

The game Rock-paper-scissors serves as a motivation to study games where the players choose their actions randomly according to some probability distribution or, in other words, where the players choose lotteries.

5.1 Mixed strategies

We recall that a *lottery* over a finite set A is a probability distribution on A, and that the set $\mathcal{L}(A)$ of all lotteries over A is a compact and convex set that can be identified with the subset

$$M = \{(x_1, x_2, \dots, x_n) \in \mathbf{R}^n_+ \mid x_1 + x_2 + \cdots + x_n = 1\}$$

of $\mathbf{R}^n$, if n is the number of elements in A.

The n vertices $(1, 0, \dots, 0)$, $(0, 1, \dots, 0)$, $\dots$, $(0, 0, \dots, 1)$ of the set M correspond to the lotteries δ_a in $\mathcal{L}(A)$ that result in the outcome $a \in A$ with probability one. As already mentioned in Section1.3, we will identify the lottery δ_a with the element a and denote it by a, except in cases where we need to be particularly pedantic. Thus, we perceive A as a subset of the set $\mathcal{L}(A)$ of lotteries.

Definition 5.1.1 A *mixed strategy* of a player i of a finite strategic game $\langle N, (A_i), (u_i)\rangle$ is a lottery over the set A_i. The set of all mixed strategies of player i, that is the set of all lotteries over A_i, is denoted by $\mathcal{L}(A_i)$.

The lotteries δ_{a_i} in $\mathcal{L}(A_i)$ that result in an outcome $a_i \in A_i$ with probability one, are called *pure strategies* of player i.

Suppose that the players of the game $\langle N, (A_i), (u_i)\rangle$ choose their actions in the sets A_i independently of each other by using the outcomes of the lotteries $p_i \in \mathcal{L}(A_i)$. The probability that the action vector

$$a = (a_1, a_2, \dots, a_n)$$

will be chosen will then be equal to

$$P(a) = p_1(a_1)p_2(a_2)\cdots p_n(a_n).$$

The function P is a probability measure, that is in our terminology a *lottery*, on the Cartesian product $A = A_1 \times A_2 \times \cdots \times A_n$, and we call this lottery the *product lottery* of the given lotteries and write

$$P = p_1 \times p_2 \times \cdots \times p_n.$$

Note that the product $\delta_{a_1} \times \delta_{a_2} \times \cdots \times \delta_{a_n}$ of pure strategies δ_{a_i} is equal to the lottery δ_a over A, which has $a = (a_1, a_2, \dots, a_n) \in A$ as certain outcome.

Given a function u on the product set A, we may of course form the expected value of u with respect to the product lottery $p_1 \times p_2 \times \cdots \times p_n$. This expected value is denoted $\tilde{u}(p_1, p_2, \dots, p_n)$, and according to the general definition of expected value,

$$\tilde{u}(p_1, p_2, \dots, p_n) = \sum_{a_1 \in A_1} \sum_{a_2 \in A_2} \dots \sum_{a_n \in A_n} u(a_1, a_2, \dots, a_n) p_1(a_1) p_2(a_2) \cdots p_n(a_n).$$

The function $(p_1, p_2, \dots, p_n) \mapsto \tilde{u}(p_1, p_2, \dots, p_n)$ is defined on the Cartesian product $\mathcal{L}(A_1) \times \cdots \times \mathcal{L}(A_n)$, which can be perceived as a convex, compact subset of $\mathbf{R}^m$ for a suitable m. The function is continuous and affine in each variable, i.e.

$$\tilde{u}(p_{-i}, \alpha p_i + \beta q_i) = \alpha \tilde{u}(p_{-i}, p_i) + \beta \tilde{u}(p_{-i}, q_i)$$

if α and β are nonnegative numbers with sum equal to 1.

The expected value $\tilde{u}(\delta_{a_1}, \dots, \delta_{a_n})$ is of course equal to $u(a_1, \dots, a_n)$, so using our convention to identify lotteries with certain outcomes with the corresponding actions, we have

$$\tilde{u}(a_1, a_2, \dots, a_n) = u(a_1, a_2, \dots, a_n).$$

A strategy combination, that we will sometimes consider, is the one where player i chooses a mixed strategy p_i while all other players k choose pure strategies a_k. The expected value $\tilde{u}(a_{-i}, p_i)$ of the function u is now

$$\tilde{u}(a_{-i}, p_i) = \sum_{a_i \in A_i} u(a_{-i}, a_i) p_i(a_i),$$

and we can express the expected value $\tilde{u}(q_{-i}, p_i)$ for an arbitrary vector q_{-i} of mixed strategies in terms of these expected values as

$$\tilde{u}(q_{-i}, p_i) = \sum_{a_{-i} \in A_{-i}} \tilde{u}(a_{-i}, p_i) q_1(a_1) \cdots \widehat{q_i(a_i)} \cdots q_n(a_n),$$

where the circumflux above $q_i(a_i)$ indicates that this term in the product should be omitted. The expected value $\tilde{u}(q_{-i}, p_i)$ is in other words a weighted average of all expected values $\tilde{u}(a_{-i}, p_i)$ that are obtained by letting a_{-i} run through A_{-i}, and this implies that

$$\tilde{u}(q_{-i}, p_i) \geq \min_{a_{-i} \in A_{-i}} \tilde{u}(a_{-i}, p_i) \tag{1}$$

for all mixed strategy vectors q_{-i}.

5.2 The mixed extension of a game

Definition 5.2.1 Let $G = \langle N, (A_i), (u_i) \rangle$ be a finite strategic game with cardinal utility functions u_i. The game $\widetilde{G} = \langle N, (\mathcal{L}(A_i)), (\tilde{u}_i) \rangle$

- with the same set N of players;
- for each player $i \in N$, the set $\mathcal{L}(A_i)$ of lotteries over A_i as set of actions;
- for each player $i \in N$, the expected utility $\tilde{u}_i$ as utility function

is called the *mixed extension of the game* G.

We thus assume, here and in the future, that the utility functions u_i in the original game G are cardinal, because otherwise it is not very meaningful to form expected values.

Using the mixed extension, we can now immediately generalize a number of concepts and results for actions in strategic games to apply to mixed strategies, which are nothing but actions in the extension $\widetilde{G}$.

Definition 5.2.2 A Nash equilibrium $p^* = (p_1^*, p_2^*, \ldots, p_n^*)$ of the mixed extension $\widetilde{G} = \langle N, (\mathcal{L}(A_i)), (\tilde{u}_i) \rangle$ of a finite strategic game $G = \langle N, (A_i), (u_i) \rangle$ is called a *mixed Nash equilibrium* of the game G.

A mixed Nash equilibrium that only consists of pure strategies is called a *pure Nash equilibrium.*

Mixed Nash equilibria can of course be characterized using best-response sets in $\widetilde{G}$. Given a vector p_{-i} of mixed strategies for all players except player i, we define player i's *best-response set* $\widetilde{B}_i(p_{-i})$ *of mixed strategies* as the player's best-response set in the mixed extension, i.e.

$$\widetilde{B}_i(p_{-i}) = \{q_i \in \mathcal{L}(A_i) \mid \tilde{u}_i(p_{-i}, q_i) \geq \tilde{u}_i(p_{-i}, r_i) \quad \text{for all } r_i \in \mathcal{L}(A_i)\}.$$

It follows immediately from Proposition 2.2.1 that a vector p^* of mixed strategies of the game $G = \langle N, (A_i), (u_i)\rangle$ is a mixed Nash equilibrium if and only if $p_i^* \in \widetilde{B}_i(p_{-i}^*)$ for all players i.

The following theorem, due to John Nash, is the fundamental theorem of strategic game theory

Proposition 5.2.1 *Every finite strategic game has a mixed Nash equilibrium.*

Proof. The proposition is a corollary to Proposition 2.3.3, because the mixed extension of a game $\langle N, (A_i), (u_i)\rangle$ meets the conditions of the proposition: the sets $\mathcal{L}(A_i)$ are convex and compact and the expected utility functions $\tilde{u}_i$ are continuous and affine in the ith variable and consequently quasiconcave. □

We will now illustrate Nash's theorem by determining all mixed Nash equilibria of some simple two-player strategic games. We will use the characterization of the Nash equilibrium in terms of best-response functions. In cases where both players have two action options, there is a simple graphical solution.

Example 5.2.1 Consider the following two-person game

	L	R
T	$(3,3)$	$(0,2)$
B	$(2,1)$	$(5,5)$

The game has two Nash equilibria, viz. (T, L) and (B, R). To determine all mixed Nash equilibria, we let $p_1 = (\alpha, 1-\alpha)$ and $p_2 = (\beta, 1-\beta)$ be two arbitrary mixed strategies for the row and the column player respectively. The corresponding expected utility functions are

$$\tilde{u}_1(p_1, p_2) = 3\alpha\beta + 2(1-\alpha)\beta + 5(1-\alpha)(1-\beta) = 5 - 3\beta + (6\beta - 5)\alpha$$

and

$$\begin{aligned}\tilde{u}_2(p_1, p_2) &= 3\alpha\beta + 2\alpha(1-\beta) + (1-\alpha)\beta + 5(1-\alpha)(1-\beta)\\ &= 5 - 3\alpha + (5\alpha - 4)\beta.\end{aligned}$$

The best-response set $\widetilde{B}_1(p_2)$ consists of the lotteries p_1 that maximize the function $\tilde{u}_1(p_1, p_2)$, and $\tilde{u}_1(p_1, p_2) = 5 - 3\beta + (6\beta - 5)\alpha$ is obviously maximized by $\alpha = 1$ if $\beta > \frac{5}{6}$, by all numbers $0 \le \alpha \le 1$ if $\beta = \frac{5}{6}$, and by $\alpha = 0$ if $\beta < \frac{5}{6}$. This means that

$$\widetilde{B}_1(\beta, 1-\beta) = \begin{cases} \{(0,1)\} & \text{if } \beta < \frac{5}{6},\\ \{(\alpha, 1-\alpha) \mid 0 \le \alpha \le 1\} & \text{if } \beta = \frac{5}{6},\\ \{(1,0)\} & \text{if } \beta > \frac{5}{6}.\end{cases}$$

Analogously we get

$$\widetilde{B}_2(\alpha, 1-\alpha) = \begin{cases} \{(0,1)\} & \text{if } \alpha < \frac{4}{5},\\ \{(\beta, 1-\beta) \mid 0 \le \beta \le 1\} & \text{if } \alpha = \frac{4}{5},\\ \{(1,0)\} & \text{if } \alpha > \frac{4}{5}.\end{cases}$$

In a coordinate system we now draw the two curves

$$\{(\alpha, \beta) \mid (\alpha, 1-\alpha) \in \widetilde{B}_1(\beta, 1-\beta)\}$$

and

$$\{(\alpha, \beta) \mid (\beta, 1-\beta) \in \widetilde{B}_2(\alpha, 1-\alpha)\}.$$

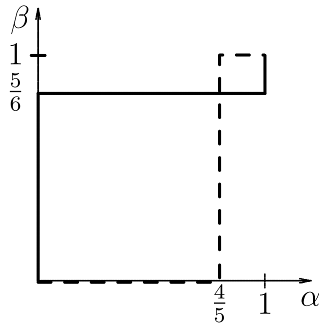

Figure 5.1. Best-response sets in Example 5.2.1.

In figure 5.1, the first curve is solid and the second curve is dashed.

The two curves' intersections (α, β) correspond to mixed strategies $p_1 = (\alpha, 1-\alpha)$ and $p_2 = (\beta, 1-\beta)$ such that $p_1 \in \widetilde{B}_1(p_2)$ and $p_2 \in \widetilde{B}_2(p_1)$, i.e. to the mixed Nash equilibria.

In the present example, the two curves intersect in three points, namely $(\alpha, \beta) = (0,0)$, $(1,1)$ and $(\frac{4}{5}, \frac{5}{6})$. There are thus three mixed Nash equilibria, $\big((0,1),(0,1)\big)$, $\big((1,0),(1,0)\big)$ and $\big((\frac{4}{5},\frac{1}{5}),(\frac{5}{6},\frac{1}{6})\big)$. The first two Nash equilibria are comprised of pure strategies, namely (δ_B, δ_R) and (δ_T, δ_L), meaning that the players should select the action vector (B, R) and the action vector (T, L), respectively.

The pure Nash equilibria give the players the following expected utilities:

$$\tilde{u}_1(B,R) = u_1(B,R) = 5 \quad \text{and} \quad \tilde{u}_2(B,R) = u_2(B,R) = 5,$$
$$\tilde{u}_1(T,L) = u_1(T,L) = 3 \quad \text{and} \quad \tilde{u}_2(T,L) = u_2(T,L) = 3.$$

In the mixed Nash equilibrium $(p_1^*, p_2^*) = \big((\frac{4}{5},\frac{1}{5}),(\frac{5}{6},\frac{1}{6})\big)$, the expected utilities are

$$\tilde{u}_1(p_1^*, p_2^*) = \tfrac{5}{2} \quad \text{and} \quad \tilde{u}_2(p_1^*, p_2^*) = \tfrac{13}{5}.$$

Both players' expected utility is thus less in the mixed Nash equilibrium than in the pure Nash equilibria. □

In the example above, the original game's two Nash equilibria survived as pure Nash equilibria in the mixed extension of the game. This is no coincidence, because we have the following general result.

Proposition 5.2.2 *The outcome $a^* = (a_1^*, a_2^*, \dots, a_n^*)$ is a Nash equilibrium of a finite strategic game, if and only if the corresponding outcome $(\delta_{a_1^*}, \delta_{a_2^*}, \dots, \delta_{a_n^*})$ is a pure Nash equilibrium of the game's mixed extension.*

Proof. In the future, we use the same notation $(a_1, a_2, \dots, a_n)$ for an outcome in the original game G and for the outcome $(\delta_{a_1}, \delta_{a_2}, \dots, \delta_{a_n})$ in the game's mixed extension $\widetilde{G}$. Thus, we also write $\tilde{u}_i(a_1, a_2, \dots, a_n)$, or shorter $\tilde{u}_i(a)$, instead of $\tilde{u}_i(\delta_{a_1}, \delta_{a_2}, \dots, \delta_{a_n})$.

Suppose first that a^* is a pure Nash equilibrium of the mixed extension. Then by definition,

$$u_i(a^*) = \tilde{u}_i(a^*_{-i}, a^*_i) \geq \tilde{u}_i(a^*_{-i}, p_i)$$

for all players i and all lotteries $p_i \in \mathcal{L}(A_i)$, and in particular this holds for all certain lotteries of the typ $p = \delta_{a_i}$, which means that

$$(2) \qquad u_i(a^*) \geq \tilde{u}_i(a^*_{-i}, a_i) = u_i(a^*_{-i}, a_i)$$

for all $a_i \in A_i$. This proves that the outcome a^* is a Nash equilibrium of the original game.

Conversely, suppose that the outcome a^* is a Nash equilibrium of G, i.e. that inequality (2) holds for all players i and all actions $a_i \in A_i$. Let p_i be an arbitrary mixed strategy of player i, i.e. an arbitrary lottery over A_i. By multiplying the inequality (2) with $p_i(a_i)$ and adding the inequalities so obtained as a_i runs through the set A_i, we get the following inequality:

$$(3) \quad \sum_{a_i \in A_i} \tilde{u}_i(a^*)p_i(a_i) \geq \sum_{a_i \in A_i} \tilde{u}_i(a^*_{-i}, a_i)p_i(a_i) = \sum_{a_i \in A_i} \tilde{u}_i(a^*_{-i}, \delta_{a_i})p_i(a_i).$$

Now

$$\sum_{a_i \in A_i} \tilde{u}_i(a^*)p_i(a_i) = \tilde{u}_i(a^*) \sum_{a_i \in A_i} p_i(a_i) = \tilde{u}_i(a^*)$$

and

$$\sum_{a_i \in A_i} \tilde{u}_i(a^*_{-i}, \delta_{a_i})p_i(a_i) = \tilde{u}_i\big(a^*_{-i}, \sum_{a_i \in A_i} p_i(a_i)\delta_{a_i}\big) = \tilde{u}_i(a^*_{-i}, p_i),$$

so inequality (3) means that $\tilde{u}_i(a^*_{-i}, a^*_i) \geq \tilde{u}_i(a^*_{-i}, p_i)$ for all lotteries p_i in the lottery set $\mathcal{L}(A_i)$. This shows that the outcome a^* is a pure Nash equilibrium of the mixed extension of the game G. □

Exercises

5.1 Find all mixed Nash equilibria of the game Prisoner's Dilemma.

5.2 Find all mixed Nash equilibria of the game Battle of Sexes.

5.3 Find all mixed Nash equilibria of the following game:

	L	R
T	$(2,2)$	$(0,2)$
B	$(2,1)$	$(6,6)$

5.3 The indifference principle

In Example 5.2.1 we computed the Nash equilibria of the game

	L	R
T	$(3,3)$	$(0,2)$
B	$(2,1)$	$(5,5)$

and found that in addition to the two pure Nash equilibria (T, L) and (B, R), there is also a mixed Nash equilibrium p^* with $p^*_1 = (\frac{4}{5}, \frac{1}{5})$ and $p^*_2 = (\frac{5}{6}, \frac{1}{6})$.

If player 2 chooses his Nash strategy p^*_2, player 1 will receive the same expected utility, regardless of whether he chooses the action T or the action B, because

$$3 \cdot \tfrac{5}{6} + 0 \cdot \tfrac{1}{6} = 2 \cdot \tfrac{5}{6} + 5 \cdot \tfrac{1}{6} = \tfrac{5}{2}.$$

Similarly, player 2 gets the same expected utility when player 1 chooses his Nash strategy, regardless of whether player 2 chooses L or R, because

$$3 \cdot \tfrac{4}{5} + 1 \cdot \tfrac{1}{5} = 2 \cdot \tfrac{4}{5} + 5 \cdot \tfrac{1}{5} = \tfrac{13}{5}.$$

This is no coincidence because of the following universal indifference principle for Nash equilibria.

Proposition 5.3.1 (The indifference principle) *Let $p^* = (p_1^*, p_2^*, \dots, p_n^*)$ be a vector of mixed strategies for the players in a finite game $\langle N, (A_i), (u_i)\rangle$.*

(i) If p^* is a mixed Nash equilibrium, then for each player i and each action $a_i \in A_i$

$$p_i^*(a_i) > 0 \;\Rightarrow\; \tilde{u}_i(p_{-i}^*, a_i) = \tilde{u}_i(p^*).$$

Thus, all actions that occur with positive probability in the player's equilibrium strategy give him the same expected utility.

(ii) Conversely, p^* is a mixed Nash equilibrium if, for each player i, there is a constant c_i such that

$$(4) \qquad \begin{cases} p_i^*(a_i) > 0 \;\Rightarrow\; \tilde{u}_i(p_{-i}^*, a_i) = c_i \\ p_i^*(a_i) = 0 \;\Rightarrow\; \tilde{u}_i(p_{-i}^*, a_i) \le c_i. \end{cases}$$

Proof. (i) Let $B_i = \{a_i \in A_i \mid p_i^*(a_i) > 0\}$. Then

$$\sum_{a_i \in B_i} p_i^*(a_i) = 1 \qquad \text{and} \qquad p_i^* = \sum_{a_i \in A_i} p_i^*(a_i)\delta_{a_i} = \sum_{a_i \in B_i} p_i^*(a_i)\delta_{a_i},$$

and using linearity, we obtain

$$(5) \quad \tilde{u}_i(p^*) = \tilde{u}_i(p_{-i}^*, p_i^*) = \sum_{a_i \in B_i} \tilde{u}_i(p_{-i}^*, \delta_{a_i}) p_i^*(a_i) = \sum_{a_i \in B_i} \tilde{u}_i(p_{-i}^*, a_i) p_i^*(a_i).$$

Since p^* is a mixed Nash equilibrium, $\tilde{u}_i(p_{-i}^*, a_i) \le \tilde{u}_i(p^*)$ for all $a_i \in B_i$, and if there is an a_i such that strict inequality prevails, then it follows by insertion into equation (5) that

$$\tilde{u}_i(p^*) < \sum_{a_i \in B_i} \tilde{u}_i(p^*) p_i^*(a_i) = \tilde{u}_i(p^*) \sum_{a_i \in B_i} p_i^*(a_i) = \tilde{u}_i(p^*),$$

which is a contradiction. Hence, $\tilde{u}_i(p_{-i}^*, a_i) = \tilde{u}_i(p^*)$ for all $a_i \in B_i$.

(ii) Conversely, assume that the condition of (ii) is satisfied, and consider player i. Since $\tilde{u}_i(p_{-i}^*, a_i) \le c_i$ for each action a_i of player i, it follows that

$$\tilde{u}_i(p_{-i}^*, a_i) q_i(a_i) \le c_i q_i(a_i)$$

for each mixed strategy q_i and each $a_i \in A_i$. Moreover, for the particular mixed strategy p_i^* it follows from the implications in (4) that

$$\tilde{u}_i(p_{-i}^*, a_i) p_i^*(a_i) = c_i p_i^*(a_i)$$

for all $a_i \in A_i$. By summing the above two inequalities and equalities over A_i, we get the following inequality

$$\begin{aligned}\tilde{u}_i(p^*_{-i}, q_i) &= \sum_{a_i \in A_i} \tilde{u}_i(p^*_{-i}, a_i) q_i(a_i) \leq \sum_{a_i \in A_i} c_i q_i(a_i) = c_i \\ &= \sum_{a_i \in A_i} c_i p^*_i(a_i) = \sum_{a_i \in A_i} \tilde{u}_i(p^*_{-i}, a_i) p^*_i(a) = \tilde{u}_i(p^*),\end{aligned}$$

which shows that p^* is a mixed Nash equilibrium. □

EXAMPLE 5.3.1 Let us use the indifference principle to determine the Nash equilibria of the game

	L	R
T	$(3,1)$	$(1,2)$
B	$(2,4)$	$(4,3)$

We first note that there are no pure Nash equilibria. Because of the indifference principle, there is also no Nash equilibrium with exactly one pure strategy. For example, (T, p_2) is not a Nash equilibrium for any choice of mixed strategy p_2 for player 2, because $u_2(T, L) = 1 \neq 2 = u_2(T, R)$, and the other three options with one pure strategy are excluded for similar reasons.

So let (p_1^*, p_2^*) be a Nash equilibrium, and write $p_1^* = (\alpha, 1-\alpha)$ and $p_2^* = (\beta, 1-\beta)$ with $0 < \alpha < 1$ and $0 < \beta < 1$. According to the indifference principle, $\tilde{u}_1(T, p_2^*) = \tilde{u}_1(B, p_2^*)$ and $\tilde{u}_2(p_1^*, L) = \tilde{u}_2(p_1^*, R)$, which gives us the system

$$\begin{aligned} 3\beta + (1-\beta) &= 2\beta + 4(1-\beta) \\ \alpha + 4(1-\alpha) &= 2\alpha + 3(1-\alpha) \end{aligned}$$

with the solution $\alpha = \frac{1}{2}$ and $\beta = \frac{3}{4}$. We conclude that the game has a unique mixed Nash equilibrium, namely $\big((\frac{1}{2}, \frac{1}{2}), (\frac{3}{4}, \frac{1}{4})\big)$. □

Exercises

5.4 Consider the game

	k_1	k_2	k_3
r_1	$(2,*)$	$(1,*)$	$(5,*)$
r_2	$(1,*)$	$(3,*)$	$(4,*)$
r_3	$(4,*)$	$(0,*)$	$(2,*)$

where the column player's payoffs have been lost. On the other hand, it is known that the mixed strategy $p_1^* = (\frac{1}{3}, \frac{1}{6}, \frac{1}{2})$ is a Nash equilibrium strategy for the row player. Use this information to determine a mixed strategy p_2^* for the column player that makes (p_1^*, p_2^*) to a Nash equilibrium.

5.5 The same question as in the previous exercise if you instead know that $p_1^* = (\frac{1}{2}, 0, \frac{1}{2})$ is a Nash equilibrium strategy for the row player.

5.4 Dominance

It is sometimes possible for a player to rule out an action option because there are other options that give him a larger payoff regardless of how the opponents play.

Example 5.4.1 Consider the game

	k_1	k_2	k_3
r_1	$(2,*)$	$(1,*)$	$(2,*)$
r_2	$(3,*)$	$(4,*)$	$(1,*)$
r_3	$(3,*)$	$(2,*)$	$(3,*)$

where the column player's payoffs are omitted because they do not matter for the reasoning. Here, $u_1(r_3, k_i) > u_1(r_1, k_i)$ for each k_i, so no matter what action the column player chooses, the r_3 option is better for the row player than the r_1 option. We express this by saying that the r_1 action is *strictly dominated* by the r_3 action.

A rational player would never choose a strictly dominated action, and such an action can not, as we will show later, be a part of any Nash equilibrium. □

EXAMPLE 5.4.2 In the game with the payoff matrix

	k_1	k_2	k_3
r_1	$(2, *)$	$(1, *)$	$(3, *)$
r_2	$(1, *)$	$(3, *)$	$(4, *)$
r_3	$(4, *)$	$(0, *)$	$(3, *)$

none of the row player's actions is strictly dominated by any other action. But

$$u_1(r_1, k_i) < \tfrac{1}{2}u_1(r_2, k_i) + \tfrac{1}{2}u_1(r_3, k_i) \tag{6}$$

for each of the column player's three options k_1, k_2, k_3, since $2 < \frac{1}{2}(1 + 4)$, $1 < \frac{1}{2}(3 + 0)$ and $3 < \frac{1}{2}(4 + 3)$.

Let $\hat{p}_1$ be the row player's mixed strategy defined by $\hat{p}_1(r_1) = 0$, $\hat{p}_1(r_2) = \hat{p}_1(r_3) = \frac{1}{2}$, and consider the row players expected utility function $\tilde{u}_1$. Since $\tilde{u}_1(\hat{p}_1, k_i) = \frac{1}{2}u_1(r_2, k_i) + \frac{1}{2}u_1(r_3, k_i)$, inequality (6) tells us that

$$u_1(r_1, k_i) < \tilde{u}_1(\hat{p}_1, k_i)$$

for all pure strategies k_i of the column player. However, each mixed strategy is a convex combination of pure strategies and expected utility functions are linear in each variable, so it follows from the inequality above that

$$\tilde{u}_1(r_1, p_2) < \tilde{u}_1(\hat{p}_1, p_2)$$

for all mixed strategies p_2 of the column player. Thus, in the mixed extension of the current game, the row player's pure strategy r_1 is strictly dominated by his mixed strategy $\hat{p}_1$. A rational player should therefore not choose the pure strategy r_1 because it gives him less expected payoff than the mixed strategy $\hat{p}_1$, no matter how the column player acts. □

The examples above serve as justification for the following general definition.

Definition 5.4.1 Let $\langle N, (A_i), (u_i)\rangle$ be a finite strategic game. The action $a_i \in A_i$ is *strictly dominated* if player i has a mixed strategy p_i that gives him a greater expected utility than the pure strategy a_i, regardless of the other players' choices of actions, i.e. if

$$u_i(x_{-i}, a_i) < \tilde{u}_i(x_{-i}, p_i)$$

for all $x_{-i} \in A_{-i}$.

Remark. If the action a_i is strictly dominated by the mixed strategy p_i, then it is also strictly dominated by the mixed strategy $\hat{p}_i$ which is defined by $\hat{p}_i(a_i) = 0$ and $\hat{p}_i(x_i) = p_i(x_i)/(1 - p_i(a_i))$ for all $x_i \in A_i \setminus \{a_i\}$. Therefore, there is no restriction to assume that $p_i(a_i) = 0$ in the definition above.

Proposition 5.4.1 *Assume that p^* is a mixed Nash equilibrium of the finite strategic game $\langle N, (A_i), (u_i)\rangle$ and that player i's action a_i is strictly dominated. Then, $p^*_i(a_i) = 0$.*

In other words, a strictly dominated action can not be included in a player's Nash equilibrium strategy with positive probability.

Proof. Since the action a_i is strictly dominated, player i has a mixed stragegy $\hat{p}_i$ such that

$$u_i(x_{-i}, a_i) < \tilde{u}_i(x_{-i}, \hat{p}_i)$$

for all $x_{-i} \in A_{-i}$. Let $p = (p_1, p_2, \dots, p_n)$ be an arbitrary strategy vector. The expected utilities $\tilde{u}_i(p_{-i}, a_i)$ and $\tilde{u}_i(p_{-i}, \hat{p}_i)$ are obtained as expected values of the functions $x_{-i} \mapsto u_i(x_{-i}, a_i)$ and $x_{-i} \mapsto \tilde{u}_i(x_{-i}, \hat{p}_i)$ with respect to the product measure on A_{-i} formed by the lotteries in p_{-i}. It therefore follows from the inequality above that

$$\tilde{u}_i(p_{-i}, a_i) < \tilde{u}_i(p_{-i}, \hat{p}_i).$$

We use this inequality when p is the mixed Nash equilibrium p^* and conclude that

$$\tilde{u}_i(p^*_{-i}, a_i) < \tilde{u}_i(p^*_{-i}, \hat{p}_i) \le \tilde{u}_i(p^*).$$

It now follows from the indifference principle that $p^*_i(a_i) = 0$, because the assumption $p^*_i(a_i) > 0$ implies that $\tilde{u}_i(p^*_{-i}, a_i) = \tilde{u}_i(p^*)$, which contradicts the inequality above. □

Definition 5.4.2 The game $G' = \langle N, (A'_i), (u'_i)\rangle$ is a *subgame* of the strategic game $G = \langle N, (A_i), (u_i)\rangle$ if $A'_i \subseteq A_i$ and the utility function u'_i is the restriction of the utility function u_i to the set $A' = A'_1 \times A'_2 \times \cdots \times A'_n$ for each player $i \in N$.

