Strategic Games

Krzysztof R. Apt

October 22, 2018

Contents

Introduction

Mathematical game theory, as launched by Von Neumann and Morgenstern in their seminal book [63], followed by Nash' contributions [42, 43], has become a standard tool in Economics for the study and description of various economic processes, including competition, cooperation, collusion, strategic behaviour and bargaining. Since then it has also been successfuly used in Biology, Political Sciences, Psychology and Sociology. With the advent of the Internet game theory became increasingly relevant in Computer Science.

One of the main areas in game theory are ***strategic games***, (sometimes also called ***non-cooperative games***), which form a simple model of interaction between profit maximizing players. In strategic games each player has a payoff function that he aims to maximize and the value of this function depends on the decisions taken *simultaneously* by all players. Such a simple description is still amenable to various interpretations, depending on the assumptions about the existence of *private information*. The purpose of these lecture notes is to provide a simple introduction to the most common concepts used in strategic games and most common types of such games.

Many books provide introductions to various areas of game theory, including strategic games. Most of them are written from the perspective of applications to Economics. In the nineties the leading textbooks were [41], [11], [22] and [48].

Moving to the next decade, [47] is an excellent, broad in its scope, undergraduate level textbook, while [49] is probably the best book on the market for the graduate level. Undeservedly less known is the short and lucid [61]. An elementary, short introduction, focusing on the concepts, is [57]. In turn, [50] is a comprehensive book on strategic games that also extensively discusses ***extensive games***, i.e., games in which the players choose actions in turn. Finally, [12] is thoroughly revised version of [11].

Several textbooks on microeconomics include introductory chapters on game theory, including strategic games. Two good examples are [35] and [27]. In turn, [45] is a recent collection of surveys and introductions to the computational aspects of game theory, with a number of articles concerned with strategic games and mechanism design.

Finally, [36] is a most recent, very comprehensive account of various areas of game theory, while [59] is an elegant introduction to the subject.

Chapter 1

Nash equilibria and social optima

1.1 Nash equilibrium

Assume a set $\{1, \ldots, n\}$ of players, where $n > 1$. A ***strategic game*** (or ***non-cooperative game***) for n players consists of

- a non-empty (possibly infinite) set S_i of ***strategies***,
- a ***payoff function*** $p_i : S_1 \times \cdots \times S_n \to \mathbb{R}$,

for each player i. We write it as $(n, S, \mathbf{p})$, where $S := S_1 \times \cdots \times S_n$ and $\mathbf{p} := p_1, \ldots, p_n$.

We study strategic games under the following basic assumptions:

- players choose their strategies *simultaneously*; subsequently each player receives a payoff from the resulting joint strategy,
- each player is ***rational***, which means that his objective is to maximize his payoff,
- players have ***common knowledge*** of the game and of each others' rationality.[1]

[1]Intuitively, common knowledge of some fact means that everybody knows it, everybody knows that everybody knows it, etc.

Here are three classic examples of strategic two-player games which we shall discuss more extensively a moment. We represent such games in the form of a matrix assuming that the players choose respectively a row or a column. Each entry represents the resulting payoffs to the row and column players. So for instance in the first example, the **Prisoner's Dilemma** game, when the row player chooses C (cooperate) and the column player chooses D (defect), then the payoff for the row player is 0 and the payoff for the column player is 3. Such a representation is called a ***bimatrix*** .

Prisoner's Dilemma

	C	D
C	$2,2$	$0,3$
D	$3,0$	$1,1$

Battle of the Sexes

	F	B
F	$2,1$	$0,0$
B	$0,0$	$1,2$

Matching Pennies

	H	T
H	$1,-1$	$-1,\ 1$
T	$-1,\ 1$	$1,-1$

We introduce now some basic notions that will allow us to discuss and analyze strategic games in a meaningful way. Fix a strategic game $(n, S, \mathbf{p})$. We call each element $s \in S$ a ***joint strategy***, sometimes also called a ***strategy profile***, denote the ith element of s by s_i, and abbreviate the sequence $(s_j)_{j \neq i}$ to s_{-i}. We write (s'_i, s_{-i}) to denote the joint strategy in which player's i strategy is s'_i and each other player's strategy is s_j. Occasionally we write (s_i, s_{-i}) instead of s. Finally, we abbreviate the Cartesian product $\times_{j \neq i} S_j$ to S_{-i}.

We call a strategy s_i of player i a ***best response*** to a joint strategy s_{-i} of his opponents if

$$\forall s'_i \in S_i : p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i}).$$

Next, we call a joint strategy s a ***Nash equilibrium*** if each s_i is a best response to s_{-i}, that is, if

$$\forall i \in \{1, \ldots, n\} : \forall s'_i \in S_i \ p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i}).$$

So a joint strategy is a Nash equilibrium if no player can achieve a higher payoff by *unilaterally* switching to another strategy. Intuitively, a Nash equilibrium is a situation in which each player is a posteriori satisfied with his choice. Let us return now the three above introduced games.

Prisoner's Dilemma

The Prisoner's Dilemma game has a unique Nash equilibrium, namely (D, D). One of the peculiarities of this game is that in its unique Nash equilibrium each player is worse off than in the outcome (C, C). We shall return to this game once we have more tools to study its characteristics.

To clarify the importance of this game we now provide a couple of simple interpretations of it. The first one,

> Each player decides whether he will receive 1000 dollars or the other will receive 2000 dollars. The decisions are simultaneous and independent.

So the entries in the bimatrix of the Prisoner's Dilemma game refer to the thousands of dollars each player will receive. For example, if the row player asks to give 2000 dollars to the other player, and the column player asks for 1000 dollar for himself, the row player gets nothing while column player gets 3000 dollars. This contingency corresponds to the 0,3 entry in the bimatrix.

The original interpretation of this game that explains its name refers to the following story.

> Two suspects are taken into custody and separated. The district attorney is certain that they are guilty of a specific crime, but he does not have adequate evidence to convict them at a trial. He points out to each prisoner that each has two alternatives: to confess to the crime the police are sure they have done (C), or not to confess (N).
>
> If they both do not confess, then the district attorney states he will book them on some very minor trumped-up charge such as petty larceny or illegal possession of weapon, and they will both

> receive minor punishment; if they both confess they will be prosecuted, but he will recommend less than the most severe sentence; but if one confesses and the other does not, then the confessor will receive lenient treatment for turning state's evidence whereas the latter will get "the book" slapped at him.

This is represented by the following bimatrix, in which each negative entry, for example -1, corresponds to the 1 year prison sentence ('the lenient treatment' referred to above):

	C	N
C	$-5, -5$	$-1, -8$
N	$-8, -1$	$-2, -2$

The negative numbers are used here to be compatible with the idea that each player is interested in maximizing his payoff, so, in this case, of receiving a lighter sentence. So for example, if the row suspect decides to confess, while the column suspect decides not to confess, the row suspect will get 1 year prison sentence (the 'lenient treatment'), the other one will get 8 years of prison (' "the book" slapped at him').

Many other natural situations can be viewed as a Prisoner's Dilemma game. This allows us to explain the underlying, undesidered phenomena.

Consider for example the arms race between two, equally strong, countries. For each of them it is beneficial not to arm instead of to arm. Yet both countries end up arming themselves. Next, suppose that two companies produce a similar product and may choose between low and high advertisement costs. Both end up heavily advertising. As a final example consider a couple seeking a divorce. Each partner can choose an inexpensive (bad) or an expensive (good) layer. In the end both partners end up choosing expensive lawyers.

Matching Pennies game

Next, consider the Matching Pennies game. This game formalizes a game that used to be played by children. Each of two children has a coin and simultaneously shows heads (H) or tails (T). If the coins match then the first child wins, otherwise the second child wins. This game has no Nash equilibrium. This corresponds to the intuition that for no outcome both players are satisfied. Indeed, in each outcome the losing player regrets his choice. Moreover, the sum of the payoffs is always 0. Such games, unsurprisingly, are called ***zero-sum games*** and we shall return to them later. Also, we shall return to this game once we have introduced *mixed strategies.*

Battle of the Sexes game

Finally, consider the Battle of the Sexes game. The interpretation of this game is as follows. A couple has to decide whether to go out for a football match (F) or a ballet (B). The man, the row player prefers a football match over the ballet, while the woman, the column player, the other way round. Moreover, each of them prefers to go out together than to end up going out separately. This game has two Nash equilibria, namely (F, F) and (B, B).

Obviously, all three games are very simplistic. They deal with two players and each player has to his disposal just two strategies. In what follows we shall introduce many interesting examples of strategic games. Some of them will deal with many players and some games will have several, sometimes an infinite number of strategies.

To close this section we consider two examples of more interesting games, one for two players and another one for an arbitrary number of players.

Example 1.1 (Traveler's dilemma) Suppose that two travellers have identical luggage, for which they both paid the same price. Their luggage is damaged (in an identical way) by an airline. The airline offers to compensate them for their luggage. They may ask for any dollar amount between \$2 and \$100. There is only one catch. If they ask for the same amount, then that is what they will both receive. However, if they ask for different amounts —say one asks for $\$m$ and the other for $\$m'$, with $m < m'$— then whoever asks for $\$m$ (the lower amount) will get $\$(m + 2)$, while the other traveller will get $\$(m - 2)$. The question is: what amount of money should each traveller ask for?

We can formalize this problem as a two-player strategic game, with the set $\{2, \ldots, 100\}$ of natural numbers as possible strategies. The following payoff function[2] formalizes the conditions of the problem:

$$p_i(s) := \begin{cases} s_i & \text{if } s_i = s_{-i} \\ s_i + 2 & \text{if } s_i < s_{-i} \\ s_{-i} - 2 & \text{otherwise} \end{cases}$$

It is easy to check that $(2, 2)$ is a Nash equilibrium. To check for other Nash equilibria consider any other combination of strategies (s_i, s_{-i}) and

[2]We denote in two-player games the opponent of player i by $-i$, instead of $3 - i$.

suppose that player i submitted a larger or equal amount, i.e., $s_i \geq s_{-i}$. Then player's i payoff is s_{-i} if $s_i = s_{-i}$ or $s_{-i} - 2$ if $s_i > s_{-i}$.

In the first case he will get a strictly higher payoff, namely $s_{-i} + 1$, if he submits instead the amount $s_{-i} - 1$. (Note that $s_i = s_{-i}$ and $(s_i, s_{-i}) \neq (2, 2)$ implies that $s_{-i} - 1 \in \{2, \ldots, 100\}$.) In turn, in the second case he will get a strictly higher payoff, namely s_{-i}, if he submits instead the amount s_{-i}.

So in each joint strategy $(s_i, s_{-i}) \neq (2, 2)$ at least one player has a strictly better alternative, i.e., his strategy is not a best response. This means that $(2, 2)$ is a unique Nash equilibrium. This is a paradoxical conclusion, if we recall that informally a Nash equilibrium is a state in which both players are satisfied with their choice. □

Example 1.2 Consider the following ***beauty contest game***. In this game there are $n > 2$ players, each with the set of strategies equal $\{1, \ldots, 100\}$, Each player submits a number and the payoff to each player is obtained by splitting 1 equally between the players whose submitted number is closest to $\frac{2}{3}$ of the average. For example, if the submissions are $29, 32, 29$, then the payoffs are respectively $\frac{1}{2}, 0, \frac{1}{2}$.

Finding Nash equilibria of this game is not completely straightforward. At this stage we only observe that the joint strategy $(1, \ldots, 1)$ is clearly a Nash equilibrium. We shall answer the question of whether there are more Nash equilibria once we introduce some tools to analyze strategic games. □

Until now we associated with each player a payoff function p_i. An alternative is to associate with each player a ***cost function*** c_i. Then the objective of each player is to minimize the cost. Consequently, when the cost functions are used, a joint strategy s is a Nash equilibrium if

$$\forall i \in \{1, \ldots, n\} \ \forall s'_i \in S_i : c_i(s_i, s_{-i}) \leq c_i(s'_i, s_{-i}).$$

It is straightforward to associate with each game that uses cost functions a customary strategic game by using

$$p_i(s) := -c_i(s). \tag{1.1}$$

We denote games with cost functions by $(n, S, \mathbf{c})$, where $S := S_1 \times \cdots \times S_n$ and $\mathbf{c} := c_1, \ldots, c_n$.

Example 1.3 Consider the following example of a game with cost functions. Assume that there are 5 drivers. Each driver chooses a road from Katowice to Gliwice. The more drivers choose the same road the larger delay results. In Figure 1.1 we use the notation 1/4/5 to denote the delays associated with the middle road, which are: 1 if one driver uses it, 4 if two drivers use it and 5 when three or more use it, and analogously for the other two roads. combinations.

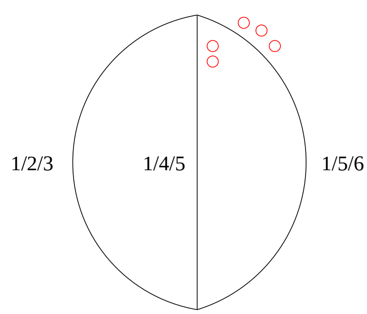

Figure 1.1: Three roads with the incurred delays

This translates into a game for 5 players, in which each player i has three strategies 1, 2 or 3 corresponding to the selected road, and the following cost function c_i:

$$c_i(s) := \begin{cases} 1 & \text{if } s_i = 1 \text{ and } |\{j \mid s_j = 1\}| = 1 \\ 2 & \text{if } s_i = 1 \text{ and } |\{j \mid s_j = 1\}| = 2 \\ 3 & \text{if } s_i = 1 \text{ and } |\{j \mid s_j = 1\}| \geq 3 \\ 1 & \text{if } s_i = 2 \text{ and } |\{j \mid s_j = 2\}| = 1 \\ 4 & \text{if } s_i = 2 \text{ and } |\{j \mid s_j = 2\}| = 2 \\ 5 & \text{if } s_i = 2 \text{ and } |\{j \mid s_j = 2\}| \geq 3 \\ 1 & \text{if } s_i = 3 \text{ and } |\{j \mid s_j = 3\}| = 1 \\ 5 & \text{if } s_i = 3 \text{ and } |\{j \mid s_j = 3\}| = 2 \\ 6 & \text{if } s_i = 3 \text{ and } |\{j \mid s_j = 3\}| \geq 3 \end{cases}$$

It is easy to see that in this game all Nash equilibria are joint strategies in which three drivers choose the first road and the remaining two drivers choose respectively the second and the third road. This game is an example of a ***congestion game***, a class of games introduced in Section 3.3 and further studied in Chapter 4. □

1.2 Social optima

To discuss strategic games in a meaningful way we need to introduce further, natural, concepts. Fix a strategic game $(n, S, \mathbf{p})$.

Given a joint strategy s we call the sum $\sum_{j=1}^{n} p_j(s)$ the ***social welfare*** of s and denote by $SW(s)$. Next, we call a joint strategy s a ***social optimum*** if its social welfare is maximal.

The concepts of a Nash equilibrium and a social optimum differ. For example in the Prisoner's Dilemma game there is only one social optimum, namely (C, C), and the unique Nash equilibrium, (D, D), is not a social optimum. In contrast in the Battle of Sexes game both concepts coincide.

The tension between Nash equilibria and social optima present in the Prisoner's Dilemma game occurs in several other natural games. It forms one of the fundamental topics in the theory of strategic games. We illustrate now this phenomenon by a number of examples and return to it in Chapter 4.

Example 1.4 (Prisoner's Dilemma for n players)
The Prisoner's Dilemma game can be easily generalized to n players as follows. It is convenient to assume that each player has two strategies, 1, representing cooperation (formerly C), and 0, representing defection (formerly D). Then, given a joint strategy s_{-i} of the opponents of player i, $\sum_{j \neq i} s_j$ denotes the number of 1-strategies in s_{-i}. Denote by $\mathbf{1}$ the joint strategy in which each strategy equals 1 and similarly with $\mathbf{0}$.

We put

$$p_i(s) := \begin{cases} 2\sum_{j \neq i} s_j + 1 & \text{if } s_i = 0 \\ 2\sum_{j \neq i} s_j & \text{if } s_i = 1 \end{cases}$$

Note that for $n = 2$ we get the original Prisoner's Dilemma game.

It is easy to check that the strategy profile $\mathbf{0}$ is the unique Nash equilibrium in this game. Indeed, in each other strategy profile a player who chose 1 (cooperate) gets a higher payoff when he switches to 0 (defect).

Finally, note that the social welfare in $\mathbf{1}$ is $2n(n-1)$, which is strictly more than n, the social welfare in $\mathbf{0}$. We now show that $\mathbf{1}$ is a unique social optimum. To this end it suffices to note that if a single player switches from 0 to 1, then his payoff decreases by 1 but the payoff of each other player increases by 2, and hence the social welfare increases. □

The next two examples deal with the depletion of ***common resources***, which in economics are goods that are not *excludable* (people cannot be prevented from using them) but are *rival* (one person's use of them diminishes another person's enjoyment of it). Examples are congested toll-free roads, fish in the ocean, or the environment. The overuse of such common resources leads to their destruction. This phenomenon is called the ***tragedy of the commons***.

One way to model it is as a Prisoner's dilemma game for n players. But such a modeling is too crude as it does not reflect the essential characteristics of the problem. We provide two more adequate modeling of it, one for the case of a binary decision (for instance, whether to use a congested road or not), and another one for the case when one decides about the intensity of using the resource (for instance on what fraction of a lake should one fish).

Example 1.5 (Tragedy of the commons I) Assume $n > 1$ players, each having to its disposal two strategies, 1 and 0 reflecting respectively that the player decides to use the common resource or not. If he does not use the resource, he gets a fixed payoff. Further, the users of the resource get the same payoff. Finally, the more users of the common resource the smaller payoff for each of them gets, and when the number of users exceeds a certain threshold it is better for the other players not to use the resource.

The following payoff function realizes these assumptions:

$$p_i(s) := \begin{cases} 0.1 & \text{if } s_i = 0 \\ \frac{F(\sum_{j=1}^{n} s_j)}{\sum_{j=1}^{n} s_j} & \text{otherwise} \end{cases}$$

where

$$F(m) := 1.1m - 0.1m^2.$$

Indeed, the function $F(m)/m$ is strictly decreasing. Moreover, $F(9)/9 = 0.2$, $F(10)/10 = 0.1$ and $F(11)/11 = 0$. So when there are already ten or more users of the resource it is indeed better for other players not to use the resource.

To find a Nash equilibrium of this game, note that given a strategy profile s with $m = \sum_{j=1}^{n} s_j$ player i profits from switching from s_i to another strategy in precisely two circumstances:

- $s_i = 0$ and $F(m+1)/(m+1) > 0.1$,
- $s_i = 1$ and $F(m)/m < 0.1$.

In the first case we have $m + 1 < 10$ and in the second case $m > 10$.

Hence when $n < 10$ the only Nash equilibrium is when all players use the common resource and when $n \geq 10$ then s is a Nash equilibrium when either 9 or 10 players use the common resource.

Assume now that $n \geq 10$. Then in a Nash equilibrium s the players who use the resource receive the payoff 0.2 (when $m = 9$) or 0.1 (when $m = 10$). So the maximum social welfare that can be achieved in a Nash equilibrium is $0.1(n - 9) + 1.8 = 0.1n + 0.9$.

To find a strategy profile in which social optimum is reached with the largest social welfare we need to find m for which the function $0.1(n - m) + F(m)$ reaches the maximum. Now, $0.1(n - m) + F(m) = 0.1n + m - 0.1m^2$ and by elementary calculus we find that $m = 5$ for which $0.1(n-m)+F(m) = 0.1n+2.5$. So the social optimum is achieved when 5 players use the common resource. □

Example 1.6 (Tragedy of the commons II) Assume $n > 1$ players, each having to its disposal an infinite set of strategies that consists of the real interval $[0, 1]$. View player's strategy as its chosen fraction of the common resource. Then the following payoff function reflects the fact that player's enjoyment of the common resource depends positively from his chosen fraction of the resource and negatively from the total fraction of the common resource used by all players:

$$p_i(s) := \begin{cases} s_i(1 - \sum_{j=1}^{n} s_j) & \text{if } \sum_{j=1}^{n} s_j \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

The second alternative reflects the phenomenon that if the total fraction of the common resource by all players exceeds a feasible level, here 1, then player's enjoyment of the resource becomes zero. We can write the payoff function in a more compact way as

$$p_i(s) := \max(0, s_i(1 - \sum_{j=1}^{n} s_j)).$$

To find a Nash equilibrium of this game, fix $i \in \{1, \ldots, n\}$ and s_{-i} and denote $\sum_{j \neq i} s_j$ by t. Then $p_i(s_i, s_{-i}) = \max(0, s_i(1 - t - s_i))$.

By elementary calculus player's i payoff becomes maximal when $s_i = \frac{1-t}{2}$. This implies that when for all $i \in \{1, \ldots, n\}$ we have

$$s_i = \frac{1 - \sum_{j \neq i} s_j}{2},$$

then s is a Nash equilibrium. One can check that this system of n linear equations has a unique solution, $s_i = \frac{1}{n+1}$ for $i \in \{1, \ldots, n\}$ (see Exercise 1.7). In this strategy profile each player's payoff is $\frac{1-n/(n+1)}{n+1} = \frac{1}{(n+1)^2}$, so its social welfare is $\frac{n}{(n+1)^2}$.

There are other Nash equilibria. Indeed, suppose that for all $i \in \{1, \ldots, n\}$ we have $\sum_{j \neq i} s_j \geq 1$, which is the case for instance when $s_i = \frac{1}{n-1}$ for $i \in \{1, \ldots, n\}$. It is straightforward to check that each such strategy profile is a Nash equilibrium in which each player's payoff is 0 and hence the social welfare is also 0. It is easy to check that no other Nash equilibria exist.