Since the utility functions of a subgame are restrictions of the utility functions of the original game, we may use the same notation for them without any risk of misunderstanding. The subgames of a game $\langle N, (A_i), (u_i)\rangle$ will thus be written as $\langle N, (A'_i), (u_i)\rangle$.

If an outcome a^* belonging to A' is a Nash equilibrium of the game G, then it is obviously also a Nash equilibrium of the subgame G'.

A mixed strategy p_i for player i in the subgame G' can also be perceived as a mixed strategy in the game G by simply extending the definition of p_i to the entire set A_i by defining $p_i(a_i) = 0$ for $a_i \in A_i \setminus A'_i$. Conversely, each mixed strategy p_i in the game G with the property that $p_i(a_i) = 0$ for all $a_i \in A_i \setminus A'_i$ can be perceived as a mixed strategy in the game G'. This allows us to perceive a player's set $\mathcal{L}(A'_i)$ of strategies in the subgame G' as a subset of the same player's strategy set $\mathcal{L}(A_i)$ in the game G.

This means that the mixed extension of a subgame G' of a game G is a subgame of the mixed extension of G. By applying our trivial observation above concerning Nash equilibria in games and subgames to the mixed extensions of the games, we therefore get the following result.

Proposition 5.4.2 *Let $p^* = (p^*_1, p^*_2, \dots, p^*_n)$ be a vector of mixed strategies in a subgame G' of the game G. If p^* is a mixed Nash equilibrium of the game G, then p^* is also a mixed Nash equilibrium of the subgame G'.*

Iterated elimination of strictly dominated actions

EXAMPLE 5.4.3 Consider the following strategic game G:

	k_1	k_2	k_3	k_4
r_1	$(2,3)$	$(2,4)$	$(2,3)$	$(4,2)$
r_2	$(4,2)$	$(3,3)$	$(0,2)$	$(2,1)$
r_3	$(1,4)$	$(1,2)$	$(0,0)$	$(3,1)$
r_4	$(1,0)$	$(2,1)$	$(5,5)$	$(3,2)$

The row player's action r_3 is strictly dominated by the action r_1 and will not be selected by any rational player. Therefore, let us delete the r_3 action, which results in the following subgame G^1:

	k_1	k_2	k_3	k_4
r_1	$(2,3)$	$(2,4)$	$(2,3)$	$(4,2)$
r_2	$(4,2)$	$(3,3)$	$(0,2)$	$(2,1)$
r_4	$(1,0)$	$(2,1)$	$(5,5)$	$(3,2)$

Now, we see that the column player's action k_1 is strictly dominated by k_2, and that k_4 is strictly dominated by k_3. Therefore, we eliminate k_1 and k_4 from the game because these actions will not be selected by any rational column player of the game G^1. This gives us the following subgame G^2 of the game G^1:

	k_2	k_3
r_1	$(2,4)$	$(2,3)$
r_2	$(3,3)$	$(0,2)$
r_4	$(2,1)$	$(5,5)$

In the game G^2 the row player's action r_1 is strictly dominated by his mixed strategy $(0, \frac{1}{2}, \frac{1}{2})$, which gives him the expected utility $\frac{5}{2}$, regardless of the column player's choice. So we eliminate the alternative r_1 and obtain the subgame G^3:

	k_2	k_3
r_2	$(3,3)$	$(0,2)$
r_4	$(2,1)$	$(5,5)$

No action is strictly dominated in this subgame.

Rational players of the game G should, if they believe that their opponent is also rational, by reasoning as we have done, come to the conclusion that they should only choose actions that occur in the subgame G_3, i.e. the row player should select r_2 or r_4 and the column player should choose k_2 or k_3. In Chapter 7, we will provide further support for this conclusion. □

The reasoning in the example above can of course be generalized and suggests the following definition.

Definition 5.4.3 Let $G = \langle N, (A_i), (u_i)\rangle$ be a strategic game. We say that the subset $B = B_1 \times B_2 \times \cdots \times B_n$ of $A = A_1 \times A_2 \times \cdots \times A_n$ *survives iterated elimination of strictly dominated actions* if there exists a finite sequence $A^t = A_1^t \times A_2^t \times \cdots \times A_n^t$, $t = 0, 1, 2, \ldots, T$, of product sets that satisfies the following conditions:

- $A^0 = A$ and $A^T = B$;
- $A^{t+1} \subseteq A^t$ for $t = 0, 1, \ldots, T-1$;
- The actions in the sets $A_i^t \setminus A_i^{t+1}$ are strictly dominated in the game $G^t = \langle N, (A_i^t), (u_i)\rangle$ for $t = 0, 1, \ldots, T-1$;
- No action in the game $G^T = \langle N, (B_i), (u_i)\rangle$ is strictly dominated.

It follows from the remark after Definition 5.4.1 that we may assume that each action in the set $A_i^t \setminus A_i^{t+1}$ is strictly dominated by a mixed strategy that belongs to the game G^{t+1}.

In Chapter 7, we will show that the set B of surviving outcomes is unique, which is not obvious as there may be several different ways to perform the eliminations of strictly dominated actions. (See Corollary 7.2.3.) Of course it may happen that $B = A$, i.e. that there are no strictly dominated actions in the original game G.

If B consists of just one outcome $\hat{a}$, we say that the game G is *solvable by repeated strict dominance* with $\hat{a}$ as the solution.

EXAMPLE 5.4.4 The outcomes of the subgame G^3 are the surviving outcomes of the game G in Example 5.4.3 with respect to iterated elimination of strictly dominated actions. □

Proposition 5.4.3 *Let $G = \langle N, (A_i), (u_i)\rangle$ be a finite strategic game, suppose that the set $B = B_1 \times B_2 \times \cdots \times B_n$ survives iterated elimination of strictly dominated actions, and let $p^* = (p_1^*, p_2^*, \ldots, p_n^*)$ be a vector of mixed strategies in the game G. The following two assertions are then equivalent:*

(i) p^* is a mixed Nash equilibrium of the game G.

(ii) $p_i^*(a_i) = 0$ for each player i and each action $a_i \in A_i \setminus B_i$, and p^* is a mixed Nash equilibrium of the subgame $H = \langle N, (B_i), (u_i)\rangle$.

Actions that disappear during iterated elimination of strictly dominated actions can thus not occur with positive probability in any Nash equilibrium strategy.

Proof. Let $G^0 = G$, G^1, G^2, …, G^{T-1}, $G^T = H$ be the chain of subgames that appears in Definition 5.4.3. Due to induction, it is sufficient to show that the following two statements are equivalent for an arbitrary strategy vector p^* in the game G^t.

(a) p^* is a mixed Nash equilibrium of the game G^t.

(b) $p_i^*(a_i) = 0$ for all $i \in N$ and all $a_i \in A_i^t \setminus A_i^{t+1}$, and p^* is a mixed Nash equilibrium of the subgame G^{t+1}.

Assume first that statement (a) applies and that a_i is an action in the set $A_i^t \setminus A_i^{t+1}$. The action a_i is then strictly dominated in the game G^t, so it follows from Proposition 5.4.1 that $p_i^*(a_i) = 0$. The strategy p_i^* can therefore be perceived as a mixed strategy in the subgame G^{t+1} for player i, and the strategy vector p^* is according to Proposition 5.4.2 a Nash equilibrium of the subgame G^{t+1}. Thus we have shown that (a) implies (b).

Conversely, suppose that statement (b) applies. Due to the indifference principle we have

$$(7) \qquad \tilde{u}_i(p^*_{-i}, a_i) \begin{cases} = \tilde{u}_i(p^*_{-i}, p^*_i) & \text{if } p^*_i(a_i) > 0, \\ \leq \tilde{u}_i(p^*_{-i}, p^*_i) & \text{if } p^*_i(a_i) = 0 \end{cases}$$

for each player i and each action $a_i \in A_i^{t+1}$.

For each action $a_i \in A_i^t \setminus A_i^{t+1}$, player i has a mixed strategy $\hat{p}_i$ in the subgame G^{t+1} that strictly dominates the action a_i, which implies that

$$(8) \qquad \tilde{u}_i(p^*_{-i}, a_i) < \tilde{u}_i(p^*_{-i}, \hat{p}_i) \leq \tilde{u}_i(p^*_{-i}, p^*_i).$$

Since moreover $p^*_i(a_i) = 0$ for all $a_i \in A_i^t \setminus A_i^{t+1}$, by assumption, we conclude from the inequalities (7) and (8), that the strategy vector p^* satisfies the indifference principle in the game G^t and hence is a Nash equilibrium of G^t. □

Example 5.4.5 The surviving actions of the game G in Example 5.4.3 are r_2, r_4 and k_2, k_3. A strategy vector (p^*_1, p^*_2) is therefore a Nash equilibrium of G if and only if $p^*_1(r_1) = p^*_1(r_3) = 0$, $p^*_2(k_1) = p^*_2(k_4) = p^*_2(k_5) = 0$ and the restriction of (p^*_1, p^*_2) to the subgame G^3 is a Nash equilibrium of the subgame. The subgame G^3 has two pure Nash equilibria, namely (r_2, k_2) and (r_4, k_3), and one mixed Nash equilibrium that is easily computed using the indifference principle, and it consists of the row player choosing r_2 and r_4 with the probabilities $\frac{4}{5}$ and $\frac{1}{5}$, respectively, and the column player choosing k_2 and k_3 with the probabilities $\frac{5}{6}$ and $\frac{1}{6}$, respectively. We conclude that the original game G has three Nash equilibria, the two pure equilibria (r_2, k_2) and (r_4, k_3), and the mixed Nash equilibrium that consists of the strategies $(0, \frac{4}{5}, 0, \frac{1}{5})$ and $(0, \frac{5}{6}, \frac{1}{6}, 0, 0)$. □

Example 5.4.6 In the game with payoff matrix

	k_1	k_2	k_3
r_1	$(2,5)$	$(1,3)$	$(2,2)$
r_2	$(4,1)$	$(3,2)$	$(0,1)$
r_3	$(1,4)$	$(0,3)$	$(1,4)$

row 3 is strictly dominated by row 1. Elimination of row 3 results in a game in which column 3 is strictly dominated by column 2. Elimination of column 3 leads to a game where row 2 strictly dominates row 1. After having eliminated the first row, there remains a subgame in which column 2 strictly dominates column 1. Iterated elimination of strictly dominated actions thus results in the trivial subgame

	k_2
r_2	$(3,2)$

We conclude that the original game has a unique Nash equilibrium, namely the pure Nash equilibrium (r_2, k_2). □

Weak dominance

Definition 5.4.4 Let $\langle N, (A_i), (u_i)\rangle$ be a finite strategic game. The action $a_i \in A_i$ is said to be *weakly dominated* if player i has a mixed stragegy p_i that gives him at least as great expected utility as the pure strategey a_i, no matter what actions the other players choose, and a greater expected utility in some case, i.e. if

$$u_i(x_{-i}, a_i) \leq \tilde{u}_i(x_{-i}, p_i)$$

for all $x_{-i} \in A_{-i}$ with strict inequality for at least one x_{-i}.

EXAMPLE 5.4.7 Consider the game

	L	R
T	$(1,2)$	$(0,3)$
M	$(2,1)$	$(0,1)$
B	$(2,4)$	$(1,2)$

The action T is weakly dominated by M, M is weakly dominated by B, and T is strictly dominated by B. □

Elimination of weakly dominated actions does not yield as satisfactory results as the elimination of strictly dominated actions. For example, Nash equilibria may disappear during elimination of weakly dominated actions. However, every Nash equilibrium of a subgame that remains after elimination of weakly dominated actions is a Nash equilibrium also of the original game.

EXAMPLE 5.4.8 The outcome (B, R) is a Nash equilibrium of the game

	L	R
T	(1,1)	(0,0)
B	(0,0)	(0,0)

in spite of the fact that the row player's action B is weakly dominated by the action T, and the column player's action R is weakly dominated by the action L. □

The set of actions that survive iterated elimination of weakly dominated actions may also depend on the order in which the actions are eliminated.

EXAMPLE 5.4.9 Consider the following game:

	L	R
T	$(1,1)$	$(0,0)$
M	$(1,1)$	$(2,1)$
B	$(0,0)$	$(2,1)$

By first eliminating T (weakly dominated by M) and then eliminating L (weakly dominated by R), we are left with the two outcomes (M, R) and (B, R), both with payoff $(2, 1)$. If instead we first eliminate B (weakly dominated by M) and then R (weakly dominated by L), we are left with the two outcomes (T, L) and (M, L), both with payoff$(1, 1)$. □

Exercises

5.6 Solve the game Prisoner's Dilemma by iterated elimination of strictly dominated actions.

5.7 Find the subset that survives iterated elimination of strictly dominated actions in the game

	k_1	k_2	k_3	k_4
r_1	$(3,1)$	$(1,4)$	$(3,2)$	$(2,2)$
r_2	$(4,0)$	$(2,3)$	$(3,1)$	$(3,2)$
r_3	$(3,2)$	$(1,0)$	$(2,2)$	$(2,2)$
r_4	$(4,1)$	$(1,1)$	$(2,2)$	$(2,2)$

Also determine all mixed Nash equilibria.

5.8 Consider the game Guess 2/3 of the average in Exercise 2.5.
a) Is there any strictly dominated action in the game?
b) What outcomes survive iterated elimination of weakly dominated actions?

5.5 Maxminimizing strategies

Definition 5.5.1 By a player's *mixed safety level* in a finite strategic game G is meant the player's safety level in the mixed extension $\widetilde{G}$ of the game. A mixed strategy in the game G is called a *mixed maxminimizing strategy* if it is a maxminimizing action in the extension $\widetilde{G}$.

A mixed maxminimizing strategy $\hat{p}_i$ for player i is in other words a mixed strategy (i.e. a lottery over his set A_i of actions) that maximizes the function

$$f_i(p_i) = \min_{q_{-i}} \tilde{u}_i(q_{-i}, p_i),$$

where the minimum is taken over all strategy vectors q_{-i} for the other players.

The problem of computing a player's mixed safety level and mixed maxminimizing strategies is a straightforward optimization problem, and we can simplify the problem by noting that

$$f_i(p_i) = \min_{a_{-i}\in A_{-i}} \tilde{u}_i(a_{-i}, p_i) = \min_{a_{-i}\in A_{-i}} \sum_{a_i\in A_i} u_i(a_{-i}, a_i)p_i(a_i),$$

because

$$\tilde{u}_i(q_{-i}, p_i) \ge \min_{a_{-i}\in A_{-i}} \tilde{u}_i(a_{-i}, p_i)$$

for all strategy vectors q_{-i}, according to inequality (1) in Section 5.1, and this implies that the minimum of $\tilde{u}_i(q_{-i}, p_i)$ is assumed when q_{-i} is a vector consisting of pure strategies.

We leave as an exercise to verify that the functions f_i are *concave*, i.e. that $f_i(\alpha p+\beta q) \ge \alpha f_i(p)+\beta f_i(q)$ for all positive numbers α and β with sum 1 and all mixed strategies p, q. The problem of determining the maxminimizing strategies, i.e. to maximize f_i over the convex set $\mathcal{L}(A_i)$, is therefore a convex optimization problem. It is even better than that because the problem can easily be reformulated into a *linear programming problem*.

The smallest number of the numbers $t_1, t_2, \dots, t_m$ is namely equal to the largest number v that satisfies the inequalities $v \le t_1$, $v \le t_2$, $\dots$, $v \le t_m$. The function value $f_i(p)$ is therefore equal to the largest of all the numbers v that satisfy all the inequalities

$$v \le \sum_{a_i \in A_i} u_i(a_{-i}, x) p(a_i)$$

that are obtained by letting a_{-i} run through the set A_{-i}. The mixed safety level ℓ_i, i.e. the maximum value of $f_i(p_i)$, and the maxminimizing strategies are thus obtained by solving the optimization problem

$$\text{(9)} \qquad \text{Maximize } v \text{ as} \quad \begin{cases} v \le \sum_{a_i \in A_i} u_i(a_{-i}, a_i) p(a_i) & \text{for all } a_{-i} \in A_{-i} \\ p_i \in \mathcal{L}(A_i) \end{cases}$$

Suppose A_i consists of m possible actions $e_1, e_2, \dots, e_m$, and let us introduce the variables x_1, x_2, $\dots$, x_m by defining $x_k = p_i(e_k)$. The above maximization problem is then a problem in the $m+1$ variables $x_1, x_2, \dots, x_m$

and v with linear inequalities and equalities as constraints, because the last condition $p_i \in \mathcal{L}(A_i)$ is equivalent to the conditions $x_k \geq 0$ for $k = 1, 2, \dots, m$ and $x_1 + x_2 + \cdots + x_m = 1$. Problem (9) is therefore a typical linear programming problem, and for such problems there are efficient solution algorithms, for example the simplex algorithm.

Example 5.5.1 We will compute the mixed maxminimization strategies for the players of the game in Example 5.2.1 with payoff table

	L	R
T	$(3,3)$	$(0,2)$
B	$(2,1)$	$(5,5)$

Let p_1 be a mixed strategy of the row player, and let $x_1 = p_1(T)$ and $x_2 = p_1(B)$. The lottery set $\mathcal{L}(A_1)$ can be identified with the line segment

$$X = \{(x_1, x_2) \in \mathbf{R}^2 \mid x_1 + x_2 = 1,\, x_1, x_2 \geq 0\}$$

between the two points $(0,1)$ and $(1,0)$ in $\mathbf{R}^2$, and the row players maxminimization problem consists in maximizing the function

$$\begin{aligned} f_1(x_1, x_2) &= \min_{a_2 \in A_2} \sum_{a_1 \in A_1} u_1(a_1, a_2)\, p_1(a_1) \\ &= \min\bigl(u_1(T,L)p_1(T) + u_1(B,L)p_1(B), u_1(T,R)p_1(T) + u_1(B,R)p_1(B)\bigr) \\ &= \min(3x_1 + 2x_2, 0x_1 + 5x_2) \end{aligned}$$

when $(x_1, x_2) \in X$. The equivalent linear programming problem has the form

$$\text{Maximize } v \text{ as } \begin{cases} 3x_1 + 2x_2 \geq v \\ 0x_1 + 5x_2 \geq v \\ x_1 + x_2 = 1 \\ x_1, x_2 \geq 0 \end{cases}$$

The easiest way to solve this simple problem is to eliminate the variable x_1 $(= 1 - x_2)$ from the function f_1 and then solve the problem graphically. It is clear from Figure 5.2 that $f_1(1 - x_2, x_2) = \min(3 - x_2, 5x_2)$ assumes its largest value $\frac{5}{2}$ for $x_2 = \frac{1}{2}$. The mixed strategy $(\frac{1}{2}, \frac{1}{2})$ is in other words player 1's maxminimizing strategy, and his mixed safety level is equal to $\frac{5}{2}$.

Similarly, the column player should maximize the function

$$f_2(y_1, y_2) = \min(3y_1 + 2y_2, y_1 + 5y_2) = \min(3 - y_2, 1 + 4y_2)$$

when $y_1 + y_2 = 1$ och $y_1, y_2 \geq 0$. Maximum is obtained for $y_2 = \frac{2}{5}$, so the player's maxminimizing strategy consists in choosing L with probability $\frac{3}{5}$ and R with probability $\frac{2}{5}$. His mixed safefy level is $\frac{13}{5}$. □

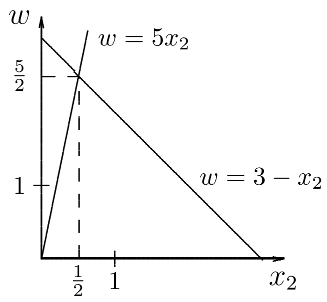

Figure 5.2. Graphical solution to the maximization problem in Example 5.5.1.

The two players' mixed safety levels $\frac{5}{2}$ and $\frac{13}{5}$ coincide with their expected utilities in the mixed Nash equilibrium in the above example. A player's expected utility in a Nash equilibrium can never be less than his safety level because of the following general result, which is an immediate consequnece of Proposition 2.4.1 applied to the mixed extension of a game.

Proposition 5.5.1 *In a mixed Nash equilibrium, each player's expected utility is greater than or equal to the player's mixed safety level.*

Exercise

5.9 Determine the mixed safety levels and mixed maxminimizing strategies for the players of the game in Exercise 5.3.

Chapter 6

Two-person Zero-sum Games

In Section 2.5, we studied strictly competitive two-person games. *Two-person zero-sum games* forms a special subclass of such games, and they are characterized by the players' preferences being given by cardinal utility functions whose sum is identical to zero. In this chapter we will study the mixed extension of finite zero-sum games and, above all, characterize their mixed Nash equilibria.

6.1 Optimal strategies and the value

A finite two-person zero-sum game is completely determined by the game's payoff matrix $A = [a_{ij}]$, which in the rest of this chapter is assumed to be a matrix with m rows and n columns, unless otherwise stated. Player 1, the row player, selects a row i in the matrix, and player 2, the column player, selects simultaneously, and unaware of the row player's choice, a column j in the matrix. Thereafter, the column player pays the amount of a_{ij} units to the row player, which in the case of $a_{ij} < 0$ means that the column player gets $-a_{ij}$ from the row player.

A two-person zero-sum game's pure Nash equilibria, if any, are characterized by the saddle point inequality (2) in Section 2.5. The pair (i, j) is a Nash equilibrium if and only if the matrix element a_{ij} is the largest in its column and the smallest in its row. However, most matrices lack saddle points, and this is a reason to study mixed extensions of zero-sum games.

The row player's set of lotteries will be denoted by X and the column player's set of lotteries will be denoted by Y. An element x in X is thus a probability distribution on the set of rows of the payoff matrix A, i.e. $x = (x_1, x_2, \ldots, x_m)$, where all entries x_i are nonnegative numbers and $\sum_{i=1}^{m} x_i = 1$. The number x_i stands for the row player's probability of choosing row i. Similarly, the elements $y \in Y$ have the form $y = (y_1, y_2, \ldots, y_n)$ with $y_j \geq 0$ for all j and $\sum_{j=1}^{n} y_j = 1$.

The sets X and Y are the two players' sets of actions in the mixed extension of the finite zero-sum game. The row player's expected payoff is his utility function in the extension; the expected payoff U is defined on the product set $X \times Y$ and is given by the formula

$$U(x, y) = \sum_{i=1}^{m} \sum_{j=1}^{n} a_{ij} x_i y_j.$$

The column player's expected payoff is of course equal to $-U$, so the mixed extension is also a zero-sum game.

We can now apply Nash's existence theorem (Proposition 5.2.1) and the characterization of Nash equilibria in strictly competitive games (Propositions 2.5.2 and 2.5.3). This immediately gives us the following proposition.

Proposition 6.1.1 (The Max-min theorem)

(i) *Every finite zero-sum game G has a mixed Nash equilibrium.*

(ii) *A mixed strategy vector (x^*, y^*) in G is a mixed Nash equilibrium if and only if x^* is a mixed maxminimizing strategy of the row player and y^* is a mixed maxminimizing strategy of the column player.*

(iii) *The row player's expected utility in a mixed Nash equilibrium equals his mixed safety level. The column player's mixed safety level is equal to the row player's mixed safety level with reversed sign.*

Definition 6.1.1 The row player's mixed safety level in a finite zero-sum game is called the *value* of the game. The game is said to be *fair* if its value is equal to zero.

The value of a zero-sum game is according to Proposition 6.1.1 equal to the row player's expected utility in an arbitrary mixed Nash equilibrium.

The Max-min theorem indicates that the Nash equilibrium is a stable and satisfactory solution concept for zero-sum games. The row player can determine a Nash equilibrium strategy x^* without worrying about his opponent's action, since x^* is a maxminimizing strategy, and by choosing such a strategy he has a guaranteed expected payoff that is at least as great as the game's value. The same applies to the column player, who can make sure on average not to lose more than the game's value. For that reason, strategies that are part of a mixed Nash equilibrium are also called *optimal strategies*.

However, note that the Max-min theorem is about expected values; in individual games, payouts to the row player may, of course, be less than the value of the game.

The indifference principle applies to zero-sum games and can be formulated as follows for such games.

Proposition 6.1.2 (Indifference principle) *In a finite zero-sum game with payoff matrix* $A = [a_{ij}]$*, the mixed strategies* x^* *and* y^* *are optimal (i.e.* (x^*, y^*) *is a mixed Nash equilibrium) and* v^* *is the game's value, if and only if the following four implications are true:*

$$x_i^* > 0 \;\Rightarrow\; \sum_j a_{ij} y_j^* = v^*$$

$$x_i^* = 0 \;\Rightarrow\; \sum_j a_{ij} y_j^* \le v^*$$

$$y_j^* > 0 \;\Rightarrow\; \sum_i a_{ij} x_i^* = v^*$$

$$y_j^* = 0 \;\Rightarrow\; \sum_i a_{ij} x_i^* \ge v^*.$$

Example 6.1.1 The game Matching Pennies with payoff table

	Head	*Tail*
Head	$(1, -1)$	$(-1, 1)$
Tail	$(-1, 1)$	$(1, -1)$

is a two-person zero-sum game with payoff matrix

$$\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}.$$

The matrix has no saddle point, so there is no pure Nash equilibrium. Any mixed Nash equilibrium (x^*, y^*) with $x_1^*, x_2^* > 0$ must, because of the indifference principle, satisfy the equation $y_1^* - y_2^* = -y_1^* + y_2^* = v^*$, and combining this with the condition $y_1^* + y_2^* = 1$, we obtain $y_1^* = y_2^* = \frac{1}{2}$ and $v^* = 0$. For simimilar reasons, $x_1^* = x_2^* = \frac{1}{2}$ if (x^*, y^*) is a mixed Nash equilibrium with $y_1^*, y_2^* > 0$. We conclude that the game has a unique mixed Nash equilibrium $\bigl((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2})\bigr)$ in which both players should choose Head and Tail with the same probability $\frac{1}{2}$, and that the game is fair. □

As an application of the indifference principle, we will give a general formula for the Nash equilibrium in two-person zero-sum games where both players have two action options. Recall that a saddle point in a matrix is a matrix element that is the smallest in its row and the largest in its column.

Proposition 6.1.3 *Consider a two-person zero-sum game in which both players have two action options, and suppose that the payoff matrix*

$$A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$$

has no saddle points. The game has then a unique mixed Nash equilibrium (x^*, y^*)*, given by*

$$\begin{aligned} x_1^* &= \frac{a_{22} - a_{21}}{s(A)}, \quad & x_2^* &= \frac{a_{11} - a_{12}}{s(A)}, \\ y_1^* &= \frac{a_{22} - a_{12}}{s(A)}, \quad & y_2^* &= \frac{a_{11} - a_{21}}{s(A)}, \end{aligned}$$

where

$$s(A) = a_{11} - a_{12} - a_{21} + a_{22}.$$

The value of the game is

$$v^* = \frac{\det A}{s(A)} = \frac{a_{11}a_{22} - a_{12}a_{21}}{s(A)}.$$

Proof. We first show that $s(A) \neq 0$ by considering three separate cases.

Case 1: $a_{11} = a_{12}$. We will prove that this contradicts the assumption that the matrix lacks saddle points. So assume that $a_{21} \geq a_{22}$. If $a_{11} \geq a_{21}$, then $(1,1)$ and $(1,2)$ are both saddle points, if $a_{21} > a_{11} \geq a_{22}$, then $(1,2)$

is a saddle point, and if finally $a_{21} > a_{22} > a_{11}$, then $(2, 2)$ is a saddle point. This is a contradiction. Hence, $a_{21} \leq a_{22}$, but this assumption leads to a contradiction in a similar way. Thus, case 1 is impossible.

Case 2: $a_{11} > a_{12}$. We now conclude that $a_{22} > a_{12}$ (because $(1, 2)$ is not a saddle point), that $a_{22} > a_{21}$ (because $(2, 2)$ is not a saddle point), and that $a_{11} > a_{21}$ (because $(2, 1)$ is not a saddle point).

Case 3: $a_{11} < a_{12}$. As in case 2, it now follows that $a_{21} > a_{11}$, $a_{21} > a_{22}$ and $a_{12} > a_{22}$.

The number $s(A)$ $(= (a_{11} - a_{12}) + (a_{22} - a_{21}))$ is positive in case 2 and negative in case 3. The four numbers x_1^*, x_2^*, y_1^* and y_2^* are positive in both cases, and the sums $x_1^* + x_2^*$ and $y_1^* + y_2^*$ are both equal to 1.