To find a strategy profile in which social optimum is reached fix a strategy profile s and let $t := \sum_{j=1}^{n} s_j$. First note that if $t > 1$, then the social welfare is 0. So assume that $t \leq 1$. Then $\sum_{j=1}^{n} p_j(s_j) = t(1-t)$. By elementary calculus this expression becomes maximal precisely when $t = \frac{1}{2}$ and then it equals $\frac{1}{4}$.

Now, for all $n > 1$ we have $\frac{n}{(n+1)^2} < \frac{1}{4}$. So the social welfare of each solution for which $\sum_{j=1}^{n} s_j = \frac{1}{2}$ is superior to the social welfare of the Nash equilibria. In particular, no such strategy profile is a Nash equilibrium.

In conclusion, the social welfare is maximal, and equals $\frac{1}{4}$, when precisely half of the common resource is used. In contrast, in the 'best' Nash equilibrium the social welfare is $\frac{n}{(n+1)^2}$ and the fraction $\frac{n}{n+1}$ of the common resource is used. So when the number of players increases, the social welfare of the best Nash equilibrium becomes arbitrarily small, while the fraction of the common resource being used becomes arbitrarily large. □

The analysis carried out in the above two examples reveals that for the adopted payoff functions the common resource will be overused, to the detriment of the players (society). The same conclusion can be drawn for a much larger of class payoff functions that properly reflect the characteristics of using a common resource.

Example 1.7 (Cournot competition) This example deals with a situation in which n companies independently decide their production levels of a given product. The price of the product is a linear function that depends negatively on the total output.

We model it by means of the following strategic game. We assume that for each player i:

- his strategy set is $\mathbb{R}_+$,

- his payoff function is defined by

$$p_i(s) := s_i\Big(a - b\sum_{j=1}^{n} s_j\Big) - cs_i$$

for some given a, b, c, where $a > c$ and $b > 0$.

Let us explain this payoff function. The price of the product is represented by the expression $a - b\sum_{j=1}^{n} s_j$, which, thanks to the assumption $b > 0$, indeed depends negatively on the total output, $\sum_{j=1}^{n} s_j$. Further, cs_i is the production cost corresponding to the production level s_i. So we assume for simplicity that the production cost functions are the same for all companies.

Further, note that if $a \leq c$, then the payoffs would be always negative or zero, since $p_i(s) = (a - c)s_i - bs_i\sum_{j=1}^{n} s_j$. This explains the assumption that $a > c$. For simplicity we do not exclude that the price of the product is negative. In Exercise 1.9 a modified payoff function is given in which the price is nonnegative. The assumption $c > 0$ is obviously meaningful but not needed.

To find a Nash equilibrium of this game fix $i \in \{1, \ldots, n\}$ and s_{-i} and denote $\sum_{j \neq i} s_j$ by t. Then $p_i(s_i, s_{-i}) = s_i(a - c - bt - bs_i)$. By elementary calculus player i's payoff becomes maximal iff

$$s_i = \frac{a-c}{2b} - \frac{t}{2}.$$

This implies that s is a Nash equilibrium iff for all $i \in \{1, \ldots, n\}$

$$s_i = \frac{a-c}{2b} - \frac{\sum_{j \neq i} s_j}{2}.$$

One can check that this system of n linear equations has a unique solution, $s_i = \frac{a-c}{(n+1)b}$ for $i \in \{1, \ldots, n\}$ (see Exercise 1.8). So this is a unique Nash equilibrium of this game.

Note that for these values of s_i the price of the product is

$$a - b\sum_{j=1}^{n} s_j = a - b\frac{n(a-c)}{(n+1)b} = \frac{a+nc}{n+1}.$$

To find the social optimum let $t := \sum_{j=1}^{n} s_j$. Then $\sum_{j=1}^{n} p_j(s) = t(a - c - bt)$. By elementary calculus this expression becomes maximal precisely when $t = \frac{a-c}{2b}$. So s is a social optimum iff $\sum_{j=1}^{n} s_j = \frac{a-c}{2b}$. The price of the product in a social optimum is $a - b\frac{a-c}{2b} = \frac{a+c}{2}$.

Now, the assumption $a > c$ implies that $\frac{a+c}{2} > \frac{a+nc}{n+1}$. So we see that the price in the social optimum is strictly higher than in the Nash equilibrium. This can be interpreted as a statement that the competition between the producers of the product drives its price down, or alternatively, that a cartel between the producers leads to higher profits for them (i.e., higher social welfare), at the cost of a higher price. So in this example reaching the social optimum is not a desirable state of affairs. The reason is that in our analysis we focussed only on the profits of the producers and omitted the customers.

Further notice that when n, i.e., the number of companies, increases, the price $\frac{a+nc}{n+1}$ in the Nash equilibrium decreases. This corresponds to the intuition that increased competition is beneficial for the customers. Note also that in the limit the price in the Nash equilibrium converges to the production cost c.

Finally, let us compare the social welfare in the unique Nash equilibrium and a social optimum. We just noted that for $t := \sum_{j=1}^{n} s_j$ we have $\sum_{j=1}^{n} p_j(s) = t(a - c - bt)$, and that for the unique Nash equilibrium s we have $s_i = \frac{a-c}{(n+1)b}$ for $i \in \{1, \ldots, n\}$. So $t = \frac{a-c}{b}\frac{n}{n+1}$. Consequently

$$\sum_{j=1}^{n} p_j(s) = \frac{a-c}{b}\frac{n}{n+1}(a - c - (a-c)\frac{n}{n+1})$$
$$= \frac{a-c}{b}\frac{n}{n+1}\frac{1}{n+1}(a-c) = \frac{(a-c)^2}{b}\frac{n}{(n+1)^2}$$

This shows that the social welfare in the unique Nash equilibrium converges to 0 when n, the number of companies, goes to infinity. This can be interpreted as a statement that the increased competition between producers results in their profits becoming arbitrary small.

In contrast, the social welfare in each social optimum remains constant. Indeed, we noted that s is a social optimum iff $t = \frac{a-c}{2b}$ where $t := \sum_{j=1}^{n} s_j$. So for each social optimum s we have

$$\sum_{j=1}^{n} p_j(s) = t(a - c - bt) = \frac{a-c}{2b}(a - c - \frac{a-c}{2}) = \frac{(a-c)^2}{4b}.$$

□

While the last two examples refer to completely different scenarios, their mathematical analysis is very similar. Their common characteristics is that in both examples the payoff functions can be written as $f(s_i, \sum_{j=1}^{n} s_j)$, where f is increasing in the first argument and decreasing in the second argument.

1.3 Exercises

Exercise 1.1 Find all Nash equilibria in the following games:

Stag hunt

	S	R
S	$2,2$	$0,1$
R	$1,0$	$1,1$

Pareto Coordination

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

Hawk-dove

	H	D
H	$0,0$	$3,1$
D	$1,3$	$2,2$

□

Exercise 1.2 A popular game played on the British TV is the Split or Steal contest. In the jackpot there is a large amount of money, for instance 100,000 £. Each of two players has two golden balls. On one it is written 'split' and on the other 'steal'. The players choose simultaneously one of the balls. If they both choose the 'split' ball, they split the jackpot equally. If one of them chooses the 'steal' ball and the other one the 'split' ball, the first player gets all the money while the other one gets nothing. Finally, if both players choose the 'steal' ball, they get nothing.

Define the underlying game. What are its Nash equilibria?

Exercise 1.3 Consider the following ***inspection game***.

There are two players: a worker and the boss. The worker can either Shirk or put an Effort, while the boss can either Inspect or Not. Finding a shirker has a benefit b while the inspection costs c, where $b > c > 0$. So if the boss carries out an inspection his benefit is $b - c > 0$ if the worker shirks and $-c < 0$ otherwise.

The worker receives 0 if he shirks and is inspected, and g if he shirks and is not found. Finally, the worker receives w, where $g > w > 0$ if he puts in the effort.

This leads to the following bimatrix:

	I	N
S	$0, b-c$	$g, 0$
E	$w, -c$	$w, 0$

Analyze the best responses in this game. What can we conclude from it about the Nash equilibria of this game?

□

Exercise 1.4 Consider the following parametrized version of the Traveler's Dilemma game from Example 1.1:

$$p_i(s) := \begin{cases} s_i & \text{if } s_i = s_{-i} \\ s_i + b & \text{if } s_i < s_{-i} \\ s_{-i} - b & \text{otherwise} \end{cases}$$

where $b > 1$ is the bonus. What are the Nash equilibria of this game? □

Exercise 1.5 Express the social welfare in the Prisoner's Dilemma game for n players from Example 1.4 as a function of the number of players who choose to cooperate. □

Exercise 1.6 Prove that in the game discussed in Example 1.6 indeed no other Nash equilibria exist apart of the mentioned ones. □

Exercise 1.7 Show that the set of equations from Example 1.6

$$s_i = \frac{1 - \sum_{j \neq i} s_j}{2},$$

where $i \in \{1, \ldots, n\}$, has a unique solution, $s_i = \frac{1}{n+1}$ for $i \in \{1, \ldots, n\}$. □

Exercise 1.8 Show that the set of equations from Example 1.7

$$s_i = \frac{a-c}{2b} - \frac{\sum_{j \neq i} s_j}{2},$$

where $i \in \{1, \ldots, n\}$, has a unique solution, $s_i = \frac{a-c}{(n+1)b}$ for $i \in \{1, \ldots, n\}$. □

Exercise 1.9 Modify the game from Example 1.7 by considering the following payoff functions:

$$p_i(s) := s_i \max\Big(0, a - b\sum_{j=1}^{n} s_j\Big) - cs_i.$$

Compute the Nash equilibria of this game. □

Exercise 1.10 A joint strategy s is called a ***Pareto efficient outcome*** if for no joint strategy s'

$$\forall i \in \{1, \ldots, n\} : p_i(s') \geq p_i(s) \text{ and } \exists i \in \{1, \ldots, n\}\ p_i(s') > p_i(s).$$

That is, a joint strategy is a Pareto efficient outcome if no joint strategy is both a weakly better outcome for all players and a strictly better outcome for some player. Clearly, if s is a social optimum, then s is Pareto efficient.

Check which joint strategies in the Prisoner's Dilemma game for n players are Pareto efficient. □

1.4 Bibliographic remarks

The concept of a Nash equilibrium was introduced in [42] in the context of mixed strategies that will be introduced in Chapter 8. Therefore our definition of a Nash equilibrium is sometimes called ***pure Nash equilibrium***.

The first interpretation of Prisoner's dilemma game is from [8, page21]. The original interpretation of this game that explains its name is from [32, page 95] who attribute it to A. W. Tucker. The Traveler's dilemma game is from [9, 10]. The beauty contest game is due to [40].

Common resources are discussed in many economic textbooks, e.g., in [33, pages 226-227]). The tragedy of the commons was introduced and extensively discussed in [25]. Example 1.5 is from [23, pages 108-111] while Example 1.6 is from [47, Exercise 63.1] and [60, pages 6-7]. Cournot competition was originally defined in [18].

Chapter 2

Dominance notions

2.1 Strict dominance

Let us return now to our analysis of an arbitrary strategic game $(n, S, \mathbf{p})$. Let s_i, s'_i be strategies of player i.

We say that s_i ***dominates*** s'_i (or equivalently, that s'_i is ***dominated by*** s_i) if

$$\forall s_{-i} \in S_{-i} : p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i}).$$

We also say that s_i is ***dominant*** if it dominates all strategies of player i.

Further, we say that s_i ***strictly dominates*** s'_i (or equivalently, that s'_i is ***strictly dominated by*** s_i) if

$$\forall s_{-i} \in S_{-i} : p_i(s_i, s_{-i}) > p_i(s'_i, s_{-i}).$$

We also say that s_i is ***strictly dominant*** if it strictly dominates all other strategies of player i.

The following observation clarifies our interest in both notions of dominance.

Note 2.1 (Dominance) *Consider a strategic game G and a joint strategy s.*

(i) Suppose that each s_i is a dominant strategy. Then it is a Nash equilibrium of G.

(ii) Suppose that each s_i is a strictly dominant strategy. Then it is a unique Nash equilibrium of G.

Proof. See Exercise 2.1. □

Clearly, a rational player will not choose a strategy that is strictly dominated. As an illustration let us return to the Prisoner's Dilemma. In this game for each player C (cooperate) is a strictly dominated strategy. So the assumption of players' rationality implies that each player will choose strategy D (defect). That is, we can predict that rational players will end up choosing the joint strategy (D, D) in spite of the fact that the social optimum (C, C) yields for each of them a strictly higher payoff.

The same holds in the Prisoner's Dilemma game for n players, where for all players i strategy 1 is strictly dominated by strategy 0, since for all $s_{-i} \in S_{-i}$ we have $p_i(0, s_{-i}) - p_i(1, s_{-i}) = 1$.

So the notion of a strictly dominated strategy clarifies why this game is indeed a dilemma.

The following example clarifies the notions of dominant and strictly dominant strategies and the use of the above Note.

Example 2.2 (Public goods game)
Consider the following game for n players. Every player chooses an amount $s_i \in [0, b]$ that he contributes to a public good, where $b \in \mathbb{R}_+$ is the **budget**. The game designer collects the individual contributions of all players, multiplies their sum by $c > 1$ and distributes the resulting amount evenly among all players. The payoff of player i is thus

$$p_i(s) := b - s_i + \frac{c}{n} \sum_{j=1}^{n} s_j.$$

Choosing 0 is called 'free riding' as then the player contributes nothing but 'reaps' the profit generated by other players. Suppose that $c \leq n$. Then every player does have an incentive to free ride as contributing 0 to the public good is then a dominant strategy. Consequently, by Dominance Note 2.1(i) $\mathbf{0}$ is a Nash equilibrium. (As before we denote by $\mathbf{0}$ the joint strategy in which each player chooses 0.)

Additionally, when $c < n$ then contributing 0 is a strictly dominant strategy, so by the Dominance Note 2.1(ii) $\mathbf{0}$ is a unique Nash equilibrium. Note that $SW(\mathbf{0}) = bn$.

In general $SW(s) = bn + (c-1)\sum_{i \in N} s_i$, so a unique social optimum of this game is $s = \mathbf{b}$ for which $p_i(s) = cb$ for every player i and $SW(s) = cbn$. □

We assumed that each player is rational. So when searching for an outcome that is optimal for all players we can safely remove strategies that are strictly dominated by some other strategy. This can be done in a number of ways. For example, we could remove all or some strictly dominated strategies simultaneously, or start removing them in a *round Robin* fashion starting with, say, player 1. To discuss this matter more rigorously we introduce the notion of a restriction of a game.

Given a game $G := (n, S, \mathbf{p})$ and (possibly empty) sets of strategies $R_1, \ldots, R_n$ such that $R_i \subseteq S_i$ for $i \in \{1, \ldots, n\}$ we say that $(n, R_1 \times \cdots \times R_n, \mathbf{p})$ is a ***restriction*** of G. Here of course we view each p_i as a function on the subset $R_1 \times \cdots \times R_n$ of S.

In what follows, given a restriction R we denote by R_i the set of strategies of player i in R. Further, given two restrictions R and R' of G we write $R' \subseteq R$ when $\forall i \in \{1, \ldots, n\} : R'_i \subseteq R_i$. We now introduce the following notion of reduction between the restrictions R and R' of G:

$$R \rightarrow_S R'$$

when $R \neq R'$, $R' \subseteq R$ and

$$\forall i \in \{1, \ldots, n\} \; \forall s_i \in R_i \setminus R'_i \; \exists s'_i \in R_i : s_i \text{ is strictly dominated in } R \text{ by } s'_i.$$

That is, $R \rightarrow_S R'$ when R' results from R by removing from it some strictly dominated strategies. The subscript $_S$ stands for 'strict dominance'. Later we shall consider analogous reduction relations for other concepts of dominance.

We now clarify whether a one-time elimination of (some) strictly dominated strategies can affect Nash equilibria. We call a game $(n, S, \mathbf{p})$ ***finite*** if S is finite and ***infinite*** otherwise.

Lemma 2.3 (Strict Elimination) *Given a strategic game G consider two restrictions R and R' of G such that $R \rightarrow_S R'$. Then*

(i) if s is a Nash equilibrium of R, then it is a Nash equilibrium of R',

(ii) if G is finite and s is a Nash equilibrium of R', then it is a Nash equilibrium of R.

We shall clarify why in *(ii)* the restriction to finite games is necessary after Example 2.8.

Proof.

(i) For each player the set of his strategies in R' is a subset of the set of his strategies in R. So to prove that s is a Nash equilibrium of R' it suffices to prove that no strategy constituting s is eliminated. Suppose otherwise. Then some s_i is eliminated, so for some $s'_i \in R_i$

$$p_i(s'_i, s''_{-i}) > p_i(s_i, s''_{-i}) \text{ for all } s''_{-i} \in R_{-i}.$$

In particular

$$p_i(s'_i, s_{-i}) > p_i(s_i, s_{-i}),$$

so s is not a Nash equilibrium of R.

(ii) Suppose s is not a Nash equilibrium of R. Then for some $i \in \{1, \ldots, n\}$ strategy s_i is not a best response of player i to s_{-i} in R.

Let $s'_i \in R_i$ be a best response of player i to s_{-i} in R (which exists since R_i is finite). The strategy s'_i is eliminated since s is a Nash equilibrium of R'. So for some $s^*_i \in R_i$

$$p_i(s^*_i, s''_{-i}) > p_i(s'_i, s''_{-i}) \text{ for all } s''_{-i} \in R_{-i}.$$

In particular

$$p_i(s^*_i, s_{-i}) > p_i(s'_i, s_{-i}),$$

which contradicts the choice of s'_i. □

In general an elimination of strictly dominated strategies is not a one step process; it is an iterative procedure. Its use is justified by the assumption of common knowledge of rationality.

Example 2.4 Consider the following game:

	L	M	R
T	$3,0$	$2,1$	$1,0$
C	$2,1$	$1,1$	$1,0$
B	$0,1$	$0,1$	$0,0$

Note that B is strictly dominated by T and R is strictly dominated by M. By eliminating these two strategies we get:

	L	M
T	$3,0$	$2,1$
C	$2,1$	$1,1$

Now C is strictly dominated by T, so we get:

	L	M
T	$3, 0$	$2, 1$

In this game L is strictly dominated by M, so we finally get:

	M
T	$2, 1$

□

This brings us to the following notion, where given a binary relation $\rightarrow$ we denote by $\rightarrow^*$ its ***transitive reflexive closure***, i.e., the outcome of zero or more iterations of it. Consider a strategic game G. Suppose that $G \rightarrow_S^* R$, i.e., R is obtained by an iterated elimination of strictly dominated strategies, in short ***IESDS***, starting with G.

- If for no restriction R' of G, $R \rightarrow_S R'$ holds, we say that R is ***an outcome of IESDS from*** G.
- If R has just one joint strategy, we say that G ***is solved by IESDS***.

The following result then clarifies the relation between the IESDS and Nash equilibrium.

Theorem 2.5 (IESDS) *Suppose that G' is an outcome of IESDS from a strategic game G.*

(i) If s is a Nash equilibrium of G, then it is a Nash equilibrium of G'.

(ii) If G is finite and s is a Nash equilibrium of G', then it is a Nash equilibrium of G.

(iii) If G is finite and solved by IESDS, then the resulting joint strategy is a unique Nash equilibrium.

Proof. By the Strict Elimination Lemma 2.3. □

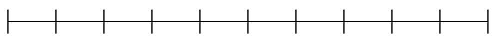

Example 2.6 A nice example of a game that is solved by IESDS is the ***location game***. Assume that that the players are two vendors who simultaneously choose a location. Then the customers choose the closest vendor. The profit for each vendor equals the number of customers it attracts.

To be more specific we assume that the vendors choose a location from the set $\{1, \ldots, n\}$ of natural numbers, viewed as points on a real line, and that at each location there is exactly one customer. For example, for $n = 11$ we have 11 locations:
and when the players choose respectively the locations 3 and 8:

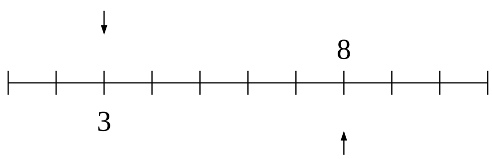

we have $p_1(3, 8) = 5$ and $p_2(3, 8) = 6$. When the vendors 'share' a customer, for instance when they both choose the location 6:

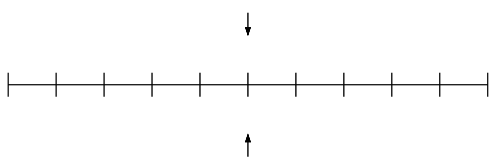

they end up with a fractional payoff, in this case $p_1(6, 6) = 5.5$ and $p_1(6, 6) = 5.5$.

In general, we have the following game:

- each set of strategies consists of the set $\{1, \ldots, n\}$,
- each payoff function p_i is defined by:

$$p_i(s_i, s_{-i}) := \begin{cases} \dfrac{s_i + s_{-i} - 1}{2} & \text{if } s_i < s_{-i} \\ n - \dfrac{s_i + s_{-i} - 1}{2} & \text{if } s_i > s_{-i} \\ \dfrac{n}{2} & \text{if } s_i = s_{-i} \end{cases}$$

It is easy to check that for $n = 2k + 1$ this game is solved by k rounds of IESDS, and that each player is left with the 'middle' strategy k. In each round both 'outer' strategies are eliminated, so first 1 and n, then 2 and $n - 1$, and so on. □

There is one more natural question that we left so far unanswered. Is the outcome of an iterated elimination of strictly dominated strategies unique, or in the game theory parlance: is strict dominance ***order independent***? The answer is positive. The proof can be found in the appendix of this chapter.

Theorem 2.7 (Order Independence I) *Given a finite strategic game all iterated eliminations of strictly dominated strategies yield the same outcome.*

The above result does not hold for infinite strategic games.