That (x^*, y^*) is a mixed Nash equilibrium follows immediately from the indifference principle, because

$$a_{11}y_1^* + a_{12}y_2^* = a_{21}y_1^* + a_{22}y_2^* = a_{11}x_1^* + a_{21}x_2^* = a_{12}x_1^* + a_{22}x_2^* = v^*.$$

For example,

$$a_{11}y_1^* + a_{12}y_2^* = \frac{a_{11}(a_{22} - a_{12}) + a_{12}(a_{11} - a_{21})}{s(A)} = \frac{a_{11}a_{22} - a_{12}a_{21}}{s(A)} = v^*.$$

It remains to show that the Nash equilibrium is unique. Therefore, suppose that $(\hat{x}, \hat{y})$ is another mixed Nash equilibrium. Then, $(x^*, \hat{y})$ is also a Nash equilibrium with the same value v^*, and it now follows from the indifference principle that $a_{11}\hat{y}_1 + a_{12}\hat{y}_2 = v^*$. The strategy $\hat{y} = (\hat{y}_1, \hat{y}_2)$ is therefore a solution to the linear system

$$\begin{cases} a_{11}y_1 + a_{12}y_2 = v^* \\ \quad\; y_1 + \quad\; y_2 = 1. \end{cases}$$

The system's determinant $a_{11} - a_{12}$ is nonzero because case 1 above can not occur. The solution of the system is therefore unique, and since the Nash solution y^* also solves the system, we conclude that $\hat{y} = y^*$. That $\hat{x} = x^*$ is of course proven analogously. □

Example 6.1.2 Consider the zero-sum game with payoff matrix

$$A = \begin{bmatrix} 2 & -3 \\ -1 & 0 \end{bmatrix}.$$

The matrix lacks saddle points, so the optimal strategies x^* and y^* are obtained from the formulas in Proposition 6.1.3, which yield $s(A) = 2 - (-3) -$

$(-1) + 0 = 6$ and

$$x_1^* = \frac{0-(-1)}{6} = \frac{1}{6}, \quad x_2^* = \frac{2-(-3)}{6} = \frac{5}{6},$$
$$y_1^* = \frac{0-(-3)}{6} = \frac{1}{2}, \quad y_2^* = \frac{2-(-1)}{6} = \frac{1}{2}.$$

The value of the game is

$$v^* = \frac{2 \cdot 0 - (-3) \cdot (-1)}{6} = -\frac{1}{2}. \qquad \square$$

Exercises

6.1 In the game *Odd or Even*, two players choose one of the numbers 1 or 2 independently of each other. If the sum of the two chosen numbers is odd, the row player wins the product of the two numbers from the opponent. If the sum is even, he loses the product of the two numbers to the opponent. Set up the game's payoff matrix, and solve the game, i.e. determine the players' optimal strategies as well as the game's value.

6.2 Solve using the indifference principle the games with the following payoff matrices.

a) $\begin{bmatrix} 2 & 3 \\ 4 & 1 \end{bmatrix}$ b) $\begin{bmatrix} 1 & -1 & -1 \\ 0 & 2 & 1 \\ 0 & 0 & 3 \end{bmatrix}$

6.3 Decide whether $\overline{x} = (\frac{1}{9}, 0, \frac{1}{18}, \frac{5}{6})$ and $\overline{y} = (0, 0, \frac{1}{3}, \frac{2}{3})$ are optimal strategies in the game with payoff matrix

$$\begin{bmatrix} 0 & 3 & -3 & 2 \\ -2 & -3 & 2 & -1 \\ -5 & 0 & -3 & 2 \\ 3 & 0 & 1 & 0 \end{bmatrix}.$$

Find the game's value.

6.4 Solve the games with payoff matrices

a) $\begin{bmatrix} 2 & 5 & 3 \\ 5 & 4 & 4 \\ 7 & 2 & 3 \end{bmatrix}$ b) $\begin{bmatrix} 2 & 3 \\ 4 & -1 \end{bmatrix}$ c) $\begin{bmatrix} 4 & 5 & 2 \\ 3 & 4 & 4 \\ 1 & 6 & 3 \end{bmatrix}$.

[Hint: In c), start by eliminating strictly dominated actions.]

6.2 Two-person zero-sum games and linear programming

In order to determine his mixed maxminimizing strategies in the game with payoff matrix $A = [a_{ij}]$, the row player has to solve the optimization problem

$$\max_{x \in X} \min_{y \in Y} U(x, y),$$

where

$$U(x, y) = \sum_{i=1}^{m} \sum_{j=1}^{n} a_{ij} x_i y_j.$$

This problem can be reformulated as a linear programming problem by noting that the expected payoff

$$U(x, y) = \sum_{j=1}^{n} y_j \Big(\sum_{i=1}^{m} a_{ij} x_i\Big)$$

for each $x \in X$ is a weighted average of the n numbers $\sum_{i=1}^{m} a_{ij} x_i$, $(j = 1, 2, \dots, n)$ with y_j as the weights. A weighted average is a number between the smallest and the largest of the given numbers, and the weighted average is equal to the smallest of the given numbers if the smallest number is given weight 1 and all other numbers weight 0. Consequently,

$$\min_{y\in Y} U(x,y) = \min\{\sum_{i=1}^{m} a_{ij}x_i \mid j = 1,2,\dots,n\}.$$

But the smallest of a number of given numbers is equal to the largest number s that is less than or equal to the given numbers. The number $\min_{y\in Y} U(x,y)$ is therefore, for each x, equal to the largest real number s that satisfies the inequalities

$$\sum_{i=1}^{m} a_{ij}x_i \geq s$$

for all column indices j.

The row player's maxminimizing problem is therefore equivalent to the optimization problem of finding the largest real number s that satisfies the above inequalities for some x belonging to X. This means that the row player's maxminimizing strategies are obtained as solutions x to the linear maximization problem

$$\text{(P)}\qquad \text{Maximize } s \text{ as} \\ \begin{cases} a_{11}x_1 + a_{21}x_2 + \dots + a_{m1}x_m \geq s \\ a_{12}x_1 + a_{22}x_2 + \dots + a_{m2}x_m \geq s \\ \qquad\vdots \\ a_{1n}x_1 + a_{2n}x_2 + \dots + a_{mn}x_m \geq s \\ x_1 + \quad x_2 + \dots + \quad x_m = 1 \\ x_1,\, x_2, \dots, x_m \geq 0. \end{cases}$$

The column player's maxminimizing problem is analogous, but to maximize the minimal expected payoff $-U(x,y)$ is equivalent to the problem of minimizing maximum of $U(x,y)$. The column player's maxminimizing strategies are therefore obtained as solutions to the optimization problem

$$\min_{y\in Y} \max_{x\in X} U(x,y).$$

Using an argument simular to that above, we find that this problem is equivalent to the linear minimization problem

$$(\mathrm{P}') \qquad \text{Minimize } t \text{ as} \quad \begin{cases} a_{11}y_1 + a_{12}y_2 + \ldots + a_{1n}y_n \le t \\ a_{21}y_1 + a_{22}y_2 + \ldots + a_{2n}y_n \le t \\ \qquad \vdots \\ a_{m1}y_1 + a_{m2}y_2 + \ldots + a_{mn}y_n \le t \\ y_1 + y_2 + \ldots + y_n = 1 \\ y_1, y_2, \ldots, y_n \ge 0. \end{cases}$$

In summary, we have shown the following proposition.

Proposition 6.2.1 *A mixed strategy vector* (x^*, y^*) *in a finite two-person zero-sum game with payoff matrix* $A = [a_{ij}]$ *is a Nash equilibrium if and only if* x^* *and* y^* *are optimal solutions to the linear programming problems* (P) *and* (P′)*, respectively. The value of the game is equal to the maximum value of* (P) *and to the minimum value of* (P′).

The two problems (P) and (P′) are instances of *dual linear programming problems*. According to an important theorem in linear programming, the so called Duality Theorem, dual problems (with feasible points) have the *same optimal values*. This means in our case that the maximum value in the problem (P) is equal to the minimum value in the problem (P′). Since the optimal maximum value is the row players mixed safety level and the optimal minimum value is the column players mixed safety value with reversed sign, we get an alternative proof from the Duality Theorem for statement (iii) in the Max-min theorem (Proposition 6.1.1) and thus also, because of Proposition 2.5.1, for the existence of mixed Nash equilibria in two-person zero-sum games.

The indifference principle is, by the way, also a consequence of a corollary to the Duality Theorem (the so called Complementary Theorem)

In cases where the game's value is positive, we can rewrite the maximization problem (P) in a simpler form. Since we only need to consider positive s values in these cases, we can after division by s write the constraints on the form

$$\begin{cases} a_{11}x_1/s + a_{21}x_2/s + \ldots + a_{m1}x_m/s \geq 1 \\ a_{12}x_1/s + a_{22}x_2/s + \ldots + a_{m2}x_m/s \geq 1 \\ \qquad \vdots \\ a_{1n}x_1/s + a_{2n}x_2/s + \ldots + a_{mn}x_m/s \geq 1 \\ x_1/s + x_2/s + \ldots + x_m/s = 1/s \\ s > 0, \ \ x_1, x_2, \ldots, x_m \geq 0. \end{cases}$$

Maximizing s is equivalent to minimizing $1/s$. So by introducing new variables $z_i = x_i/s$ and by using the constraint equality for $1/s$, we can replace the maximization problem (P) by the minimization problem

$$\text{Minimize } \ z_1 + z_2 + \cdots + z_m \ \text{ as}$$
$$\begin{cases} a_{11}z_1 + a_{21}z_2 + \ldots + a_{m1}z_m \geq 1 \\ a_{12}z_1 + a_{22}z_2 + \ldots + a_{m2}z_m \geq 1 \\ \qquad \vdots \\ a_{1n}z_1 + a_{2n}z_2 + \ldots + a_{mn}z_m \geq 1 \\ z_1, z_2, \ldots, z_m \geq 0. \end{cases}$$

Let c be the minimum value of this problem and let z^* be a minimum point. The game's value v^* is then equal to c^{-1}, and $c^{-1}z^*$ is an optimal strategy for the row player.

The assumption that the value of the game should be positive is not a serious restriction. We can always accomplish this by adding a sufficiently large positive constant K to all matrix elements in the payoff matrix A so that all of these will be nonnegative. This increases the game's value by the same constant K but does not affect the players' optimal strategies.

Example 6.2.1 Consider the game in Example 6.1.2 with payoff matrix

$$A = \begin{bmatrix} 2 & -3 \\ -1 & 0 \end{bmatrix}.$$

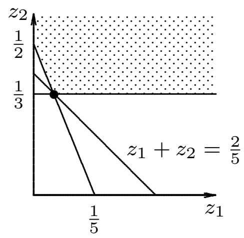

Figure 6.1. A graphical solution to Example 6.2.1.

To obtain a game with a positive value we add 3 to all matrix elements and get a new game with the matrix

$$\begin{bmatrix} 5 & 0 \\ 2 & 3 \end{bmatrix}.$$

In order to determine the row player's maxminimization strategies we solve the optimization problem

$$\text{Minimize } z_1 + z_2 \text{ as}$$
$$\begin{cases} 5z_1 + 2z_2 \geq 1 \\ \qquad\;\; 3z_2 \geq 1 \\ \;\; z_1,\, z_2 \geq 0. \end{cases}$$

graphically. See figure 6.1. The minimum is assumed at the intersection $z^* = (\frac{1}{15}, \frac{1}{3})$ of the lines $5z_1 + 2z_2 = 1$ and $3z_2 = 1$, and it is equal to $\frac{2}{5}$. This means that the modified game's value is equal to $\frac{5}{2}$, and that the row player's optimal strategy in both the modified and the original game is $x^* = \frac{5}{2}z^* = (\frac{1}{6}, \frac{5}{6})$. The original game's value is $\frac{5}{2} - 3 = -\frac{1}{2}$.

The column player's optimal strategy (y_1^*, y_2^*) is most easily calculated using the indifference principle according to which

$$\begin{cases} 2y_1^* - 3y_2^* = -\frac{1}{2} \\ -y_1^* \qquad\quad = -\frac{1}{2}. \end{cases}$$

It follows that $y_1^* = y_2^* = \frac{1}{2}$. □

Exercise

6.5 Charlie and Rick have three playing cards each. Both have ace of diamonds and ace of spades. Charlie also has 2 of diamonds and Rick has 2 of spades. The players simultaneously play a card each. Charlie wins if both of these cards have the same color and loses in the opposite case. The winner receives in payment the value of his winning card from the loser, with ace counting as 1. Write down the payoff matrix for this two-person game, and formulate column player Charlie's problem to optimize the expected payoff as an LP problem.

Chapter 7

Rationalizability

7.1 Beliefs

Consider a strategic game $\langle N, (A_i), (u_i)\rangle$, and suppose player i knows what action vector a_{-i} the other players will choose. The player's problem is then a pure decision problem, and if he is rational, he should of course choose an action that maximizes the function $x_i \mapsto u_i(a_{-i}, x_i)$.

Suppose the player instead believes that the other players choose their actions in the set A_{-i} randomly according to some probability distribution μ. The probability that the other players select the action vector a_{-i} is, in other words, $\mu(a_{-i})$, and player i's expected utility if he chooses the option $x_i \in A_i$ is

$$U_i(x_i;\mu) = \sum_{a_{-i}\in A_{-i}} u_i(a_{-i}, x_i)\mu(a_{-i}).$$

Given that player i wants to maximize his expected utility, his problem is still a pure decision problem. Now he has to maximize the function $x_i \mapsto U_i(x_i;\mu)$.

A player can of course not be sure that the other players play according to a certain probability distribution, so we use the following terminology.

Definition 7.1.1 A probability distribution on the set A_{-i} is called a *belief* of player i. An action $a_i \in A_i$ is said to be *supported by the belief* μ if it maximizes the function $x_i \mapsto U_i(x_i;\mu)$.

Note that we do not rule out beliefs in which the opponents cooperate and coordinate their actions. In other words, player i's belief does not have to be a product measure on A_{-i}.

Example 7.1.1 Let us determine some beliefs that support various actions in the two-person game with payoff matrix

	L	M	R
T	$(4,12)$	$(6,4)$	$(2,5)$
N	$(8,3)$	$(2,6)$	$(4,5)$
B	$(6,5)$	$(5,9)$	$(3,8)$

The row player's belief δ_M, i.e. the belief that the column player chooses his actions with the probabilities 0, 1 and 0 respectively, obviously supports the action T, while the belief δ_L instead supports the action N. The belief $\mu = 0.1\delta_L + 0.4\delta_M + 0.5\delta_R$ supports the action B, because the calculation

$$\begin{aligned} U_1(T;\mu) &= 0.1 \cdot 4 + 0.4 \cdot 6 + 0.5 \cdot 2 = 3.8, \\ U_1(N;\mu) &= 0.1 \cdot 8 + 0.4 \cdot 2 + 0.5 \cdot 4 = 3.6, \\ U_1(B;\mu) &= 0.1 \cdot 6 + 0.4 \cdot 5 + 0.5 \cdot 3 = 4.1, \end{aligned}$$

shows that $U_1(B;\mu)$ is the largest.

The column player has no belief $\nu = \alpha_1\delta_T + \alpha_2\delta_N + \alpha_3\delta_B$ that supports his action R, because by eliminating α_1 $(= 1 - \alpha_2 - \alpha_3)$ we get

$$\begin{aligned} U_2(L;\nu) &= 12\alpha_1 + 3\alpha_2 + 5\alpha_3 = 12 - 9\alpha_2 - 7\alpha_3, \\ U_2(M;\nu) &= 4\alpha_1 + 6\alpha_2 + 9\alpha_3 = 4 + 2\alpha_2 + 5\alpha_3, \\ U_2(R;\nu) &= 5\alpha_1 + 5\alpha_2 + 8\alpha_3 = 5 + 3\alpha_3, \end{aligned}$$

and it is easy to verify that $U_2(M;\nu) > U_2(R;\nu)$ if $\alpha_2 + \alpha_3 > \frac{1}{2}$, and that $U_2(L;\nu) > U_2(R;\nu)$ if $\alpha_2 + \alpha_3 \le \frac{1}{2}$. In other words, there is no belief ν that makes $U_2(R;\nu)$ greatest. □

Our next proposition shows that actions without support of any belief can not be included with positive probability in a player's Nash equilibrium strategy.

Proposition 7.1.1 *Let p^* be a mixed Nash equilibrium of the strategic game $\langle N, (A_i), (u_i)\rangle$. Each player i has a belief that supports all actions $a_i \in A_i$ such that $p_i^*(a_i) > 0$.*

Proof. Let μ be the product lottery on A_{-i} that is obtained by multiplying the lotteries $p_1^*, \dots, p_{i-1}^*, p_{i+1}^*, \dots, p_n^*$ that form the strategy vector p_{-i}^*. Then

$$U_i(x_i;\mu) = \sum_{a_{-i}\in A_{-i}} u_i(a_{-i}, x_i)\mu(a_{-i}) = \tilde{u}_i(p_{-i}^*, x_i)$$

for all $x_i \in A_i$. If $a_i \in A_i$ is an action with $p_i^*(a_i) > 0$, then $\tilde{u}_i(p_{-i}^*, a_i) = \tilde{u}_i(p^*)$ because of the indifference principle, and it follows that

$$U_i(a_i;\mu) = \tilde{u}_i(p_{-i}^*, a_i) = \tilde{u}_i(p^*) \ge \tilde{u}_i(p_{-i}^*, x_i) = U_i(x_i;\mu)$$

for all $x_i \in A_i$. The action a_i is thus supported by the belief μ. □

What will be the outcome of a game with rational participants? If a player's action is not supported by any belief, the player should not reasonably choose the action in question. A natural common sense requirement is that each player makes his choice among the actions that are supported by a belief.

Therefore, it seems natural to start by characterizing the set of actions that lack support in beliefs.

Proposition 7.1.2 *A player's action in a game* $G = \langle N, (A_i), (u_i) \rangle$ *lacks beliefs that support the action, if and only if the action is strictly dominated.*

Proof. By renumbering the players we may without loss of generality restrict our attention to player 1. So let $\hat{a}_1$ be an action of player 1. We will prove that the player has no belief that supports $\hat{a}_1$ if and only if the action is strictly dominated. To achieve this we consider the two-person zero-sum game $G' = \langle \{1, 2\}, (B_i), (v_i) \rangle$ with $B_1 = A_1 \setminus \{\hat{a}_1\}$ as player 1's set of actions, $B_2 = A_{-1}$ as player 2's set of actions, and the following function v as player 1's utility function:

$$v(b_1, b_2) = u_1(b_1, b_2) - u_1(\hat{a}_1, b_2).$$

The function v is well-defined since (b_1, b_2) and $(\hat{a}_1, b_2)$ are outcomes in the game G. Player 2's utility function is of course $-v$.

A mixed strategy of player 2 in the game G' is by definition a probability distribution on the set A_{-1}, i.e. a belief μ of player 1 in the game G, and a mixed strategy p_1 of player 1 in the game G' can be perceived as a mixed strategy of player 1 in the game G with the property that $p_1(\hat{a}_1) = 0$. Let P denote the set of all mixed strategies of player 1, and let M denote the set of all mixed strategies of his opponent player 2 in the game G'.

Player 1's expected utility of his mixed strategy p_1 when player 2 chooses the mixed strategy μ is

$$\tilde{v}(p, \mu) = \sum_{b_1 \in B_1} \sum_{b_2 \in B_2} v(b_1, b_2) p(b_1) \mu(b_2).$$

If in particular $p_1 = \delta_{a_1}$ is a pure strategy for player 1 (i.e. a_1 is an action in the set B_1) and $\mu \in M$ is an arbitrary mixed strategy of player 2, then

$$\begin{aligned} \text{(1)} \qquad \tilde{v}(\delta_{a_1}, \mu) &= \sum_{b_2 \in B_2} v(a_1, b_2) \mu(b_2) \\ &= \sum_{b_2 \in A_{-1}} u_1(a_1, b_2) \mu(b_2) - \sum_{b_2 \in A_{-1}} u_1(\hat{a}_1, b_2) \mu(b_2) \\ &= U_1(a_1; \mu) - U_1(\hat{a}_1; \mu). \end{aligned}$$

By definition, there is no belief that supports the action $\hat{a}_1$ in the game G if and only if for each probability distribution μ on A_{-1} there exists an action a_1 of player 1 such that $U_1(a_1; \mu) > U_1(\hat{a}_1; \mu)$, and using equation (1) we conclude that there is no belief that supports $\hat{a}_1$ if and only if for each mixed strategy $\mu \in M$ in the game G' there is a pure strategy $\delta_{a_1} \in P$ such that $\tilde{v}(\delta_{a_1}, \mu) > 0$.

The expected value $\tilde{v}(p_1, \mu)$ is for each mixed strategy $p_1 \in P$ a convex combination of the expected values $\tilde{v}(\delta_{a_1}, \mu)$, where a_1 belongs to the set B_1. Therefore, there exists, given $\mu \in M$, an action $a_1 \in B_1$ such that $\tilde{v}(\delta_{a_1}, \mu) > 0$ if and only if there exists a mixed strategy $p_1 \in P$ such that $\tilde{v}(p_1, \mu) > 0$.

To summarize, we have shown that there is no belief μ that supports $\hat{a}_1$ if and only if

$$\max_{p_1 \in P} \tilde{v}(p_1, \mu) > 0 \quad \text{for all } \mu \in M,$$

that is if and only if

$$\min_{\mu \in M} \max_{p_1 \in P} \tilde{v}(p_1, \mu) > 0.$$

But

$$\max_{p_1 \in P} \min_{\mu \in M} \tilde{v}(p_1, \mu) = \min_{\mu \in M} \max_{p_1 \in P} \tilde{v}(p_1, \mu)$$

according to the Max-min theorem for zero-sum games (Proposition 6.1.1). That there is no belief that supports $\hat{a}_1$ is therefore equivalent to the condition

$$\max_{p_1 \in P} \min_{\mu \in M} \tilde{v}(p_1, \mu) > 0,$$

i.e. to the condition that there is a mixed strategy p_1^* of player 1 in the game G' such that $\tilde{v}(p_1^*, \mu) > 0$ for all beliefs μ of player 1 in the game G.

Since $\tilde{v}(p_1^*, \mu)$ is linear function with respect to the second variable and since every belief μ is a convex combination of the safe lotteries $\delta_{a_{-1}}$ over the set A_{-1}, $\tilde{v}(p_1^*, \mu) > 0$ for all beliefs μ if and only if

$$\sum_{a_1 \in B_1} v(a_1, a_{-1}) p_1^*(a_1) = \tilde{v}(p_1^*, \delta_{a_{-1}}) > 0$$

for all action vectors a_{-1} in A_{-1}, i.e. if and only if

$$\begin{aligned} \tilde{u}_1(p_1^*, a_{-1}) - u_1(\hat{a}_1, a_{-1}) &= \sum_{a_1 \in B_1} \big(u_1(a_1, a_{-1}) - u_1(\hat{a}_1, a_{-1})\big) p_1^*(a_1) \\ &= \sum_{a_1 \in B_1} v(a_1, a_{-1}) p_1^*(a_1) > 0 \end{aligned}$$

for all $a_{-1} \in A_{-1}$, which means that the action $\hat{a}_1$ is strictly dominated by the mixed strategy p_1^*.

Thus, we have shown that there are no beliefs that support the action $\hat{a}_1$ if and only if the action is strictly dominated. □

EXAMPLE 7.1.2 The column player in the game in Example 7.1.1 has no belief that supports the action R, which is thus strictly dominated. Indeed, the mixed strategy $p_2 = 0.2\delta_L + 0.8\delta_M$, which means that the action L is selected with 20% probability and the action M is selected with 80% probability, dominates R strictly, since

$$\begin{aligned} \tilde{u}_2(T, p_2) &= 0.2 \cdot 12 + 0.8 \cdot 4 = 5.6 > 5 = u_2(T, R) \\ \tilde{u}_2(N, p_2) &= 0.2 \cdot 3 + 0.8 \cdot 6 = 5.4 > 5 = u_2(N, R) \\ \tilde{u}_2(B, p_2) &= 0.2 \cdot 5 + 0.8 \cdot 9 = 8.2 > 8 = u_2(B, R). \end{aligned}$$

On the other hand, none of the row player's actions are strictly dominated because everyone is supported by beliefs. □

Exercise

7.1 a) Show for the game in Example 7.1.1 that $\big((\frac{1}{3}, 0, \frac{2}{3}), (\frac{1}{3}, \frac{2}{3}, 0)\big)$ is a mixed Nash equilibrium.

b) Verify that the row player's belief $\frac{1}{3}\delta_L + \frac{2}{3}\delta_M$ supports the actions T och B, and that the column player's belief $\frac{1}{3}\delta_T + \frac{2}{3}\delta_B$ supports the actions L and M.

7.2 Rationalizability

We have argued that each player should only choose actions that he can support by beliefs. But all beliefs are not rational. For example, it does not seem sensible for a player to believe that another player will choose one for him strictly dominated alternative with positive probability. Each player should, as he forms his beliefs, assume that his opponents only will choose actions that they in turn can support by beliefs.

Example 7.2.1 Consider a game G with the following payoff table:

	L	M	R
T	$(2,2)$	$(1,1)$	$(4,0)$
B	$(1,2)$	$(4,1)$	$(3,5)$

The column player's action M is strictly dominated by L, and by eliminating the latter alternative we obtain the subgame G^1 with payoff table

	L	R
T	$(2,2)$	$(4,0)$
B	$(1,2)$	$(3,5)$

The row player's action B is strictly dominated in this game, so we eliminate B and obtain the subgame G^2:

	L	R
T	$(2, 2)$	$(4, 0)$

The column player's action R is now strictly dominated, and a final elimination leads to the trivial game G^3:

	L
T	$(2, 2)$

with just one outcome (T, L), which we recognize as the Nash equilibrium of the original game.

Now let us see what arguments are needed for the players to reach this conclusion.

1. A rational column player does not select the M action because it is strictly dominated by L, and according to Proposition 7.1.2 there is no belief under which M is the best choice.
2. If the row player knows that the column player is rational, then he realizes that the column player does not select the M action. He can therefore eliminate it from the discussion and is thus in the game G^1. If he is also rational himself, he does not choose B because that action is strictly dominated by T and therefore has no support in any belief.
3. If the column player is rational, knows that the row of player is rational and that the row player knows that the column player is rational, then the column player knows that the row player will select the T action. Therefore, only the game G^2 remains, and in that game, R is strictly dominated by L. Thus, the column player should select the L action, and thus the end result is that the players select the action vector (T, L). □

It may be worth noting what kind of knowledge the two players must possess for the above reasoning to be valid. The row player must know that the column player is rational. The column player must know that the row player knows that the column player is rational, which is a "higher level" of knowledge. It is not enough for a player to know that the opponent is rational, but he must also be sure that the opponent knows that the former player is rational. There are, of course, even higher levels of knowledge. I may know that my opponent is rational and that he knows I am. But he may not know I know he knows.

We will now clarify the above reasoning by introducing the concept of rationalization. In the game $G = \langle N, (A_i), (u_i)\rangle$, let us provisionally call player i's belief μ *reasonable* if $\mu(a_{-i}) = 0$ for all action vectors a_{-i} of the opponents that contain at least one action a_k that is not supported by any reasonable belief of player k. Also, let us call player i's action a_i *rationalizable* if it is supported by a reasonable belief of the player.

The recursive definition of the concept of reasonable belief may seem problematic and, of course, need to be specified. We can do that by first considering player 1, defining B_1 as the set of all actions of the player that are supported by beliefs. In the subgame G^1 with $A_{-1} \times B_1$ as set of outcomes, we define B_2 to be the subset of A_2 whose actions are supported by beliefs of player 2, and G^2 to be the subgame of G^2 with $B_1 \times B_2 \times A_3 \times \cdots \times A_n$ as set of outcomes. After n steps we have obtaine a subgame G^n with $B_1 \times B_2 \times \cdots \times B_n$ as set of outcomes, where B_n consists of the actions that are supported in the subgame G^{n-1} by beliefs of player n. Now we start again with player 1, defining C_1 to be the subset of B_1 that consists of the actions of player 1 that are supported by beliefs in the subgame G^n. After a number of laps we end up in a subgame G' with outcome set $Z_1 \times Z_2 \times \cdots \times Z_n$, in which, for each player k, every action $a_k \in Z_k$ is supported by some belief. These beliefs are the *reasonable* beliefs, and the actions belonging to Z_k are player k's rationalizable actions.

Instead of defining rationalizable actions in terms of reasonable convictions, however, it is easier to define the concept of rationalizability directly as follows.

Definition 7.2.1 Player i's action a_i in the game $\langle N, (A_i), (u_i)\rangle$ is *rationalizable* if for each $k \in N$ there is a subset Z_k with the following two properties:

(i) $a_i \in Z_i$.

(ii) For each $k \in N$ and each action $a_k \in Z_k$, player k has a belief μ that supports a_k and is zero outside the set Z_{-k}, i.e. $\mu(a_{-k}) = 0$ for all $a_{-k} \notin Z_{-k}$.

Note that if the sets $Z_1, Z_2, \dots, Z_n$ meet the conditions of the definition, then all actions in Z_k are rationalizable for each player k.

Are there always rationalizable actions? The affirmative answer is given by the next proposition.

Proposition 7.2.1 *In every finite game $\langle N, (A_i), (u_i)\rangle$ there is a unique largest nonempty subset $Z = Z_1 \times Z_2 \times \cdots \times Z_n$ of $A = A_1 \times A_2 \times \cdots \times A_n$ with the following property:*

For each player i and each action $a_i \in Z_i$, there exists a belief μ that supports a_i and is zero outside Z_{-i}.

Proof. We first show that there is a nonempty product set Z with the stated property and then that there is a largest such set.

Therefore, let p^* be a mixed Nash equilibrium of the game, and define the product set Z by for each player i letting

$$Z_i = \{a_i \in A_i \mid p_i^*(a_i) > 0\}.$$

It immediately follows from the proof of Proposition 7.1.1 that each action in Z_i is supported by the product measure μ on A_{-i} that is formed by the mixed strategies in p_{-i}^*, and μ is zero outside Z_{-i}.

To prove that there is a unique largest set Z with the stated property we consider all product sets that meet the condition of the proposition. There are of course only finitely many such sets $Z^1, Z^2, \dots, Z^m$.

Let $Z = Z_1 \times Z_2 \times \cdots \times Z_n$ be the product set obtained by defining $Z_i = Z_i^1 \cup Z_i^2 \cup \cdots \cup Z_i^m$ for each player i. This product set meets the condition of the proposition, because each action a_i in Z_i belongs to the set Z_i^k for some k and is therefore supported by some belief μ that is zero outside Z_{-i}^k and then apriori zero outside the larger set Z_{-i}. The set Z is obviously the unique largest set that meets the condition of the proposition since it contains all the sets Z^k as subsets. □

EXAMPLE 7.2.2 The product set of rationalizable actions in the game

	L	M	R
T	$(4, 12)$	$(6, 4)$	$(2, 5)$
N	$(8, 3)$	$(2, 6)$	$(4, 5)$
B	$(7, 5)$	$(1, 9)$	$(3, 8)$

is the set $Z = \{T, N\} \times \{L, M\}$. The row player's action T is supported by the belief δ_M and the action N is supported by the belief δ_L, and both beliefs are zero outside $\{L, M\}$. Similarly, the column player's action L is supported by the belief δ_T and the action M is supported by the belief δ_N. It is also easy to see that there is no larger set of rationalizable actions than Z. □

The proof of Proposition 7.2.1 shows that all actions that are included in a mixed Nash equilibrium with positive probability are rationalizable. It also follows from Proposition 7.1.2 that no strictly dominated action can be rationalizable. The precise connection between rationalizability and strict dominance is given by the following proposition.

Proposition 7.2.2 *An action in a finite strategic game survives iterated elimination of strictly dominated actions if and only if the action is rationalizable.*

Proof. Let Z be the product set of rationalizable actions in the game $G = \langle N, (A_i), (u_i)\rangle$, and let B be a product set of actions that survive iterated elimination of strictly dominated actions. Let further $G^t = \langle N, (A_i^t), (u_i)\rangle$, $t = 0, 1, \ldots, T$, denote the chain of subgames that through iterated elimination of strictly dominated actions leads from $G^0 = G$ to the subgame $G^T = \langle N, (B_i), (u_i)\rangle$ (compare with Definition 5.4.3). We will prove that $Z_i = B_i$ for each player i.

We begin by inductively proving that Z is a subset of the set $A^t = A_1^t \times \cdots \times A_n^t$ for $t = 0, 1, \ldots, T$. The initial step $t = 0$ is trivial since $A^0 = A$, so assume that $Z \subseteq A^t$ for some $t < T$.

All actions in Z_i are by definition supported by beliefs that are zero outside the set Z_{-i} and therefore also, because of the induction assumtion, zero outside the larger set A_{-i}^t. This means that all actions in Z_i are supported by beliefs in the game G^t, so it follows from Proposition 7.1.2 that no action in Z_i is strictly dominated when considered as an action in the subgame G^t. Thus, all actions in Z_i survive the elimination of strictly dominated actions that leads from G^t to the subgame G^{t+1}, which means that $Z_i \subseteq A_i^{t+1}$ and that consequently Z is a subset of A^{t+1}.

This completes the induction set, and at the endpoint $t = T$ we get the desired inclusion $Z \subseteq A^T = B$.

It remains to prove the converse $B \subseteq Z$, and for this purpose, it is sufficient to prove that the product set $B = A^T$ satisfies the condition in Proposition 7.2.1, i.e. that every action $a_i \in A_i^T$ is supported by some belief that is zero outside A_{-i}^T, because Z is the largest product set with this property.

Let a_i be an arbitrary action in A_i^T. Since a_i is not strictly dominated by any strategy in the subgame G^T, the action a_i is (according to Proposition 7.1.2 applied to the subgame G^T) supported by some belief μ of player i that is zero outside the set A_{-i}^T. This means that a_i maximizes the player's expected utility $U_i(x_i; \mu)$ as x_i ranges over the set A_i^T. The problem is that we have to show that the belief μ also supports the action a_i in the original game $G = G^0$, i.e. that a_i maximizes the expected utility $U_i(x_i; \mu)$ for all $x_i \in A_i$.

Suppose that is not the case. Then there is a $t < T$ such that a_i maximizes the expected utility $U_i(x_i; \mu)$ for all $x_i \in A_i^{t+1}$ but not for all $x_i \in A_i^t$. Let b_i an action in A_i^t that gives maximum; the action b_i is in other words an action in the game G^t that is supported by the belief μ. Since $U_i(b_i; \mu) > U_i(a_i; \mu)$, b_i can not belong to the set A_i^{t+1}, and hence $b_i \in A_i^t \setminus A_i^{t+1}$. This means that the action b_i has been eliminated at the transition from the subgame G^t to the subgame G^{t+1}, and it is therefore strictly dominated by some mixed strategy in the game G^t. This contradicts Proposition 7.1.2 since b_i is supported by the belief μ in the same game.

Our contradiction shows that the action a_i is supported by the belief μ also considered as an action in the game G, and this proves that B satisfies the condition of Proposition 7.2.1. □

Corollary 7.2.3 *The product set of actions that survives iterated elimination of strictly dominated actions is uniquely determined.*

Proof. The corollary follows immediately from Proposition 7.2.2, because the player's rationalizable actions are uniquely determined according to Proposition 7.2.1, □

Example 7.2.3 In the game in Example 7.2.2, the row player's action B is strictly dominated by the action N, and by eliminating B we obtain a subgame where the column player's action R is strictly dominated by the mixed strategy $0.25\delta_L + 0.75\delta_M$. Iterated elimination of strictly dominated actions thus leads to the subgame

	L	M
T	$(4, 12)$	$(6, 4)$
N	$(8, 3)$	$(2, 6)$

This is consistent with the fact that T and N are the row player's rationalizable actions, and L and M are the column player's rationalizable actions. $\square$

Chapter 8

Extensive Games with Perfect Information

The strategic form describes games in a very compact way, but it is still possible to analyze many interesting questions. However, the assumption that all players choose their action plans simultaneously and once and for all is a limitation that causes certain important aspects to be lost. In many conflict situations and games, participants can consider and modify their action plans as the situation evolves and not only at the beginning. The extensive form is a way to model such situations.

8.1 Game trees

Chess is a board game played by two players, White and Black, who alternate moves. White moves first and has 20 legal first moves to choose from. Then it is Black's turn to move a piece, and Black has also 20 possible first moves. After that, it is White's turn, and the number of legal moves now depends on the position of the pieces on the chessboard. After a number of moves, a position can arise in which the player to move does not have any legal moves, either because the player's king is placed in check and there is no move the player can make to escape check; in that case the player's king is checkmated and the opponent has won, or because the player is not in check but has no legal move; this situation is called a stalemate and the game is declared a draw. The game is also a draw when there is no possibility of checkmate for either side with any series of legal moves due to insufficient material, or if the same board position has occurred three times with the same player to move and all pieces having the same rights to move, or if there has been no capture or pawn move in the last fifty moves by each player. A game can

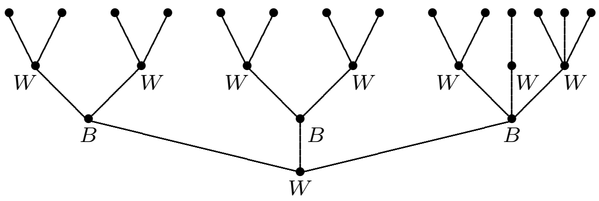

Figure 8.1. A game tree

also be terminated prematurely by one of the players resigning in a hopeless position or by the players agreeing on a draw. All chess games therefore end after a finite number of moves.

We can illustrate chess using a game tree like that in Figure 8.1. The root of the tree symbolizes the starting position of the game, and the branches from the root are White's possible first moves (twenty but in the figure we have just drawn three). After each of these we have come to a node in the tree (a position) where it is Black's turn to move, and Black has a number of possible moves each symbolized by a branch (in the figure, we have drawn a variable number despite the fact that Black in reality has twenty possible moves), etc. The tree in the figure, of course, shows only a small part of the enormously large game tree of chess, which is so big and complex that it is not possible to draw it in practice.

Each node in the game tree corresponds to a chess position, but note that there is no one-to-one correspondance between nodes and configurations of pieces on the board, because the same configuration may occur after different series of moves, and the legal moves at a certain position of the game are not entirely determined by the position of the pieces but may also depend on the previous moves. (For example, castling is not allowed if the king or the rook has been previously moved.) However, the branches that emanate from each node of the game tree are unambiguously defined by the chess rules.

The final positions of chess games, i.e. positions where one of the players has won or where the game has ended in a draw, corespond to nodes from which no branches emanate. Each path in the tree from the root up to a final position corresponds to a specific game of chess.

With the above description of chess as a model, we now make the following general definition.

Definition 8.1.1 A *game tree* $\langle P, p_0, s\rangle$ is a structure consisting of the following components:

- a nonempty set P of *positions*;
- a special element $p_0 \in P$, the *initial position*;
- a function $s\colon P \to \mathcal{P}(P)$ from the position set P to the set $\mathcal{P}(P)$ of all

subsets of P with the following property:

For each position $p \neq p_0$ there is a unique sequence $(p_k)_{k=1}^{n}$ of positions such that $p_n = p$ and $p_{k+1} \in s(p_k)$ for $k = 0, 1, \ldots, n-1$.

The positions belonging to $s(p)$ are called the *successors* of position p, and the set-valued function s is the *successor function*. Positions without successors are called *terminal positions*. The set of all nonterminal positions is denoted P°.

A pair (p, q) consisting of a nonterminal position p and a successor q of p is called a *move*.

A *move sequence* of *length* $n-1$ from p_1 is a sequences $(p_k)_{k=1}^{n}$ of positions such that p_{k+1} is a successor of p_k for $k = 1, 2, \ldots, n-1$, or equivalently, a sequence of $n-1$ moves (p_k, p_{k+1}), $k = 1, 2, \ldots, n-1$. An *infinite move sequence* from p_1 is an infinite sequence $(p_k)_{k=1}^{\infty}$ of positions such that p_{k+1} succeeds p_k for all positive integers k.

A move sequence that starts in the game tree's initial position p_0 and ends in a terminal position or continues indefinitely is called a *play*. The set of all plays in a game tree is denoted Π.

We will often assign names to the moves in a game tree in order to be able to refer to them comfortably. In addition, we will usually specify move sequences by listing the successive moves instead of the successive positions.

Example 8.1.1 We illustrate game trees by drawing them as "trees" with the initial position at the bottom as root and with line segments as "branches" from each position to each succeeding position. Figure 8.2 shows a game tree with nine positions. Position p_0 is the initial position, p_0, p_1, p_2 and p_3 are nonterminal positions, and p_4, p_5, p_6, p_7 and p_8 are terminal positions. We have given the moves letter names so that for example the move (p_0, p_1) is called A. The move sequence p_0, p_1, p_3, p_8 is a play of length 3 that can also be described as ACH. There are a total of five plays, namely ACG, ACH, AD, BE and BF. □

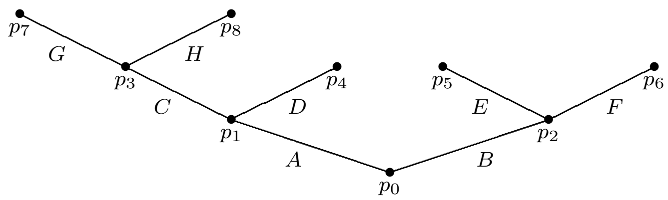

Figure 8.2. The game tree in Example 8.1.1.

Definition 8.1.2 A game tree is *finite* if its position set P is finite. A game tree is of *finite height* if there is a number c such that no play has length greater than c, and the *height* of the game tree is the length of the longest play in the tree.

All finite game trees are of course of finite height, but the converse does not hold.

Exercises

8.1 Give an example of a non-finite game tree of finite height.

8.2 Give an example of a game tree of infinite height where all plays have finite length.

8.2 Extensive form games

We have introduced the concept of game tree $\langle P, p_0, s \rangle$, but to make a game of the tree we also need players and rules that determine who to perform the moves in the nonterminal positions. Let $N = \{1, 2, \ldots, n\}$ be the set of players. We assume that in each nonterminal position p of the game tree only one player is allowed to make a move, and if we denote that player $Pl(p)$, we have thereby also defined a function Pl from the set P° of nonterminal positions to the set N of players. Obviously, it is unnecessary to decide who should act at a terminal position, as there are no moves to do there.

In order to give the players incentives to act, we also need to specify their preferences for the different outcomes of the game. In chess, each player prefers a win in front of a draw and a draw in front of a loss, which we can of course also indicate with a utility function that allocates 1 to wins, $\frac{1}{2}$ to draws and 0 to losses.

In an arbitrary game, we assume similarly that each player i has a preference relation $\succeq_i$ (or a utility function u_i) that is defined on the set Π of all plays. If there are no infinite plays, i.e. if each play ends in a terminal position, then there is a one-to-one correspondance between plays and terminal positions, and we may as well assume that the preference relations are defined on the set of all terminal positions in the game tree.

We are now ready for the formal definition of an extensive game.

Definition 8.2.1 An n-person game $\langle N, P, p_0, s, Pl, (\succeq_i)\rangle$ in *extensive form with perfect information* consists of

- a set $N = \{1, 2, \dots, n\}$ of *players*;
- a *game tree* $T = \langle P, p_0, s\rangle$;
- a function $Pl\colon P^\circ \to N$, *the player function*;
- for each player $i \in N$ a preference relation $\succeq_i$ on the set Π of all plays.

The game has a *finite horizon* if the game tree T has a finite height, and the game is *finite* if the game tree is finite.

Remark. The above definition does not require all players in N to participate by making moves. In other words, the player function Pl need not be surjective. The reason for allowing passive players is that we thereby simplify some inductive arguments that we will have to do later. For the same reason, we do not require that there be any nonterminal positions. A game with no nonterminal position is of course trivial − it only consists of the initial position, and the player function is empty.

The add "perfect information" is motivated by the fact that the players each time they have to perform a move know all previous moves and the exact position of the game. In the next chapter we will study games where this information is not available.

When we illustrate extensive games in figures, we describe the player function and the utility functions by entering the player value $Pl(p)$ at each nonterminal position p and the players' utility values $(u_1(p), \dots, u_n(p))$ at each terminal position p. However, we omit the utility values in cases where these do not matter in the current discussion.

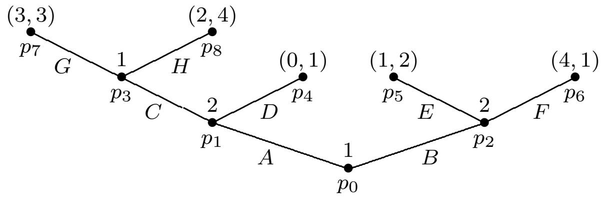

Figure 8.3. The game in Example 8.2.1.

In a specific instance of the extensive game, player $Pl(p_0)$ begins by selecting one of the possible moves at the initial position p_0, i.e. the player chooses a position $p_1 \in s(p_0)$. Then player $Pl(p_1)$, the player at position p_1, makes a move by selecting a position p_2 that succeeds p_1. Then player $Pl(p_2)$ chooses a successor p_3 of position p_2, and so on. The game is over as soon as any of the succeeding positions, p_k say, is a terminal position with the play $p_0, p_1, \dots, p_k$ as a result. An infinite play is obtained if the players constantly choose nonterminal positions.

Example 8.2.1 Figure 8.3 shows an extensive two-person game with the tree in Figure 8.2 as game tree. Player 1 moves in the initial position p_0 and in p_3, and player 2 moves in positions p_1 and p_2.

The move sequence ACH corresponds to the play $p_0p_1p_3p_8$, which results in 2 utility units for player 1 and 4 utility units for player 2. Therefore, both players prefer this play in front of for example the play BE. On the other hand, the play BF is better for player 1. □

Strategies

We will now define the concept of strategy for games in extensive form so that it corresponds to actions or pure strategies for games in strategic form.

A strategy for a player is a plan that describes what the player should do in **each** position of the game, where it is the player's turn to perform a move, regardless of whether the player can reach that position or not if he follows his plan.

This differs from the normal use of the concept of strategy in, for example, chess. No chess player can have a game plan that describes what to do in all possible positions, but the game plan develops during the game. Good chess players usually follow pre-planned plans during the opening game, let's say during the first ten moves, plans that mean that they know in advance how to respond to all sensible moves of the opponent. The first moves are therefore performed quickly and follow traditional courses.

However, theoretically we can imagine that both players have plans for everything (including positions that can not occur if a player follows his plan). In that case, it is of course unnecessary that the players themselves play the game. Instead, they can give their respective plans to a third party, a referee who can then play the game according to their intentions and see how it ends.

The corresponding applies to all extensive games if the players have complete plans for all positions in the game – the outcome is then completely determined by the plans.

We are now ready for the formal definition.

Definition 8.2.2 Let P_i denote the set of all positions in the extensive game $\langle N, P, p_0, s, Pl, (\succeq_i)\rangle$ where player i is supposed to make a move, i.e.

$$P_i = \{p \in P^\circ \mid Pl(p) = i\}.$$

A *strategy* σ for player i is a function $\sigma\colon P_i \to P$ with the property that $\sigma(p)$ is a successor of p for each $p \in P_i$.

The set of all strategies of player i is denoted by Σ_i, and we put

$$\Sigma = \Sigma_1 \times \Sigma_2 \times \cdots \times \Sigma_n.$$

Remark. If player i does not participate actively in the game, then P_i is the empty set and i's only strategy is the empty function This is of course compatible with the player not doing anything.

If σ is a strategy of player i and p is a position belonging to P_i, then $(p, \sigma(p))$ is by definition one of the player's possible moves at p. So instead of defining strategies as functions, we might as well define a strategy for player i as a *set of moves* that contains exactly one move (p, p') for each position $p \in P_i$. We will switch freely between these two equivalent interpretations and use the variant that is simplest in the current situation. In our examples and figures we often assign names to the moves but not to the positions, and then it is most natural to describe strategies by listing the moves that are contained in them.

EXAMPLE 8.2.2 Consider the game in Example 8.2.1. Player 1 has four strategies, which written as sets of moves are $\{A, G\}$, $\{A, H\}$, $\{B, G\}$ and $\{B, H\}$. From now on we will drop the set brackets and describe strategies by just listing their moves in an arbitrary order. Player 1 strategies can thus be described briefly as AG, AH, BG, and BH. Player 2 also has four strategies, and with the corresponding notation, these can be written as CE, CF, DE and DF.

The strategy AG means that player 1 chooses the move A in the initial position p_0 and the move G if the play reaches the position p_3.

The strategy BG means that player 1 chooses the move B in the initial position p_0 and that he should choose the move G in the position p_3. But if the player starts with the move B then it is inpossible for the play to reach the position p_3. In other words, the strategy describes an action option for a situation that can not occur if the player follows his strategy. This is a consequence of our definition of strategy, because the definition requires the player to specify a move in each position that belongs to him.

Is it possible to ease this requirement? The Nash equilibrium concept, which we will immediately study, can be formulated with a more intuitive strategy concept that only requires the strategies to be defined in positions that can arise if the strategies are followed. However, the Nash equilibrium concept is not particularly compelling for extensive games, so we will introduce an intuitively better equilibrium concept, and then we need the strategy concept in the version we have given above.

One way to understand the strategy BG, which does not make it so unreasonable, is to see it as a plan for what the player would do in the position p_3 if he by mistake would have chosen the move A instead of B in the initial position. □

Each strategy vector $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n) \in \Sigma$ determines uniquely a play $\pi(\sigma)$ in the game, i.e. a sequence of moves that starts in the initial position and either ends in a terminal position or goes on forever. The map $\sigma \mapsto \pi(\sigma)$ is in other words a function

$$\pi\colon \Sigma \to \Pi$$

from the product $\Sigma = \Sigma_1 \times \Sigma_2 \times \cdots \times \Sigma_n$ of the players' strategy sets to the set Π of all plays.

The formal definition of the play $\pi(\sigma)$, which is called the *outcome of the strategy vector* σ, is straightforward and looks like this:

Definition 8.2.3 Let $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n)$ be a strategy vector in the extensive game $\langle N, P, p_0, s, Pl, (\succeq_i)\rangle$.

Assume inductively that the position sequence $(p_j)_{j=0}^{k}$ has already been defined. Stop if p_k is a terminal position. Otherwise, let $i_k = Pl(p_k)$ be the player in turn to make a move in position p_k and set $p_{k+1} = \sigma_{i_k}(p_k)$. The position p_{k+1} is thus the position that succeeds the p_k position if player $Pl(p_k)$ follows his strategy.

The procedure will either end with a finite play $p_0, p_1, \dots, p_k$ or give rise to an infinite play, and the play obtained is the outcome $\pi(\sigma)$ of σ.

Example 8.2.3 We continue our study of the game in Example 8.2.1, which is again shown in Figure 8.4. Let σ_1 be player 1's strategy AH and σ_2 player 2's strategy CF. This means that player 1 will start with the move A leading to position p_1, where player 2 continues with the move C leading to position p_3, where player 1 does the move H leading to the terminal position p_8. The outcome $\pi(\sigma_1, \sigma_2)$ is thus the play $p_0p_1p_3p_8$, for which we also use the notation ACH.

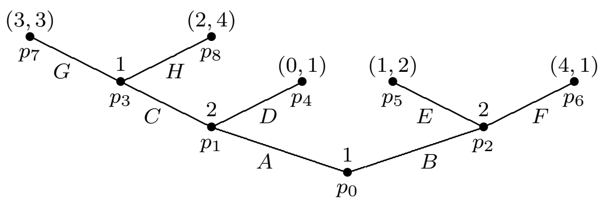

Figure 8.4

Player 1's strategy σ_1 and player 2's strategy $\sigma_2' = DE$ results correspondingly in the outcome $\pi(\sigma_1, \sigma_2') = AD$. A table of outcomes for all strategy combinations looks like this:

	CE	CF	DE	DF
AG	ACG	ACG	AD	AD
AH	ACH	ACH	AD	AD
BG	BE	BF	BE	BF
BH	BE	BF	BE	BF

□

Reduction from extensive to strategic form

The players' preference relations $\succeq_i$, or utility functions u_i, are apriori defined on the set Π of all plays. However, they give rise to preference relations $\succeq'_i$, or utility functions u'_i respectively, on the product Σ of strategy sets via the following definition:

$$\sigma \succeq'_i \tau \;\Leftrightarrow\; \pi(\sigma) \succeq_i \pi(\tau)$$

or

$$u'_i(\sigma) = u_i(\pi(\sigma)).$$

It is immediately verified that the relations $\succeq'_i$ are preference relations, respectively that the functions u'_i are utility functions, on Σ. This makes it possible to transfer a game in extensive form $\langle N, P, p_0, s, Pl, (\succeq_i)\rangle$ to a game in strategic form, namely to the strategic game $\langle N, (\Sigma_i), (\succeq_i')\rangle$. The latter game is called the *reduced strategic form* of the extensive game.

EXAMPLE 8.2.4 In the previous example, we determined the outcomes of all strategy combinations for the game in Example 8.2.1. In Figure 8.4, the payoffs to the players are listed at the terminal positions, and by entering these payoffs in the table of outcomes, we obtain the following strategic form of the game, where player 1 is the row player as usual:

	CE	CF	DE	DF
AG	$(3,3)$	$(3,3)$	$(0,1)$	$(0,1)$
AH	$(2,4)$	$(2,4)$	$(0,1)$	$(0,1)$
BG	$(1,2)$	$(4,1)$	$(1,2)$	$(4,1)$
BH	$(1,2)$	$(4,1)$	$(1,2)$	$(4,1)$

□

Nash equilibrium

We assume that every player of an extensive game tries to choose his strategy so that his outcome is as good as possible, but just like in strategic games there is generally no outcome that is best for all players. We must therefore focus on defining and studying different kinds of equilibrium solutions. The following equilibrium definition is natural, given the corresponding definition for strategic games.

Definition 8.2.4 A strategy vector $\sigma^* \in \Sigma = \Sigma_1 \times \Sigma_2 \times \cdots \times \Sigma_n$ in an extensive game is called a *Nash equilibrium* if, for each player i and each strategy $\tau_i \in \Sigma_i$,

$$\pi(\sigma^*_{-i}, \sigma^*_i) \succeq_i \pi(\sigma^*_{-i}, \tau_i).$$

The following proposition is a trivial consequence of the definition of the reduced strategic form of an extensive game.

Proposition 8.2.1 *A strategy vector $\sigma^* = (\sigma^*_1, \sigma^*_2, \dots, \sigma^*_n)$ is a Nash equilibrium of an extensive game if and only if it is a Nash equilibrium of the reduced strategic form of the game.*

EXAMPLE 8.2.5 The game in Example 8.2.1 has three Nash equilibria, namely (AG, CE), (BG, DE) and (BH, DE). This follows immediately from the payoff table in Example 8.2.4.

A Nash equilibrium should be a stable and trustworthy solution, but there is a problem with the Nash equilibrium (BG, DE) (and similarly with

(BH, DE)) which makes it difficult to take it seriously. The problem does not appear in the strategic formulation of the game, but it occurs as soon as you take into account that the players of extensive games choose their moves in succession.

The reason why the strategy vector (BG, DE) is a Nash equilibrium is the fact that player 2 has a potential threat that could prevent player 1 from choosing the strategy AG, namely the move D, which in that case would result in the play AD with payoff 0 to player 1. But this threat is not very credible, because if player 1 starts with the move A, then player 2 is guaranteed at least 3 utility units by the move C, which is better than 1 utility unit for the move D. Under such circumstances, a rational player should not choose D.

To get around this problem, we will introduce a more stringent equilibrium concept in the next section. □

Exercises

8.3 Consider the extensive game in Figure 8.5.

a) List the players' strategies.
b) Determine all Nash equilibria.
c) Reduce the game to strategic form.

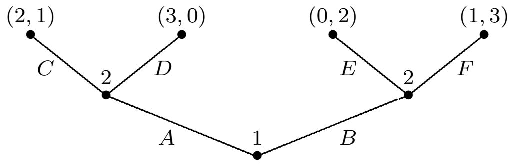

Figure 8.5

8.4 Two people should choose one of three options A, B and C by successively excluding one of them. First, person 1 excludes one of the options, and then person 2 excludes one of the two remaining options. The option remaining after these two rounds becomes the choice of the people. Formulate the procedure as an extensive game and determine the Nash equilibria in cases where the two persons' preferences for the alternatives are given by

a) $A \succ_1 B \succ_1 C$ and $C \succ_2 B \succ_2 A$ b) $A \succ_1 B \succ_1 C$ and $C \succ_2 A \succ_2 B$.

8.5 Charlie and Lisa use the following method to split $\$\,100$. Charlie offers Lisa $\$\,x$, where x is an arbitrary integer in the range of $0 \leq x \leq 100$. If Lisa accepts the offer, Charlie will retain the remaining $\$\,(100 - x)$. However, if Lisa rejects the offer, none of them will receive any money. Both Charlie and Lisa only care about the amount they receive and prefer to get as much as possible. Formulate the procedure as an extensive game and determine the Nash equilibria. (The game is usually called the *ultimatum game*.)

8.6 Consider the same situation as in the previous exercise, but assume that the amount offered may be any real number x in the interval $[0, 100]$. Formulate this situation as an extensive game, and determine the Nash equilibria. The new game is of course no longer a finite game.

8.3 Subgame perfect equilibria

By "cutting the branch" that leads to a position p in a game tree we get a new game tree with p as the initial position. We call this tree a subgame tree. The formal definition reads as follows:

Definition 8.3.1 Let p be a position in the game tree $T = \langle P, p_0, s\rangle$, and let P' be the set of positions that consists of p and all positions $q \in P$ that can be reached from p by a sequence of moves in the game tree. Let moreover s' denote the restriction of the successor function s to the set P', i.e. $s'(q) = s(q)$ for all $q \in P'$. Then $T_p = \langle P', p, s'\rangle$ is a game tree with p as initial position, and it is called a *subgame tree* of T.