Example 2.8 Consider a game in which the set of strategies for each player is the set of natural numbers. The payoff to each player is the number (strategy) he selected.

Note that in this game every strategy is strictly dominated. Consider now three ways of using IESDS:

- by removing in one step all strategies that are strictly dominated,
- by removing in one step all strategies different from 0 that are strictly dominated,
- by removing in each step exactly one strategy, for instance the least even strategy.

In the first case we obtain the restriction with the empty strategy sets, in the second one we end up with the restriction in which each player has just one strategy, 0, and in the third case we obtain an infinite sequence of reductions. □

The above example also shows that in the limit of an infinite sequence of reductions different outcomes can be reached. So for infinite games the definition of the order independence has to be modified.

The above example also shows that in the Strict Elimination Lemma 2.3(*ii*) and the IESDS Theorem 2.5(*ii*) and (*iii*) we cannot drop the assumption that the game is finite. Indeed, the above infinite game has no Nash equilibria, while the game in which each player has exactly one strategy has a Nash equilibrium.

2.2 Weak dominance

Consider now again an arbitrary strategic game $G := (n, S, \mathbf{p})$. Let s_i, s'_i be strategies of player i.

We say that s_i ***weakly dominates*** s'_i (or equivalently, that s'_i is ***weakly dominated by*** s_i) if

$$\forall s_{-i} \in S_{-i} : p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i}) \text{ and } \exists s_{-i} \in S_{-i} : p_i(s_i, s_{-i}) > p_i(s'_i, s_{-i}).$$

Further, we say that s_i is ***weakly dominant*** if it weakly dominates all other strategies of player i.

The following is a special case of the Dominance Note 2.1(i).

Note 2.9 (Weak Dominance) *Consider a strategic game G and a joint strategy s.*

Suppose that each s_i is a weakly dominant strategy. Then it is a Nash equilibrium of G.

Note that we do not claim here that s is a unique Nash equilibrium of G. In fact, such a stronger claim does not hold. Indeed, consider the game

	L	R
T	$1, 1$	$1, 1$
B	$1, 1$	$0, 0$

Here T is a weakly dominant strategy for the player 1, L is a weakly dominant strategy for player 2 and, as prescribed by the above Note, (T, L), is a Nash equilibrium. However, this game has two other Nash equilibria, (T, R) and (B, L).

Analogous considerations to the ones concerning strict dominance can be carried out for the elimination of weakly dominated strategies. To this end we consider the reduction relation $\rightarrow_W$ on the restrictions of G, defined by

$$R \rightarrow_W R'$$

when $R \neq R'$, $R' \subseteq R$ and

$$\forall i \in \{1, \ldots, n\}\ \forall s_i \in R_i \setminus R'_i\ \exists s'_i \in R_i : s_i \text{ is weakly dominated in } R \text{ by } s'_i.$$

We now define the iterated elimination of weakly dominated strategies (in short ***IEWDS***) and the concepts of ***an outcome of IEWDS*** from an initial game G and of a game that ***is solved by IEWDS*** analogously as in the case of IESDS.

However, now some complications arise. To illustrate them consider the following game that results from equipping each player in the Matching Pennies game with a third strategy E (for Edge):

	H	T	E
H	$1, -1$	$-1, \ \ 1$	$-1, -1$
T	$-1, \ \ 1$	$1, -1$	$-1, -1$
E	$-1, -1$	$-1, -1$	$-1, -1$

Note that

- (E, E) is its only Nash equilibrium,
- for each player E is the only strategy that is weakly dominated.

Any form of elimination of these two E strategies, simultaneous or iterated, yields the same outcome, namely the Matching Pennies game, that, as we have already noticed, has no Nash equilibrium. So during this eliminating process we 'lost' the only Nash equilibrium. In other words, part (i) of the IESDS Theorem 2.5 does not hold when reformulated for weak dominance.

On the other hand, some partial results are still valid here. As before we prove first a lemma that clarifies the situation.

Lemma 2.10 (Weak Elimination) *Given a finite strategic game G consider two restrictions R and R' of G such that $R \rightarrow_W R'$. If s is a Nash equilibrium of R', then it is a Nash equilibrium of R.*

Proof. Suppose s is a Nash equilibrium of R' but not a Nash equilibrium of R. Then for some $i \in \{1, \ldots, n\}$ the set

$$A := \{s'_i \in R_i \mid p_i(s'_i, s_{-i}) > p_i(s)\}$$

is non-empty.

Weak dominance is a strict partial ordering (i.e. an irreflexive transitive relation) and A is finite, so some strategy s'_i in A is not weakly dominated in R by any strategy in A. But each strategy in A is eliminated in the reduction

$R \rightarrow_W R'$ since s is a Nash equilibrium of R'. So some strategy $s_i^* \in R_i$ weakly dominates s_i' in R. Consequently

$$p_i(s_i^*, s_{-i}) \geq p_i(s_i', s_{-i})$$

and as a result $s_i^* \in A$. But this contradicts the choice of s_i'. □

This brings us directly to the following result.

Theorem 2.11 (IEWDS) *Suppose that G is a finite strategic game.*

(i) If G' is an outcome of IEWDS from G and s is a Nash equilibrium of G', then s is a Nash equilibrium of G.

(ii) If G is solved by IEWDS, then the resulting joint strategy is a Nash equilibrium of G.

Proof. By the Weak Elimination Lemma 2.10. □

In contrast to the IESDS Theorem 2.5 we cannot claim in part (*ii*) of the IEWDS Theorem 2.11 that the resulting joint strategy is a *unique* Nash equilibrium. Further, in contrast to strict dominance, an iterated elimination of weakly dominated strategies can yield several outcomes.

The following example reveals even more peculiarities of this procedure.

Example 2.12 Consider the following game:

	L	M	R
T	$0,1$	$1,0$	$0,0$
B	$0,0$	$0,0$	$1,0$

It has three Nash equilibria, (T, L), (B, L) and (B, R). This game can be solved by IEWDS but only if in the first round we do not eliminate all weakly dominated strategies, which are M and R. If we eliminate only R, then we reach the game

	L	M
T	$0,1$	$1,0$
B	$0,0$	$0,0$

that is solved by IEWDS by eliminating B and M. This yields

	L
T	$0, 1$

So not only IEWDS is not order independent; in some games it is advantageous *not* to proceed with the deletion of the weakly dominated strategies 'at full speed'. One can also check that the second Nash equilibrium, (B, L), can be found using IEWDS, as well, but not the third one, (B, R). $\Box$

To summarize, the iterated elimination of weakly dominated strategies

- can lead to a deletion of Nash equilibria,
- does not need to yield a unique outcome,
- can be too restrictive if we stipulate that in each round all weakly dominated strategies are eliminated.

Finally, note that the above IEWDS Theorem 2.11 does not hold for infinite games. Indeed, Example 2.8 applies here, as well.

2.3 Never best responses

Iterated elimination of strictly or weakly dominated strategies allow us to solve various games, i.e., reduce them to a single joint strategy that is a Nash equilibrium. However, several simple games cannot be solved using them.

For example, consider the following game:

	X	Y
A	$2, 1$	$0, 0$
B	$0, 1$	$2, 0$
C	$1, 1$	$1, 2$

Here no strategy is strictly or weakly dominated. On the other hand C is a ***never best response***, that is, it is not a best response to any strategy of the opponent. Indeed, A is a unique best response to X and B is a unique best response to Y. Clearly, the above game is solved by an iterated elimination of never best responses. So this procedure can be stronger than IESDS and IEWDS.

Formally, we introduce the following reduction notion between the restrictions R and R' of a given strategic game G:

$$R \rightarrow_N R'$$

when $R \neq R'$, $R' \subseteq R$ and

$$\forall i \in \{1, \ldots, n\} \ \forall s_i \in R_i \setminus R'_i \ \neg \exists s_{-i} \in R_{-i} : s_i \text{ is a best response to } s_{-i} \text{ in } R.$$

That is, $R \rightarrow_N R'$ when R' results from R by removing from it some strategies that are never best responses. Note that in the case of strict and weak dominance a strategy was eliminated because another one was 'better'. Here in contrast there is no 'better' strategy that accounts for a removal of a given strategy.

We now focus on the iterated elimination of never best responses, in short ***IENBR***, obtained using the $\rightarrow_N^*$ relation in analogy to $\rightarrow_S^*$ and IESDS. The following counterpart of the IESDS Theorem 2.5 holds.

Theorem 2.13 (IENBR) *Suppose that G' is an outcome of IENBR from a strategic game G.*

(i) If s is a Nash equilibrium of G, then it is a Nash equilibrium of G'.

(ii) If G is finite and s is a Nash equilibrium of G', then it is a Nash equilibrium of G.

(iii) If G is finite and solved by IENBR, then the resulting joint strategy is a unique Nash equilibrium.

Proof. See Exercise 2.6. □

Further, we have the following analogue of the Hereditarity I Lemma 2.19.

Lemma 2.14 (Hereditarity II) *The relation of never being a best response is hereditary on the set of restrictions of a given finite game.*

Proof. Suppose a strategy $s_i \in R'_i$ is a never best response in R and $R \rightarrow_N R'$. Assume by contradiction that for some $s_{-i} \in R'_{-i}$, s_i is a best response to s_{-i} in R', i.e.,

$$\forall s'_i \in R'_i : p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i}).$$

The initial game is finite, so there exists a best response s'_i to s_{-i} in R. Then s'_i is not eliminated in the step $R \rightarrow_N R'$ and hence is a strategy in R'_i. But s_i is not a best response to s_{-i} in R, so

$$p_i(s'_i, s_{-i}) > p_i(s_i, s_{-i}),$$

so we reached a contradiction. □

This leads us to the following analogue of the Order Independence I Theorem 2.7.

Theorem 2.15 (Order Independence II) *Given a finite strategic game all iterated eliminations of never best responses yield the same outcome.*

Proof. By Theorem 2.17 and the Hereditarity II Lemma 2.14. □

In the case of infinite games we encounter the same problems as in the case of IESDS. Indeed, Example 2.8 readily applies to IENBR, as well, since in this game no strategy is a best response. In particular, this example shows that if we solve an infinite game by IENBR we cannot claim that we obtained a Nash equilibrium. Still, IENBR can be useful in such cases.

Example 2.16 Consider the following infinite variant of the location game considered in Example 2.6. We assume that the players choose their strategies from the open interval $(0, 100)$ and that at each real in $(0, 100)$ there resides one customer. We have then the following payoffs that correspond to the intuition that the customers choose the closest vendor:

$$p_i(s_i, s_{-i}) := \begin{cases} \dfrac{s_i + s_{-i}}{2} & \text{if } s_i < s_{-i} \\ 100 - \dfrac{s_i + s_{-i}}{2} & \text{if } s_i > s_{-i} \\ 50 & \text{if } s_i = s_{-i} \end{cases}$$

In this game each strategy 50 is a best response (namely to strategy 50 of the opponent) and no other strategies are best responses. So this game is solved by IENBR, in one step.

We cannot claim automatically that the resulting joint strategy $(50, 50)$ is a Nash equilibrium, but it is clearly so since each strategy 50 is a best response to the 'other' strategy 50. Moreover, by the IENBR Theorem 2.13(i) we know that this is a unique Nash equilibrium. □

Appendix

We provide here the proof of the Order Independence I Theorem 2.7. Conceptually it is useful to carry out these considerations in a more general setting. We assume an initial strategic game

$$G := (n, S, \mathbf{p}).$$

By a ***dominance relation*** D we mean a function that assigns to each restriction R of G a restriction D_R of R. We say then that the strategy s_i of player i is ***D-dominated*** if $s_i \in D_i$, where $D_R := (n, D_1 \times \cdots \times D_n, \mathbf{p})$. Intuitively, the dominance relation assigns to each restriction R the set $\bigcup_{i=1}^{n} D_i$ of D-dominated strategies.

Given two restrictions R and R' of G we write $R \rightarrow_D R'$ when $R \neq R', R' \subseteq R$ and

$$\forall i \in \{1, \ldots, n\} \; \forall s_i \in R_i \setminus R'_i : s_i \text{ is } D\text{-dominated in } R.$$

Clearly being strictly dominated by another strategy is an example of a dominance relation and $\rightarrow_S$ is an instance of $\rightarrow_D$.

An ***outcome*** of an iteration of $\rightarrow_D$ starting in a game G is a restriction R that can be reached from G using $\rightarrow_D$ in finitely many steps and such that for no R', $R \rightarrow_D R'$ holds.

We call a dominance relation D

- ***order independent*** if for all initial finite games G all iterations of $\rightarrow_D$ starting in G yield the same outcome,

- ***hereditary*** if for all initial games G, all restrictions R and R' such that $R \rightarrow_D R'$ and a strategy s_i in R'

$$s_i \text{ is } D\text{-dominated in } R \text{ implies that } s_i \text{ is } D\text{-dominated in } R'.$$

We now establish the following general result.

Theorem 2.17 *Every hereditary dominance relation is order independent.*

To prove it we introduce the notion of an ***abstract reduction system***. It is simply a pair $(A, \rightarrow)$ where A is a set and $\rightarrow$ is a binary relation on A. Recall that $\rightarrow^*$ denotes the transitive reflexive closure of $\rightarrow$.

- We say that b is a $\to$*-**normal form of*** a if $a \to^* b$ and no c exists such that $b \to c$, and omit the reference to $\to$ if it is clear from the context. If every element of A has a unique normal form, we say that $(A, \to)$ (or just $\to$ if A is clear from the context) satisfies the ***unique normal form property***.

- We say that $\to$ is ***weakly confluent*** if for all $a, b, c \in A$

$$\begin{array}{ccc} & a & \\ \swarrow & & \searrow \\ b & & c \end{array}$$

  implies that for some $d \in A$

$$\begin{array}{ccc} b & & c \\ \searrow^{*} & & {}^{*}\swarrow \\ & d & \end{array}$$

We need the following crucial lemma.

Lemma 2.18 (Newman) *Consider an abstract reduction system* $(A, \to)$ *such that*

- *no infinite* $\to$ *sequences exist,*

- $\to$ *is weakly confluent.*

Then $\to$ *satisfies the unique normal form property.*

Proof. By the first assumption every element of A has a normal form. To prove uniqueness call an element a *ambiguous* if it has at least two different normal forms. We show that for every ambiguous a some ambiguous b exists such that $a \to b$.

So suppose that some element a has two distinct normal forms n_1 and n_2. Then for some b, c we have $a \to b \to^* n_1$ and $a \to c \to^* n_2$. By weak confluence some d exists such that $b \to^* d$ and $c \to^* d$. Let n_3 be a normal form of d. It is also a normal form of b and of c. Moreover $n_3 \neq n_1$ or $n_3 \neq n_2$. If $n_3 \neq n_1$, then b is ambiguous and $a \to b$. And if $n_3 \neq n_2$, then c is ambiguous and $a \to c$.

This means that if an ambiguous element exists, then an infinite $\rightarrow$ sequence exists. We conclude that no ambiguous elements exist. $\Box$

Proof of Theorem 2.17.

Take a hereditary dominance relation D. Consider a restriction R. Suppose that $R \rightarrow_D R'$ for some restriction R'. Let R'' be the restriction of R obtained by removing all strategies that are D-dominated in R.

We have $R'' \subseteq R'$. Assume that $R' \neq R''$. Choose an arbitrary strategy s_i such that $s_i \in R'_i \setminus R''_i$. So s_i is D-dominated in R. By the hereditarity of D, s_i is also D-dominated in R'. This shows that $R' \rightarrow_D R''$.

So we proved that either $R' = R''$ or $R' \rightarrow_D R''$, i.e., that $R' \rightarrow_D^* R''$. This implies that $\rightarrow_D$ is weakly confluent. It suffices now to apply Newman's Lemma 2.18. $\Box$

To apply this result to strict dominance we establish the following fact.

Lemma 2.19 (Hereditarity I) *The relation of being strictly dominated is hereditary on the set of restrictions of a given finite game.*

Proof. Suppose a strategy $s_i \in R'_i$ is strictly dominated in R and $R \rightarrow_S R'$. The initial game is finite, so there exists in R_i a strategy s'_i that strictly dominates s_i in R and is not strictly dominated in R. Then s'_i is not eliminated in the step $R \rightarrow_S R'$ and hence is a strategy in R'_i. But $R' \subseteq R$, so s'_i also strictly dominates s_i in R'. $\Box$

The promised proof is now immediate.

Proof of the Order Independence I Theorem 2.7.

By Theorem 2.17 and the Hereditarity I Lemma 2.19. $\Box$

2.4 Exercises

Exercise 2.1 Prove the Dominance Note 2.1. $\Box$

Exercise 2.2

(i) What is the outcome of IESDS in the location game with an even number of locations?

(ii) Modify the location game from Example 2.6 to a game for three players. Exhibit the Nash equilibria when $n \geq 5$. Prove that no Nash equilibria exist when $n > 5$.

(iii) Define a modification of the above game for three players to the case when the set of possible locations (both for the vendors and the customers) forms all points of a circle. (So the set of strategies is infinite.) Find the set of Nash equilibria. □

Exercise 2.3 Show that the beauty contest game from Example 1.2 is solved by IEWDS. What is the outcome?

This allows us to conclude by the IEWDS Theorem 2.11 that this game has a Nash equilibrium, though not necessarily a unique one. □

Exercise 2.4 Show that in the location game from Example 2.16 no strategy is strictly or weakly dominant. □

Exercise 2.5 Show that the relation of being weakly dominated is not hereditary. □

Exercise 2.6 Prove the IENBR Theorem 2.13. □

2.5 Bibliographic remarks

The location game was introduced in [26]. The Order Independence I Theorem 2.7 was established in [24] and [58]. The proof given here is from [4]. Newman's Lemma is due to [44].

An iterated elimination of strictly dominated strategies in infinite games is considered in [19] and [3], where two different options are proposed and some limited order independence results are established. Example 2.8 has been noted in [19]. Theorem 2.17 and the Order Independence II Theorem 2.15 are from [4].

Chapter 3

Games with a potential

3.1 Best response dynamics

The existence of a Nash equilibrium is clearly a desirable property of a strategic game. In this chapter we discuss some natural classes of games that do have a Nash equilibrium. First, notice the following obvious nondeterministic algorithm, called ***best response dynamics***, to find a Nash equilibrium (NE) in a game $(n, S, \mathbf{p})$:

choose $s \in S$;
while s is not a NE
  choose $i \in \{1, \ldots, n\}$ such that s_i is not a best response to s_{-i};
  $s_i :=$ a best response to s_{-i}

By definition if the best response dynamics terminates, it yields a Nash equilibrium. However, this procedure does not need to terminate if a Nash equilibrium exists. Take for instance the following extension of the Matching Pennies game already considered in Section 2.2:

	H	T	E
H	$1, -1$	$-1, 1$	$-1, -1$
T	$-1, 1$	$1, -1$	$-1, -1$
E	$-1, -1$	$-1, -1$	$-1, -1$

Then an execution of the best response dynamics may end up in a cycle repeatedly going through the sequence

$$((H, H),\ (H, T),\ (T, T),\ (T, H)).$$

However, for various games, for instance the Prisoner's Dilemma game (also for n players) and the Battle of the Sexes game, all executions of the best response dynamics terminate. This is a consequence of a general approach that forms a topic of this chapter. First, note the following simple observation to which we shall return later in the chapter.

Note 3.1 (Best Response Dynamics) *Consider a strategic game for n players. Suppose that every player has a strictly dominant strategy. Then all executions of the best response dynamics terminate after at most n steps and their outcome is unique.*

Proof. Each strictly dominant strategy is a unique best response to each joint strategy of the opponents, so in each execution of the best response dynamics every player can modify his strategy at most once. □

3.2 Potentials

In this section we introduce the main concept of this chapter. Given a game $G := (n, S, \mathbf{p})$ we call the function $P : S \to \mathbb{R}$ a ***exact potential*** (in short, just a ***potential***) for G if

$$\begin{aligned}&\forall i \in \{1, \ldots, n\}\ \forall s_{-i} \in S_{-i}\ \forall s_i, s'_i \in S_i :\\ &p_i(s'_i, s_{-i}) - p_i(s_i, s_{-i}) = P(s'_i, s_{-i}) - P(s_i, s_{-i}).\end{aligned}$$

We call then a game that has a potential a ***potential game***.

The intuition behind the potential is that it tracks the changes in the payoff when some player deviates, without taking into account which one. The following observation explains the interest in the potential games.

Note 3.2 (Potential) *For finite potential games all executions of the best response dynamics terminate. A fortiori all finite potential games have a Nash equilibrium.*

Proof. The domain of the potential function is then finite and at each step of each execution of the best response dynamics the value of the potential strictly increases. □

Consequently, each finite potential game has a Nash equilibrium. This is also a consequence of the fact that by definition each maximum of a potential is a Nash equilibrium.

A number of games that we introduced in the earlier chapters are potential games. Take for instance the Prisoner's Dilemma game for n players from Example 1.4. Indeed, we noted already that in this game we have $p_i(0, s_{-i}) - p_i(1, s_{-i}) = 1$. This shows that $P(s) := n - \sum_{j=1}^{n} s_j$ is a potential function. Intuitively, this potential counts the number of players who selected 0, i.e., the defect strategy.

Also, the Battle of the Sexes is a potential game. We present the game and its potential in Figure 3.1.

	F	B
F	$2,1$	$0,0$
B	$0,0$	$1,2$

	F	B
F	2	1
B	0	2

Figure 3.1: The Battle of the Sexes game (left) and its potential (right)

In the example below we use the following classic result.

Theorem 3.3 (Fundamental Theorem of Calculus) *If $f : [a, b] \to \mathbb{R}$ is a continuous function defined on a real interval $[a, b]$ and F is an antiderivative of f, then*

$$F(b) - F(a) = \int_a^b f(t)\, dt.$$

Example 3.4 It is less trivial to show that the Cournot competition game from Example 1.7 is a potential game. Recall that the set of strategies for each player is $\mathbb{R}_+$ and payoff for each player i is defined by

$$p_i(s) := s_i(a - b\sum_{j=1}^{n} s_j) - cs_i$$

for some given a, b, c, where (these conditions play no role here) $a > c$ and $b > 0$.

We prove that

$$P(s) := a\sum_{i=1}^{n} s_i - b\sum_{i=1}^{n} s_i^2 - b\sum_{1 \le i < j \le n} s_i s_j - \sum_{i=1}^{n} cs_i$$

is a potential.

Applying the Fundamental Theorem of Calculus to the functions p_i and P we get that for all $i \in \{1, \ldots, n\}$, $s_{-i} \in S_{-i}$ and $s_i, s'_i \in S_i$ such that $s'_i > s_i$

$$p_i(s'_i, s_{-i}) - p_i(s_i, s_{-i}) = \int_{s_i}^{s'_i} \frac{\partial p_i}{\partial s_i}(t, s_{-i})\, dt$$

and

$$P(s'_i, s_{-i}) - P(s_i, s_{-i}) = \int_{s_i}^{s_i} \frac{\partial P}{\partial s_i}(t, s_{-i})\, dt.$$

So to prove that P is a potential it suffices to show that for all $i \in \{1, \ldots, n\}$

$$\frac{\partial p_i}{\partial s_i} = \frac{\partial P}{\partial s_i}.$$

But for all $i \in \{1, \ldots, n\}$ and $s \in S$

$$\frac{\partial p_i}{\partial s_i}(s) = \Big(a - b\sum_{j=1}^{n} s_j\Big) - bs_i - c = a - 2bs_i - b \sum_{j\in\{1,\ldots,n\}\setminus\{j\}} s_j - c = \frac{\partial P}{\partial s_i}(s).$$

Note that the fact that Cournot competion is a potential game does not automatically imply that it has a Nash equilibrium. Indeed, the set of strategies is infinite, so the Potential Note 3.2 does not apply. So the proof of existence of a Nash equilibrium given in Example 1.7 is not superfluous. □

The potential tracks the precise changes in the payoff function. We can relax this requirement and only track the sign of the changes of the payoff function. This leads us to the following notion.

Given a game $G := (n, S, \mathbf{p})$ we call the function $P : S \to \mathbb{R}$ an ***ordinal potential*** for G if

$\forall i \in \{1, \ldots, n\}\ \forall s_{-i} \in S_{-i}\ \forall s_i, s'_i \in S_i :$
$p_i(s'_i, s_{-i}) - p_i(s_i, s_{-i}) > 0$ iff $P(s'_i, s_{-i}) - P(s_i, s_{-i}) > 0$.

Clearly, every potential is an ordinal potential. The converse does not hold. Indeed, consider a modification of the Prisoner's Dilemma game and its ordinal potential given in Figure 3.2.

Note that this game has no potential. Indeed every potential has to satisfy the following conditions:

$$\begin{aligned} P(C,C) - P(D,C) &= -1, \\ P(D,C) - P(D,D) &= -2, \\ P(D,D) - P(C,D) &= 1, \\ P(C,D) - P(C,C) &= 1, \end{aligned}$$

	C	D
C	$2,2$	$0,3$
D	$3,0$	$1,2$

	C	D
C	0	1
D	1	2

Figure 3.2: A game (left) and its ordinal potential (right)

which implies $0 = -1$.

An even more general notion is the following one. Given a game $G := (n, S, \mathbf{p})$ we call the function $P : S \to \mathbb{R}$ a ***generalized ordinal potential*** for G if

$$\forall i \in \{1, \ldots, n\}\ \forall s_{-i} \in S_{-i}\ \forall s_i, s_i' \in S_i :$$
$$p_i(s_i', s_{-i}) - p_i(s_i, s_{-i}) > 0 \text{ implies } P(s_i', s_{-i}) - P(s_i, s_{-i}) > 0.$$

Clearly, every ordinal potential is a generalized ordinal potential. The converse does not hold. As an example consider the game and its generalized ordinal potential given in Figure 3.3.

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

	L	R
T	0	3
B	1	2

Figure 3.3: A game (left) and its generalized ordinal potential (right)

It is easy to check that this game has no ordinal potential. Indeed, every ordinal potential has to satisfy

$$P(T, L) < P(B, L) < P(B, R) < P(T, R).$$

But $p_2(T, L) = p_2(T, R)$, so $P(T, L) = P(T, R)$.

We now characterize the finite games that have a generalized ordinal potential. We first introduce the used concepts.

Fix a strategic game $G := (n, S, \mathbf{p})$. By a ***profitable deviation*** we mean a pair (s, s') of joint strategies, written as $s \to s'$, such that for some $i \in \{1, \ldots, n\}$ and $s_i' \in S_i$ we have $s' = (s_i', s_{-i})$ and $p_i(s') > p_i(s)$. We say then that player i can ***profitably deviate*** to s_i' and that s_i' is player's i ***better response*** (not necessarily unique) to the joint strategy s.

An ***improvement path*** is a maximal sequence (i.e., a sequence that cannot be extended) of joint strategies such that each consecutive pair of joint strategies is a profitable deviation. Clearly, if an improvement path is finite, then its last element is a Nash equilibrium. Moreover, if s is a Nash equilibrium, then s is also an improvement path.

We say that G has the ***finite improvement property*** (***FIP*** in short), if every improvement path is finite. Finally, by an ***improvement sequence*** we mean a prefix of an improvement path. Obviously, if G has the FIP, then it has a Nash equilibrium. We can now state the announced result.

Theorem 3.5 (FIP) *A finite game has a generalized ordinal potential iff it has the FIP.*

In the proof below we use the following classic result. We say here that a tree is ***finitely branching*** is each node has only finitely many successors.

Lemma 3.6 (König's Lemma) *Any finitely branching tree is either finite or it has an infinite path.* □

Proof. Consider an infinite, but finitely branching tree T. We construct an infinite path in T, that is, an infinite sequence

$$\xi : n_0 \; n_1 \; n_2 \; \ldots$$

of nodes such that, for each $i \geq 0$, n_{i+1} is a child of n_i. We define ξ inductively such that every n_i is the root of an infinite subtree of T. As n_0 we take the root of T. Suppose now that $n_0, \ldots, n_i$ are already constructed. By induction hypothesis, n_i is the root of an infinite subtree of T. Since T is finitely branching, there are only finitely many children $m_1, \ldots, m_n$ of n_i. At least one of these children is a root of an infinite subtree of T, so we take n_{i+1} to be such a child of n_i. This completes the inductive definition of ξ. □

Proof of the FIP Theorem 3.5.
($\Rightarrow$) Let P be a generalized ordinal potential. Suppose by contradiction that an infinite improvement path exists. Then the corresponding values of P form a strictly increasing infinite sequence. This is a contradiction, since there are only finitely many joint strategies.

($\Leftarrow$) Consider the following branching tree. Its root is some anonymous node, while all other nodes are joint strategies. The successors of the root are all joint strategies. Further, a node s' is a successor of a node s if $s \to s'$. So the branches of this tree are all the improvement paths. Because the game is finite this tree is finitely branching.

The game has the FIP, so this tree has no infinite paths. Consequently by König's Lemma this tree is finite and hence the number of improvement sequences is finite. Given a joint strategy s define $P(s)$ to be the number of improvement sequences that terminate in s. Then in the considered game $(n, S, \mathbf{p})$

$$\forall i \in \{1, \ldots, n\}\ \forall s_{-i} \in S_{-i}\ \forall s_i, s_i' \in S_i :$$
$$p_i(s_i', s_{-i}) - p_i(s_i, s_{-i}) > 0 \text{ implies } P(s_i', s_{-i}) - P(s_i, s_{-i}) = 1,$$

Indeed, the premise implies that each improvement sequence that terminates in (s_i, s_{-i}) can be extended by the profitable deviation $(s_i, s_{-i}) \to P(s_i', s_{-i})$ to form an improvement sequence that terminates in (s_i', s_{-i}). Additionally we have the improvement sequence that consists of just this profitable deviation. So P is a generalized ordinal potential. □

3.3 Congestion games

We now study an important class of games that have a potential. These are games with cost functions, called ***congestion games***, defined as follows. Consider n players and assume a non-empty finite set E of ***facilities***, for example road segments. Given a player i his strategy is a non-empty subset of E, i.e., his set of strategies is the set of non-empty subsets of E. That is, $S_i \subseteq \mathcal{P}(E) \setminus \{\{\emptyset\}\}$, where $\mathcal{P}(E)$ denotes the power set of set E.

We define the cost functions c_i as follows. First, we introduce the ***delay function*** $d_j : \{1, \ldots, n\} \to \mathbb{R}$ for using facility $j \in E$; $d_j(k)$ is the *delay* for using facility j when there are k users of j. Next, we define a function $u_j : S \to \{1, \ldots, n\}$ by

$$u_j(s) := |\{r \in \{1, \ldots, n\} \mid j \in s_r\}|.$$

So $u_j(s)$ is the number of users of facility j given the joint strategy s. Finally, we define the cost function by

$$c_i(s) := \sum_{j \in s_i} d_j(u_j(s)).$$

So $c_i(s)$ is the aggregate delay incurred by player i when each player j selected the set of facilities s_j. Note that we imposed no conditions on the delay functions. In fact, depending on the class of games they can be increasing or decreasing functions.

The following important result clarifies our interest in the congestion games.

Theorem 3.7 (Congestion) *Every congestion game is a potential game.*

Proof. First, we extend each function d_j to $\{0, 1, \ldots, n\}$ by putting $d_j(0) := 0$. We now define

$$P(s) := \sum_{j \in E} \sum_{k=0}^{u_j(s)} d_j(k). \tag{3.1}$$

Intuitively, $P(s)$ is the sum of "accumulated delays" aggregated over all facilities. We prove that P is indeed a potential.

Below χ_A denotes the set characteristic function for the set A, i.e., $\chi_A(j) = 1$ if $j \in A$ and 0 otherwise. We have then for $i \in \{1, \ldots, n\}$

$$c_i(s) = \sum_{j \in E} d_i(u_j(s)) \chi_{s_i}(j). \tag{3.2}$$

From (3.1) and (3.2) it follows that for $i \in \{1, \ldots, n\}$

$$P(s) - c_i(s) = \sum_{j \in E} \sum_{k=0}^{u_j(s) - \chi_{s_i}(j)} d_j(k) \tag{3.3}$$

and

$$P(s'_i, s_{-i}) - c_i(s'_i, s_{-i}) = \sum_{j \in E} \sum_{k=0}^{u_j(s'_i, s_{-i}) - \chi_{s'_i}(j)} d_j(k). \tag{3.4}$$

But for $j \in E$ we have $u_j(s) - \chi_{s_i}(j) = u_j(s'_i, s_{-i}) - \chi_{s'_i}(j)$, so from (3.3) and (3.4) it follows that P is a potential. □

Corollary 3.8 *Every congestion game has a Nash equilibrium.*

Proof. By the Potential Note 3.2. □

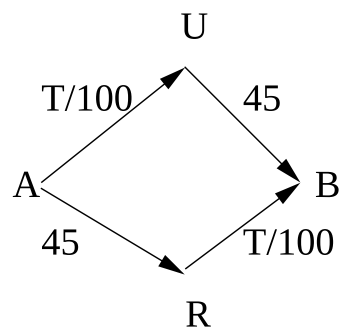

Figure 3.4: A road network

Example 3.9 We now discuss the so-called ***Braess paradox*** showing that adding new roads to a road network can lead to an increased travel time. To discuss it we use the game theoretic concepts of a Nash equilibrium, strictly dominant strategies and a social welfare. Consider the road network given in Figure 3.4.

Each label presents a delay for the corresponding road segment: either it is a constant (here 45) or it is a function (here $T/100$ of the number of the users (T).

Assume that there are 4000 players (drivers), travelling from A to B, where n is even. Each of them has two strategies consisting of a road A - U - B or A - R - B.

It is easy to see that a joint strategy is a Nash equilibrium iff the drivers evenly split among the two possible roads, that is 2000 players choose one strategy and 2000 the other strategy. The resulting cost (travel time) for each player (driver) equals $2000/100 + 45 = 45 + 2000/100 = 65$.

Suppose now that a new, fast, road from U to R is added to the network with delay 0, see Figure 3.5.

Now each player (driver) has three possible strategies (routes): A - U - B, A - R - B, and A - U - R - B. It is easy to see (see Exercise 3.6) that in this new congestion game for each player A - U - R - B is a strictly dominant strategy. Consequently, by the Strict Dominance Note 2.1(i) this new game has a unique Nash equilibrium that consists of all players choosing this strictly dominant strategy. Moreover, by the Best Response Dynamics Note 3.1, all executions of the best response dynamics terminate in this unique Nash equilibrium.

Now, the resulting travel time for each driver equals $4000/100+4000/100 =$

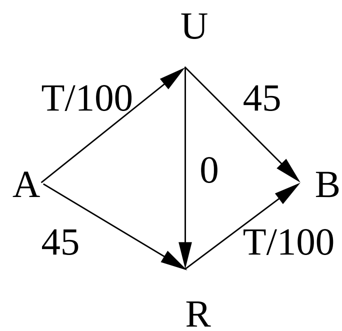

Figure 3.5: An augmented road network

80, so it increased. This shows that adding the new road segment, in this case U - R, can result in a longer travel time. It is easy to check that this paradox remains to hold as long as the delay for using U - R is smaller than 5. □

A special case of congestion games are the ***fair cost sharing games***. In these games each facility $j \in E$ has a cost $c_j \in \mathbb{R}$ associated with it. Then the delay function for a facility is obtained by dividing its cost equally between the users. So we use

$$d_j(u_j(s)) := \frac{c_j}{u_j(s)}$$

in the definition of the congestion game. Consequently

$$c_i(s) := \sum_{j \in s_i} \frac{c_j}{u_j(s)}.$$

In this context the delay function should be viewed as the charge for the use of the facility. Since this charge decreases with the number of the users of the facility, in this class of games the delay functions are decreasing functions. So fair cost sharing games form a class of congestion games in which the costs decrease when the number of users of the shared facilities increases.

3.4 Weakly acyclic games

In the FIP Theorem 3.5 we identified the class of games that have the finite improvement property (FIP). They are obviously of interest, since they have a Nash equilibrium that can be reached from any initial joint strategy by means of an improvement path. However, FIP is a very strong property and for several games only a weaker property holds.

To illustrate this point in the remainder of this section we consider the following class of games. Consider a finite directed graph in which we view each node as a player. Assume that each player has a finite set of strategies that we call *colours*. The payoff to each player is the number of (in)neighbours who chose the same colour.

More precisely, given a directed graph G, let N_j denote the set of all neighbours of node j in G. Then each payoff function is defined by

$$p_i(s) := |\{j \in N_i \mid s_i = s_j\}|.$$

We call such games ***graph games***.

Example 3.10 Consider the directed graph and the colour assignment depicted in Figure 3.6.

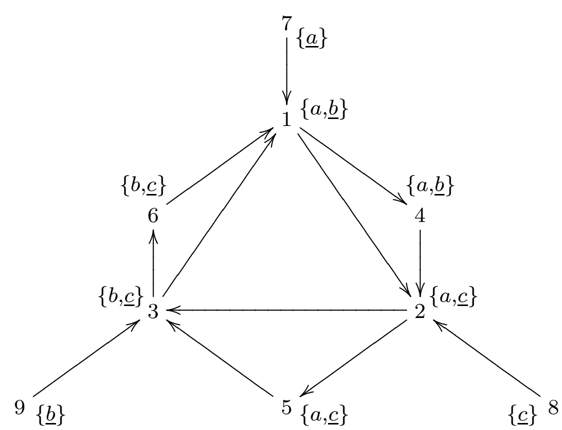

Figure 3.6: A directed graph with a colour assignment.

Take the joint strategy s that consists of the underlined strategies, so

- node 7 selects a,
- nodes 1, 4 and 9 select b,
- nodes 2, 3, 5, 6 and 8 select c.

Then the payoffs are as follows:

- 0 for the nodes 1, 7, 8 and 9,
- 1 for the nodes 2, 4, 5, 6,
- 2 for the node 3.

Note that the above joint strategy is not a Nash equilibrium. For example, node 1 can profitably deviate to colour a. □

Let us focus now on another example.

Example 3.11 Consider a graph game on a simple cycle $1 \to 2 \to \ldots \to n \to 1$, where $n \geq 3$ and such that the nodes share at least two colours, say a and b. Take the initial colouring $(a, b, \ldots, b)$. Then both $(a, \underline{b}, b, \ldots, b) \to (a, a, b, \ldots, b)$ and $(\underline{a}, a, b, \ldots, b) \to (b, a, b, \ldots, b)$ are profitable deviations. We use here the '$\to$' notation introduced in Section 3.2 and to increase readability we underlined the strategies that were modified. After these two steps we obtain a colouring that is a rotation of the first one. Iterating we obtain an infinite improvement path. Hence this game does not have the FIP. □

On the other hand a weaker property holds for the above game. We call a strategic game ***weakly acyclic*** if for any joint strategy there exists a finite improvement path that starts at it. The following result holds.

Theorem 3.12 *Every graph game on a simple cycle is weakly acyclic.*

Proof. As in Example 3.11 we denote the considered cycle by $1 \to 2 \to \ldots \to n \to 1$, where $n \geq 2$.

Given a joint strategy s, let

$$NBR(s) := \{i \in N \mid s_i \text{ is not a best response to } s_{-i}\},$$

where $N = \{1, \ldots, n\}$. Consider a function $f : S \to N \cup \{0\}$, defined as follows:

$$f(s) := \begin{cases} \min(NBR(s)) & \text{if } NBR(s) \neq \emptyset, \\ 0 & \text{otherwise} \end{cases}$$

So at each joint strategy that is not a Nash equilibrium, the function f specifies the player who is supposed to profitably deviate: it is the player with the smallest index who can do so. For each Nash equilibrium the function yields 0.

Consider an improvement path $\rho = s^1, s^2, \ldots$ which satisfies the condition that for all $k \geq 1$, if s^k is not a Nash equilibrium, then $s^k \to s^{k+1}$ is a profitable deviation for the player $f(s^k)$. We argue that ρ is finite, i.e., that it reaches a Nash equilibrium. In the proof below we repeatedly rely on the fact that the payoff of each player depends only on his and his $\to$-predecessor's strategies.

Let l be the maximum that f reaches over the set $\{s^1, s^2, \ldots\}$. Three cases arise.

Case 1. $l = 0$.

Then s^1 is a Nash equilibrium.

Case 2. $0 < l < n$.

Take the smallest m such that $f(s^m) = l$. Then s^{m+1} is well-defined. By the definition of the game all players in $\{1, \ldots, l-1\}$ play a best response in s^m and hence all players in $\{2, \ldots, l\}$ play a best response in s^{m+1}. By assumption $l \neq n$, so player l is not the $\to$-predecessor of player 1. This means that also player 1 plays a best response in s^{m+1}.

Further, by the choice of l and m all players in $\{l+1, \ldots, n\}$ play a best response in s^{m+1}, as well. Hence s^{m+1} is a Nash equilibrium.

Case 3. $l = n$.

Take the smallest m such that $f(s^m) = n$. Then s^{m+1} is well-defined. In s^m exactly $n-1$ players play a best response. By the definition of the game this property continues to hold for all the consecutive strategies of ρ until a Nash equilibrium is reached. So it suffices to show that eventually a player is reached whose strategy modification does not result in his successor losing the property of playing a best response.

Let $s_n^{m+1} = c$. If $c \in \bigcap_{i=1}^n S_i$, then the improvement path ρ terminates when all players switched to c.

Otherwise, let n_0 be the least number in N such that $c \notin S_{n_0}$. Then the improvement path ρ terminates when all players i such that $i \in \{1, \ldots, n_0 - 1\}$, where we identify $1-1$ with n, switched to c. $\Box$

As in the case of the games that have the FIP we can characterize finite weakly acyclic games by means of appropriate potentials. Given a game

$G := (n, S, \mathbf{p})$ we call the function $P : S \to \mathbb{R}$ a ***weak potential*** for G if

> $\forall s$: if s is not a Nash equilibrium, then for some profitable deviation $s \to s'$, $P(s) < P(s')$.

The following is a natural counterpart of Theorem 3.5.

Theorem 3.13 (Weakly Acyclic) *A finite game has a weak potential iff it is weakly acyclic.*

Proof.
$(\Rightarrow)$ Let P be a weak potential. Take a joint strategy s. Suppose that s is not a Nash equilibrium. Then for some profitable deviation $s \to s'$ we have $P(s) < P(s')$. We set s, s' to be the prefix of an improvement path that starts with s. By iterating this process we construct an improvement path. This path cannot be infinite since the corresponding values of P form a strictly increasing sequence and there are only finitely many joint strategies.

$(\Leftarrow)$ Given a joint strategy s, let $P(s)$ be the negated length of the shortest finite improvement path that starts with s. To prove that P is a weak potential consider a joint strategy s that is not a Nash equilibrium. Let $s, s^1, s^2, \ldots, s^k$ be a shortest finite improvement path that starts with s. Then $P(s) = -(k+1)$, $s \to s^1$ is a profitable deviation, and $P(s^1) = -k$. So P is indeed a weak potential. □

In Section 2.3 we considered games that can be solved by IENBR, the iterated elimination of never best responses. We now relate them to the weakly acyclic games.

Theorem 3.14 (Connection) *If a finite game can be solved by IENBR, then it is weakly acyclic.*

Proof. Suppose that a finite game G is solved by IENBR. Let $R^1, \ldots, R^m$ be the corresponding sequence of restrictions, that is, $G = R^1$, $R^i \to_N R^{i+1}$ for $i \in \{1, \ldots, m-1\}$ and R^m has just one joint strategy.