Each play π' in a subgame tree T_p with p as the initial position can be uniquely extended to a play π in the original game tree T by starting with the unique sequence of moves that leads to p from the initial position p_0 in T, and then continuing with the play π'. The assignment $\pi' \mapsto \pi$ is obviously an injective map from the set Π_p of all plays in the subgame tree T_p to the set Π of all plays in T.

It is now obvious what should be meant by a subgame of an extensive game; a subgame is formed by starting in an arbitrary position, forgetting what has happened before. The exact definition looks like this:

Definition 8.3.2 Let $G = \langle N, P, p_0, s, Pl, (\succeq_i)\rangle$ be an extensive game, and let p be an arbitrary position in the game. The game $G_p = \langle N, T_p, Pl', (\succeq'_i)\rangle$, where

- T_p is the subgame tree of $T = \langle P, p_0, s\rangle$ with p as initial position;
- the player function Pl' is the restriction of the player function Pl to the set of all nonterminal positions in T_p;
- the preference relations $\succeq'_i$ are defined on the set Π_p of all plays π' in T_p by declaring $\pi'_1 \succeq'_i \pi'_2$ if and only if $\pi_1 \succeq_i \pi_2$, where π_1 and π_2 are the unique extensions in T of the plays π'_1 and π'_2;

is called a *subgame* of G.

Each strategy σ_i for player i in the game G gives naturally rise to a strategy σ'_i in the subgame G_p; the function σ'_i is simply the restriction of the strategy σ_i to those positions of the subgame G_p where it is the player's turn to perform a move.

Example 8.3.1 Figure 8.6 shows the game tree T of a game G that we have already met in Example 8.2.1 and three subgame trees T_{p_1}, T_{p_2} and T_{p_3}. Each play in the subgame T_{p_1} can be extended uniquely to a play in the game tree T; for example, the play ACH is the extension of the play CH.

In the game G, AH is a strategy for player 1. The restriction of this strategy to the three subgames G_{p_1}, G_{p_2} and G_{p_3} are the strategy H, the empty strategy, and the strategy H, respectively. Player 2 has CE as one of his strategies in the game G; the restriction of this strategy to the three subgames are in turn C, E and the empty strategy.

The game G has a natural solution, which we will now determine. Suppose we for some reason have come to the position p_3. This means that we are in the initial position of the G_{p_3} subgame whose game tree has height 1. For the player in turn to move, i.e. player 1, there is only one rational move, the move G, because this move gives him more payoff than the other move H. In other words, player 1's optimal strategy in the subgame is the G strategy.

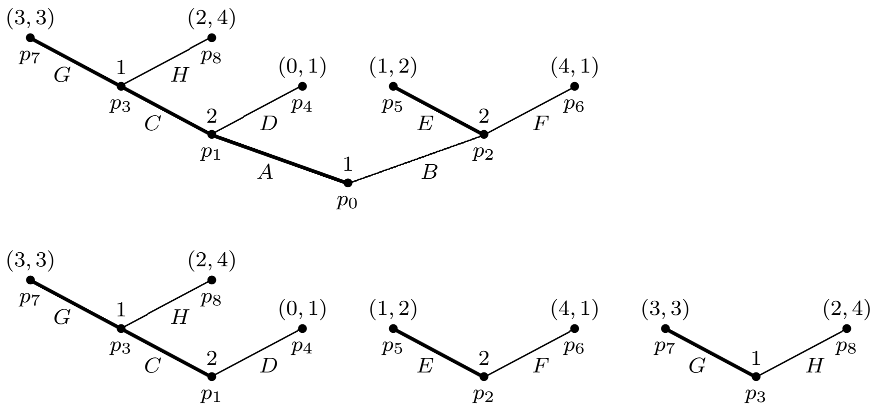

Figure 8.6. At the top a game G (and the corresponding game tree T), at the bottom the three subgames G_{p_1}, G_{p_2} and G_{p_3} (and the corresponding game trees).

In Figure 8.6 we have highlighted this by thickening the edge G. Player 2, of course, can not affect the outcome of the subgame G_{p_3}.

Similarly, player 2's optimal strategy in the subgame G_{p_2} is the strategy E, which benefits him more than the other strategy F.

Suppose now that the game has progressed to the p_1 position, which is the initial position of the subgame G_{p_1}. Player 2 is to move, and he can now figure out that player 1's optimal strategy in position p_3 consists in selecting the G move which gives 3 utility units to player 2, which is more than player 2 receives through the move D. Player 2's optimal strategy in the game G_{p_1} is thus the C strategy, which we have highlighted by a thick edge. Note that the strategy vector (G, C) is a Nash equilibrium in the game G_{p_1}, because player 1 loses by unilaterally switching from G to H, and player 2 loses by unilaterally changing from C to D.

Consider now the initial position p_0 of the original game G. If player 1 selects the A move, the optimal continuation of the two players in the subgame G_{p_1} will result in the play ACG with 3 utility units to player 1. If player 1 instead chooses B, the optimal continuation of the subgame G_{p_2} will result in the play BE with only 1 utility unit to player 1. Player 1's optimal move in the starting position is thus A, and his optimal strategy in the original game is the strategy AG, while player 2's optimal strategy is CE.

The strategy vector (AG, CE) is a Nash equilibrium of the game G with the special feature that the restriction of the vector to each subgame is a Nash equilibrium of the subgame. This property distinguishes it from the two other less credible Nash equilibria that we found when we studied the game G in Example 8.2.5. □

The characteristics of the Nash equilibrium in the previous example are intuitively very appealing and justify the following general definition.

Definition 8.3.3 A strategy vector $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n)$ in an extensive game G is called a *subgame perfect equilibrium* if for each subgame of G the restriction of the strategy vector σ to the subgame is a Nash equilibrium of the subgame.

Example 8.3.2 The strategy vector (AG, CE) is a subgame perfect equilibrium of the game in the previous example. □

The way to argue in Example 8.3.1 can be generalized and leads to what is commonly called *backward induction*.

Proposition 8.3.1 *Let $G = \langle N, P, p_0, s, Pl, (\succeq_i)\rangle$ be an extensive game with initial position p_0, and consider the subgames G_p for each successor p of p_0, i.e. for each $p \in s(p_0)$.*

Assume that each such subgame G_p has a subgame perfect equilibrium

$$\sigma^*_p = (\sigma^*_{p1}, \sigma^*_{p2}, \dots, \sigma^*_{pn}),$$

and let π^*_p denote the corresponding play in G_p. Let Π^* denote the set of all plays that are obtained by extending these plays to plays in G, i.e. all plays *of the type p_0, π^*_p that starts with the move (p_0, p) and then continues with the play π^*_p.*

Let i_0 be the player in turn to make a move at the initial position p_0, and suppose there is a first move $(p_0, \overline{p})$ such that the play $p_0, \pi^*_{\overline{p}}$ is a maximal element in the set Π^* with respect to the player's preference relation $\succeq_{i_0}$. (There is certainly such a maximal element if the set $s(p_0)$ of successors to the initial position is finite.)

Now extend each player j's strategies σ^*_{pj} to a strategy σ^*_j in the game G by defining σ^*_j as the uniquely determined strategy whose restrictions to the subgames G_p are equal to σ^*_{pj}, in the case $j = i_0$ supplemented with the definition $\sigma^*_{i_0}(p_0) = \overline{p}$.

Then

$$\sigma^* = (\sigma^*_1, \sigma^*_2, \dots, \sigma^*_n)$$

is a subgame perfect equilibrium of the game G, and all subgame perfect equilibria of G are formed in this way.

Proof. The restrictions of the strategies σ^*_i to the subgames G_p are by definition subgame perfect equilibrium strategies of these subgames, and they are therefore Nash equilibrium strategies of every proper subgame of G. So it suffices to show that σ^* is a Nash equilibrium of the game G.

First consider an arbitrary strategy τ_{i_0} for player i_0, and let (p_0, p) be the first move of this strategy. The remaining part of the play will then take part in the subgame G_p. Since the restrictions to the subgame G_p of the players' strategies σ^*_j are Nash equilibrium strategies of this game, the strategy vector $(\sigma^*_{-i_0}, \tau_{i_0})$ will result in a play with an outcome that is not preferred by player i_0 when compared with the outcome of the Nash equilibrium σ^*_p of the subgame G_p, and this outcome, in turn, is not better than the outcome of σ^* because of the definition of the first move $(p_0, \overline{p})$. This proves that player i_0 does not profit from unilaterally switching from strategy $\sigma^*_{i_0}$ to any other strategy τ_{i_0}.

Now, let τ_j be an arbitrary strategy for any player j other than player i_0, and suppose that all players except j stick to their strategies in σ^*. In particular, player i_0 does so, which means that the strategy vector (σ^*_{-j}, τ_j) results in a play that after one move takes place in the subgame $G_{\overline{p}}$ and then is identical to the play obtained when all players except player j use their Nash equilibrium strategies $\sigma^*_{\overline{p}}$ and player j uses the restriction of τ_j to the subgame $G_{\overline{p}}$. For player j, this play can not be better than the play given by the Nash equilibrium $\sigma^*_{\overline{p}}$ of $G_{\overline{p}}$, which apart from the first move is the play in G given by the strategy vector σ^*.

This concludes the proof of the proposition. □

The construction in Proposition 8.3.1 gives rise to the following algorithm.

Algorithm for determining subgame perfect equilibria of extensive games G with finite horizon

1. Set $k = 0$. The only subgame perfect strategy of subgames G_p, when p is a terminal position, is the empty strategy.
2. Stop if the height of G's game tree is equal to k. Use otherwise Proposition 8.3.1 to determine all subgame perfect equilibria of all subgames with game trees of height $k + 1$.
3. Set $k := k + 1$ and return to step 2.

For the induction step 2 to work, there must exist optimal moves, which obviously is the case if the game is finite.

As a corollary of Proposition 8.3.1 and the algorithm we get the following result.

Proposition 8.3.2 *Each finite extensive game with perfect information has a subgame perfect equilibrium.*

Corollary 8.3.3 *Each finite extensive game with perfect information has a Nash equilibrium.*

Example 8.3.3 Chess is a finite extensive game and therefore has a Nash equilibrium. The game is furthermore strictly competitive, so White has the same payoff in each Nash equilibrium (if there are several), according to Proposition 2.5.3. If the value of the game is 1, then White has a strategy that guarantees him to win no matter how Black plays, if the value is $\frac{1}{2}$, both players have strategies that guarantee them at least a draw, and finally if the value is 0, then Black has a strategy that guarantees him to win regardless of how White plays. Theoretically, chess is therefore an uninteresting game. But the game's value is not known, and even less are the Nash equilibria, and because of the great complexity of chess, it is unlikely that they ever will be calculated. □

Example 8.3.4 An extensive game can have more than one subgame perfect equilibrium. Consider the game in Figure 8.7. Using backward induction we get two subgame perfect equilibria, namely (B, CE) and (A, DE). The first one results in the play BE with payoff $(3, 4)$. The second equilibrium vector results in the play AD with payoff $(5, 2)$, which is better for player 1 and worse for player 2. In this game there is no way for player 1 to rationalize the choice of strategy B instead of A or vice versa, because in the initial position he can not predict if player 2 will choose C or D if he chooses the move A, □

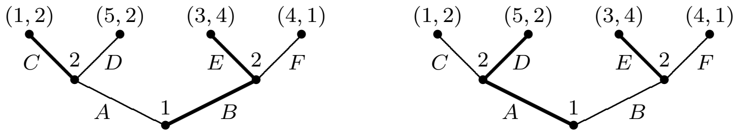

Figure 8.7. A game with two subgame perfect equilibria.

The problem in Example 8.3.4 can not arise in games where the optimal choice of strategy is unique in each subgame.

Exercises

8.7 Determine all subgame perfect equilibria of the game in Exercise 8.3.

8.8 Determine all subgame perfect equilibria of the game in Exercise 8.4.

8.9 Determine all subgame perfect equilibria of the game in Exercise 8.5.

8.10 Determine all subgame perfect equilibria of the game in Exercise 8.6.

8.11 Two people use the following method to divide a cake. Person 1 divides the cake into two pieces, after which person 2 chooses one of the parts and person 1 gets the remaining part. They only care about the size of their own cake piece. Formulate the method as an extensive game. How is the cake divided by a subgame perfect equilibrium?

8.12 Find all subgame perfect equilibria of the game in Figure 8.8.

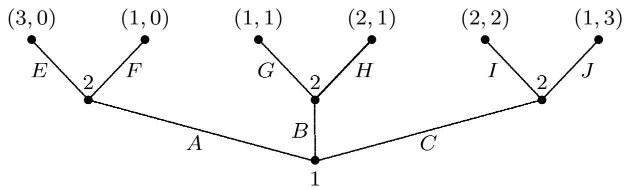

Figure 8.8

8.13 The game in Figure 8.9 is called the *centipede game*. The game has 100 nonterminal positions that are located at the integer points 1, 2, ..., 100 on the line. From each nonterminal position it is possible to move either one step to the right or one step downwards. The game starts in position 1 with a move by player 1, who will then move at all odd positions $2k-1$, while it is player 2's turn to move at the even positions $2k$. The payoff to the two players at a terminal position under an odd number $2k-1$ is given by the pair (k, k), and under an even number $2k$ by the pair $(k-1, k+2)$.

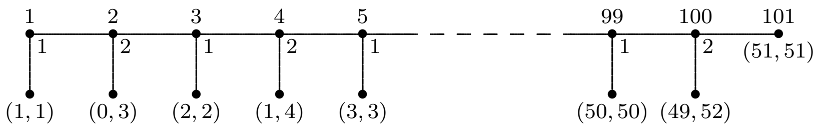

Figure 8.9

a) Determine the subgame perfect equilibria.
b) Is it likely that the players would choose this solution?
c) Are there any other Nash equilibria?
d) Is there any Nash equilibrium with a payoff that differs from the payoff at the subgame perfect equilibrium?

8.4 Stackelberg duopoly

In this section we will study a market model for two firms that produce and sell an identical product. Firm i's cost of producing q_i units of the product is $C_i(q_i)$, and just as in the Cournot model, the price of the product depends on the total output – if this is $Q = q_1 + q_2$, the product is sold at the price $P(Q)$

per unit. It is the size of the output, which is the firm's strategic variable, but contrary to the Cournot model, the firms do not take their strategic decisions independently, but the market leader, firm 1, first determines its output, and the other firm knows this decision when it decides its output. We assume that the two firms can choose any non-negative number as size of their outputs.

The problem of deciding the optimal outputs can be formulated as an extensive two-person game, and the game has been named after economist von Stackelberg who studied a similar asymmetric situation.

The players in the Stackelberg game are of course the two firms. Firm 1 begins by selecting any non-negative number q_1, whereupon firm 2 continues by also selecting an arbitrary non-negative number q_2. Then the game is over. This gives rise to an infinite game tree, as indicated by Figure 8.10.

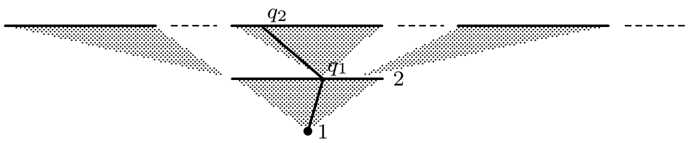

Figure 8.10. Stackelberg's duopoly game. The intervals symbolize $\mathbf{R}_+$.

A play in the game is in other words an ordered pair (q_1, q_2) of nonnegative numbers, and the set of all plays coincides with the Cartesian product $\mathbf{R}_+ \times \mathbf{R}_+$. The two players value their plays using their respective profit functions $V_i(q_1, q_2) = q_i P(q_1 + q_2) - C_i(q_i)$.

Since the game has a finite horizon, we can determine a subgame perfect equilibrium using backward induction:

1. Given that firm 1 has chosen the output q_1, firm 2 should maximize the function $q_2 \mapsto V_2(q_1, q_2)$. Let us assume that the maximum exists and is assumed in a unique point $m(q_1)$.
2. Firm 1's problem now consists in maximizing its profit given that firm 2 acts in accordance with paragraph 1. The output q_1 results in the profit
$$V_1(q_1, m(q_1)) = q_1 P(q_1 + m(q_1)) - C_1(q_1),$$
so the problem consists in maximizing this function. Let us assume that the maximum exists and is assumed in a unique point $\overline{q}_1$.
3. Under these conditions, the game has a unique subgame perfect equilibrium, namely $(\overline{q}_1, \overline{q}_2)$, where $\overline{q}_2 = m(\overline{q}_1)$. This equilibrium is called the *Stackelberg equilibrium solution.*

Remark. If there are several maximizing points in step 1, then it is of course necessary to carry out the investigation in step 2 for each maximum point in order to obtain all equilibrium solutions. The same applies if the maximum point in step 2 is not unique.

EXAMPLE 8.4.1 Let us determine the Stackelberg equilibrium solution in the case of a linear inverse demand function

$$P(Q) = \begin{cases} a - Q & \text{if } 0 \le Q \le a, \\ 0 & \text{if } Q > a \end{cases}$$

and the same constant piece cost c for both companies, where $0 < c < a$.

Firm 2's profit $V_2(q_1, q_2)$ is for $0 \le q_1 \le a - c$ given by the second degree polynomial $q_2(a - c - q_1 - q_2)$ in the interval $0 \le q_2 \le a - c - q_1$, while the profit is negativ outside that interval. The maximum is attained for $q_2 = \frac{1}{2}(a - c - q_1)$. Firm 2's profit is equal to $-cq_2$ if $q_1 > a - c$, so the maximum is in this case assumed for $q_2 = 0$. This means that

$$m(q_1) = \begin{cases} \frac{1}{2}(a - c - q_1) & \text{if } 0 \le q_1 \le a - c, \\ 0 & \text{if } q_1 > a - c. \end{cases}$$

Hence,

$$q_1 + m(q_1) = \begin{cases} \frac{1}{2}(a - c + q_1) & \text{if } 0 \le q_1 \le a - c, \\ q_1 & \text{if } q_1 > a - c, \end{cases}$$

and the profit of firm 1 is therefore

$$V(q_1, m(q_1)) = \begin{cases} \frac{1}{2}q_1(a-c-q_1) & \text{if } 0 \le q_1 \le a-c, \\ q_1(a-c-q_1) & \text{if } a-c < q_1 \le a, \\ -cq_1 & \text{if } q_1 > a. \end{cases}$$

The maximum is attained at the point $\frac{1}{2}(a-c)$. Stackelberg's equilibrium solution $(\overline{q}_1, \overline{q}_2)$ is therefore given by

$$\overline{q}_1 = \frac{1}{2}(a-c) \quad \text{and} \quad \overline{q}_2 = m(\overline{q}_1) = \frac{1}{4}(a-c).$$

The total supply at the equilibrium point is equal to $\frac{3}{4}(a-c)$, the price is $\frac{1}{4}(a+3c)$, and firm 1's profit is $\frac{1}{8}(a-c)^2$, while firm 2's profit is half that big, that is $\frac{1}{16}(a-c)^2$.

Compare with the Cournot model (see Example 3.1.1), where the total equilibrium supply is $\frac{2}{3}(a-c)$ to the price $\frac{1}{3}(a+2c)$ with a profit for each company equal to $\frac{1}{9}(a-c)^2$. Stackelberg's model gives a greater supply, a lower price and a lower overall profit than the Cournot model. □

Exercises

8.14 Determine the Stackelberg equilibrium solution if $P(Q) = 64(Q+2)^{-1}$ and both companies have the same linear cost function $C(q) = 9q$.

8.15 Consider a duopoly with the sole assumption regarding the two firms' profit functions $V_1(q_1, q_2)$ and $V_2(q_1, q_2)$ that there is a Cournot equilibrium and a Stackelberg equilibrium. Prove that the market leading firm, firm 1, receives a higher profit from the Stackelberg solution than from the Cournot solution.

8.5 Chance moves

The extensive games we have studied so far have been deterministic, i.e. when the players have chosen their strategies and play according to them, the outcome is entirely given. But there are also many sequential games in which there are situations where the continuation is due to chance. This is true not least in many popular parlor games, where you are forced to draw a playing card or perform a roll of the dice, and where the continuation of the game then depends on the color or the value of the drawn card or the number of eyes on the dice. In many more serious games, players, when deciding on their action plans, must also take into account future situations that they themselves can not control, situations that depend on a whimsical nature or chance, and where at best it is possible to identify a number of possible alternatives that can be assigned probabilities.

EXAMPLE 8.5.1 Instead of starting with a general definition, we begin with an example. Figure 8.11 illustrates an extensive two-person game with two nonterminal positions marked with a c. The letter c stands for "chance", and none of the players should make a choice there but the continuation is due to chance. The probabilities of the different random or chance moves that lead from such a position have been marked along the respective moves. The two chance moves x and y are equally probable because the probability of each of them is $\frac{1}{2}$. In the second chance position, the probability of the move z is equal to $\frac{1}{4}$ and the probability of the move w equals $\frac{3}{4}$.

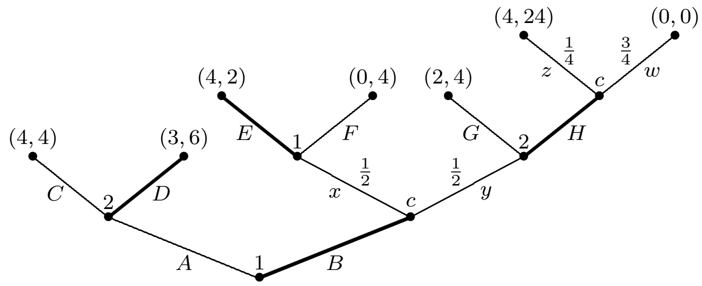

Figure 8.11

Assume that player 1 has decided to make the moves B and E, and that player 2 has decided for the moves D and H − the thick marked edges in the figure. The outcome of the game is then not given but depends on which random moves will actually be made in the positions where this is applicable. If chance results in the moves x and z, which occurs with probability $\frac{1}{8}$, the players' strategies lead to the play BxE with payoff $(4, 2)$ to the two players. If instead the chance moves are y and w, which occurs with probability $\frac{3}{8}$, then we get the play $ByHw$ with payoff $(0, 0)$.

The following table summarizes all possibilities:

Random moves	Probability	Play	Payoffs
x, z	$\frac{1}{8}$	BxE	$(4, 2)$
y, z	$\frac{1}{8}$	$ByHz$	$(4, 24)$
x, w	$\frac{3}{8}$	BxE	$(4, 2)$
y, w	$\frac{3}{8}$	$ByHw$	$(0, 0)$

Using the above table we can now compute the two players' *expected* *utilities* of the strategy combinations BE and DH. They are

$$U_1(BE,DH) = \frac{1}{8}\cdot 4 + \frac{1}{8}\cdot 4 + \frac{3}{8}\cdot 4 + \frac{3}{8}\cdot 0 = \frac{5}{2},$$
$$U_2(BE,DH) = \frac{1}{8}\cdot 2 + \frac{1}{8}\cdot 24 + \frac{3}{8}\cdot 2 + \frac{3}{8}\cdot 0 = 4.$$

Similarly, we could calculate the players' expected utilities for all strategy vectors, and it should now be obvious what should be meant by a Nash equilibrium – it is a strategy vector with the property that no player can increase his expected utility by unilaterally changing strategy.

In this example, it is easy to see that (BE, DH) is not a Nash equilibrium. Player 1 profits from unilaterally switching to the AE strategy, which results in the play AD with the expected payoff (in this case, sure payoff) of 3 units to player 1. □

The concept of subgame perfect equilibrium can be naturally generalized to games with random moves, and the method of backward induction works for games with finite horizon as in the game above. For each position, we can "backwards" gradually calculate the players' expected payoffs for the subgame starting from the position in question, given that the players choose their best possible strategies in the subgame.

Example 8.5.2 Figure 8.12 shows the result of the successive computations needed to compute the subgame perfect equilibrium of the game in Example 8.5.1. In turn, we consider the subgames that start in the positions marked p_5, p_4, p_3, p_2, p_1 and p_0.

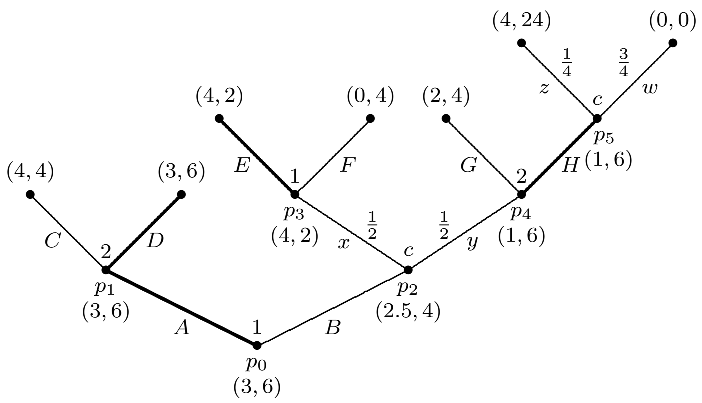

Figure 8.12. Computation of the subgame perfect equilibrium of the game in Example 8.5.1 and Figure 8.11 using backward induction.

The move at position p_5 is a random move, and the two players' expected utilities after this move are given by the pair $(1, 6)$ $(= \frac{1}{4}(4, 24) + \frac{3}{4}(0, 0))$. We note this in the game tree by writing the number pair $(1, 6)$ below the p_5 node.

Consider now the subgame that starts in p_4. The move that gives player 2 the highest expected payoff is the move H, giving him the expected utility 6 (and player 1 expected utility 1). We highlight this in the figure with a thick edge H and by writing the number pair $(1, 6)$ below the node p_4.

In the subgame starting at p_3, E is the best move for player 1; this move results in a play with the (expected) payoffs $(4, 2)$. We thicken the corresponding edge and write the expected payoffs below the p_3 node.

The expected payoffs for the subgame starting at position p_2, where a random move is performed, are $(2.5, 4)$ $(= \frac{1}{2}(4, 2) + \frac{1}{2}(1, 6))$, provided that the players follow their subgame strategies calculated so far.

The D strategy is the best strategy for player 2 in the subgame starting at p_1, and it gives the players the payoffs $(3, 6)$.

Therefore, at the initial position p_0, it is best for player 1 to select the A move, as it gives the player the expected utility 3. We have thus found that the game has a unique subgame perfect equilibrium, namely the strategy pair (AE, DH), which of course is also a Nash equilibrium. □

With the example above in mind, it is now easy to provide formal definitions of the concept of extensive game with chance moves and the concepts of strategy, Nash equilibrium and subgame perfect equilibrium for such games.

Definition 8.5.1 An n-person game in *extensive form with perfect information and chance moves* consists of

- a set $N = \{1, 2, \dots, n\}$ of players;
- a game tree $T = \langle P, p_0, s\rangle$;
- a function $Pl\colon P^\circ \to N \cup \{c\}$, *the player function*, defined on the set P° of all nonterminal positions in the game;
- for each position p such that $Pl(p) = c$ a probability measure μ_p on the set of $s(p)$ of all successors of p (or, equivalently, on the set of all moves at p);
- for each player $i \in N$ a cardinal utility function u_i on the set Π of all plays.

We define

$$P_i = \{p \in P^\circ \mid Pl(p) = i\}$$

for $i \in N \cup \{c\}$. The set P_c consists of those positions where the moves are determined by chance. The probability of a certain move at a chance position p is given by the probability measure μ_p. We assume that the probability distributions at different chance positions are *independent* of each other.

The concept of strategy can now be defined for games with chance moves in exactly the same way as for games without any chance moves.

Definition 8.5.2 A *strategy* σ for player i in a game in extensive form with perfect information and chance moves is a function $\sigma\colon P_i \to P$ such that $\sigma(p)$ is a successor of p for all positions $p \in P_i$.

The set of all strategies of player i will be denoted by Σ_i as before, and we set $\Sigma = \Sigma_1 \times \Sigma_2 \times \cdots \times \Sigma_n$.

Unlike earlier, the outcome of the game now no longer depends on the players' choice of strategy vector $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n) \in \Sigma$ only, but also on the random moves performed at the chance positions in P_c. Therefore, let Σ_c denote the set of all functions $\sigma_c\colon P_c \to P$ such that $\sigma_c(p)$ is a successor of p for all $p \in P_c$. As before, we will not make any distinction between succeeding positions and moves when describing the functions σ_c. To describe a specific function σ_c therefore becomes equivalent to describing a move at each chance position.

The independent probability measures μ_p induce a probability measure μ on the set Σ_c, the product measure of the μ_p:s.