We define the *height* of a strategy s_i from G as the largest $l \in \{1, \ldots, m\}$ such that $s_i \in R_i^l$, where R_i^l is the strategy set of player i in the restriction R^l. For a joint strategy s from G we then define

$$P(s) := \sum_{i=1}^{n} height(s_i).$$

We now prove that P is a weak potential. Suppose that s is not a Nash equilibrium. Take i such that $height(s_i)$ is minimal. By the IENBR Theorem 2.13(iii), s is not in R^m. So $P(s) < m \cdot n$ and consequently $height(s_i) < m$. Suppose $height(s_i) = l$. So $s_i \in R_i^l$ and $s_i \notin R_i^{l+1}$. That is, s_i is a never best response in R^l. In particular, s_i is not a best response to s_{-i} in R^l.

Let s'_i be a best response to s_{-i} in R^l. Put $s' := (s'_i, s_{-i})$. Then $p_i(s) < p_i(s')$, i.e., $s \to s'$ is a profitable deviation. Also $s'_i \in R_i^{l+1}$, so $height(s'_i) > l = height(s_i)$. Hence $P(s) < P(s')$, so by the Weakly Acyclic Theorem 3.13 G is weakly acyclic. □

3.5 Exercises

Exercise 3.1 Prove that two potentials for a given game are the same modulo a constant. More precisely, suppose that P and P' are two potentials for a game G. Prove that these potentials satisfy the following property: there exists a constant c such that for all joint strategies s

$$P(s) - P'(s) = c.$$

□

Exercise 3.2 Find a potential game that has no Nash equilibrium.
Hint. Analyze the game from Example 2.8. □

Exercise 3.3 Suppose that P_1 and P_2 are potentials for some game G. Prove that there exists a constant c such that for every joint strategy s we have $P_1(s) - P_2(s) = c$. □

Exercise 3.4 Prove that

$$P(s) := s_1 s_2 \ldots s_n (a - b \sum_{j=1}^{n} s_j - c)$$

is an ordinal potential for the Cournot competition game introduced in Example 1.7 and analyzed in Example 3.4. □

Exercise 3.5 Generalize the best response dynamics to the ***better response dynamics*** by replacing the qualification 'best' by 'better'. So the executions of the better response dynamics are exactly all improvement paths.

(i) Find a game in which all executions of the best response dynamics terminate while some executions of the better response dynamics don't.

(ii) Find a game in which all executions of the best response dynamics diverge while some executions of the better response dynamics terminate. □

Exercise 3.6 Prove that in Example 3.9 for each player the route A - U - R - B is indeed a strictly dominant strategy. □

Exercise 3.7 What is the minimum number of players for which the conclusions of Example 3.9 hold? □

Exercise 3.8 Prove that the game given in Example 3.10 has no Nash equilibrium. □

Exercise 3.9 Find a weak potential for the game considered in Theorem 3.12. □

3.6 Bibliographic remarks

Section 3.2 is based on [39]. König's Lemma 3.6 is from [29]. The proof of the FIP Theorem 3.5 is from [37]. The Congestion Theorem 3.7 is due to [51]. The presentation of the proof is from [61]. Braess paradox is from [14]. Fair congestion games were introduced in [2].

Weakly acyclic games were introduced in [64] and [37]. In the last paper it is shown that the congestion games in which the strategies are singleton sets of facilities and the delay functions are player specific, are weakly acyclic. The notion of a weak potential and the Weakly Acyclic Theorem 3.13 are from [38]. The Connection Theorem 3.14 is from [6], while the proof is from [38]. Weakly acyclic games were further studied in [34], [21], and [20]. Graph games were studied in [7], where in particular Theorem 3.12 was established.

Chapter 4

Efficiency of equilibria

WRITTEN BY Bart de Keijzer.

It is a natural question to ask, given a game or a class of games, how good the social welfare of the Nash equilibria are in comparison to the the social welfare of a social optimum. We have seen an analysis of this type already in Chapter 1. In the present chapter, we investigate this question further, by introducing the formal notions of ***price of anarchy*** and ***price of stability***, which are widely used to measure the efficiency of equilibria in this way.

4.1 The price of anarchy and price of stability

We defined in Chapter 1 the social welfare as the sum of payoffs of all players, and a social optimum as a joint strategy that maximizes the social welfare. Given a game, we shall refer to the social welfare of a social optimum as the ***optimal social welfare***.

Informally, the ***price of anarchy*** of a game is defined as *the worst-case factor by which the optimal social welfare exceeds the social welfare of a Nash equilibrium*. Formally, given a game G, the ***price of anarchy of*** G, denoted by $PoA(G)$, is given by

$$\mathrm{PoA}(G) := \sup \left\{ \frac{SW(s^*)}{SW(s)} \;\middle|\; s \in \mathrm{NE}(G) \right\},$$

where s^* is a social optimum of G, and $\mathrm{NE}(G)$ is the set of all Nash equilibria of G. We define division by zero and the supremum of the empty set as infinity. For a class of games $\mathcal{G}$, the ***price of anarchy of*** $\mathcal{G}$ is then defined as

$$\text{PoA}(\mathcal{G}) := \sup\{\text{PoA}(G) \mid G \in \mathcal{G}\}, \tag{4.1}$$

i.e., the highest price of anarchy of all games in $\mathcal{G}$. Intuitively, for a given game, we look at the performance of the worst possible Nash equilibrium in the game. So the price of anarchy is essentially a worst case measure.

We can define a best case measure as well, which we obtain by taking the best possible social welfare of the equilibria in the game instead of the worst one. This yields the ***price of stability***, defined by

$$\text{PoS}(G) := \inf\left\{\frac{SW(s^*)}{SW(s)} \;\middle|\; s \in NE(G)\right\},$$

and for a class of games $\mathcal{G}$ defined as the highest price of stability of all games in the class:

$$\text{PoS}(\mathcal{G}) := \sup\{\text{PoS}(G) \mid G \in \mathcal{G}\}. \tag{4.2}$$

Note that these definitions only make sense in case there exists at least one Nash equilibrium. Therefore, prior to analyzing the price of anarchy or the price of stability for a given class of games, one should first check that the games in this class have a Nash equilibrium.

Example 4.1 We start off by reformulating some of the observations of Chapter 1 in terms of price of anarchy and price of stability.

(i) In Example 1.4 we studied a generalization of the Prisoner's Dilemma game to $n > 2$ players. We established for this class of games that

- the optimal social welfare is $2n(n-1)$,
- there is a unique Nash equilibrium,
- and that the social welfare of this unique Nash equilibrium is n.

It follows that both the price of anarchy and the price of stability for this class of games is

$$\frac{2n(n-1)}{n} = 2n - 2.$$

(ii) The public goods game of Example 2.2 has a unique Nash equilibrium if $c < n$. The social welfare of this Nash equilibrium is bn, whereas the optimal social welfare is cbn, so the price of anarchy and price of stability are both $cbn/bn = c$.

Further, if $c > n$, it is a strongly dominant strategy for each player to contribute his full budget b. This results in a unique Nash equilibrium that is also a social optimum. Therefore, in this case both the price of anarchy and the price of stability are 1.

Lastly, if $c = n$, it can be seen that each player is indifferent towards which strategy he plays; every possible contribution that the player can make is a weakly dominant strategy. So by Note 2.9 every joint strategy is a Nash equilibrium. This means that the price of anarchy is c (attained when all players contribute 0), and the price of stability is 1 (attained when all players contribute b).

(iii) For the class of games studied in Example 1.5, we saw that for $n \geq 10$ the optimal social welfare is $0.1n + 2.5$ and that a Nash equilibrium is achieved when either 9 or 10 players use the common resource. When 9 players use the common resource, the social welfare is $0.1n + 0.9$. When 10 players use the common resource, the social welfare is $0.1n$. Therefore, if $n \geq 10$ the price of anarchy is $(0.1n + 2.5)/0.1n = 1 + 25/n$ and the price of stability is $(0.1n + 2.5)/(0.1n + 0.9) = 1 + 16/(n + 9)$.

In case $5 \leq n < 10$ the unique Nash equilibrium consists of all players using the common resource, resulting in a social welfare of $n(1.1 - 0.1n)$. Moreover, the social optimum is attained when 5 players choose the common resource. In case $n < 5$ the unique Nash equilibrium and the social optimum both consist of all players using the common resource.

Therefore, the price of stability and price of anarchy are both 1 when $1 \leq n \leq 5$, and are both $(0.1n + 2.5)/(n(1.1 - 0.1n))$ when $6 \leq n \leq 9$. It can be checked by hand that the price of stability is maximized for $n = 9$ and the price of anarchy is maximized for $n = 10$. It follows that for the complete class of games (i.e., when n is not fixed) the price of anarchy is 3.5 and the price of stability is $3.4/1.8 = 17/9 \approx 1.889$. □

In the remainder of this chapter, we establish bounds on the price of anarchy and the price of stability of two important classes of congestion games. These are games in which the players are assumed to minimize their cost instead of to maximize their profit. Although we defined the price of anarchy only for the latter type of games, these definitions are easily adapted to cost minimization games. Let $G = (n, S, \mathbf{c})$ be such a cost minization game, and let $SC(s) := \sum_{i=1}^{n} c_i(s)$. We refer to SC as the ***social cost*** function of the game. A social optimum is now defined as a joint strategy s^* that minimizes SC. The price of anarchy of G is then defined as

$$\text{PoA}(G) := \sup \left\{ \frac{SC(s)}{SC(s^*)} \;\middle|\; s \in NE(G) \right\}.$$

Similarly, the price of stability of G is defined as

$$\text{PoS}(G) := \inf \left\{ \frac{SC(s)}{SC(s^*)} \;\middle|\; s \in NE(G) \right\}.$$

For *classes* of cost minimization games, the price of anarchy and the price of stability are defined as before, so by (4.1) and (4.2) respectively.

4.2 Affine congestion games

We established in the previous chapter, in Section 3.3, that congestion games have a potential, and therefore, by Note 3.2 a Nash equilibrium is guaranteed to exist in such games. In this section we study the price of anarchy and the price of stability of the subclass of ***affine congestion games***, defined as follows. By an ***affine delay function*** we mean a function $d : \{0, 1, \ldots, n\} \to \mathbb{R}$ for which there exist rational numbers a and b such that $d(0) = 0$ and $d(x) = ax + b$ for all $x \in \{1, \ldots, n\}$. The assumption that a and b are rational (rather than real) is made only to simplify the technical content that follows. Let E be the set of facilities of a given congestion game G. Then G is called ***affine*** iff for every facility $j \in E$ the associated delay function d_j is an affine delay function. The congestion game given in Example 3.9 is an example of an affine congestion game. For every $j \in E$ we let a_j and b_j be the numbers such that $d_j(k) = a_j k + b_j$.

Note that the affine congestion game in that example can be represented as a network, i.e., the facilities and strategy sets of the players correspond to the paths of a graph. In general, there need not be such a network structure, such as in the following example of an affine congestion game with the price of anarchy equal to 5/2.

Example 4.2 Consider a congestion game with 6 facilities and 3 players. Partition the facilities into two sets E_1 and E_2 of three facilities each. Let h_0, h_1 and h_2 be the facilities in E_1, and let e_0, e_1 and e_2 be the facilities in E_2. Each player $i \in \{1, 2, 3\}$ has two strategies in his strategy set: the strategy

$s_i = \{h_{i-1}, e_{i-1}\}$ and the strategy $t_i = \{h_{i \bmod 3}, h_{(i+1) \bmod 3}, e_{i \bmod 3)}\}$. All the delay functions are of the form $k \mapsto k$. Observe that $s = (s_1, s_2, s_3)$ is a social optimum, since each facility is used then by exactly one player, so $\text{SC}(s) = 6$.

Moreover, if the players select the joint strategy $t = (t_1, t_2, t_3)$, then all facilities in E_1 are used by 2 players, and all facilities in E_2 are used by 1 player. Therefore, $\text{SC}(t) = 3 \cdot 2 \cdot 2 + 3 \cdot 1 = 15$, and every player i has a cost $c_i(t) = 5$.

We now check that t is a Nash equilibrium. By symmetry we only need to verify that the cost of player 1 does not decrease if he switches from t_1 to s_1. If he switches, then he shares facility s_0 with players 2 and 3, and facility t_0 with player 3. So his cost becomes 5, i.e., his cost does not change. It follows that t is a Nash equilibrium, and hence the price of anarchy for this game is at least $15/6 = 5/2$. □

This shows that the price of anarchy of the class of affine congestion games is at least 5/2, because the congestion game we just presented belongs to that class. It turns out that this is actually a worst-case example. We prove now that *all* affine congestion games have a price of anarchy of at most 5/2. We shall need the following two lemmas, the proofs of which can be found in the appendix of this chapter.

Lemma 4.3 *Let G be an affine congestion game. There exists an affine congestion game G' such that all delay functions are of the form $d(x) = x$ for $x \in \{0, 1, \ldots, n\}$ and $PoA(G) = PoA(G')$.*

Lemma 4.4 *For all $\alpha, \beta \in \mathbb{N}$ it holds that $\alpha(\beta + 1) \leq (5/3)\alpha^2 + (1/3)\beta^2$.*

Theorem 4.5 *The price of anarchy of the class of affine congestion games is at most* 5/2*.*

Proof. Consider an affine congestion game for n players. Let $E = \{1, \ldots, m\}$ be its set of facilities, where d_j is the delay function of facility $j \in E$. By Lemma 4.3, we can assume that $d_j(k) = k$ for all $j \in E, k \in \{0, 1, \ldots, n\}$. We use the notation $u_j(s) := \{i \in \{1, \ldots, n\} \mid j \in s_i\}$, for $j \in E$ and $s \in S$, as defined in Section 3.3.

Now, let s be a Nash equilibrium and let s^* be a social optimum of this game. Then we can bound $C(s)$ as follows.

$$C(s) = \sum_{i=1}^{n} c_i(s) \leq \sum_{i=1}^{n} c_i(s_i^*, s_{-i}),$$

where the equality follows from the definition of C and the inequality holds by the definition of a Nash equilibrium. In turn, we can upper bound the right hand side as follows:

$$\begin{aligned}\sum_{i=1}^{n} c_i(s_i^*, s_{-i}) &= \sum_{i=1}^{n} \sum_{j \in s_i^*} d_j(u_j(s_i^*, s_{-i})) \\ &= \sum_{i=1}^{n} \sum_{j \in s_i^*} u_j(s_i^*, s_{-i}) \\ &\leq \sum_{i=1}^{n} \sum_{j \in s_i^*} (u_j(s) + 1),\end{aligned}$$

where the equality holds by the assumption about the delay functions and the inequality follows because for all $i \in \{1, \ldots, n\}$ it holds that in (s_i^*, s_{-i}), the number of players on each facility is at most one more than in s, as only a single player deviates. We proceed to bound the right hand side by rewriting it as a single summation over the facilities. Observe that in the last summation for each facility $j \in E$ the term $u_j(s) + 1$ is counted $u_j(s^*)$ times (i.e., once for each player i such that $j \in s_i^*$), so

$$\sum_{i=1}^{n} \sum_{j \in s_i^*} (u_j(s) + 1) \leq \sum_{j \in E} u_j(s^*)(u_j(s) + 1).$$

Next, we apply Lemma 4.4, where we take $u_j(s^*)$ for α and $u_j(s)$ for β.

$$\sum_{j \in E} u_j(s^*)(u_j(s) + 1) \leq \sum_{j \in E} \frac{5}{3} u_j(s^*)^2 + \frac{1}{3} u_j(s)^2 = \frac{5}{3} C(s^*)^2 + \frac{1}{3} C(s)^2.$$

The equality holds because definition, $u_j(s')$ is the cost that facility j causes to each player choosing it in the joint strategy $s' \in S$, and there are $u_j(s')$ such players, so $\sum_{j \in E} u_j(s')^2 = C(s')$. Combining all the above derivations we obtain that $C(s) \leq (5/3)C(s^*) + (1/3)C(s)$. This can be rewritten as $C(s)/C(s^*) \leq 5/2$ which proves the claim, as s is an arbitrary Nash equilibrium. □

From Example 4.2 and Theorem 4.5 we can now derive the exact price of anarchy of affine congestion games.

Corollary 4.6 *The price of anarchy of the class of affine congestion games is* 5/2.

Proof. By Theorem 4.5 all affine congestion games have a price of anarchy of at most 5/2. By Example 4.2 there exists a congestion game with a price of anarchy of at least 5/2. The claim therefore follows by the definition of the price of anarchy of a class of games. □

For the price of stability we use the following lemma that is applicable to all potential games. It essentially tells us that a bound on the value of the potential of a given game automatically yields a bound on the price of stability. We state this lemma for cost minimization games, but there is an obvious analogue of it for payoff maximization games.

Lemma 4.7 *Let* $G = (n, S, \mathbf{c})$ *be a cost minimization game that has a potential* P*. If there exist* $\alpha, \beta \in \mathbb{R}$ *such that for all* $s \in S$

$$\frac{C(s)}{\alpha} \leq P(s) \leq \beta C(s),$$

then the price of stability of G *is at most* $\alpha\beta$*.*

Proof. Let $s \in S$ be a joint strategy that minimizes P. By the definition of a potential s is a Nash equilibrium. Let s^* be a social optimum. Then we have

$$C(s) \leq \alpha P(s) \leq \alpha P(s^*) \leq \alpha\beta C(s^*),$$

where the first and last inequalities follow by our assumed bounds on P and the second inequality holds because s minimizes P. □

Theorem 4.8 *The price of stability of the class of affine congestion games is at most* 2*.*

Proof. Consider an affine congestion game for n players. Let E be its set of facilities, where d_j is the delay function of facility $j \in E$. By Lemma 4.3 we can assume that $d_j(k) = k$ for all $j \in E, k \in \{1, \ldots, n\}$. Recall from the proof of Theorem 3.7 that $P(s) = \sum_{j \in E} \sum_{k=1}^{u_j(s)} k$ is a potential for this game. Let $s \in S$ be a joint strategy. We obtain the following expression for $P(s)$:

$$P(s) = \sum_{j \in E} \sum_{k=1}^{u_j(s)} k = \sum_{j \in E} \frac{1}{2} u_j(s)(u_j(s) + 1) = \frac{1}{2} C(s) + \frac{1}{2} \sum_{j \in E} u_j(s).$$

The latter value lies in between $\frac{1}{2}C(s)$ and $C(s)$. We may therefore apply Lemma 4.7 with $\alpha = 2$ and $\beta = 1$ and conclude that the price of stability of the game is at most 2. □

4.3 Fair cost sharing games

We now study the price of anarchy and the price of stability of another class of congestion games: the fair cost sharing games defined at the end of Chapter 3. Unfortunately, for this class of games the price of anarchy and the price of stability are not constant. Consider the two fair cost sharing games depicted in Figure 4.1 as networks. The arcs represent the facilities and the labels on the arcs are the facility costs.

Example 4.9 In the game G_1 all players have to get from node s to t. There are n players and two facilities. The cost of Facility 1 is $c_1 = 1$ and the cost of Facility 2 is $c_2 = n$. Every player can choose between either facility, i.e., for every $i \in \{1, \dots, n\}$, $S_i = \{\{1\}, \{2\}\}$. If everyone chooses Facility 2, the cost n is shared evenly among all n players, so the cost of every player is 1. If a player deviated to Facility 1, he would still have a cost of 1, so it is a Nash equilibrium when both players choose Facility 2. Its social cost is n. A social optimum arises when everyone instead chooses 1. This shows that for all $n \in \mathbb{N}$, the price of anarchy of the class of n player fair cost sharing games is at least n. □

Example 4.10 Next, consider the fair cost sharing game G_2. This game has n players and $2n+1$ facilities. Each player $i \in \{1, \dots, n\}$ has to get from s_i to t. The costs of the facilities are as follows: for $j \in \{1, \dots, n\}$, $c_j = 1/j$. For $j \in \{n+1, \dots, 2n\}$, $c_j = 0$. Finally, $c_{2n+1} = 1 + \epsilon$. The strategy sets are as follows: for all $i \in \{1, \dots, n\}$, $S_i = \{\{i\}, \{n+1, 2n+1\}\}$.

Consider the joint strategy s in which every player i chooses strategy $\{i\}$. It is easy to see that s is a Nash equilibrium.

We show that s is a *unique* Nash equilibrium of G_2. Consider some other joint strategy s', and let i be the highest-numbered player who does *not* choose strategy $\{i\}$ (i.e., the highest numbered player i choosing strategy $\{n+i, 2n+1\}$. We first compute the cost of i under s'. Since $c_{n+i} = 0$, facility $n+i$ does not induce any cost to player i. Since $c_{2n+1} = 1 + \epsilon$, and there are at most i players who are using facility $2n+1$ under s', it follows

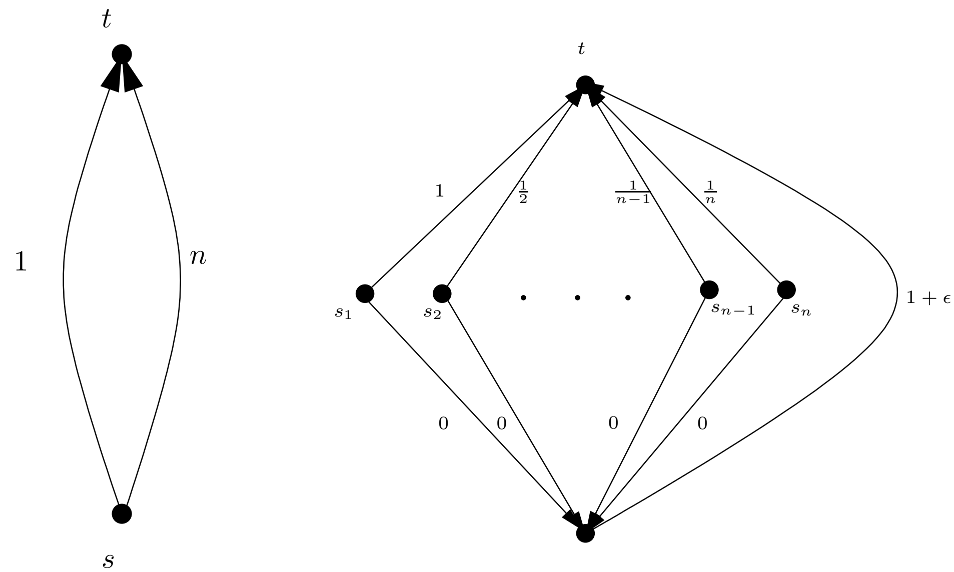

Figure 4.1: The fair cost sharing games G_1 (left) and G_2 (right).

that player i's cost is strictly higher than $1/i$. If i deviated to strategy $\{i\}$, then his cost would improve to $1/i$, so s is not a Nash equilibrium.