Example 8.5.3 In example 8.5.1, the set Σ_c consists of four functions σ_c^i,

$i = 1, 2, 3, 4$, which (in terms of moves) are defined by

$$\sigma_c^1(p_2) = x,\ \ \sigma_c^1(p_5) = z; \qquad \sigma_c^2(p_2) = x,\ \ \sigma_c^2(p_5) = w$$
$$\sigma_c^3(p_2) = y,\ \ \sigma_c^3(p_5) = z; \qquad \sigma_c^4(p_2) = y,\ \ \sigma_c^4(p_5) = w.$$

The induced probability distribution μ on Σ_c is defined by

$$\begin{aligned} \mu(\sigma_c^1) &= \tfrac{1}{2} \cdot \tfrac{1}{4} = \tfrac{1}{8} \\ \mu(\sigma_c^2) &= \tfrac{1}{2} \cdot \tfrac{3}{4} = \tfrac{3}{8} \\ \mu(\sigma_c^3) &= \tfrac{1}{2} \cdot \tfrac{1}{4} = \tfrac{1}{8} \\ \mu(\sigma_c^4) &= \tfrac{1}{2} \cdot \tfrac{3}{4} = \tfrac{3}{8}. \end{aligned}$$

□

Each combination $(\sigma, \sigma_c) \in \Sigma \times \Sigma_c$ gives rise to a unique play $\pi(\sigma, \sigma_c)$ in the game and thereby also to a uniquely determined payoff $u_i(\pi(\sigma, \sigma_c))$ to player i. The outcome is random, insofar as it is due to the random moves σ_c that have been executed. More precisely, the outcome is a stochastic variable on the probability space Σ_c with μ as probability measure. We denote by $U_i(\sigma)$ players i's expected utility from the strategy vector $\sigma \in \Sigma$. The expected utility is (if Σ_c is a finite probability space) given by the formula

$$U_i(\sigma) = \sum_{\sigma_c \in \Sigma_c} u_i(\pi(\sigma, \sigma_c))\mu(\sigma_c).$$

We assume that the players use these expected utilities when they evaluate the value of a given strategic vector in Σ. It is now easy to generalize the concepts of Nash equilibrium and subgame perfect equilibrium to games with chance moves.

Definition 8.5.3 A strategy vector $\sigma^* = (\sigma_1^*, \sigma_2^*, \dots, \sigma_n^*)$ in an extensive game with chance moves is called

- a *Nash equilibrium* if for each player i and each strategy $\tau_i \in \Sigma_i$

$$U_i(\sigma_{-i}^*, \sigma_i^*) \geq U_i(\sigma_{-i}^*, \tau_i);$$

- a *subgame perfect equilibrium* if the restriction of the strategy vector to each subgame is a Nash equilibrium of the subgame.

Backward induction works equally well in extensive games with chance moves as in games without such moves, so the proof of the following proposition requires only minor modifications of the proof of Proposition 8.3.1 and is therefore left to the reader.

Proposition 8.5.1 *Each finite game in extensive form with perfect information and chance moves has a subgame perfect equilibrium and consequently a Nash equilibrium.*

Exercise

8.16 Determine the subgame perfect equilibrium and the players expected utility in the equilibrium for the game in Figure 8.13.

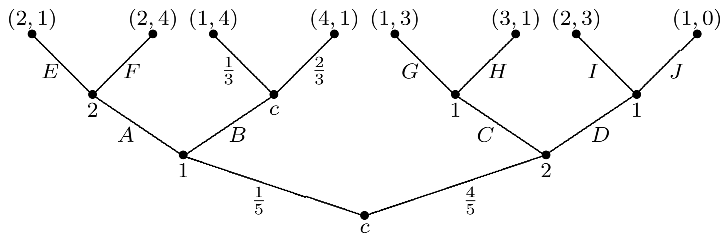

Figure 8.13

Chapter 9

Extensive Games with Imperfect Information

In the extensive games we have studied so far, each player, when performing a move, has complete information about all players' previous moves including the outcome of any chance moves, and the player knows exactly his actual position in the game tree. But in many games, the players do not have that information. For example, in many card games, a player usually does not know which cards the opponents have.

In this chapter we will study extensive games where the players' information about the situation is limited − when a player is to perform a move, he may be in one of several possible positions. We model this by dividing the player's set of positions into pairwise disjoint subsets that we call the player's information sets. During the play, the player always knows in what information set the play is, but he does not know the exact position in the information set. However, we still assume that all players know the game tree, what moves are possible in the information sets, and last but not least, each other's preferences.

9.1 Basic Endgame

Poker is a classic example of a game where the players do not have full information, and where winning game modes contain moments of bluffing. The game has therefore interested many prominent game theorists who have analyzed simplified poker models. We will study a model introduced by W.H. Cutler and commonly called *Basic Endgame*. A detailed analysis of the model and different extensions of this has been done by Tom and Chris Ferguson, father and son, the first professor of mathematics, the latter professional poker player.

Basic Endgame adequately describes the considerations a player faces in the final phase of poker, when only two players remain. The two players start by placing \$1 each in the pot. Player 1 then receives a card from a deck of cards. The card is a winning card with a probability of 1/3 and a losing card with a probability of 2/3. The player sees the card but keeps it hidden for the opponent who does not get a card.

Player 1 can now choose between passing or betting. If he passes with a winning card he gets the pot and then wins \$1. If he passes with a losing card, he loses the pot and in that way he has lost \$1. If player 1 bets he must add another \$2 into the pot. Player 2, who does not know which card player 1 has, can now choose between calling or folding. If he folds, player 1 gets the whole pot regardless of which card he has and thus has won \$1. If player 2 calls, he adds \$2 to the pot, after which player 1 shows his card. If player 1 has a winning card, he will win the whole pot and has won \$3. If not, player 2 gets the whole pot and player 1 has lost \$3.

The game tree of Basic Endgame is shown in Figure 9.1. The game starts with a chance move. Player 1 then knows his position in the game tree and has four options, namely

1. to pass with both winning and losing cards: pass_{w}-pass_{l},
2. to pass with a winning card and bet with a losing card: pass_{w}-bet_{l},
3. to bet with a winning card and pass with a losing card: bet_{w}-pass_{l},
4. to bet with both winning and losing cards: bet_{w}-bet_{l}.

The second option may seem bizarre but is still an option. Option 3 is the player's honest strategy, and option 4 is his bluffing strategy.

The problem for player 2 is that he does not know if player 1 has a winning or a losing card, if player 1 has chosen to bet. In other words, player 2 does not know if he is in the left or the right part of the game tree. Player 2 therefore has only two options, namely to call or to fold regardless of player 1's card. We indicate player 2's lack of information in the drawing of the game tree by connecting his two positions with a dotted line, and

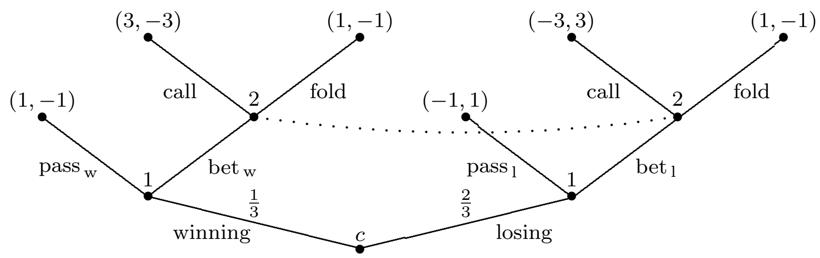

Figure 9.1. The game Basic Endgame

by marking a move from each position with ”call” and a move from each position with ”fold”. The move that is then actually performed in the game if player 2 chooses ”call” depends on player 1’s previous move, and similarly for ”fold”.

The two positions that have been connected by a dotted line constitute an *information set* for player 2, who knows that he is supposed to perform a move from a position in the information set but does not know from which position.

We can analyze the game by reducing it to strategic form. The expected payoff matrix with player 1 as the row player is calculated as follows:

If player 1 selects the option pass_{w}-pass_{l}, his expected payoff will be $\frac{1}{3} \cdot 1 + \frac{2}{3} \cdot (-1) = -\frac{1}{3}$, regardless of whether player 2 had chosen to call or to fold.

If player 1 chooses the option pass_{w}-bet_{l}, his expected payoff will be $\frac{1}{3} \cdot 1 + \frac{2}{3} \cdot (-3) = -\frac{5}{3}$ if player 2 calls, and $\frac{1}{3} \cdot 1 + \frac{2}{3} \cdot 1 = 1$ if player 2 folds, etc.

The full expected payoff matrix looks like this:

	call	fold
$\text{pass}_{\,\text{w}}$-$\text{pass}_{\,\text{l}}$	$-\frac{1}{3}$	$-\frac{1}{3}$
$\text{pass}_{\,\text{w}}$-$\text{bet}_{\,\text{l}}$	$-\frac{5}{3}$	1
$\text{bet}_{\,\text{w}}$-$\text{pass}_{\,\text{l}}$	$\frac{1}{3}$	$-\frac{1}{3}$
$\text{bet}_{\,\text{w}}$-$\text{bet}_{\,\text{l}}$	-1	1

Row 1 is weakly dominated by row 3, and row 2 is weakly dominated by row 4. By eliminating the two top rows we get the matrix

	call	fold
$\text{bet}_{\,\text{w}}$-$\text{pass}_{\,\text{l}}$	$\frac{1}{3}$	$-\frac{1}{3}$
$\text{bet}_{\,\text{w}}$-$\text{bet}_{\,\text{l}}$	-1	1

Since the matrix has no saddle point, there is no pure Nash equilibrium in the game, but there is of course a mixed Nash equilibrium. The row player's optimal mixed strategy consists in choosing $\text{bet}_{\,\text{w}}$-$\text{pass}_{\,\text{l}}$ with probability $\frac{3}{4}$ and $\text{bet}_{\,\text{w}}$-$\text{bet}_{\,\text{l}}$ with probability $\frac{1}{4}$, while player 2's optimal mixed strategy consists in calling and folding with equal probability $\frac{1}{2}$. The game is fair because its value is equal to 0.

In order to play optimally in Basic Endgame, player 1 should always bet when he has a winning card and bluff by betting on average once by four with a losing card.

Exercise

9.1 Solve the game Basic Endgame when the probability that player 1 will get a winning card is

a) $\frac{1}{4}$ b) p, where $0 < p < 1$.

9.2 Extensive games with incomplete information

We will now give a general definition of extensive games in which players, like in Basic Endgame, have limited information about their positions in the game tree. We do this by using the concepts of *information set* and *label*.

The set P_i of positions that belong to player i is partitioned into pairwise disjoint subsets P_{ij}, $j = 1, 2, \ldots, n_i$, and each such subset is called an information set. A player in turn to perform a move always knows what information set the play has come to, but he does not know the exact position in the information set unless the information set happens to consist of just one position. In order for the player to lack information about the exact position, it is necessary that all positions in an information set have the same number of succeeding positions, i.e. that there are as many possible moves to perform from each position in an information set.

The player must also somehow be able to indicate how he will continue the game from an information set P_{ij}. Since he does not know his position $p \in P_{ij}$, he can not do this by specifying a move, i.e. a succeeding position $q \in e(p)$.

We solve this problem by demanding from the rules of the game that the set of moves from all positions in a given information set P_{ij} be grouped into pairwise disjoint subsets with the property that each such subset E contains exactly one move from each position in the given information set. This means that all positions in the information set P_{ij} must have the same number of successors and that the number of subsets is equal to this common number. To keep track of the moves in such a subset E, we then put a unique common name, a *label*, on all the moves in the subset E.

In Basic Endgame (see Figure 9.1) player 2 has one information set consisting of two positions which in the figure are connected by a dotted line. There are two moves from each position, and two subsets are formed by the four moves − one consists of the two moves labeled "call", the other of the two moves labeled "'fold". Both subsets consist of exactly one move from each position in the information set.

Player 1 has two information sets in Basic Endgame, but both of them are singleton sets, i.e. consist of only one position, and in such cases, we can use the names of the actual moves as labels.

We are now ready for the general definition of games in extensive form that includes the possibility of chance moves and of incomplete information.

Definition 9.2.1 An n-person game in *extensive form* consists of

- a set $N = \{1, 2, \dots, n\}$ of players;
- a game tree $T = \langle P, p_0, s\rangle$;
- a function $Pl\colon P^\circ \to N \cup \{c\}$, *the player function*, defined on the set P° of all nonterminal positions in the game;
- for each position p with $Pl(p) = c$ a probability measure μ_p on the set $s(p)$ of all successors of p (or, equivalently, on the set of all moves at p);

- for each player i a partition of the player's nonterminal positions

$$P_i = \{p \in P^\circ \mid Pl(p) = i\}$$

  into subsets P_{ij}, the player's *information sets*, with the following property: All positions in a given information set P_{ij} have the same number of successors;
- for each information set P_{ij} a partition of the set of moves from the positions belonging to P_{ij} into pairwise disjoint subsets called *labels* with the following property: Each label contains exactly one move from each position in the information set;
- for each player $i \in N$ a cardinal utility function u_i defined on the set Π of all plays.

In our drawings of games we connect positions belonging to the same information set by a dotted line and write the name of the player at the line. Labels are indicated by giving all the moves that are contained in a label a common name.

A concrete instance of an extensive game progresses as follows. When the play reaches a position p, the player who owns the information set to which p belongs selects a label from the information set. The selected label contains a move that leads from p to a succeeding position p'. The player to whom position p' belongs, then continues. If the play comes to a chance position p, then the following position is selected according to the probability distribution that belongs to p.

Example 9.2.1 Figure 9.2 shows an extensive two-person game with chance moves. We have omitted the players' utility functions.

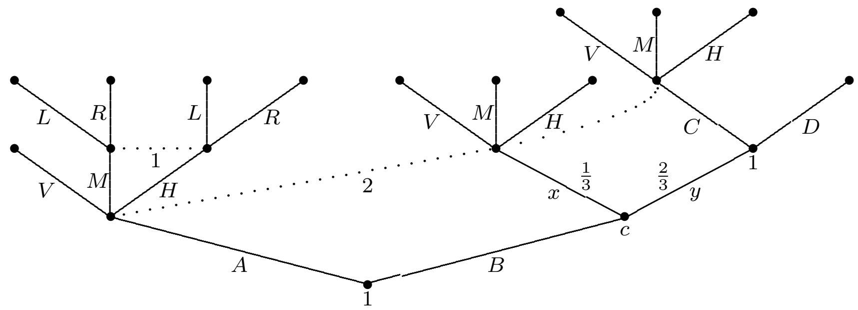

Figure 9.2. An extensive game with chance moves and imperfect information.

Player 1 has three information sets, but two of them consist of just one position. The third one consists of two positions and there are two moves at each position, labeled L and R.

Player 2 has only one information set that consists of three positions. In each position, the player has three moves to choose from and therefore three labels are required: V, M and H.

If player 1 decides to use the moves labeled A, C and R, and player 2 decides to use the M label, then we get the play AMR.

If player 1 instead uses B, C and R, and player 2 still has chosen M, the outcome will depend on chance; we get the play BxM with probability 1/3 and the play $ByCM$ with probability 2/3. □

The general definition of extensive games is very flexible. Games with perfect information are of course special cases; in them, all information sets are singleton sets and each label corresponds to a unique move.

Games where the players make their moves simultaneously and independently can be described as extensive games with imperfect information. Suppose, for example, that we have two players 1 and 2, that player 1 selects one of the options $a_1, a_2, \ldots, a_k$, that player 2 selects one of the options $b_1, b_2, \ldots, b_m$, and that the pair (a_i, b_j) results in the payoff vector (u_{ij}, v_{ij}). In Chapter 2, we expressed this situation as a strategic two-person game, but we can also formulate it as an extensive game with

- an initial position p_0 that belongs to player 1;

- k positions $p_1, p_2, \dots, p_k$ that belong to player 2 and form his only information set;
- km terminal positions $p_{11}, \dots, p_{1m}, p_{21}, \dots, p_{2m}, \dots,\ p_{k1}, \dots, p_{km}$;
- for each i a move from p_0 to p_i with the label a_i, and for each pair i, j a move from p_i to p_{ij} with the label b_j;
- utility functions u and v that are defined at the terminal positions by $u(p_{ij}) = u_{ij}$ and $v(p_{ij}) = v_{ij}$.

Figure 9.3 illustrates the game when $k = 3$ and $m = 2$. Each n-person game in strategic form can of course similarly be written as an n-person game in extensive form with imperfect information.

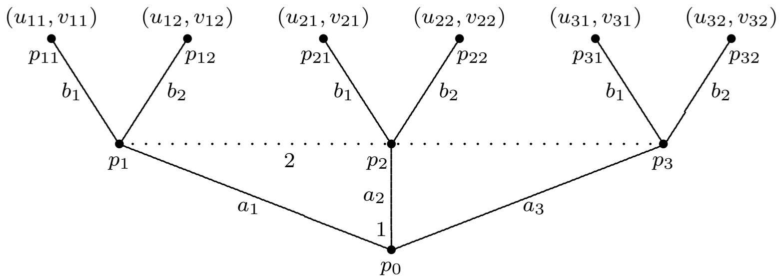

Figure 9.3. Extensive form of a strategic two-person game.

Pure strategies

The concept of strategy was defined for extensive games with perfect information in Sections 8.2 and 8.5. It is now easy to generalize this concept to games with imperfect information.

Definition 9.2.2 A *(pure) strategy* for a player in a general extensive n-person game is a set of labels that consists of one label from each information set that belongs to the player.

We denote the set of all pure strategies belonging to player i by Σ_i and put $\Sigma = \Sigma_1 \times \Sigma_2 \times \cdots \times \Sigma_n$.

A strategy is thus a set $\{A_1, A_2, \dots, A_k\}$ of labels, but in order to simplify the notation we often specify strategies by listing their labels in an arbitrary order, for example as $A_1 A_2 \dots A_k$.

If player i has ℓ information sets $P_{i1}, P_{i2}, \dots, P_{i\ell}$, and m_j is the number of labels in the information set P_{ij}, then the player's number of pure strategies is obviously equal to $m_1 m_2 \cdots m_\ell$.

EXAMPLE 9.2.2 In the game in Example 9.2.1 (Figure 9.2), ACR is a pure strategy for player 1 and M is a pure strategy for player 2. Player 1 has 8 $(= 2 \cdot 2 \cdot 2)$ pure strategies, and player 2 has 3 pure strategies. □

The outcome of an extensive game with no chance moves is completely determined when each player i has choosen a strategy $\sigma_i \in \Sigma_i$. Let $\pi(\sigma)$ denote the play obtained from the strategy vector $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n)$. This gives us a function $\pi\colon \Sigma \to \Pi$, where as previously Π denotes the set of all plays.

In the presence of chance moves, the play is no longer unambiguously determined by the players' strategy choices but random. Given the strategy choice $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n)$, the play π occurs with a certain probability that we denote $\lambda_\sigma(\pi)$. This probability depends, of course, on the probabilities of the chance moves that occur in the game. In the general case, λ_σ is a *probability measure on the set* Π *of all plays.*

EXAMPLE 9.2.3 Consider the game Basic Endgame from previous section (see Figure 9.1), and let σ_1 denote player 1's strategy "bet_{w}-pass_{l}" and σ_2 denote player 2's strategy "fold". The strategy pair $\sigma = (\sigma_1, \sigma_2)$ results in the following probability distribution λ_σ for the plays of the game:

Play	Probability
winning–pass$_{\text{w}}$	0
winning–bet$_{\text{w}}$–call	0
winning–bet$_{\text{w}}$–fold	1/3
losing–pass$_{\text{l}}$	2/3
losing–bet$_{\text{l}}$–call	0
losing–bet$_{\text{l}}$–fold	0

□

Reduction to strategic form

Each strategic game can, as we have seen, be perceived as an extensive game with imperfect information. The converse is also true – each extensive n-person game can be transformed into a strategic game as follows.

Let in the extensive game

- Σ_i denote player i's set of pure strategies;
- $u_i\colon \Pi \to \mathbf{R}$ denote player i's utility function, where Π is the set of all plays;
- λ_σ be the probability measure on Π that is induced by the strategy choice $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n) \in \Sigma = \Sigma_1 \times \Sigma_2 \times \cdots \times \Sigma_n$;
- $U_i(\sigma)$ denote the expected value of the utility function u_i with respect to this probability measure λ_σ, which means that

$$U_i(\sigma) = \sum_{\pi \in \Pi} u_i(\pi)\lambda_\sigma(\pi),$$

  if the game is finite.

We now have all the ingredients of a strategic game $\langle N, (\Sigma_i), (U_i)\rangle$, and this is the *reduced strategic form* of the given extensive game.

An action vector $\sigma^* \in \Sigma$ is by definition a Nash equilibrium in the strategic game if the inequality $U_i(\sigma^*_{-i}, \sigma^*_i) \geq U_i(\sigma^*_{-i}, \tau_i)$ holds for all players i and all actions $\tau_i \in \Sigma_i$. We can of course formulate this condition directly for the extensive game, and we will then get the following definition.

Definition 9.2.3 A vector σ^* of pure strategies in a finite extensive game is a *(pure) Nash equilibrium* if

$$\sum_{\pi \in \Pi} u_i(\pi)\lambda_{(\sigma^*_{-i},\sigma^*_i)}(\pi) \geq \sum_{\pi \in \Pi} u_i(\pi)\lambda_{(\sigma^*_{-i},\tau_i)}(\pi)$$

for all players i and all pure strategies $\tau_i \in \Sigma_i$.

In Section 9.1, we reduced the game Basic Endgame to strategic form and found that it has no pure Nash equilibrium. So extensive games with imperfect information can lack Nash equilibria.

Games with perfect recall

We can use the extensive form with information sets to describe games in which players forget moves they have previously made. For example, such games as bridge can be modeled in this way. Bridge is played by four players who form two teams with two players in each team, but since it is the teams that play against each other, bridge should be perceived as a two-person game. During the bidding phase, the players in a team bid alternately on their hands, and each player only sees his own hand. Each team can therefore be perceived as **one** player who alternately remembers and forgets a part of what he previously knew. The corresponding applies when the cards are played and the players in the non-playing team only know their own cards, the cards on the table and the cards that have already been played.

Example 9.2.4 Figure 9.4 shows a game where player 1 after two moves in the game has forgotten his opening move. If he had remembered that he started with L (or R), he would know the game's position after two moves, and his information set after two moves would then consist of one instead of two positions. □

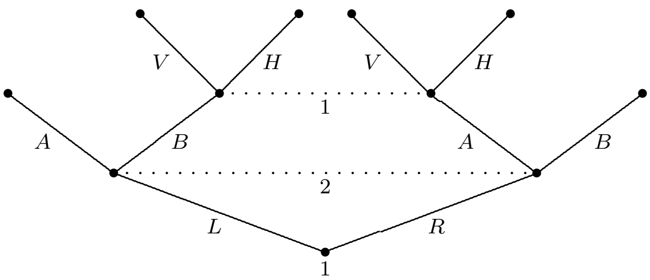

Figure 9.4. A game where player 1 has forgotten a move he has made previously.

Example 9.2.5 In the game to the left in Figure 9.5, player 1 has forgotten his initial move when the game comes to position p_2.

Note that the definition of a pure strategy requires the player to choose **one** label from each of his information sets. This means that a player must choose the same label no matter how many times the game comes to a certain information set. If player 1's strategy consists of choosing the V label in the

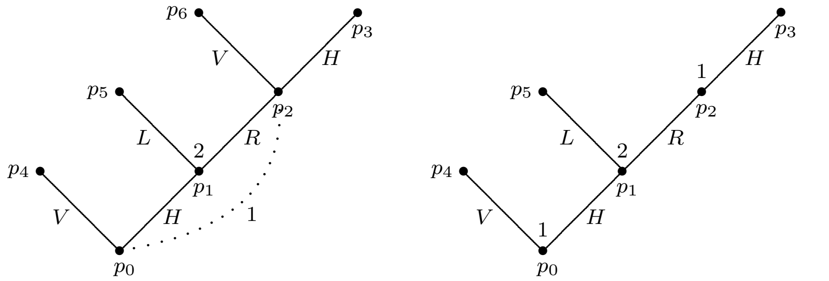

Figure 9.5. In the game to the left, player 1 forgets his initial move. The game is strategically equivalent to the right game.

information set $\{p_0, p_2\}$, then the game ends at p_4, and if the strategy instead consists of choosing H, then the game will end in p_5 or p_3. So the terminal position p_6 is unreachable and may as well be deleted from the game, resulting in the game to the right in the figure. □

Similar situations arise as soon as there is an information set containing two positions p and q that are connected by some move sequence. (In the example above, there is a move sequence from p_0 to p_2.) Some authors include in the definition of extensive games that there shall be no such positions in any information set, but we have not done so.

Extensive games with imperfect information are, as we have seen, able to model situations where a player forgets moves that he has previously made. Games in which all players remember all their previous moves are called games with *perfect recall*, and we will now provide a precise definition of this term.

Consider an arbitrary extensive game, and let p be a position in the game where player i is expected to make a move. In turn, let $\mathcal{I}_1$, $\mathcal{I}_2,\dots$, $\mathcal{I}_k$ be the information sets which belong to player i and have been passed by during the move sequence that leads from the intial position p_0 up to the p position, and denote by L_j the label chosen by player i when the game was in the information set $\mathcal{I}_j$. The sequence $\mathcal{I}_1$, L_1, $\mathcal{I}_2$, $L_2,\dots$, $\mathcal{I}_k$, L_k is called the player's *memory list* in position p, and it is denoted $\mathcal{M}_i(p)$ for the rest of this section. The memory list $\mathcal{M}_i(p)$ is of course empty if player i has not made any moves before position p.

Example 9.2.6 To exemplify the concept of memory list, we consider the game in Figure 9.6. (We have refrained from naming the final positions and from entering the payoffs because they do not matter.)

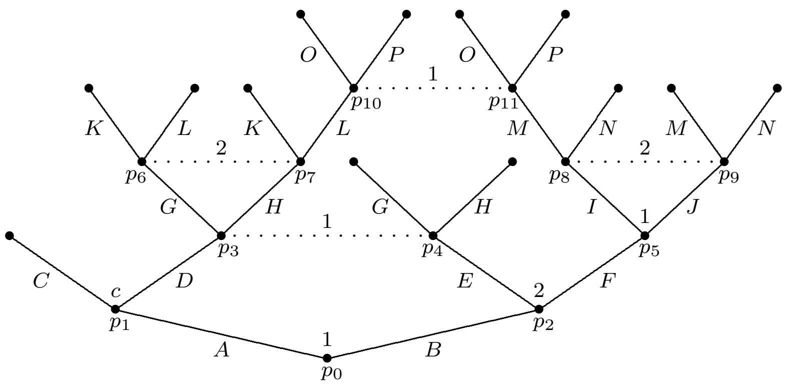

Figure 9.6. The game in Example 9.2.6.

The two players memory lists at some different positions look like this:

$\mathcal{M}_1(p_{10})$: $\{p_0\}, A, \{p_3, p_4\}, H$ $\mathcal{M}_1(p_{11})$: $\{p_0\}, B, \{p_5\}, I$

$\mathcal{M}_2(p_6)$: $\emptyset$ $\mathcal{M}_2(p_8)$: $\{p_2\}, F$ $\mathcal{M}_2(p_9)$: $\{p_2\}, F$ □

A player i with perfect recall should of course remember his memory lists, so if two memory lists $\mathcal{M}_i(p)$ and $\mathcal{M}_i(p')$ differ, then the player has different information about the two positions p and p', which therefore can not belong to the same information set. This observation justifies the following definition.

Definition 9.2.4 A player of an extensive game has *perfect recall* if the player's memory lists at positions in the same information set are identical for all information sets of the player.

The extensive game has *perfect recall* if all players in the game have perfect recall.

EXAMPLE 9.2.7 Player 2 has perfect recall in the game in Figure 9.6. However, player 1 does not have perfect recall, because the player's memory lists in the positions p_{10} and p_{11} are different. So the game does not have perfect recall.

The game Basic Endgame and the games in Figures 9.2 and 9.3 have perfect recall. □

A player with perfect recall can not have any information set $\mathcal{I}$ that contains two positions p and p' which are connected by a move sequence from p to p', because the memory list at p' must then be equal to the memory list at p extended with at least the information set $\mathcal{I}$ and a label from $\mathcal{I}$. Player 1 therefore does not have perfect recall in the left game in Figure 9.5.

Exercises

9.2 Figure 9.7 shows a two-person zero-sum game in extensive form. The numbers at the terminal positions are the payoffs to player 1.

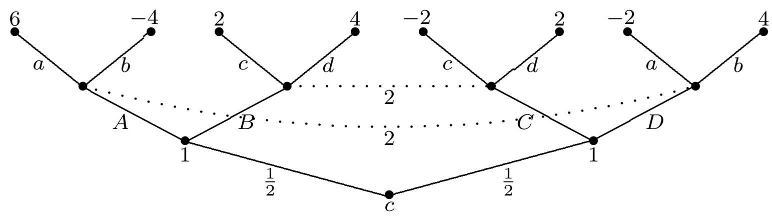

Figure 9.7

a) Does the game have perfect recall?

b) Determine the equivalent strategic form of the game.

c) Solve the game, i.e. determine a pair of optimal mixed strategies for the two players and compute the game's value.

[Hint: Eliminate dominated strategies in order to get a manageable problem.]

9.3 Player 2 chooses one of two rooms where he hides a coin. Player 1, who does not know in which room the coin is hidden, selects one of the rooms to search for the coin. If he is looking in room A and the coin is there, he will find it with 50 % probability, but if he searches in room B and the coin is there he

will find it with only 25 % probability. If he is looking in the wrong room, he will of course not find the coin. Player 1 keeps the coin if he finds it, otherwise it is returned to player 2.
Formulate the situation as an extensive game, then reduce it to strategic form and determine the optimal mixed strategies and the game's value.

9.4 Consider the game in the previous exercise and suppose player 1 gets a second chance to find the coin if he fails in the first attempt. In other words, he may choose a room again, either the same as in the first attempt or the second room, and is then allowed to search for the coin in the selected room with the same probabilities of finding it as before. Draw the game tree for this extensive game, reduce it to strategic form and determine the optimal mixed strategies and the game's value.

9.5 Determine the equivalent strategic form of the extensive zero-sum game in Figure 9.8 and solve the game.

9.6 Coin A is a fair coin with equal probability for heads and tails, while coin B is fake and has probability 1/3 of heads and 2/3 of tails. Player 1 begins by predicting "heads" or "tails". If he predicts heads, coin A is tossed, and if he predicts tails, coin B is tossed. Player 2 is informed whether player 1's prediction was right or wrong, but not about the prediction or about which coin was used, and he is then supposed to guess which of the two coins was

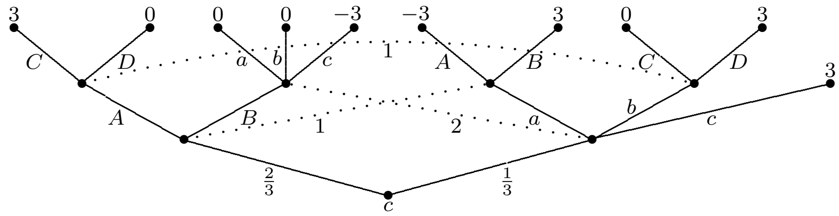

Figure 9.8

tossed. If player 2 guesses correctly, he wins 1 dollar from player 1, but if he guesses wrong and player 1's prediction was correct, he loses 2 dollar to the opponent. If both are wrong, no payment is made.
Draw the game tree, reduce the game to strategic form and solve it.

9.3 Mixed strategies and behavior strategies

By reducing extensive games to strategic form we can easily translate many concepts and results for strategic games to coresponding concepts and results for extensive games. Here follows an example.

Definition 9.3.1 A *mixed strategy* for a player i in an extensive game is a lottery, i.e. a probability distribution, on the set Σ_i of the player's strategies.

To clarify the difference between mixed strategies and strategies in Σ_i, we will often call the latter the player's *pure* strategies.

The concept of mixed strategy is natural for games in strategic form, but does not feel particularly appealing for extensive games. It does not seem natural for a player to make a random choice of strategy once and for all and then follow the chosen strategy regardless of how the opponents do their moves during the game. Instead, it seems more natural to choose a move randomly every time it is the player's turn to make a move. Such thinking leads to the concept of *behavior strategy.*

Definition 9.3.2 Let $\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_m$ be the information sets of player i in an extensive game. A *behavior strategy* λ for player i is a set $\lambda = \{\lambda^1, \lambda^2, \dots, \lambda^m\}$ of *independent* probability measures, where each λ^j is a probability measure on the set of labels in the information set $\mathcal{I}_j$.

EXAMPLE 9.3.1 Consider the game in Figure 9.8. Player 1 has four pure strategies, namely AC, AD, BC and BD. A mixed strategy for player 1 could be to choose these pure strategies with the probabilities $\frac{1}{7}$, $\frac{2}{7}$, $\frac{3}{7}$ and $\frac{1}{7}$, and the set of all his mixed strategies can obviously be identified with the subset $\{x \in \mathbf{R}^4 \mid x_1 + x_2 + x_3 + x_4 = 1,\ \ x_1, x_2, x_3, x_4 \geq 0\}$ of $\mathbf{R}^4$.

A behavior strategy for the same player has the following form: "Choose the label A with probability α and the label B with probability $(1-\alpha)$, the label C with probability β and the label D with probability $(1-\beta)$", where α and β are arbitrary real numbers in the interval $[0,1]$. The set of all behavior strategies can evidently be identified with the subset $[0,1] \times [0,1]$ of $\mathbf{R}^2$. □

Now consider an arbitrary extensive game and suppose that each player i has chosen a mixed or a behavior stratey λ_i. The outcome of a game that is played using these strategies will of course be stochastic; the probability of a certain play π depends on the strategy vector $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$ (and on the chance moves, if any). In other words, the strategy vector λ gives rise to a probability distribution P_λ on the set Π of all plays, and we will now describe how to determine this distribution.

Let π be a play of the game. A *pure* strategy of player i is said to be *compatible with the play* π if all moves that the player makes in the play have labels that belong to the strategy. The set of all pure strategies for player i that are compatible with the play π will be denoted $\Sigma_i(\pi)$.

If a player uses a strategy that is incompatible with the play π, then this play can not occur, regardless of how the opponents play and of the outcome of chance moves.

EXAMPLE 9.3.2 Figure 9.9 shows a piece of an extensive game. Let π denote the play that ends in the position p. During the path from the initial position p_0, the play π passes through three information sets that belong to player 1, and we have indicated these by dotted lines. The player has three additional information sets that are not visited by the play, and they are also indicated by dotted lines.

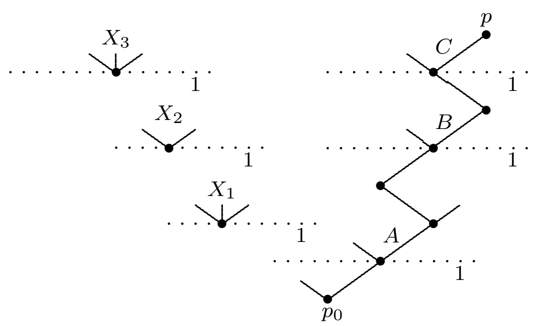

Figure 9.9. The game in Example 9.3.2.

A pure strategy of player 1 is a set of labels with one label from each information set. The strategy is compatible with the play π if and only if it contains the labels A, B and C from the three information sets in the play. The set $\Sigma_1(\pi)$ of all pure strategies that are compatible with the play thus consists of all pure strategies of the form $ABCX_1X_2X_3$, where the labels X_1, X_2 and X_3 can be choosen arbitrarily from their respective information sets. □

Now, let λ_i be a *mixed strategy* of player i in an extensive game, and let π be an arbitrary play in the game. The probability $P_i(\lambda_i;\pi)$ that the mixed strategy λ_i will result in the play π, *given that the other players and chance choose moves that belong to the play*, is equal to the probability that the lottery λ_i will result in an outcome belonging to the set $\Sigma_1(\pi)$ of strategies that are compatible with the play, and it is consequently given by the formula

$$(1) \qquad P_i(\lambda_i;\pi) = \sum_{\sigma\in\Sigma_i(\pi)} \lambda_i(\sigma).$$

We now consider the corresponding probability for behavior strategies. So let λ_i be a *behavior strategy* of player i. This means that λ_i is a family of independent probability measures $\lambda_i^{\mathcal{I}}$ with one probability measure $\lambda_i^{\mathcal{I}}$ for each information set $\mathcal{I}$ that belongs to the player.

Given an arbitrary play π, let $E_1(\pi), E_2(\pi), \dots, E_k(\pi)$ be the labels of the moves in the play that are performed by player i in the positions that belong to the player. The corresponding information sets are denoted $\mathcal{I}_1(\pi)$, $\mathcal{I}_2(\pi)$, ..., $\mathcal{I}_k(\pi)$. (Note that different positions in π may belong to the same information set, i.e. it may happen that $\mathcal{I}_\ell(\pi) = \mathcal{I}_m(\pi)$ for $\ell \neq m$, but then necessarily $E_\ell(\pi) = E_m(\pi)$. See Figure 9.5.) The number

$$P_i(\lambda_i; \pi) = \prod_{j=1}^{k} \lambda_i^{\mathcal{I}_j}(E_j(\pi)) \tag{2}$$

gives the probability that the play π will occur if player i follows his behavior strategy λ_i and all other players and chance choose the "right" moves. (Here we have to use our assumption that the probability measures $\lambda_i^{\mathcal{I}_j}$ are independent of each other, i.e. that the choice of a label in an information set is made independently of the choice of labels in the other information sets.)

We also need to take into account possible chance moves during a play. Therefore, let $p_1, p_2, \dots, p_\ell$ be those positions in the play π where chance moves are performed, and define $c(\pi)$ as the probability that all these chance moves are moves in the play π. The number $c(\pi)$ is then a product with ℓ factors, where factor no. j is the probability that the chance move in position p_j is a move in the play π.

Using the quantities introduced above, we now get the following formula for the probability that a play π will occur when the players have chosen their mixed or behavior strategies.

Proposition 9.3.1 *Suppose that each player i has chosen a mixed or a behavior strategy λ_i, and let $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$. The probability $P_\lambda(\pi)$ for the play π is given by the formula*

$$P_\lambda(\pi) = c(\pi) \cdot \prod_{i \in N} P_i(\lambda_i; \pi).$$

Proof. The play π occurs if and only if all players and chance choose moves that belong to the play, and these choices are independent events. The probability for π to occur is therefore equal to the product of the probabilities of these events. □

Each vector $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$ of mixed or behavior strategies of the n players thus generates a probability measure P_λ on the set Π of all plays. A pure strategy can of course be perceived as a special mixed strategy and as a special behavior strategy.

Let i be a player with utility function u_i. The player's expected utility $\tilde{u}_i(\lambda)$ of the strategy vector $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$, where the λ_j are mixed or behavior strategies, is given by the formula

$$\tilde{u}_i(\lambda) = \sum_{\pi \in \Pi} u_i(\pi) P_\lambda(\pi).$$

It is now clear how to define the concept of Nash equilibrium for both mixed strategies and behavior strategies.

Definition 9.3.3 An n-tuple $\lambda^* = (\lambda_1^*, \lambda_2^*, \dots, \lambda_n^*)$ of mixed strategies of the players in an extensive game is a *mixed Nash equilibrium* if

$$\tilde{u}_i(\lambda_{-i}^*, \lambda_i^*) \geq \tilde{u}_i(\lambda_{-i}^*, \tau_i)$$

for each player i and all mixed strategies τ_i of player i.

An n-tuple λ^* of behavior strategies of the players in an extensive game is a *behavior Nash equilibrium* if the same inequality holds for each player i and all behavior strategies τ_i of player i.

A mixed strategy vector λ^* in an extensive game is obviously a mixed Nash equilibrium if and only if the same vector is a mixed Nash equilibrium in the reduced strategic form of the game. The following result is therefore an immediate consequence of Proposition 5.2.1.

Proposition 9.3.2 *Every finite extensive game has a mixed Nash equilibrium.*

A natural question in this context is now whether all that a player can accomplish with mixed strategies can also be accomplished using behavior strategies, and vice versa. To clarify what we mean, we need the following definition:

Definition 9.3.4 In an extensive game, two strategies λ_i and τ_i of player i (both mixed or both behavior or one of each kind) are called *outcome equivalent* if the two probability measures $P_{(\sigma_{-i}, \lambda_i)}$ and $P_{(\sigma_{-i}, \tau_i)}$ are identical for all other players' choices of pure strategies $\sigma_1, \sigma_2, \dots, \sigma_n$.

Suppose that a player has m information sets and that the number of labels in these are respectively $\nu_1, \nu_2, \dots, \nu_m$. The number of pure strategies will then be $\nu = \nu_1 \nu_2 \cdots \nu_m$. The set of all mixed strategies of the player can therefore be perceived as a subset of the space $\mathbf{R}^\nu$, and the dimension of this set is $\nu - 1$, because one can choose the probability of each pure strategy as an arbitrary nonnegative number as long as the sum of all numbers is 1.

Each component λ^j of a behavior strategy λ is a probability measure on the labels in the corresponding information set, and the set of all such probability measures is a subset of $\mathbf{R}^{\nu_j}$ of dimension $\nu_j - 1$. The set of all behavior strategies is therefore a product set of dimension $\nu' - m$, where $\nu' = \nu_1 + \nu_2 + \cdots + \nu_m$.

It is easy to see that $\nu' - m \leq \nu - 1$ with strict inequality except in the trivial case when $\nu_j = 1$ for all indices j but possibly one. This means that there are generally "more" mixed strategies than behavior strategies. Therefore, one should not expect the two concepts to be equivalent in the sense that for each mixed strategy there is an outcome equivalent behavior strategy and vice versa. The concepts are indeed not equivalent for all games, as the following two examples show.

Example 9.3.3 Figure 9.10 shows a game with 4 plays, called π_1, π_2, π_3 and π_4; the names are written at the final positions of the games.

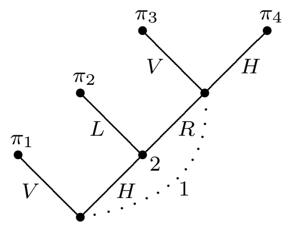

Figure 9.10. A game with a behavior strategy that is not outcome equivalent to any mixed strategy.

The two players have two pure strategies each, player 1 the strategies V and H, and player 2 the strategies L and R. The combination (V, R) results in the play π_1, while the combination (H, R) gives π_4.

Let λ_1 be player 1's mixed strategy that consists in choosing V with probability α and H with probability $1-\alpha$. The mixed strategy λ_1 combined with the pure strategy R results in the play π_1 with probability α and in the play π_4 with probability $1 - \alpha$.

The probability distribution $P_{(\lambda_1,R)}$ on the set Π of all four possible plays is thus defined by

$$P_{(\lambda_1,R)}(\pi_1) = \alpha, \;\; P_{(\lambda_1,R)}(\pi_2) = P_{(\lambda_1,R)}(\pi_3) = 0, \;\; P_{(\lambda_1,R)}(\pi_4) = 1 - \alpha.$$

Let now τ_1 be player 1's behavior strategy that consists in choosing V with probability β and consequently H with probability $1 - \beta$. When combined with player 2's pure strategy R, τ_1 gives rise to the play π_1 with probability β, the play π_3 with probability $(1 - \beta)\beta$ and the play π_4 with probability

$(1-\beta)^2$. The probability distribution $P_{(\tau_1,R)}$ is thus given by

$$P_{(\tau_1,R)}(\pi_1) = \beta,\;\; P_{(\tau_1,R)}(\pi_3) = (1-\beta)\beta,\;\; P_{(\tau_1,R)}(\pi_4) = (1-\beta)^2,$$

while $P_{(\tau_1,R)}(\pi_2) = 0$. So the play π_3 occurs with positive probability for the behavior strategy if $0 < \beta < 1$ but for no mixed strategy. In other words, there is no mixed strategy that is outcome equivalent to the behavior strategy τ_1 if for example $\beta = \frac{1}{2}$. □

Example 9.3.4 Figure 9.11 shows a game with six plays, whose names are written at the respective terminal positions.

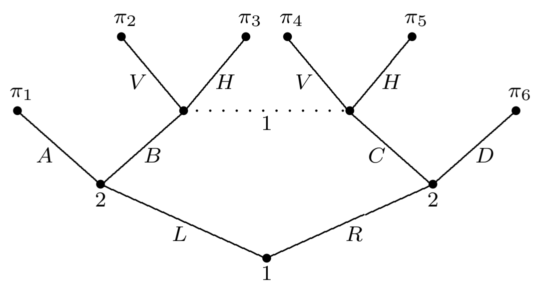

Figure 9.11. A game with a mixed strategy that is not outcome equivalent to any behavior strategy.

Player 1 has four pure strategies: LV, LH, RV and RH. Let λ_1 denote the mixed strategy which assigns these pure strategies the probabilities 0, $\frac{1}{2}$, $\frac{1}{2}$ and 0. Player 1's pure strategies in the given order combined with player 2's pure strategy BC result in the plays π_2, π_3, π_4 and π_5. The combination of the mixed strategy λ_1 and the pure strategy BC therefore leads to the two plays π_3 and π_4 with equal probability $\frac{1}{2}$ and to the other plays with probability 0. In other words, $P_{(\lambda_1,BC)}(\pi_3) = P_{(\lambda_1,BC)}(\pi_4) = \frac{1}{2}$, and $P_{(\lambda_1,BC)}(\pi) = 0$ for the other four plays π.

An arbitrary behavior strategy $\tau_1 = (\tau_1^1, \tau_1^2)$ of player 1 is defined by $\tau_1^1(L) = \alpha$, $\tau_1^1(R) = 1-\alpha$ and $\tau_1^2(V) = \beta$, $\tau_1^2(H) = 1-\beta$, where $0 \leq \alpha, \beta \leq 1$. The strategy τ_1 combined with player 2's pure strategy BC results in the plays π_2, π_3, π_4 and π_5 with probability $\alpha\beta$, $\alpha(1-\beta)$, $(1-\alpha)\beta$ and $(1-\alpha)(1-\beta)$, respectively .

In order for the strategy τ_1 to be outcome equivalent to the mixed strategy λ_1, the following system has to be satisfied

$$\begin{cases} \alpha\beta = 0 \\ \alpha(1-\beta) = \frac{1}{2} \\ (1-\alpha)\beta = \frac{1}{2} \\ (1-\alpha)(1-\beta) = 0\,. \end{cases}$$

The first equation implies that $\alpha = 0$ or $\beta = 0$, but this contradicts the second and the third equation, respectively. So the system is inconsistent, and that means that no behavior strategy is outcome equivalent to the mixed strategy λ_1. □

The games in the two examples 9.3.3 and 9.3.4 are games without perfect recall. This is no coincidence because we have the following result due to Kuhn.

Proposition 9.3.3 *Consider a player in a finite extensive game, and assume that the player has no information set that contains two positions that are connected by a move sequence; this assumption is in particular satisfied if the player has perfect recall. Then, each of the player's behavior strategy is outcome equivalent to a mixed strategy.*

Proposition 9.3.4 *For a player with perfect recall in a finite extensive game, each mixed strategy is outcome equivalent to a behavior strategy.*

Corollary 9.3.5 *For players with perfect recall in finite extensive games, each mixed strategy is outcome equivalent to a behavior strategy and vice versa.*

Proof of Proposition 9.3.3. Let player 1 be the player under consideration, and let $\mathcal{I}_1, \mathcal{I}_2, \ldots, \mathcal{I}_m$ be his information sets. To each of player 1's behavior strategies τ_1 we have to associate a mixed strategy λ_1 that is outcome equivalent to τ_1, i.e. has the property that the two probabilities $P_{(\tau_1,\sigma_2,\ldots,\sigma_n)}(\pi)$ and $P_{(\lambda_1,\sigma_2,\ldots,\sigma_n)}(\pi)$ coincide for all plays π and all pure strategies $\sigma_2, \ldots, \sigma_n$ of the other players, and because of the formula in Proposition 9.3.1 it suffices to prove that λ_1 satisfies the equality

$$P_1(\tau_1;\pi) = P_1(\lambda_1;\pi) \tag{3}$$

for all plays π.

We recall that player 1's behavior strategies have the form

$$\tau_1 = \{\tau_1^1, \tau_1^2, \ldots, \tau_1^m\},$$

where τ_1^j for each j is a probability measure on the information set $\mathcal{I}_j$. The player's mixed strategies are probability measures on the set Σ_1 that consists of the player's pure strategies, and a pure strategy σ_1 is just a list of labels with one label from each information set belonging to the player.

Given the behavior strategy τ_1 we now define the mixed strategy λ_1 by, for pure strategies σ_1 of the form $\sigma_1 = E_1E_2 \ldots E_m$ with labels E_j belonging to the information set $\mathcal{I}_j$, putting

$$\lambda_1(\sigma_1) = \prod_{j=1}^{m} \tau_1^j(E_j).$$

Then $\lambda_1(\sigma_1) \geq 0$, and $\sum_{\sigma_1\in\Sigma_1} \lambda_1(\sigma_1) = 1$, so λ_1 is indeed a probability measure on Σ_1, i.e. a mixed strategy of player 1.

It remains to show that the two strategies λ_1 and τ_1 are outcome equivalent. Therefore, let π be an arbitrary play in the game, and renumber player 1's information sets so that $\mathcal{I}_1, \mathcal{I}_2, \ldots, \mathcal{I}_k$ are the informations sets that contain the player's positions in the play π, and let $E'_1, E'_2, \ldots, E'_k$, where E'_j belongs to $\mathcal{I}_j$, denote the labels of the player's moves in the play π. It follows from the assumptions of Proposition 9.3.3 that none of the information sets contain more than one position from the play, so the information sets $\mathcal{I}_1, \mathcal{I}_2, \ldots, \mathcal{I}_k$ and hence also the labels $E'_1, E'_2, \ldots, E'_k$ are therefore certainly different. By definition,

$$P_1(\tau_1;\pi) = \prod_{j=1}^{k} \tau_1^j(E'_j).$$

To compute $P_1(\lambda_1;\pi)$, we begin by noting that the set $\Sigma_1(\pi)$ of player 1's pure strategies that are compatible with the play π consists of all strategies σ_1

of the form $\sigma_1 = E_1', E_2' \dots E_k', X_{k+1} \dots, X_m$, where the labels $X_{k+1}, \dots, X_m$ are arbitrary in their respective information sets $\mathcal{I}_{k+1}, \dots, \mathcal{I}_m$, because the choice of label is irrelevant in information sets that do not contain a position from π, whereas there is a unique choice of label in each information set that contains a position from π.

Therefore

$$
\begin{aligned}
(4) \qquad P_1(\lambda_1; \pi) &= \sum_{\sigma_1 \in \Sigma_1(\pi)} \lambda_1(\sigma_1) = \sum{}' \lambda_1(E_1' E_2' \dots E_k' X_{k+1} \dots X_m) \\
&= \sum{}' \tau_1^1(E_1')\tau_1^2(E_2') \cdots \tau_1^k(E_k')\tau_1^{k+1}(X_{k+1}) \cdots \tau_1^m(X_m) \\
&= P_1(\tau_1; \pi) \cdot \sum{}' \tau_1^{k+1}(X_{k+1}) \cdots \tau_1^m(X_m),
\end{aligned}
$$

where the symbol $\sum'$ means that the summation is taken over all choices of labels X_{k+1} in $\mathcal{I}_{k+1}$, X_{k+2} in $\mathcal{I}_{k+2}$, etc. Since each τ_1^j is a probability measure on the set of labels in $\mathcal{I}_j$,

$$
\begin{aligned}
\sum{}' \tau_1^{k+1}(X_{k+1}) \cdots \tau_1^m(X_m) &= \sum_{\text{all } X_{k+1} \in \mathcal{I}_{k+1}} \tau_1^{k+1}(X_{k+1}) \cdots \sum_{\text{all } X_m \in \mathcal{I}_m} \tau_1^m(X_m) \\
&= 1 \cdots 1 = 1,
\end{aligned}
$$

and by inserting this into equation (4) we obtain the requested equality (3). This proves the proposition. □

Example 9.3.5 We illustrate Proposition 9.3.3 with the game in Figure 9.12.

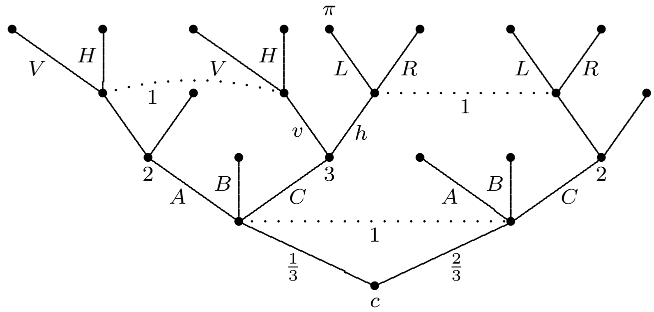

Figure 9.12. A game that illustrates Proposition 9.3.3.

Let us define a behavior strategy $\tau_1 = \{\tau_1^1, \tau_1^2, \tau_1^3\}$ by putting

$$\begin{aligned}&\tau_1^1(A) = \tfrac{1}{6},\ \ \tau_1^1(B) = \tfrac{1}{3},\ \ \tau_1^1(C) = \tfrac{1}{2}\\ &\tau_1^2(V) = \tfrac{1}{5},\ \ \tau_1^2(H) = \tfrac{4}{5}\\ &\tau_1^3(L) = \tfrac{1}{4},\ \ \tau_1^2(R) = \tfrac{3}{4}.\end{aligned}$$

Player 1 has $3 \cdot 2 \cdot 2 = 12$ pure strategies that are listed in table 9.1, which also contains the mixed strategy λ_1 that is outcome equivalent to τ_1. The probabilities $\lambda_1(\sigma_1)$ have been computed using the method in the proof of Proposition 9.3.3. For example,

$$\lambda_1(AVL) = \tau_1^1(A)\tau_1^2(V)\tau_1^3(L) = \tfrac{1}{6} \cdot \tfrac{1}{5} \cdot \tfrac{1}{4} = \tfrac{1}{120}.$$

σ_1	AVL	AVR	AHL	AHR	BVL	BVR	BHL	BHR	CVL	CVR	CHL	CHR
$\lambda_1(\sigma_1)$	$\frac{1}{120}$	$\frac{1}{40}$	$\frac{1}{30}$	$\frac{1}{10}$	$\frac{1}{60}$	$\frac{1}{20}$	$\frac{1}{15}$	$\frac{1}{5}$	$\frac{1}{40}$	$\frac{3}{40}$	$\frac{1}{10}$	$\frac{3}{10}$

Table 9.1. Player 1's pure strategies σ_1 and the mixed strategy λ_1 that is outcome equivalent to τ_1.

Let π denote the play that ends in the position which have been marked π i Figure 9.12. We will compute the two probabilities $P_{(\tau_1,\sigma_2,\sigma_3)}(\pi)$ and $P_{(\lambda_1,\sigma_2,\sigma_3)}(\pi)$ for all possible choices of pure strategies σ_2 and σ_3 of the two other players in order to show that they coincide.

The play π is impossible if player 3 chooses the strategy v, so both probabilities are equal to zero in that case.

If $\sigma_3 = h$ we get, regardless of player 2's choice σ_2,

$$P_{(\tau_1,\sigma_2,h)}(\pi) = \tfrac{1}{3} \cdot \tau_1^1(C)\tau_1^3(L) = \tfrac{1}{3} \cdot \tfrac{1}{2} \cdot \tfrac{1}{4} = \tfrac{1}{24}.$$

Player 1's pure strategies CVL and CHL are the only ones that can result in the play π, and the probabilities of those under the mixed strategy λ_1 are $\frac{1}{40}$ and $\frac{1}{10}$, respectively. Hence,

$$P_{(\lambda_1,\sigma_2,h)}(\pi) = \tfrac{1}{3}\cdot\tfrac{1}{40} + \tfrac{1}{3}\cdot\tfrac{1}{10} = \tfrac{1}{24}.$$

The two probability distributions are thus equal. □

Proof of Proposition 9.3.4. Let λ_1 be a mixed strategy of player 1. In order to define an outcome equivalent behavior strategy τ_1, we introduce a partial order $\prec$ on the set of information sets of player 1 by writing $\mathcal{I}_1 \prec \mathcal{I}_2$ if there is a move sequence from a position in the information set $\mathcal{I}_1$ to a position in the information set $\mathcal{I}_2$.

To define the behavior strategy τ_1 we have to define a probability measure $\tau_1^{\mathcal{I}}$ on the set of labels for each information set $\mathcal{I}$ that belongs to player 1. We do so recursively and begin by defining $\tau_1^{\mathcal{I}}$ for information sets $\mathcal{I}$ that are not preceeded by any other information set under the partial order $\prec$. Let $\mathcal{I}$ be such an information set; for labels E belonging to $\mathcal{I}$ we let

$$\tau_1^{\mathcal{I}}(E) = \sum_{E\in\sigma_1} \lambda_1(\sigma_1), \tag{5}$$

where the summation is over player 1's pure strategies σ_1 that contain the label E. In other words, $\tau_1^{\mathcal{I}}(E)$ is the probability that the mixed strategy λ_1 assigns the set of pure strategies that contain the label E.

Let now $\mathcal{I}$ be an arbitrary information set, and assume that the probability measures $\tau_1^{\mathcal{J}}$ have already been defined for all information sets $\mathcal{J}$ that preceed $\mathcal{I}$ with respect to the order $\prec$. Since player 1 has perfect recall, it is possible to order these sets in a chain of the form $\mathcal{I}_1 \prec \mathcal{I}_2 \prec \cdots \prec \mathcal{I}_m \prec \mathcal{I}$. The player's memory list at an arbitrary position in the information set $\mathcal{I}$ has the form $\mathcal{I}_1, A_1, \mathcal{I}_2, A_2, \dots, \mathcal{I}_m, A_m$, where each A_j is a label belonging to the information set $\mathcal{I}_j$.

Let us now suppose that the definitions are made in such a way that the following equality holds for every choice of label B_m belonging to the information set $\mathcal{I}_m$:

$$\prod_{j=1}^{m-1} \tau_1^{\mathcal{I}_j}(A_j)\cdot \tau_1^{\mathcal{I}_m}(B_m) = \sum_{A_1\dots A_{m-1}B_m\in\sigma_1} \lambda_1(\sigma_1), \tag{6}$$

where the sum is taken over all pure strategies σ_1 that contain the labels A_1, …, A_{m-1}, B_m. (The equality holds for $m = 1$ because of equality (5).)