We have $C(s) = \sum_{i=1}^{n} c_i = \sum_{i=1}^{n} 1/i$, so $C(s) = H(n)$, the nth *harmonic number*. Recall that $\lim_{n\to\infty} H(n) = \infty$. The joint strategy s^* in which every player i chooses $\{n+i, 2n+1\}$ is a social optimum with $C(s^*) = 1+\epsilon$. This shows that for all $n \in \mathbb{N}$ the price of stability of the class of n player fair cost sharing games is at least $H(n)$. □

It turns out that the lower bounds shown in the above examples are tight.

Theorem 4.11 *Let $n \in \mathbb{N}$. The price of anarchy of the class of n player fair cost sharing games is n. The price of stability of the class of n player fair cost sharing games is $H(n)$.*

Proof. The lower bounds follow from Examples 4.9 and 4.10. It remains to show the upper bounds.

Let G be a fair cost sharing game for n players with facility set E and cost c_j for $j \in E$. Let s be a Nash equilibrium and let s^* be a social optimum of G. Then we derive:

$$\begin{aligned} C(s) &= \sum_{i=1}^{n} c_i(s) \leq \sum_{i=1}^{n} c_i(s_i^*, s_{-i}) \\ &= \sum_{i=1}^{n} \sum_{j \in s_i^*} \frac{c_j}{u_j(s_i^*, s_{-i})} \\ &\leq \sum_{i=1}^{n} \sum_{j \in s_i^*} \frac{nc_j}{u_j(s^*)} \\ &= nC(s^*), \end{aligned}$$

where the first inequality holds because s is a Nash equilibrium, and the second inequality holds because for each term we have $c_j/u_j(s_i^*, s_{-i}) \leq c_j \leq (n/u_j(s^*))c_j$. This proves our claim about the price of anarchy.

For the price of stability, recall from the proof of Theorem 3.7 that $P(s) = \sum_{j \in E} \sum_{k=1}^{u_j(s)} c_j/k = \sum_{j \in E} H(u_j(s))c_j$ is a potential function of G. Because $H(u_j(s)) \geq 1$ for all facilities $j \in s_1 \cup \cdots \cup s_n$, we have $C(s) \leq P(s) \leq H(n)C(s)$ for all joint strategies $s \in S$. Therefore we can apply Lemma 4.7 with $\alpha = 1$ and $\beta = H(n)$ and conclude that the price of stability is at most $H(n)$. □

Appendix

We provide here the proofs of the lemmas that were omitted in the main body of this chapter.

Proof of Lemma 4.3. Let $E = \{1, \ldots, m\}$ be the set of facilities of the considered affine congestion game. For every $j \in E$, the delay function d_j is of the form $d_j(x) = a_j x + b_j$, where $a_j = a_j^n/a_j^d$ and $b_j = b_j^n/b_j^d$ are rational numbers. If we multiply all the coefficients a_j and b_j by $K = a_1^d \cdots a_n^d \cdot b_1^d \cdots b_n^d$, we obtain a congestion game where all a_j and b_j are natural numbers, with the same set of equilibria and the same social optimum, and therefore the same price of anarchy. So we may assume without loss of generality that a_j and b_j are natural numbers.

Next, take a facility $j \in E$. For each player i, introduce b_j new facilities $j^i(1), \ldots, j^i(b_j)$, each with a delay function in the desired form $d(x) = x$. Add the facilities $j^i(1), \ldots, j^i(b_j)$ to all strategies of player i that contain facility j. Also, introduce a_j new facilities $j'(1), \ldots, j'(a_j)$, each with a delay function of the form $d(x) = x$. Add the facilities $j'(1), \ldots, j'(a_j)$ to all strategies that contain facility j. Finally, for each player i, remove the considered facility j from each strategy that contains j.

The above transformation results in a game in which the number of facilities that have delay functions not of the form $d(x) = x$ reduces by 1. By repeatedly applying this transformation step we transform the original congestion game into one in which all delay functions are of the desired form. It therefore suffices to show that the price of anarchy does not change after one single application of this transformation step. Let j be the facility for which the transformation step has been applied. For each player $i \in \{1, \ldots, n\}$ let S_i be the his strategy set *before* applying the transformation and let S'_i be his strategy set *after* applying the transformation. Let $S = S_1 \times \cdots \times S_n$ and let $S' = S'_1 \times \cdots \times S'_n$. There is a natural bijection f between the joint strategies in S and the joint strategies in S' such that for all $i \in \{1, \ldots, n\}$ and $s \in S$,

$$f(s)_i = \begin{cases} s_i & \text{if } j \notin s_i \\ s_i \cup \{j'(1), \ldots, j'(a_j)\} \cup \{j^i(1), \ldots, j^i(b_j)\} \setminus \{j\} & \text{if } j \in s_i. \end{cases}$$

By the definition of this bijection it is clear that $c_i(s) = c_i(f(s))$ if $s_i = f(s)_i$, as the number of players who choose a facility $k \in E \setminus \{j\}$ remains unaffected by f. Also, for any $\ell \in \{1, \ldots, a_j\}$ the number of players that choose $j'(\ell)$ in $f(s)$ equals the number of players that choose j in s. Lastly, a player chooses j in s if and only if he chooses all facilities $j^i(1), \ldots, j^i(b_j)$, which no other player is able to choose. The cost $c_i(f(s))$ of a player i that chooses j in s therefore equals

$$\begin{aligned} c_i(f(s)) = \sum_{k \in f(s)_i} d_k(u_k(f(s))) &= \sum_{k \in s_i \setminus \{j\}} d_k(u_k(s)) + \sum_{\ell=1}^{a_j} u_{j'(\ell)}(f(s)) + \sum_{\ell=1}^{b_j} 1 \\ &= \sum_{k \in s_i \setminus \{j\}} d_k(u_k(s)) + a_j u_{j'(\ell)}(f(s)) + b_j \\ &= \sum_{k \in s_i \setminus \{j\}} d_k(u_k(s)) + a_j u_j(s) + b_j \\ &= \sum_{k \in s_i} d_k(u_k(s)) = c_i(s), \end{aligned}$$

so the costs of all players are preserved under f. Therefore, the social optimum is also preserved and s is a Nash equilibrium if and only if $f(s)$ is a Nash equilibrium. Therefore, the price of anarchy is also preserved under this transformation step. $\square$

Proof of Lemma 4.4. First, we rewrite the inequality we need to prove as

$$\alpha\beta + \alpha - \frac{5}{3}\alpha^2 - \frac{1}{3}\beta^2 \leq 0. \tag{4.3}$$

Suppose that $\alpha \geq \beta$. Then we may write $\alpha = \beta + \gamma$, where $\gamma \geq 0$, and we can write the left hand side of (4.3) as

$$\begin{aligned}
&\beta^2 + \beta\gamma + \beta + \gamma - \frac{5}{3}\beta^2 - \frac{10}{3}\beta\gamma - \gamma^2 - \frac{1}{3}\beta^2 \\
&\quad = \beta - \beta^2 + \gamma - \frac{7}{3}\beta\gamma - \frac{5}{3}\gamma^2 \\
&\quad = \beta(1-\beta) + \gamma\left(1 - \frac{7}{3}\beta - \frac{5}{3}\gamma\right).
\end{aligned}$$

Both of the quantities between the parentheses are non-positive, because $\beta \geq 1$ and $\gamma \geq 0$, which means that the above expression is non-positive as we wanted to show.

Suppose that $\alpha < \beta$. Then we may write $\beta = \alpha + \gamma$, where $\gamma \geq 1$ (recall that $\alpha, \beta \in \mathbb{N}$), and we can write the left hand side of (4.3) as

$$\begin{aligned}
&\alpha^2 + \alpha\gamma + \alpha - \frac{5}{3}\alpha^2 - \frac{1}{3}\alpha^2 - \frac{2}{3}\alpha\gamma - \frac{1}{3}\gamma^2 \\
&\quad = \frac{1}{3}\alpha\gamma - \alpha^2 - \frac{1}{3}\gamma^2 + \alpha \\
&\quad = \alpha(1-\alpha) + \frac{\gamma}{3}(\alpha - \gamma).
\end{aligned}$$

If $\gamma \geq \alpha$, then both quantities between parentheses are again non-positive, so the above expression is non-positive. If $\gamma < \alpha$, we can upper bound the first term by $\gamma(1-\alpha)$, so in that case, the left hand side of (4.3) is at most

$$\gamma\left(1 - \frac{2}{3}\alpha - \frac{1}{3}\gamma\right),$$

which is non-positive, because α and γ are at least 1. $\square$

4.4 Exercises

Exercise 4.1 Compute the price of stability and the price of anarchy of the class of games of Example 1.6, and for the Cournot competition game in Example 1.7. Derive these for n fixed, as well as for the complete classes of games. □

Exercise 4.2 Consider the congestion game represented as a network in Figure 4.2. The delay functions on the arcs/facilities are either constant (4 or 5) or equal to the number of players who choose the arc (denoted by T).

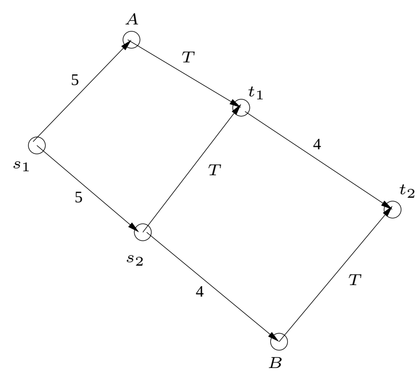

Figure 4.2: A network

There are 6 players who need to choose a path from s_1 to t_1 and 6 players who need to choose a path from s_2 to t_2. So each of the drivers in the first set has two strategies, corresponding respectively to the paths $s_1 \to A \to t_1$ and $s_1 \to s_2 \to t_1$, while each of the drivers in the second set has two strategies, corresponding respectively to the paths $s_2 \to t_1 \to t_2$ and $s_2 \to B \to t_2$.

Consider a joint strategy. Denote by

- T_1 the number of players who chose the path $s_1 \to A \to t_1$,
- T_2 the number of players who chose the path $s_1 \to s_2 \to t_1$,
- T_3 the number of players who chose the path $s_2 \to t_1 \to t_2$,
- T_4 the number of players who chose the path $s_2 \to B \to t_2$.

By assumption we have

$$T_1 + T_2 = 6, \; T_3 + T_4 = 6.$$

(i) Write the conditions on T_1, T_2, T_3, T_4 that determine that a joint strategy is a Nash equilibrium.

(ii) Write the expression in T_1, T_2, T_3, T_4 that defines the social cost of a joint strategy.

(iii) Compute the price of anarchy and the price of stability for this game. □

Exercise 4.3 Consider the congestion game of Exercise 4.2 and suppose now that one adds to the network an arc $t_1 \to B$ with delay 0. The resulting network is drawn in Figure 4.3.

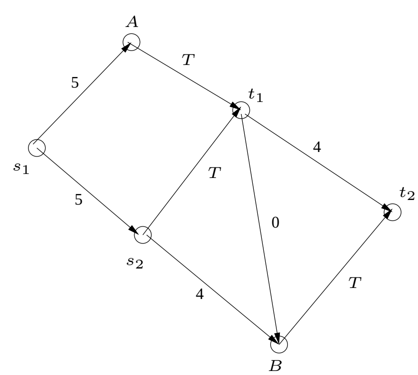

Figure 4.3: The new network

The drivers who need to choose a path from s_2 to t_2 have then three strategies. Given a joint strategy we denote now by

- T_5 the number of playerswho took the road $s_2 \to t_1 \to B \to t_2$,

and define T_1, T_2, T_3 and T_4 as in Exercise 4.2.

(i) Write the conditions on T_1, T_2, T_3, T_4, T_5 that determine that a joint strategy is a Nash equilibrium.

(ii) Write the expression in T_1, T_2, T_3, T_4, T_5 that defines the social cost of a joint strategy.

(iii) Compute the price of anarchy and the price of stability for this game. □

Exercise 4.4 Let $G = (n, S, \mathbf{c})$ be a cost minimization game, and let $s^* \in S$ be a social optimum of G. Suppose that there exist two numbers $\lambda, \mu \in$

$\mathbb{R}, \mu < 1$, such that for all $s \in S$ it holds that

$$\sum_{i=1}^{n} p_i(s_i^*, s_{-i}) \geq \lambda C(s^*) + \mu C(s). \tag{4.4}$$

(i) Prove that the price of anarchy of G is at most $\lambda/(1-\mu)$.

(ii) Prove that for affine congestion games there are two parameters $\lambda, \mu \in \mathbb{R}, \mu < 1$, such that (4.4) holds and $\lambda/(1-\mu) = 2.5$.

Hint: Inspect carefully the proof of Theorem 4.5.

(iii) Formulate an analogous condition that can be used to prove bounds on the price of anarchy of payoff maximization games. □

Exercise 4.5 Consider the following 'War of attrition' game. This game has two players, each strategy set is $\mathbb{R}_+$, and the payoff function p_i for player i defined as follows, where $v > 0$

$$p_i(s) := \begin{cases} -s_i & \text{if } s_i \leq s_{-i} \\ v - s_{-i} & \text{otherwise} \end{cases}$$

The interpretation of this game is as follows. Two players compete for a resource of value $v > 0$. A strategy is the period of time one persists. The player who persists longests wins the resource but incurs the cost equal the duration of the contest. The player who loses incurs the cost equal the period of time he persisted. If both players quit at the same time, then none of them wins the resource and both incur the cost of the duration of the contest. For example, $p_1(2,3) = -2$, $p_2(2,3) = v - 2$, while $p_1(3,3) = -p_2(3,3) = -3$. Compute the price of stability and the price of anarchy for this game. □

Exercise 4.6 Consider the following game for an odd number of $2k+1$ players. Each player has two strategies L (left) and R (right). The players who made the same choice as the majority get the payoff 1, while the others get the payoff 0.

Compute the price of stability and the price of anarchy for this game. □

4.5 Bibliographic Remarks

The price of anarchy was introduced in [30, 31]. The price of stability was defined later. Its first conceptual use was in [56] where the price of stability was analyzed for a variation of congestion games with a continuum of players each of which have an infinitesimally small impact on the delay of a facility, known as the *Wardrop model*. The price of stability was first named as such in [1], which is also the first paper to explicitly study fair cost sharing games. These papers initiated a popular line of research that aims at characterizing the price of anarchy and price of stability of many well-studied classes of games.

One of the first important classes of games for which the price of anarchy was studied is a class of affine routing games in the Wardrop model. Such games are similar to affine congestion games. The main difference is that one assumes a continuum of players, each having an infinitesimally small impact on the delays of the facilities. The price of anarchy was shown to be $4/3$ for this variation of affine congestion games, in [54].

The bound of 2.5 on the price of anarchy of affine congestion games was established in [17]. Contrary to the upper bound of 2.5 for the price of anarchy, the bound of 2 that we gave for the price of stability of affine congestion games is *not* tight. It can be improved further to $1+\sqrt{3}/3 \approx 1.577$, which was done in [16, 15].

Our definition of the social welfare as the sum of all individual players' payoffs is meaningful and standard. However, one is free to choose other objective functions and carry out analysis of the price of anarchy for such alternative objective functions. For example, one could take as the objective function the maximum or minimum payoff among the players, instead of the sum of payoffs. For the class of affine congestion games analysis of the price of anarchy under the latter objective function is carried out in [17]. The choice of the objective function may alternatively be tailored to the context of the studied game. For example, in Chapter 20 of [46], the price of anarchy of the class of *selfish load balancing games* was studied with respect to *makespan* objective function instead of the social cost. In such load balancing games, jobs have to be scheduled on machines. Jobs are considered players who are to choose a machine and the makespan of a joint strategy is the maximum among the loads on the machines.

The condition studied in Exercise 4.4 is known in the literature as (λ, μ)-*smoothness*. It was introduced in [52, 53] and has since been used as a tool for proving upper bounds on the price of anarchy of many classes of games. The smoothness framework has also been extended in various directions.

Analysis of the price of anarchy can also be carried out with respect to other types of equilibria, such as the so-called *mixed Nash equilibria* and *correlated equilibria*, which we shall discuss in Chapters 8 and 9, respectively.

Chapter 5

Sealed-bid auctions

An ***auction*** is a procedure used for selling and buying items by offering them up for bid. Auctions are often used to sell objects that have a variable price (for example oil) or an undetermined price (for example radio frequencies). There are several types of auctions. In its most general form they can involve multiple buyers and multiple sellers with multiple items being offered for sale, possibly in succession. Moreover, some items can be sold in fractions, for example oil.

Here we shall limit our attention to a simple situation in which only one seller exists and offers one object for sale that has to be sold in its entirety (for example a painting). So in this case an auction is a procedure that involves

- one seller who offers an object for sale,
- n bidders, each bidder i having a valuation $v_i \geq 0$ of the object.

The procedure we discuss here involves submission of ***sealed bids***. More precisely, the bidders simultaneously submit their bids in closed envelopes and the object is allocated, in exchange for a payment, to the bidder who submitted the highest bid (the ***winner***). Such an auction is called a ***sealed-bid auction***. We still need to determine what to do in the case of a *tie*. To keep things simple we stipulate that when more than one bidder submitted the highest bid, the object is allocated to the highest bidder with the lowest index.

To formulate a sealed-bid auction as a strategic game we consider each bidder as a player. Then we view each bid of player i as his possible strategy.

We allow any nonnegative real as a bid, that is the set of strategies of each player is $\mathbb{R}_+$.

We assume that the valuations v_i are fixed and known to all players. This is an unrealistic assumption to which we shall return in a later chapter. However, this assumption is necessary, since the valuations are used in the definition of the payoff functions and by assumption the players have common knowledge of the game and hence of each others' payoff functions. When defining the payoff functions we consider two options, each being determined by the underlying payment procedure.

Given a sequence $b := (b_1, \ldots, b_n)$ of reals, we denote the least l such that $b_l = \max_{k \in \{1,\ldots,n\}} b_k$ by $\text{argsmax}\, b$. That is, $\text{argsmax}\, b$ is the smallest index l such that b_l is a largest element in the sequence b. For example, $\text{argsmax}\,(6, 7, 7, 5) = 2$.

5.1 First-price auction

The most commonly used rule in a sealed-bid auction is that the winner i pays to the seller the amount equal to his bid. The resulting mechanism, extended with some rule to resolve the ties, is called the ***first-price auction***.

Assume the winner is bidder i, whose bid is b_i. Since his value for the sold object is v_i, his payoff (profit) is $v_i - b_i$. For the other players the payoff (profit) is 0. Note that the winner's profit can be negative. This happens when he wins the object by ***overbidding***, i.e., submitting a bid higher than his valuation of the object being sold. Such a situation is called the ***winner's curse***.

To summarize, the payoff function p_i of player i in the game associated with the first-price auction is defined as follows, where b is the vector of the submitted bids:

$$p_i(b) := \begin{cases} v_i - b_i & \text{if } i = \text{argsmax}\, b \\ 0 & \text{otherwise} \end{cases}$$

Let us now analyze the resulting game. The following theorem provides a complete characterization of its Nash equilibria.

Theorem 5.1 (Characterization I) *Consider the game associated with the first-price auction with the players' valuations v. Then b is a Nash equilibrium iff for $i = \text{argsmax}\, b$*

(i) $b_i \leq v_i$

(the winner does not suffer from the winner's curse),

(ii) $\max_{j \neq i} v_j \leq b_i$

(the winner submitted a sufficiently high bid),

(iii) $b_i = \max_{j \neq i} b_j$

(another player submitted the same bid as player i).

These three conditions can be compressed into the single statement

$$\max_{j \neq i} v_j \leq \max_{j \neq i} b_j = b_i \leq v_i,$$

where $i = \text{argsmax} b$. Also note that (i) and (ii) imply that $v_i = \max v$, which means that in every Nash equilibrium a player with the highest valuation is the winner.

Proof. $(\Rightarrow)$ (i) If $b_i > v_i$, then player's i payoff is negative and it increases to 0 if he submits the bid equal to v_i.

(ii) If $\max_{j \neq i} v_j > b_i$, then player j such that $v_j > b_i$ can win the object by submitting a bid in the open interval (b_i, v_j), say $v_j - \epsilon$, where $\epsilon \in (0, v_j - b_i)$. Then his payoff increases from 0 to ϵ.

(iii) If $b_i > \max_{j \neq i} b_j$, then player i can increase his payoff by submitting a bid in the open interval $(\max_{j \neq i} b_j, b_i)$, say $b_i - \epsilon$, where $\epsilon \in (0, b_i - \max_{j \neq i} b_j)$. Then his payoff increases from $v_i - b_i$ to $v_i - b_i + \epsilon$.

So if any of the conditions $(i) - (iii)$ is violated, then b is not a Nash equilibrium.

$(\Leftarrow)$ Suppose that a vector of bids b satisfies $(i) - (iii)$. Player i is the winner and by (i) his payoff is non-negative. His payoff can increase only if he bids less, but then by (iii) another player (the one who initially submitted the same bid as player i) becomes the winner, while player i's payoff becomes 0.

The payoff of any other player j is 0 and can increase only if he becomes the winner. This can happen only if he bids at least b_i (if $j < i$) or more than b_i (if $j > i$). But then by (ii), $\max_{j \neq i} v_j \leq b_j$, so his payoff remains 0 or becomes negative.

So b is a Nash equilibrium. □

As an illustration of the above theorem suppose that the vector of the valuations is $(1, 6, 5, 2)$. Then the vectors of bids $(1, 5, 5, 2)$ and $(1, 5, 2, 5)$ satisfy the above three conditions and are both Nash equilibria. The first vector of bids shows that player 2 can secure the object by bidding the second highest valuation. In the second vector of bids player 4 overbids but his payoff is 0 since he is not the winner.

By the ***truthful bidding*** we mean the vector b of bids, such $b = v$, i.e., each player bids his valuation. Note that by the Characterization Theorem 5.1 truthful bidding is a Nash equilibrium iff the two highest valuations coincide.

Further, note that for no player i such that $v_i > 0$ his truthful bidding is a dominant strategy. Indeed, truthful bidding by player i always results in payoff 0. However, if all other players bid 0, then player i can increase his payoff by submitting a lower, positive bid.

Observe that the above analysis does not allow us to conclude that in each Nash equilibrium the winner is the player who wins in the case of truthful bidding. Indeed, suppose that the vector of valuations is $(0, 5, 5, 5)$, so that in the case of truthful bidding by all players player 2 is the winner. Then the vector of bids $(0, 4, 5, 5)$ is a Nash equilibrium with player 3 being the winner.

Finally, notice the following strange consequence of the above theorem: in no Nash equilibrium the last player, n, is a winner. The reason is that we resolved the ties in the favour of a bidder with the lowest index. Indeed, by item (*iii*) in every Nash equilibrium b we have $\text{argsmax}\, b < n$.

5.2 Second-price auction

We consider now an auction with the following payment rule. As before the winner is the bidder who submitted the highest bid (with a tie broken, as before, to the advantage of the bidder with the smallest index), but now he pays to the seller the amount equal to the *second* highest bid. This sealed-bid auction is called the ***second-price auction***. It was proposed by W. Vickrey and is alternatively called the ***Vickrey auction***. So in the absence of ties in this auction the winner pays to the seller a lower price than in the first-price auction.

Let us formalize this auction as a game. The payoffs are now defined as follows:

$$p_i(b) := \begin{cases} v_i - \max_{j \neq i} b_j & \text{if } i = \text{argsmax}\, b \\ 0 & \text{otherwise} \end{cases}$$

Note that bidding v_i always yields a non-negative payoff but can now lead to a strictly positive payoff, which happens when v_i is a unique winning bid. However, when the highest two bids coincide, the payoffs are still the same as in the first-price auction, since then for $i = \text{argsmax} b$ we have $b_i = \max_{j \neq i} b_j$. Finally, note that the winner's curse still can occur here, namely when $v_i < b_i$ and some other bid is in the open interval (v_i, b_i).

The analysis of the second-price auction as a game leads to different conclusions that for the first-price auction. The following theorem provides a complete characterization of the Nash equilibria of the corresponding game.

Theorem 5.2 (Characterization II) *Consider the game associated with the second-price auction with the players' valuations v. Then b is a Nash equilibrium iff for $i = argsmax\, b$*

(i) $\max_{j \neq i} v_j \leq b_i$

(the winner submitted a sufficiently high bid),

(ii) $\max_{j \neq i} b_j \leq v_i$

(the winner's valuation is sufficiently high).

Proof. ($\Rightarrow$) *(i)* If $\max_{j \neq i} v_j > b_i$, then player j such that $v_j > b_i$ can win the object by submitting a bid in the open interval (b_i, v_j). Then his payoff increases from 0 to $v_j - b_i$.

(ii) If $\max_{j \neq i} b_j > v_i$, then player's i payoff is negative, namely $v_i - \max_{j \neq i} b_j$, and can increase to 0 if player i submits a losing bid.

($\Leftarrow$) Suppose that a vector of bids b satisfies *(i)* and *(ii)*. Player i is the winner and by *(ii)* his payoff is non-negative. By submitting another bid he either remains a winner, with the same payoff, or becomes a loser with the payoff 0.

The payoff of any other player j is 0 and can increase only if he becomes the winner. This can happen only if he bids at least b_i (if $j < i$) or more than b_i (if $j > i$). But then his payoff becomes $v_j - b_i$, so by *(i)* it remains 0 or becomes negative.

So b is a Nash equilibrium. □

This characterization result shows that several Nash equilibria exist. We now exhibit three specific ones that are of particular interest. In each case it is straightforward to check that conditions (i) and (ii) of the above theorem hold.

Truthful bidding

Recall from Section 5.1 that in the case of the first-price auction truthful bidding is a Nash equilibrium iff for the considered sequence of valuations the auction coincides with the second-price auction. In contrast, in the second-price auction truthful bidding is always a Nash equilibrium.

Wolf and sheep Nash equilibrium

Suppose that $i = \text{argsmax}\, v$, i.e., player i is the winner in the case of truthful bidding. Consider the strategy profile in which player i bids v_i and everybody else bids 0. This Nash equilibrium is called ***wolf and sheep***, where player i plays the role of a wolf by bidding aggressively and scaring the sheep being the other players who submit their minimal bids.

Yet another Nash equilibrium

Consider now a situation in which a single player, say i, has the highest valuation. Characterization Theorem 5.1 implies that in the first-price auction player i is then the winner in every Nash equilibrium.

This does not need to be the case for the second-price auction. Indeed, suppose that the two highest valuations are v_j and v_i, where $v_j > v_i > 0$ and $i < j$. Then the strategy profile in which player i bids $b_i = v_j$, player j bids $b_j = v_i$ and everybody else bids 0 is a Nash equilibrium.

5.3 Incentive compatibility

So far we discussed two examples of sealed-bid auctions. A general form of such an auction is determined by fixing for each bidder i the payment procedure pay_i. Given a sequence b of bids, pay_i determines the payment to bidder i in case he is the winner.

In the resulting game, that we denote by $G_{pay,v}$, the payoff function is defined by

$$p_i(b) := \begin{cases} v_i - pay_i(b) & \text{if } i = \text{argsmax}\, b \\ 0 & \text{otherwise} \end{cases}$$

Intuitively, bidding 0 means that the bidder is not interested in the object. So if all players bid 0 then none of them is interested in the object. According to our definition the object is then allocated to the first bidder. We assume that then his payment is 0. That is, we additionally stipulate that $pay_1(0, \ldots, 0) = 0$.

When designing a sealed-bid auction it is natural to try to induce the bidders to bid their valuations. This leads to the following notion.

We call a sealed-bid auction with the payment procedures $pay_1, \ldots, pay_n$ ***incentive compatible*** if for all sequences v of players' valuations each v_i is a dominant strategy of player i in the corresponding game $G_{pay,v}$.

While dominance of a strategy does not guarantee that a player will choose it, it ensures that deviating from it is not profitable. So dominance of each valuation v_i can be viewed as a statement that in the considered auction lying does not pay off.

We now show that the condition of incentive compatibility fully characterizes the corresponding auction. More precisely, the following result holds.

Theorem 5.3 (Second-price auction) *A sealed-bid auction is incentive compatible iff it is the second-price auction.*

Proof. Fix a sequence of the payment procedures $pay_1, \ldots, pay_n$ that determines the considered sealed-bid auction.

($\Rightarrow$) Choose an arbitrary sequence of bids that for the clarity of the argument we denote by (v_i, b_{-i}). Suppose that $i = \text{argsmax}\,(v_i, b_{-i})$. We establish the following four claims.

Claim 1. $pay_i(v_i, b_{-i}) \leq v_i$.

Proof. Suppose by contradiction that $pay_i(v_i, b_{-i}) > v_i$. Then in the corresponding game $G_{pay,v}$ we have $p_i(v_i, b_{-i}) < 0$. On the other hand $p_i(0, b_{-i}) \geq 0$. Indeed, if $i \neq \text{argsmax}\,(0, b_{-i})$, then $p_i(0, b_{-i}) = 0$. Otherwise all bids in b_{-i} are 0 and $i = 1$, and hence $p_i(0, b_{-i}) = v_i$, since by assumption $pay_1(0, \ldots, 0) = 0$.

This contradicts the assumption that v_i is a dominant strategy in the corresponding game $G_{pay,v}$. □

Claim 2. For all $b_i \in (\max_{j \neq i} b_j, v_i)$ we have $pay_i(v_i, b_{-i}) \leq pay_i(b_i, b_{-i})$.
Proof. Suppose by contradiction that for some $b_i \in (\max_{j \neq i} b_j, v_i)$ we have $pay_i(v_i, b_{-i}) > pay_i(b_i, b_{-i})$. Then $i = \text{argsmax}\,(b_i, b_{-i})$ so

$$p_i(v_i, b_{-i}) = v_i - pay_i(v_i, b_{-i}) < v_i - pay_i(b_i, b_{-i}) = p_i(b_i, b_{-i}).$$

This contradicts the assumption that v_i is a dominant strategy in the corresponding game $G_{pay,v}$. □

Claim 3. $pay_i(v_i, b_{-i}) \leq \max_{j \neq i} b_j$.
Proof. Suppose by contradiction that $pay_i(v_i, b_{-i}) > \max_{j \neq i} b_j$. Take some $v'_i \in (\max_{j \neq i} b_j, pay_i(v_i, b_{-i}))$. By Claim 1 $v'_i < v_i$, so by Claim 2 $pay_i(v_i, b_{-i}) \leq pay_i(v'_i, b_{-i})$. Further, by Claim 1 for the sequence (v'_i, v_{-i}) of valuations we have $pay_i(v'_i, b_{-i}) \leq v'_i$.

So $pay_i(v_i, b_{-i}) \leq v'_i$, which contradicts the choice of v'_i. □

Claim 4. $pay_i(v_i, b_{-i}) \geq \max_{j \neq i} b_j$.
Proof. Suppose by contradiction that $pay_i(v_i, b_{-i}) < \max_{j \neq i} b_j$. Take an arbitrary $v'_i \in (pay_i(v_i, b_{-i}), \max_{j \neq i} b_j)$. Then $p_i(v'_i, b_{-i}) = 0$, while

$$p_i(v_i, b_{-i}) = v_i - pay_i(v_i, b_{-i}) > v_i - \max_{j \neq i} b_j \geq 0.$$

So $p_i(v_i, b_{-i}) > p_i(v'_i, b_{-i})$. This contradicts the assumption that v'_i is a dominant strategy in the corresponding game $G_{pay,(v'_i, v_{-i})}$. □

So we proved that for $i = \text{argsmax}\,(v_i, b_{-i})$ we have $pay_i(v_i, b_{-i}) = \max_{j \neq i} b_j$, which shows that the considered sealed-bid auction is second-price.

($\Leftarrow$) We actually prove a stronger claim, namely that for all sequences of valuations v, each v_i is a weakly dominant strategy for player i.

To this end take a vector b of bids. By definition $p_i(b_i, b_{-i}) = 0$ or $p_i(b_i, b_{-i}) = v_i - \max_{j \neq i} b_j \leq p_i(v_i, b_{-i})$. But $0 \leq p_i(v_i, b_{-i})$, so

$$p_i(b_i, b_{-i}) \leq p_i(v_i, b_{-i}).$$

Consider now a bid $b_i \neq v_i$. If $b_i < v_i$, then take b_{-i} such that each element of it lies in the open interval (b_i, v_i). Then b_i is a losing bid and v_i is a winning bid and

$$p_i(b_i, b_{-i}) = 0 < v_i - \max_{j \neq i} b_j = p_i(v_i, b_{-i}).$$

If $b_i > v_i$, then take b_{-i} such that each element of it lies in the open interval (v_i, b_i). Then b_i is a winning bid and v_i is a losing bid and

$$p_i(b_i, b_{-i}) = v_i - \max_{j \neq i} b_j < 0 = p_i(v_i, b_{-i}).$$

So we proved that each strategy $b_i \neq v_i$ is weakly dominated by v_i, i.e., that v_i is a weakly dominant strategy. As an aside, recall that each weakly dominant strategy is unique, so we characterized bidding one's valuation in the second-price auction in game theoretic terms. □

5.4 Exercises

Exercise 5.1 Consider the game G associated with the first-price auction with the players' valuations v. Prove that G has no Nash equilibrium iff v_n is the unique highest valuation. □

Exercise 5.2 Prove the counterparts of the Characterization Theorems 5.1 and 5.2 when for each player the set of possible strategies is the set $\mathbb{N} \cup \{0\}$ of natural numbers augmented with zero. □

5.5 Bibliographic remarks

Winner's curse

Wolf and sheep Nash equilibrium

The second-price auction was proposed in [62].

Chapter 6

Regret Minimization and Security Strategies

Until now we implicitly adopted a view that a Nash equilibrium is a desirable outcome of a strategic game. In this chapter we consider two alternative views that help us to understand reasoning of players who either want to avoid costly 'mistakes' or 'fear' a bad outcome. Both concepts can be rigorously formalized.

6.1 Regret minimization

Consider the following game:

	L	R
T	$100, 100$	$0, \quad 0$
B	$0, \quad 0$	$1, \quad 1$

This is an example of a coordination problem, in which there are two satisfactory outcomes (read Nash equilibria), (T, L) and (B, R), of which one is obviously better for both players. In this game no strategy strictly or weakly dominates the other and each strategy is a best response to some other strategy. So using the concepts we introduced so far we cannot explain how come that rational players would end up choosing the Nash equilibrium (T, L). In this section we explain how this choice can be justified using the concept of ***regret minimization***.

With each finite strategic game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$ we first associate a ***regret-recording game*** $(S_1, \ldots, S_n, r_1, \ldots, r_n)$ in which each payoff function r_i is defined by

$$r_i(s_i, s_{-i}) := p_i(s_i^*, s_{-i}) - p_i(s_i, s_{-i}),$$

where s_i^* is player's i best response to s_{-i}. We call then $r_i(s_i, s_{-i})$ player's i ***regret of choosing*** s_i ***against*** s_{-i}. Note that by definition for all s we have $r_i(s) \geq 0$.

For example, for the above game the corresponding regret-recording game is

	L	R
T	$0, \quad 0$	$1, 100$
B	$100, \quad 1$	$0, \quad 0$

Indeed, $r_1(B, L) := p_1(T, L) - p_1(B, L) = 100$, and similarly for the other seven entries.

Let now

$$regret_i(s_i) := \max_{s_{-i} \in S_{-i}} r_i(s_i, s_{-i}).$$

So $regret_i(s_i)$ is the maximal regret player i can have from choosing s_i. We call then any strategy s_i^* for which the function $regret_i$ attains the minimum, i.e., one such that $regret_i(s_i^*) = \min_{s_i \in S_i} regret_i(s_i)$, a ***regret minimization strategy*** for player i.

In other words, s_i^* is a regret minimization strategy for player i if

$$\max_{s_{-i} \in S_{-i}} r_i(s_i^*, s_{-i}) = \min_{s_i \in S_i} \max_{s_{-i} \in S_{-i}} r_i(s_i, s_{-i}).$$

The following intuition is helpful here. Suppose the opponents of player i are able to perfectly anticipate which strategy player i is about to play (for example by being informed through a third party what strategy player i has just selected and is about to play). Suppose further that they aim at inflicting at player i the maximum damage in the form of maximal regret and that player i is aware of these circumstances. Then to miminize his regret player i should select a regret minimization strategy. We could say that a regret minimization strategy will be chosen by a player who wants to avoid making a costly 'mistake', where by a mistake we mean a choice of a strategy that is not a best response to the joint strategy of the opponents.

To clarify this notion let us return to our example of the coordination game. To visualize the outcomes of the functions $regret_1$ and $regret_2$ we put the results in an additional row and column:

	L	R	$regret_1$
T	$0, \quad 0$	$1, 100$	1
B	$100, \quad 1$	$0, \quad 0$	100
$regret_2$	1	100	

So T is the minimum of $regret_1$ and L is the minimum of $regret_2$. Hence (T, L) is the unique pair of regret minimization strategies. This shows that using the concept of regret minimization we succeeded to single out the preferred Nash equilibrium in the considered coordination game.

It is important to note that the concept of regret minimization does not allow us to solve all coordination problems. For example, it does not help us in selecting a Nash equilibrium in symmetric situations, for instance in the game

	L	R
T	$1, 1$	$0, 0$
B	$0, 0$	$1, 1$

Indeed, in this case the regret of each strategy is 1, so regret minimization does not allow us to distinguish between the strategies. Analogous considerations hold for the Battle of Sexes game from Chapter 1.

Regret minimization is based on different intuitions than strict and weak dominance or the notion of a never best response. As a result these notions are incomparable. Further, regret minimization does not necessarily lead to a selection of a Nash equilibrium for the simple reason that some finite games have no Nash equilibria. In general, only the following limited observation holds. Recall that the notion of a dominant strategy was introduced in Exercise 2.1 on page 35.

Note 6.1 (Regret Minimization) *Consider a finite game. Every dominant strategy is a regret minimization strategy.*

Proof. Fix a finite game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$. Note that each dominant strategy s_i of player i is a best response to each $s_{-i} \in S_{-i}$. So by the definition of the regret-recording game for all $s_{-i} \in S_{-i}$ we have $r_i(s_i, s_{-i}) = 0$. Hence

s_i is a regret minimization strategy for player i, since for all joint strategies s we have $r_i(s) \geq 0$. □

The process of removing strategies that do not achieve regret minimization can be iterated. We call this process the ***iterated regret minimization***. The example of the coordination game we analyzed shows that the process of regret minimization may yield to a loss of some Nash equilibria. In fact, as we shall see in a moment, during this process all Nash equilibria can be lost. On the other hand, as recently suggested by J. Halpern and R. Pass, in some games the iterated regret minimization yields a more intuitive outcome. As an example let us return to the Traveler's Dilemma game considered in Example 1.1.

Example 6.2 (Traveler's dilemma revisited)

Let us first determine in this game the regret minimization strategies for each player. Take a joint strategy s.

Case 1. $s_{-i} = 2$.

Then player's i regret of choosing s_i against s_{-i} is 0 if $s_i = s_{-i}$ and 2 if $s_i > s_{-i}$, so it is at most 2.

Case 2. $s_{-i} > 2$.

If $s_{-i} < s_i$, then $p_i(s) = s_{-i} - 2$, while the best response to s_{-i}, namely $s_{-i} - 1$, yields the payoff $s_{-i} + 1$. So player's i regret of choosing s_i against s_{-i} is in this case 3.

If $s_{-i} = s_i$, then $p_i(s) = s_{-i}$, while the best response to s_{-i}, namely $s_{-i} - 1$, yields the payoff $s_{-i} + 1$. So player's i regret of choosing s_i against s_{-i} is in this case 1.

Finally, if $s_{-i} > s_i$, then $p_i(s) = s_i + 2$, while the best response to s_{-i}, namely $s_{-i} - 1$, yields the payoff $s_{-i} + 1$. So player's i regret of choosing s_i against s_{-i} is in this case $s_{-i} - s_i - 1$.

To summarize, we have

$$regret_i(s_i) = \max(3, \max_{s_{-i} \in S_{-i}} s_{-i} - s_i - 1) = \max(3, 99 - s_i).$$

So the minimal regret is achieved when $99 - s_i \leq 3$, i.e., when the strategy s_i is in the interval $[96, 100]$. Hence removing all strategies that do not achieve regret minimization yields a game in which each player has the strategies in the interval $[96, 100]$. In particular, we 'lost' in this way the unique Nash equilibrium of this game, (2,2).

We now repeat this elimination procedure. To compute the outcome we consider again two, though now different, cases.

Case 1. $s_i = 97$.

The following table then summarizes player's i regret of choosing s_i against a strategy s_{-i} of player i:

strategy of player $-i$	best response of player i	regret of player i
96	96	2
97	96	1
98	97	0
99	98	1
100	99	2

Case 2. $s_i \neq 97$.

The following table then summarizes player's i regret of choosing s_i, where for each strategy of player i we list a strategy of player $-i$ for which player's i regret is maximal:

strategy of player i	relevant strategy of player $-i$	regret of player i
96	100	3
98	97	3
99	98	3
100	99	3

So each strategy of player i different from 97 has regret 3, while 97 has regret 2. This means that the second round of elimination of the strategies that do not achieve regret minimization yields a game in which each player has just one strategy, namely 97. □

Recall again that the unique Nash equilibrium in the Traveler's Dilemma game is (2,2). So the iterated regret minimization yields here a radically different outcome than the analysis based on Nash equilibria. Interestingly, this outcome, (97,97), has been confirmed by empirical studies.

Exercise 6.1 Show that regret minimization as a strategy elimination procedure is not order independent.
Hint. Consider the game

	L	R
T	$2,1$	$0,3$
B	$0,2$	$1,1$

□

6.2 Security strategies

Consider the following game:

	L	R
T	$0,\ \ 0$	$101,\ \ 1$
B	$1,101$	$100,100$

This is an extreme form of a ***Chicken game***, sometimes also called a ***Hawk-dove game*** or a ***Snowdrift game***.

This game models two drivers driving towards each other on a narrow road. If neither driver swerves ('chickens'), the result is a crash. The best option for each driver is to stay straight while the other swerves. This yields a situation in which each driver, in attempting to realize his best outcome, risks a crash.

The description of this game as a snowdrift game stresses advantages of a cooperation. The game models two drivers who are trapped on the opposite sides of a snowdrift. Each has the option of staying in the car or shoveling snow to clear a path. Letting the other driver do all the work is the best option, but being exploited by shoveling while the other driver sits in the car is still better than doing nothing.

Note that this game has two Nash equilibria, (T, R) and (B, L). However, there seems to be no reason in selecting any Nash equilibrium as each Nash equilibrium is grossly unfair to the player who will receive only 1.