We will now define the probability measure $\tau_1^{\mathcal{I}}$. Put

$$\alpha = \prod_{j=1}^{m} \tau_1^{\mathcal{I}_j}(A_j).$$

The number α is nonnegative, of course.

If $\alpha = 0$, then we can choose an arbitrary probability measure on the label set of $\mathcal{I}$ as $\tau_1^{\mathcal{I}}$.

If $\alpha > 0$ and E is an arbitrary label belonging to $\mathcal{I}$, we define

(7) $$\tau_1^{\mathcal{I}}(E) = \alpha^{-1} \cdot \sum_{A_1 \ldots A_m E \in \sigma_1} \lambda_1(\sigma_1).$$

Then $\tau_1^{\mathcal{I}}$ is a probability measure on the information set $\mathcal{I}$, because

$$\sum_E \tau_1^{\mathcal{I}}(E) = \alpha^{-1} \cdot \sum_E \sum_{A_1 \ldots A_m E \in \sigma_1} \lambda_1(\sigma_1) = \alpha^{-1} \cdot \sum_{A_1 \ldots A_m \in \sigma_1} \lambda_1(\sigma_1) = \alpha^{-1} \cdot \alpha = 1,$$

where we have used equation (6) with B_m replaced by A_m and the definition of α in order to get the last but one equality above. (Anyone who has studied probability theory has certainly noted that $\tau_1^{\mathcal{I}}(E)$ is defined as a conditional probability.)

Note that equation (7) implies that equation (6) holds when m is replaced by $m+1$, and this means that we have given a recursive definition of the probability measures $\tau_1^{\mathcal{I}}$ that satisfies equation (6).

It now only remains to show that the strategy τ_1 is outcome equivalent to the mixed strategy λ_1. So let π be an arbitrary play. If $\mathcal{I}_1 \prec \mathcal{I}_2 \prec \cdots \prec \mathcal{I}_m$ are the information sets that contain player 1's positions in the play and $A_1, A_2, \ldots, A_m$ are the the corresponding labels, then $\Sigma_1(\pi)$ consistes of all mixed strategies σ_1 that contains the labels $A_1, A_2, \ldots, A_m$. According to (1) and (2),

$$P_1(\lambda_1;\pi) = \sum_{A_1 \ldots A_{m-1}A_m \in \sigma_1} \lambda_1(\sigma_1)$$

and

$$P_1(\tau_1;\pi) = \prod_{j=1}^{m} \tau_1^{\mathcal{I}_j}(A_j).$$

Therefore, it follows from (6) that $P_1(\lambda_1;\pi) = P_1(\tau_1;\pi)$, and this implies that the two strategies are outcome equivalent. □

Example 9.3.6 Player 1 has perfect recall in the game in Figure 9.13, so each mixed strategy λ_1 is outcome equivalent to some behavior strategy τ_1.

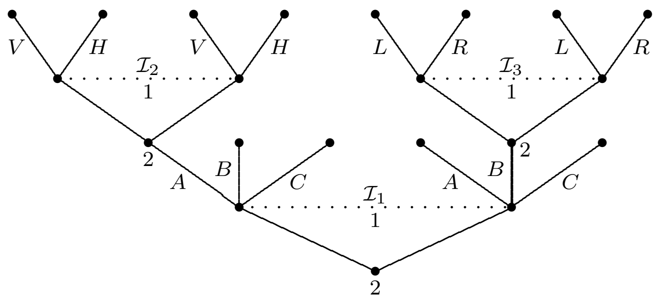

Figure 9.13. The game in Example 9.3.6.

Let us compute τ_1 when the mixed strategy λ_1 is given by table 9.2. Player 1 has three information sets, $\mathcal{I}_1$, $\mathcal{I}_2$ and $\mathcal{I}_3$, so his behavior strategy τ_1 has the form $\{\tau_1^1, \tau_1^2, \tau_1^3\}$, where each τ_1^j is a probability measure on the set of labels in $\mathcal{I}_j$. To calculate these probability measures we use the technique in the proof of Proposition 9.3.4. We first calculate

$$\tau_1^1(X) = \sum_{Y,Z} \lambda_1(XYZ) = \lambda_1(XVL) + \lambda_1(XVR) + \lambda_1(XHL) + \lambda_1(XHR)$$

σ_1	AVL	AVR	AHL	AHR	BVL	BVR	BHL	BHR	CVL	CVR	CHL	CHR
$\lambda_1(\sigma_1)$	$\frac{7}{50}$	$\frac{2}{50}$	$\frac{6}{50}$	$\frac{3}{50}$	$\frac{1}{50}$	0	$\frac{9}{50}$	$\frac{5}{50}$	$\frac{4}{50}$	$\frac{8}{50}$	$\frac{2}{50}$	$\frac{3}{50}$

Table 9.2. Player 1's mixed strategy λ_1.

for $X = A$, B and C, and the values in the table give us

$$\tau_1^1(A) = \tfrac{7}{50} + \tfrac{2}{50} + \tfrac{6}{50} + \tfrac{3}{50} = \tfrac{18}{50}$$
$$\tau_1^1(B) = \tfrac{1}{50} + 0 + \tfrac{9}{50} + \tfrac{5}{50} = \tfrac{15}{50}$$
$$\tau_1^1(C) = \tfrac{4}{50} + \tfrac{8}{50} + \tfrac{2}{50} + \tfrac{3}{50} = \tfrac{17}{50}.$$

We can then calculate $\tau_1^2(Y)$ for $Y = V$ and $Y = H$ using the formula

$$\tau_1^2(Y) = \frac{1}{\alpha}\sum_Z \lambda_1(AYZ) = \frac{1}{\alpha}\big(\lambda_1(AYL) + \lambda_1(AYR)\big), \quad \text{where} \quad \alpha = \tau_1^1(A),$$

and obtain

$$\tau_1^2(V) = \tfrac{50}{18}\cdot(\tfrac{7}{50} + \tfrac{2}{50}) = \tfrac{1}{2}, \quad \tau_1^2(H) = \tfrac{50}{18}\cdot(\tfrac{6}{50} + \tfrac{3}{50}) = \tfrac{1}{2}.$$

Analogously, $\tau_1^3(Z)$ is for $Z = L$ and $Y = R$ given by the formula

$$\tau_1^3(Z) = \frac{1}{\alpha}\sum_Y \lambda_1(BYZ) = \frac{1}{\alpha}\big(\lambda_1(BVZ) + \lambda_1(BHZ)\big) \quad \text{with} \quad \alpha = \tau_1^1(B),$$

and we get

$$\tau_1^3(L) = \tfrac{50}{15}\cdot(\tfrac{1}{50} + \tfrac{9}{50}) = \tfrac{2}{3}, \quad \tau_1^3(R) = \tfrac{50}{15}\cdot(0 + \tfrac{5}{50}) = \tfrac{1}{3}. \qquad \square$$

Exercise

9.7 Compute for the game in Figure 9.13 a behavior strategy τ_1 for player 1 that is outcome equivalen to the mixed strategyλ_1 defined by the table

σ_1	AVL	AVR	AHL	AHR	BVL	BVR	BHL	BHR	CVL	CVR	CHL	CHR
$\lambda_1(\sigma_1)$	$\frac{1}{8}$	$\frac{1}{8}$	$\frac{1}{16}$	$\frac{1}{4}$	0	0	0	0	$\frac{1}{16}$	$\frac{3}{16}$	$\frac{1}{8}$	$\frac{1}{16}$

Answers and hints for the exercises

Chapter 1

1.3 No, the completeness axiom is not satisfied.

1.4 b) No c) No

1.5 a) and c) Let x_1, x_2, x_3, x_4 be four points in the domain of a strictly concave function u with $x_1 < x_2$, $x_3 < x_4$ and $x_1 < x_3$. The slope of the chord between the points $(x_1, u(x_1))$ and $(x_2, u(x_2))$ on the curve $y = u(x)$ is then greater than the slope of the chord between the points $(x_3, u(x_3))$ and $(x_4, u(x_4))$. For a risk averse utility function u, this means that $(u(a+h) - u(a))/h > (u(b+h) - u(b))/h$ if $a < b$ and $h > 0$. In other words, the utility increase is greatest at the least wealth.

b) The utility increase of a fixed increase in wealth is the same at all wealths.

1.6 3000 SEK

1.7 No, $(\frac{1}{n}, -1) \succeq (0,0)$ for all n, but $(0,-1) = \lim(\frac{1}{n}, -1) \not\succeq (0.0)$.

1.8 To prove that a preference relation that satisfies the two implications is continuous, you can copy the proof of Lemma 1.5.4.

1.9 a) 2.9 million SEK b) 1.0427 (when the unit of x is million SEK) c) Yes d) 2 837 000 SEK

1.10 Choose a sequence of numbers α_n in the interval $]0,1[$ that converges to 1 as $n \to \infty$. If $\alpha_n p + (1-\alpha_n)r \preceq q$ for all n, then $p \preceq q$ by the continuity axiom, and this is a contradiction.

1.11 With A consisting of two elements and $\mathcal{L}(A) = \{(x, 1-x) \mid 0 \le x \le 1\}$, we get a utility function u on the lottery set by defining

$$u(x, 1-x) = \begin{cases} x & \text{if } 0 \le x < 1, \\ 0 & \text{if } x = 1. \end{cases}$$

The corresponding preference relation satisfies neither axiom 3 nor axiom 4 and is therefore not a vNM-preference relation.

Chapter 2

2.1 Payoff table:

	rock	*paper*	*scissors*
rock	$(0, 0)$	$(-1, 1)$	$(1, -1)$
paper	$(1, -1)$	$(0, 0)$	$(-1, 1)$
scissors	$(-1, 1)$	$(1, -1)$	$(0, 0)$

No Nash equilibrium.

2.2 Prisoner's dilemma: $B_1(Deny) = B_1(Confess) = \{Confess\}$, $B_2(Deny) = B_2(Confess) = \{Confess\}$.

Matching Pennies: $B_1(Head) = \{Head\}$, $B_1(Tail) = \{Tail\}$, $B_2(Head) = \{Tail\}$, $B_2(Kl) = \{Head\}$.

Stag Hunt: $B_1(X_2, \dots, X_n) = \begin{cases} \{Stag\} & \text{if all } X_i = Stag, \\ \{Hare\} & \text{otherwise.} \end{cases}$

2.3 (r_2, k_2)

2.4 (T, L) is a Nash equilibrium but not a strict one.

2.6 All players writing down the number 0 is a Nash equilibrium for all n. All players writing down the number 1 is a Nash equilibrium if $n \geq 4$.

2.7 Yes

2.8 a) Closed b) Not closed c) Not closed d) Closed e) Not closed f) Closed.

2.9 Nash equilibrium: (r_3, k_4). The row player's maxminimizing action is r_4, and his safety level is 2. The column player's maxminimizing action is k_4, and his safety level is also 2.

2.10 Both players have two maxminimizing actions, namely to choose the number 5 or the number 6. The safety level is 5.

2.11 Follows immediately from Proposition 2.4.1.

2.12 a) The row player should choose row 3, the column player column 2.
b) No Nash equilibrium.

2.13 For example the matrix $\begin{bmatrix} 1 & 1 & 2 \\ 1 & 1 & 2 \\ 0 & 0 & 3 \end{bmatrix}$.

Chapter 3

3.1 $q_i^* = (\sqrt{n^2+1} - 1)/n^2$ for $i = 1, 2, \ldots, n$.

3.2 $q^* = (2.5, 2.5)$

3.3 a) $\frac{1}{3}(a + c_2 - 2c_1, a + c_1 - 2c_2)$ if $a \geq 2c_2 - c_1$;
$\frac{1}{2}(a - c_1, 0)$ if $a < 2c_2 - c_1$.
b) $\frac{1}{3}(a - c, a - c)$ if $b \leq \frac{1}{9}(a - c)^2$;
$\frac{1}{2}(a - c, 0)$ and $\frac{1}{2}(0, a - c)$ if $\frac{1}{16}(a - c)^2 \leq b \leq \frac{1}{4}(a - c)^2$;
$(0, 0)$ if $b \geq \frac{1}{4}(a - c)^2$.

Chapter 4

4.1 (x, x, y), (x, y, x) and (y, x, x).

4.2 a) Yes. $\Phi(C, C) = 0$, $\Phi(C, F) = -1$, $\Phi(F, C) = -2$, $\Phi(F, F) = 0$ is a potential function.
b) Yes. $\Phi(H, H) = \Phi(D, D) = 0$, $\Phi(H, D) = \Phi(D, H) = 1$ is a potential function.

4.3 Every improvement path of maximal length ends at (T, R). If Φ is an ordinal potential function, then

$$\Phi(T, L) < \Phi(B, L) < \Phi(B, R) < \Phi(T, R) = \Phi(T, L),$$

which is contradictory.

Chapter 5

5.1 No mixed Nash equilibria in addition to the pure (*Confess*, *Confess*).

5.2 A total of three mixed Nash equilibria − the pure (*Consert*, *Consert*) and (*Football*, *Football*), and the mixed where player 1 chooses *Concert* with probability $\frac{2}{3}$ and player 2 chooses *Football* with probability $\frac{2}{3}$.

5.3 Only the two pure Nash equilibria (T, L) and (B, R).

5.4 $p_2^* = (\frac{7}{15}, \frac{5}{15}, \frac{3}{15})$

5.5 $p_2^* = \frac{1}{5}(3 - 2t, 5t, 2 - 3t)$, where $0 \le t \le \frac{1}{3}$.

5.7 Only (r_2, k_2) survives, and this pair is therefore the game's unique Nash equilibrium.

5.8 a) No b) Only the Nash equilibrium which consists of all players writing the number 1.

5.9 Player 1's mixed maxminimizing strategies: $(t, 1-t)$, where $0 \le t \le \frac{2}{3}$.
Player 2's mixed maxminimizing strategies: $(t, 1-t)$, where $0 \le t \le \frac{4}{5}$.
Both players' safety level: 2.

Chapter 6

6.1 The game is defined by the matrix

	1	2
1	−1	2
2	2	−4

Both players have the same optimal strategy, namely to choose the number 1 with probability $\frac{2}{3}$ and the number 2 with probability $\frac{1}{3}$. The value of the game is 0.

6.2 a) Row player: $(\frac{3}{4}, \frac{1}{4})$, Column player: $(\frac{1}{2}, \frac{1}{2})$. Value: $\frac{5}{2}$.
b) Row player: $(\frac{3}{7}, \frac{3}{7}, \frac{1}{7})$, Column player: $(\frac{5}{7}, \frac{1}{7}, \frac{1}{7})$. Value: $\frac{3}{7}$.

6.3 The strategies are optimal (the indifference principle!). The game's value: $\frac{1}{3}$.

6.4 a) Row player: $(0, 1, 0)$, Column player: $(0, 0, 1)$. Value: 4.
b) Row player: $(\frac{5}{6}, \frac{1}{6})$, Column player: $(\frac{2}{3}, \frac{1}{3})$. Value: $\frac{7}{3}$.
c) Row player: $(\frac{1}{3}, \frac{2}{3}, 0)$, Column player: $(\frac{2}{3}, 0, \frac{1}{3})$. Value: $\frac{10}{3}$.

6.5 The Row player has the following payoff matrix:

	♢A	♠A	♢2
♢A	−1	1	−2
♠A	1	−1	1
♠2	2	−1	2

The column player should solve the linear programming problem

$$\text{Minimize } t \text{ as}$$
$$\begin{cases} -y_1 + y_2 - 2y_3 \le t \\ y_1 - y_2 + y_3 \le t \\ 2y_1 - y_2 + 2y_3 \le t \\ y_1 + y_2 + y_3 = 1 \\ y_1, y_2, y_3 \ge 0 \end{cases}$$

(The minimum is equal to 0 and is obtaind for $y = (0, \frac{2}{3}, \frac{1}{3})$. The row player's optimal strategy is $(\frac{1}{2}, 0, \frac{1}{2})$.)

Chapter 8

8.1 Let $P = \{p_0, p_1, p_2, p_3, \dots\}$, $s(p_0) = \{p_1, p_2, p_3, \dots\}$ and $e(p_n) = \emptyset$ for $n = 1, 2, 3, \dots$. The game tree $\langle P, p_0, s\rangle$ is non-finite and has height equal to 1.

8.2 Let $P = \{p_0\} \cup \{p_{ik} \mid 1 \le k \le i,\ i = 1, 2, 3, \dots\}$, and define a successor function s by putting $s(p_0) = \{p_{i1} \mid i = 1, 2, 3, \dots\}$, $s(p_{ik}) = \{p_{i,k+1}\}$ for $1 \le k < i$, and $s(p_{ii}) = \emptyset$. Every play in $\langle P, p_0, s\rangle$ is of the form $p_0, p_{i1}, p_{i2}, \dots, p_{ii}$ and has finite length.

8.3 a) Player 1: A, B. Player 2: CE, CF, DE, DF.
b) (A, CE) and (A, CF)
c)

	CE	CF	DE	DF
A	$(2,1)$	$(2,1)$	$(3,0)$	$(3,0)$
B	$(0,2)$	$(1,3)$	$(0,2)$	$(1,3)$

8.4 a) With $u_1(A) = 2$, $u_1(B) = 1$, $u_1(C) = 0$ and $u_2(A) = 0$, $u_2(B) = 1$, $u_2(C) = 2$ the game tree looks like this:

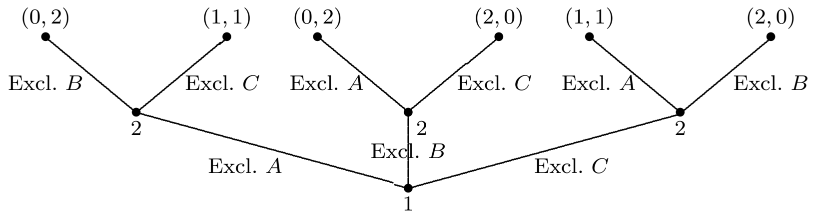

The game tree in exercise 8.4.

There are two Nash equilibria. In both of them, player 1 excludes option C. In one of them, player 2 excludes B if his opponent excludes A, and A if the opponent excludes B or C. In the other one, player 2 excludes C if his opponent excludes A, and A if the opponent excludes B or C.

b) Four Nash equilibria. In all of them, player 1 excludes alternative C. Player 2's response if player 1 excludes A, B and C, respectively is (B, A, B), (B, C, B), (C, A, B) or (C, C, B)

8.5 Nash equilibria: All combinations where player 1 offers $\$\,x$, and player 2 accepts x, declines all lower offers and accepts or declines all higher offers. In addition the combination where player 1 offers $\$\,0$ and player 2 declines all offers, and the combination where player 1 offers $\$\,0$ and player 2 declines all offers except $\$\,100$. (The number of Nash equilibria is $2^{100} + 1$.)

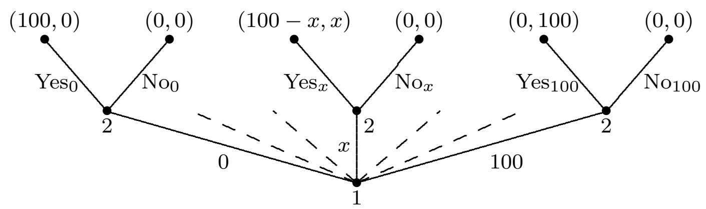

The game tree in Exercise 8.5.

8.6 Analogous to Exercise 8.5, but x is now an arbitrary real number in the interval $[0, 100]$.

8.7 (A, CF)

8.8 a) The combination where player 1 removes option C, and player 2 removes B, A and A, respectively, if player 1 removes A, B and C, respectively. (The subgame perfect equilibrium solution thus results in B remaining.)

b) The combination where player 1 removes option C, and player 2 removes B, A and B, respectively, if player 1 removes A, B and C, respectively. (The subgame perfect equilibrium solution thus results in A remaining.)

8.9 Two subgame perfect equilibria, namely that player 1 offers \$ 0 and player 2 accepts all offers (with a payoff of \$ 100 to player 1), and that player 1 offers \$ 1 and player 2 declines the offer \$ 0 and accepts all other offers (with a payoff of \$ 99 to player 1).

8.10 One subgame perfect equilibrium: Player 1 offers \$ 0 and player 2 accepts all offers.

8.11 Game tree, where x represents the share of the cake that player 1 offers player 2:

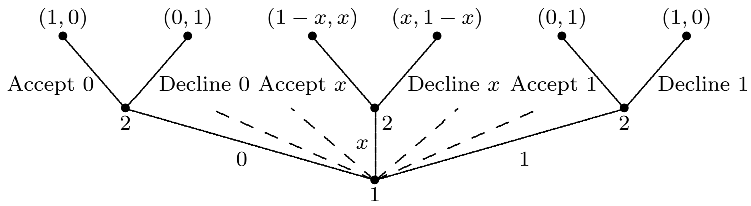

In the subgame perfect equilibrium solution, player 1 offers exactly half the cake and player 2 always takes the biggest part of the cake.

8.12 (A, EGJ), (A, EHJ), (A, FGJ), (B, FGJ), (C, FGJ) and (B, FHJ) with payoffs $(3, 0)$, $(3, 0)$, $(1, 0)$, $(1, 1)$, $(1, 3)$ and $(2, 1)$, respectively.

8.13 a) The centipede game has a unique subgame perfect equilibrium. At each non-terminal position the player in turn to move should take a step downwards. Outcome: $(1, 1)$.

c) All alternatives where the two players take a step downwards in their first moves are Nash equilibria.

d) No.

8.14 $(2, \frac{4}{3})$

8.15 Let (q_1^*, q_2^*) be the Cournot equilibrium. Player 1's Stackelberg solution $\overline{q}_1$ is obtained by maximizing the function $V_1(q_1, m(q_1))$, where $m(q_1)$ denotes is a maximum point of $V_2(q_1, q_2)$, considered as a function of q_2. We have $m(q_1^*) = q_2^*$, because $V_2(q_1^*, q_2^*) \geq V_2(q_1^*, q_2)$ for all q_2, and hence $V_1(\overline{q}_1, m(\overline{q}_1)) \geq V_1(q_1^*, m(q_1^*)) = V_1(q_1^*, q_2^*)$.

8.16 (BHI, DF). Expected utility: $(2.2, 2.8)$.

Chapter 9

9.1 a) Player 1 should with probability $\frac{5}{6}$ bet with winning cards and pass with losing cards and with probability $\frac{1}{6}$ bet with both winning and losing cards. Player 2 should call and fold with equal probability $\frac{1}{2}$. Value: $-\frac{1}{4}$.

b) If $0 < p \leq \frac{2}{3}$, player 1 should with probability $(2-3p)/(2-2p)$ bet with winning cards and pass with losing cards, and with probability $p/(2-2p)$ bet with both winning and losing cards. Player 2 should call and fold with equal probability $\frac{1}{2}$. The value of the game is $3p-1$.
If $p \geq \frac{2}{3}$, player 1 should always bet and player 2 should alwas fold, and the value of the game is 1.

9.2 a) Yes.

b) Strategic form:

	ac	ad	bc	bd
AC	2	4	-3	-1
AD	2	2	0	0
BC	0	3	0	3
BD	0	1	3	4

c) Player 1 should choose AD and BD with probability $\frac{3}{5}$ and $\frac{2}{5}$, respectively, and player 2 should choose ac and bc with probability $\frac{3}{5}$ and $\frac{2}{5}$, respectively. Value: $\frac{6}{5}$.

9.3 Game tree:

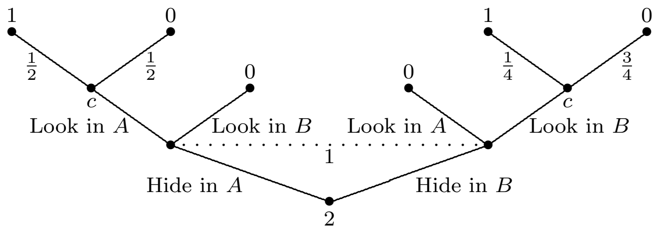

Strategic form:

	Hide in A	Hide in B
Look in A	$\frac{1}{2}$	0
Look in B	0	$\frac{1}{4}$

Player 2's optimal strategy is to hide in room A with probability $\frac{1}{3}$ and in room B with probability $\frac{2}{3}$. Player 1's optimal strategy is to look in A with probability $\frac{1}{3}$ and in B with probability $\frac{2}{3}$. Value: $\frac{1}{6}$.

9.4 Game tree:

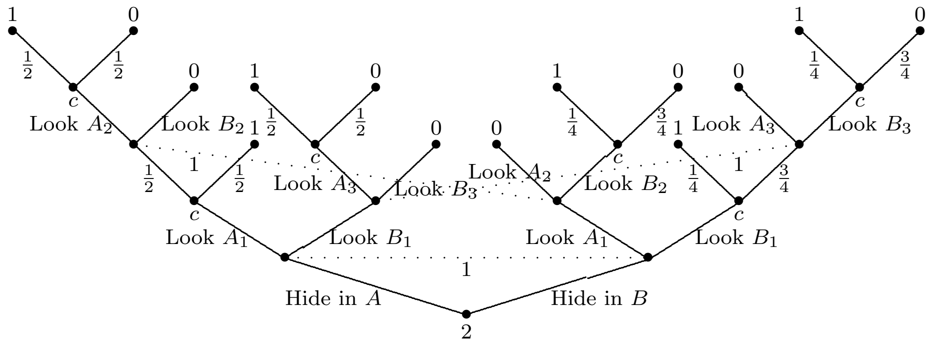

Strategic form:

	Hide in A	Hide in B
Look in A_1, Look in A_2, Look in A_3	$\frac{3}{4}$	0
Look in A_1, Look in A_2, Look in B_3	$\frac{3}{4}$	0
Look in A_1, Look in B_2, Look in A_3	$\frac{1}{2}$	$\frac{1}{4}$
Look in A_1, Look in B_2, Look in B_3	$\frac{1}{2}$	$\frac{1}{4}$
Look in B_1, Look in A_2, Look in A_3	$\frac{1}{2}$	$\frac{1}{4}$
Look in B_1, Look in A_2, Look in B_3	0	$\frac{7}{16}$
Look in B_1, Look in B_2, Look in A_3	$\frac{1}{2}$	$\frac{1}{4}$
Look in B_1, Look in B_2, Look in B_3	0	$\frac{7}{16}$

Player 2's optimal strategy is to hide the coin in room A with probability $\frac{3}{11}$ and in room B with probability $\frac{8}{11}$.
Player 1's optimal strategy:
Start the search in room B and continue in the same room if the coin is not found in the first try, with probability $\frac{4}{11}$;
Start the search in room A and continue in room B if the coin is not found, with probability $\frac{7}{11}$, or vice versa, i.e. start in room B and continue in room A with probability $\frac{7}{11}$.
The value of the game is $\frac{7}{22}$.

9.5 Strategic form:

	a	b	c
AC	1	2	3
AD	-1	1	1
BC	1	0	-1
BD	1	1	-1

Optimal strategies: Choose AC and a, respectively with probability 1.

9.6 Game tree:

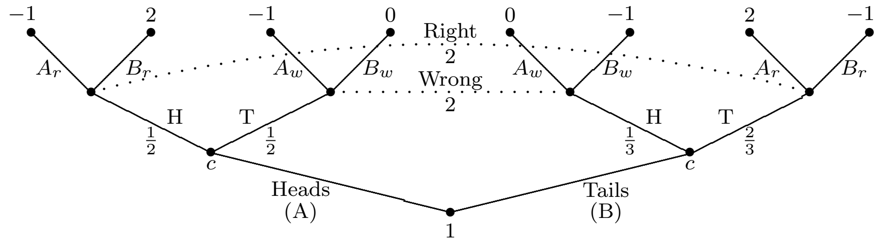

Strategic form:

	A_rA_w	A_rB_w	B_rA_w	B_rB_w
Heads	-1	$-\frac{1}{2}$	$\frac{1}{2}$	1
Tails	$\frac{4}{3}$	1	$-\frac{2}{3}$	-1

Player 1's optimal strategy is to predict Heads with probability $\frac{4}{7}$ and Tails with probability $\frac{3}{7}$. Player 2 should with probability $\frac{1}{3}$ guess coin A regardless of whether the opponent's prediction was correct or not (the strategy A_rA_w), and with probability $\frac{2}{3}$ guess coin B if the prediction was correct and coin A if the prediction was wrong (the strategy B_rA_w). The game is fair because its value is 0.

9.7 $\tau_1^1(A) = \frac{9}{16}$, $\tau_1^1(B) = 0$, $\tau_1^1(C) = \frac{7}{16}$; $\tau_1^2(V) = \frac{4}{9}$, $\tau_1^2(H) = \frac{5}{9}$; $\tau_1^3(L) = p$, $\tau_1^3(R) = 1 - p$, where $0 \le p \le 1$ is arbitrary.

Index

To see Part II download
Cooperative Games: An Introduction to Game Theory – Part II