In contrast, (B, R), which is not a Nash equilibrium, looks like a most reasonable outcome. Each player receives in it a payoff close to the one he receives in the Nash equilibrium of his preference. Also, why should a player risk the payoff 0 in his attempt to secure the payoff 101 that is only a fraction bigger than his payoff 100 in (B, R)?

Note that in this game no strategy strictly or weakly dominates the other and each strategy is a best response to some other strategy. Moreover, the regret minimization for each strategy is 1. So these concepts are useless in analyzing this game.

We now introduce the concept of a ***security strategy*** that allows us to single out the joint strategy (B, R) as the most reasonable outcome for both players.

Fix a, not necessarily finite, strategic game $G := (S_1, \ldots, S_n, p_1, \ldots, p_n)$. Player i, when considering which strategy s_i to select, has to take into account which strategies his opponents will choose. A 'worst case scenario' for player i is that, given his choice of s_i, his opponents choose a joint strategy for which player's i payoff is the lowest[1]. For each strategy s_i of player i once this lowest payoff can be identified a strategy can be selected that leads to a 'minimum damage'.

To formalize this concept for each $i \in \{1, \ldots, n\}$ we consider the function[2] $f_i : S_i \to \mathbb{R}$ defined by

$$f_i(s_i) := \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i}).$$

We call any strategy s_i^* for which the function f_i attains the maximum, i.e., one such that $f_i(s_i^*) = \max_{s_i \in S_i} f_i(s_i)$, a ***security strategy*** or a ***maxminimizer*** for player i. We denote this maximum, so

$$\max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i}),$$

by $maxmin_i$ and call it the ***security payoff*** of player i.

In other words, s_i^* is a security strategy for player i if

$$\min_{s_{-i} \in S_{-i}} p_i(s_i^*, s_{-i}) = maxmin_i.$$

Note that $f_i(s_i)$ is the minimum payoff player i is guaranteed to secure for himself when he selects strategy s_i. In turn, the security payoff $maxmin_i$

[1] We assume here that such s_i exists.

[2] In what follows we assume that all considered minima and maxima always exist. This assumption is obviously satisfied in finite games. In a later chapter we shall discuss a natural class of infinite games for which this assumption is satisfied, as well.

of player i is the minimum payoff he is guaranteed to secure for himself in general. To achieve at least this payoff he just needs to select any security strategy.

The following intuition is helpful here. Suppose the opponents of player i are able to perfectly anticipate which strategy player i is about to play. Suppose further that they aim at inflicting at player i the maximum damage (in the form of the lowest payoff) and that player i is aware of these circumstances. Then player i should select a strategy that causes the minimum damage for him. Such a strategy is exactly a security strategy and it guarantees him at least the $maxmin_i$ payoff. We could say that a security strategy will be chosen by a 'pessimist' player, i.e., one who fears the worst outcome for himself.

To clarify this notion let us return to our example of the chicken game. Clearly, both B and R are the only security strategies in this game. Indeed, we have $f_1(T) = f_2(L) = 0$ and $f_1(B) = f_2(R) = 1$. So we succeeded to single out in this game the outcome (B, R) using the concept of a security strategy.

The following counterpart of the Regret Minimization Note 6.1 holds.

Note 6.3 (Security) *Consider a finite game. Every dominant strategy is a security strategy.*

Proof. Fix a game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$ and suppose that s_i^* is a dominant strategy of player i. For all joint strategies s

$$p_i(s_i^*, s_{-i}) \geq p_i(s_i, s_{-i}),$$

so for all strategies s_i of player i

$$\min_{s_{-i} \in S_{-i}} p_i(s_i^*, s_{-i}) \geq \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i}).$$

Hence

$$\min_{s_{-i} \in S_{-i}} p_i(s_i^*, s_{-i}) \geq \max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i}).$$

This concludes the proof. □

Next, we introduce a dual notion to the security payoff $maxmin_i$. It is not needed for the analysis of security strategies but it will turn out to be relevant in the next two chapters.

With each $i \in \{1, \ldots, n\}$ we consider the function $F_i : S_{-i} \to \mathbb{R}$ defined by

$$F_i(s_{-i}) := \max_{s_i \in S_i} p_i(s_i, s_{-i}).$$

Then we denote the value $\min_{s_{-i} \in S_{-i}} F_i(s_{-i})$, i.e.,

$$\min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} p_i(s_i, s_{-i}),$$

by $minmax_i$.

The following intuition is helpful here. Suppose that now player i is able to perfectly anticipate which strategies his opponents are about to play. Using this information player i can compute the minimum payoff he is guaranteed to achieve in such circumstances: it is $minmax_i$. This lowest payoff for player i can be enforced by his opponents if they choose any joint strategy s^*_{-i} for which the function F_i attains the minimum, i.e., one such that $F_i(s^*_{-i}) = \min_{s_{-i} \in S_{-i}} F_i(s_{-i})$.

To clarify the notions of $maxmin_i$ and $minmax_i$ consider an example.

Example 6.4 Consider the following two-player game:

	L	M	R
T	3, –	4, –	5, –
B	6, –	2, –	1, –

where we omit the payoffs of the second, i.e., column, player.

To visualize the outcomes of the functions f_1 and F_1 we put the results in an additional row and column:

	L	M	R	f_1
T	3, –	4, –	5, –	3
B	6, –	2, –	1, –	1
F_1	6	4	5	

In the f_1 column we list for each row its minimum and in the F_1 row we list for each column its maximum.

Since $f_1(T) = 3$ and $f_1(B) = 1$ we conclude that $maxmin_1 = 3$. So the security payoff of the row player is 3 and T is a unique security strategy of the row player. In other words, the row player can secure for himself at least the payment 3 and achieves this by choosing strategy T.

Next, since $F_1(L) = 6, F_1(M) = 4$ and $F_1(R) = 5$ we get $minmax_1 = 4$. In other words, if the row player knows which strategy the column player is to play, he can secure for himself at least the payment 4. □

In the above example $maxmin_1 < minmax_1$. In general the following observation holds. From now on, to simplify the notation we assume that s_i and s_{-i} range over, respectively, S_i and S_{-i}.

Lemma 6.5 (Lower Bound)

(i) For all $i \in \{1, \ldots, n\}$ we have $maxmin_i \leq minmax_i$.

(ii) If s is a Nash equilibrium of G, then for all $i \in \{1, \ldots, n\}$ we have $minmax_i \leq p_i(s)$.

Given the above intuitions behind the definitions of $maxmin_i$ and $minmax_i$ we can say that item (i) formalizes the intuition that one can take a better decision when more information is available (in this case about which strategies the opponents are about to play). Item (ii) provides a lower bound on the payoff in each Nash equilibrium, which explains the name of the lemma.

Proof.
(i) Fix i. Let s_i^* be such that $\min_{s_{-i}} p_i(s_i^*, s_{-i}) = maxmin_i$ and s_{-i}^* such that $\max_{s_i} p_i(s_i, s_{-i}^*) = minmax_i$. We have then the following string of equalities and inequalities:

$$maxmin_i = min_{s_{-i}} p_i(s_i^*, s_{-i}) \leq p_i(s_i^*, s_{-i}^*) \leq max_{s_i} p_i(s_i, s_{-i}^*) = minmax_i.$$

(ii) Fix i. For each Nash equilibrium (s_i^*, s_{-i}^*) of G we have

$$min_{s_{-i}} max_{s_i} p_i(s_i, s_{-i}) \leq max_{s_i} p_i(s_i, s_{-i}^*) = p_i(s_i^*, s_{-i}^*).$$

□

The concepts of the regret minimization and security strategies bear no relation to each other. Indeed, consider the following variant of a coordination game:

	L	R
T	$100, 100$	$0, 1$
B	$1, 0$	$1, 1$

In this game the regret minimization strategies form one Nash equilibrium, (T, L), while the security strategies form the other Nash equilibrium, (B, R).

Exercise 6.2 Characterize Nash equilibria in the security strategies in the games associated with the first-price and second-price auctions by adding an appropriate condition to the ones given in the Characterization Theorems 5.1 and 5.2. □

Exercise 6.3
(i) Find a two-player game with a Nash equilibrium such that $maxmin_1 < minmax_1$.
(ii) Find a two-player game with no Nash equilibrium such that $maxmin_i = minmax_i$ for $i = 1, 2$. □

This exercise shows that in general there is no relation between the equalities $maxmin_i = minmax_i$, where $i = 1, 2$, and an existence of a Nash equilibrium. In the next chapter we shall discuss a class of two-player games for which these two properties are equivalent.

Chapter 7

Strictly Competitive Games

In this chapter we discuss a special class of two-player games for which stronger results concerning Nash equilibria can be established. To study them we shall crucially rely on the notions introduced in Section 6.2, namely security strategies and $maxmin_i$ and $minmax_i$.

More specifically, we introduce a natural class of two-player games for which the equalities between the $maxmin_i$ and $minmax_i$ values for $i = 1, 2$ constitute a necessary and sufficient condition for the existence of a Nash equilibrium. In these games any Nash equilibrium consists of a pair of security strategies.

A ***strictly competitive game*** is a two-player strategic game (S_1, S_2, p_1, p_2) in which for $i = 1, 2$ and any two joint strategies s and s'

$$p_i(s) \geq p_i(s') \text{ iff } p_{-i}(s) \leq p_{-i}(s').$$

That is, a joint strategy that is better for one player is worse for the other player. This formalizes the intuition that the interests of both players are diametrically opposed and explains the terminology.

By negating both sides of the above equivalence we get

$$p_i(s) < p_i(s') \text{ iff } p_{-i}(s) > p_{-i}(s').$$

So an alternative way of defining a strictly competitive game is by stating that this is a two-player game in which every joint strategy is a Pareto efficient outcome.

To illustrate this concept let us fill in the game considered in Example 6.4 the payoffs for the column player in such a way that the game becomes strictly competitive:

	L	M	R
T	$3,4$	$4,3$	$5,2$
B	$6,0$	$2,5$	$1,6$

Exercise 7.1 Is the Traveler's Dilemma game considered in Example 1.1 strictly competitive? □

Canonic examples of strictly competitive games are ***zero-sum games***. These are two-player games in which for each joint strategy s we have

$$p_1(s) + p_2(s) = 0.$$

So a zero-sum game is an extreme form of a strictly competitive game in which whatever one player 'wins', the other one 'loses'. A simple example is the Matching Pennies game from Chapter 1.

Another well-known zero-sum game is the ***Rock, Paper, Scissors game***. In this game, often played by children, both players simultaneously make a sign with a hand that identifies one of these three objects. If both players make the same sign, the game is a draw. Otherwise one player wins, say, 1 Euro from the other player according to the following rules:

- the rock defeats (breaks) scissors,
- scissors defeat (cut) the paper,
- the paper defeats (wraps) the rock.

Since in a zero-sum game the payoff for the second player is just the negative of the payoff for the first player, each zero-sum game can be represented in a simplified form, called ***reward matrix***. It is simply the matrix that represents only the payoffs for the first player. So the reward matrix for the Rock, Paper, Scissors game looks as follows:

	R	P	S
R	0	-1	1
P	1	0	-1
S	-1	1	0

For the strictly competitive games, so a fortiori the zero-sum games, the following counterpart of the Lower Bound Lemma 6.5 holds.

Lemma 7.1 (Upper Bound) *Consider a strictly competitive game* $G := (S_1, S_2, p_1, p_2)$*. If* (s_1^*, s_2^*) *is a Nash equilibrium of* G*, then for* $i = 1, 2$

(i) $p_i(s_i^*, s_{-i}^*) \leq \min_{s_{-i}} p_i(s_i^*, s_{-i})$*,*

(ii) $p_i(s_i^*, s_{-i}^*) \leq maxmin_i$*.*

Both items provide an upper bound on the payoff in each Nash equilibrium, which explains the name of the lemma.

Proof.

(i) Fix i. Suppose that (s_i^*, s_{-i}^*) is a Nash equilibrium of G. Fix s_{-i}. By the definition of Nash equilibrium

$$p_{-i}(s_i^*, s_{-i}^*) \geq p_{-i}(s_i^*, s_{-i}),$$

so, since G is strictly competitive,

$$p_i(s_i^*, s_{-i}^*) \leq p_i(s_i^*, s_{-i}).$$

But s_{-i} was arbitrary, so

$$p_i(s_i^*, s_{-i}^*) \leq \min_{s_{-i}} p_i(s_i^*, s_{-i}).$$

(ii) By definition

$$\min_{s_{-i}} p_i(s_i^*, s_{-i}) \leq \max_{s_i} \min_{s_{-i}} p_i(s_i, s_{-i}),$$

so by (i)

$$p_i(s_i^*, s_{-i}^*) \leq \max_{s_i} \min_{s_{-i}} p_i(s_i, s_{-i}).$$

□

Combining the Lower Bound Lemma 6.5 and the Upper Bound Lemma 7.1 we can draw the following conclusions about strictly competitive games.

Theorem 7.2 (Strictly Competitive Games) *Consider a strictly competitive game* G*.*

(i) If for $i = 1, 2$ *we have* $maxmin_i = minmax_i$*, then* G *has a Nash equilibrium.*

(ii) *If* G *has a Nash equilibrium, then for* $i = 1, 2$ *we have* $maxmin_i = minmax_i$*.*

(iii) *All Nash equilibria of* G *yield the same payoff, namely* $maxmin_i$ *for player* i*.*

(iv) *All Nash equilibria of* G *are of the form* (s_1^*, s_2^*) *where each* s_i^* *is a security strategy for player* i*.*

Proof. Suppose $G = (S_1, S_2, p_1, p_2)$.
(i) Fix i. Let s_i^* be a security strategy for player i, i.e., such that $\min_{s_{-i}} p_i(s_i^*, s_{-i}) = maxmin_i$, and let s_{-i}^* be such that $\max_{s_i} p_i(s_i, s_{-i}^*) = minmax_i$. We show that (s_i^*, s_{-i}^*) is a Nash equilibrium of G.

We already noted in the proof of the Lower Bound Lemma 6.5(i) that

$$maxmin_i = \min_{s_{-i}} p_i(s_i^*, s_{-i}) \le p_i(s_i^*, s_{-i}^*) \le \max_{s_i} p_i(s_i, s_{-i}^*) = minmax_i.$$

But now $maxmin_i = minmax_i$, so all these values are equal. In particular

$$p_i(s_i^*, s_{-i}^*) = \max_{s_i} p_i(s_i, s_{-i}^*) \tag{7.1}$$

and

$$p_i(s_i^*, s_{-i}^*) = \min_{s_{-i}} p_i(s_i^*, s_{-i}).$$

Fix now s_{-i}. By the last equality

$$p_i(s_i^*, s_{-i}^*) \le p_i(s_i^*, s_{-i}),$$

so, since G is strictly competitive,

$$p_{-i}(s_i^*, s_{-i}^*) \ge p_{-i}(s_i^*, s_{-i}).$$

But s_{-i} was arbitrary, so

$$p_{-i}(s_i^*, s_{-i}^*) = \max_{s_{-i}} p_{-i}(s_i^*, s_{-i}). \tag{7.2}$$

Now (7.1) and (7.2) mean that indeed (s_i^*, s_{-i}^*) is a Nash equilibrium of G.

(*ii*) and (*iii*) If s is a Nash equilibrium of G, by the Lower Bound Lemma 6.5(*i*) and (*ii*) and the Upper Bound Lemma 7.1(*ii*) we have for $i = 1, 2$

$$maxmin_i \leq minmax_i \leq p_i(s) \leq maxmin_i.$$

So all these values are equal.

(*iv*) Fix i. Take a Nash equilibrium (s_i^*, s_{-i}^*) of G. We always have

$$\min_{s_{-i}} p_i(s_i^*, s_{-i}) \leq p_i(s_i^*, s_{-i}^*)$$

and by the Upper Bound Lemma 7.1(*i*) we also have

$$p_i(s_i^*, s_{-i}^*) \leq \min_{s_{-i}} p_i(s_i^*, s_{-i}).$$

So

$$\min_{s_{-i}} p_i(s_i^*, s_{-i}) = p_i(s_i^*, s_{-i}^*) = maxmin_i,$$

where the last equality holds by (*iii*). So s_i^* is a security strategy for player i. □

Combining (*i*) and (*ii*) we see that a strictly competitive game has a Nash equilibrium iff for $i = 1, 2$ we have $maxmin_i = minmax_i$. So in a strictly competitive game each player can determine whether a Nash equilibrium exists without knowing the payoff of the other player. All what he needs to know is that the game is strictly competitive. Indeed, each player i then just needs to check whether his $maxmin_i$ and $minmax_i$ values are equal.

Morever, by (*iv*), each player can select on his own a strategy that forms a part of a Nash equilibrium: it is simply any of his security strategies.

There is another characterization of Nash equilibria in strictly competitive games in terms of the following notion. Given a function $f : \mathbb{R}^2 \to \mathbb{R}$ we call a pair $(x^*, y^*) \in \mathbb{R}^2$ a ***saddle point*** of f if

$$\forall x \forall y \; f(x, y^*) \leq f(x^*, y^*) \leq f(x^*, y).$$

Note 7.3 *Consider a strictly competitive game G. Then s is a Nash equilibrium iff it is a saddle point of any of the payoff functions.*

Proof. Suppose $G = (S_1, S_2, p_1, p_2)$. Fix i. By the definition of a strictly competitive game s is a Nash equilibrium iff

$$\forall s_i' \in S_i \; p_i(s_i, s_{-i}) \geq p_i(s_i', s_{-i})$$

and

$$\forall s'_{-i} \in S_{-i} \; p_i(s_i, s_{-i}) \leq p_i(s_i, s'_{-i}).$$

But the last two inequalities simply state that s is a saddle point of p_i. □

7.1 Zero-sum games

Let us focus now on the special case of zero-sum games. We first show that for zero-sum games the $maxmin_i$ and $minmax_i$ values for one player can be directly computed from the corresponding values for the other player.

Theorem 7.4 (Zero-sum) *Consider a zero-sum game* (S_1, S_2, p_1, p_2). *For* $i = 1, 2$ *we have*

$$maxmin_i = -minmax_{-i}$$

and

$$minmax_i = -maxmin_{-i}.$$

Proof. Fix i. For each joint strategy (s_i, s_{-i})

$$p_i(s_i, s_{-i}) = -p_{-i}(s_i, s_{-i}),$$

so

$$\max_{s_i} \min_{s_{-i}} p_i(s_i, s_{-i}) = \max_{s_i}(\min_{s_{-i}} -p_{-i}(s_i, s_{-i})) = -\min_{s_i} \max_{s_{-i}} p_{-i}(s_i, s_{-i}).$$

This proves the first equality. By interchanging i and $-i$ we get the second equality. □

It follows by the Strictly Competitive Games Theorem 7.2(i) that for zero-sum games a Nash equilibrium exists iff $maxmin_1 = minmax_1$. When this equality holds in a zero-sum game, the common value of $maxmin_1$ and $minmax_1$ is called the ***value*** of the game.

Example 7.5 Consider the zero-sum game represented by the following reward matrix:

	L	M	R
T	4	3	5
B	6	2	1

To compute $maxmin_1$ and $minmax_1$, as in Example 6.4, we extend the matrix with an additional row and column and fill in the minima of the rows and the maxima of the columns:

	L	M	R	f_1
T	4	3	5	3
B	6	2	1	1
F_1	6	3	5	

We see that $maxmin_1 = minmax_1 = 3$. So 3 is the value of this game. □

The above result does not hold for arbitrary strictly competitive games. To see it notice that in any two-player game a multiplication of the payoffs of player i by 2 leads to the doubling of the value of $maxmin_i$ and it does not affect the value of $minmax_{-i}$. Moreover, this multiplication procedure does not affect the property that a game is strictly competitive.

In an arbitrary strategic game with multiple Nash equilibria, for example the Battle of the Sexes game, the players face the following coordination problem. Suppose that each of them chooses a strategy from a Nash equilibrium. Then it can happen that this way they selected a joint strategy that is not a Nash equilibrium. For instance, in the Battle of the Sexes game the players can choose respectively F and B. The following result shows that in a zero-sum game such a coordination problem does not exist.

Theorem 7.6 (Interchangeability) *Consider a zero-sum game G.*

(i) Suppose that a Nash equilibrium of G exists. Then any joint strategy (s_1^*, s_2^*) such that each s_i^* is a security strategy for player i is a Nash equilibrium of G.

(ii) Suppose that (s_1^*, s_2^*) and (t_1^*, t_2^*) are Nash equilibria of G. Then so are (s_1^*, t_2^*) and (t_1^*, s_2^*).

Proof.
(i) Let (s_1^*, s_2^*) be a pair of security strategies for players 1 and 2. Fix i. By definition

$$\min_{s_i} p_{-i}(s_i, s_{-i}^*) = maxmin_{-i}. \tag{7.3}$$

But

$$\min_{s_i} p_{-i}(s_i, s^*_{-i}) = \min_{s_i} -p_i(s_i, s^*_{-i}) = -\max_{s_i} p_i(s_i, s^*_{-i})$$

and by the Zero-sum Theorem 7.4

$$maxmin_{-i} = -minmax_i.$$

So (7.3) implies

$$\max_{s_i} p_i(s_i, s^*_{-i}) = minmax_i. \tag{7.4}$$

We now rely on the Strictly Competitive Games Theorem 7.2. By item (ii) for $j = 1, 2$ we have $maxmin_j = minmax_j$, so by the proof of item (i) and (7.4) we conclude that (s^*_i, s^*_{-i}) is a Nash equilibrium.

(ii) By (i) and the Strictly Competitive Games Theorem 7.2(iv). □

The assumption that a Nash equilibrium exists is obviously necessary in item (i) of the above theorem. Indeed, in the finite zero-sum games security strategies always exist, in contrast to the Nash equilibrium.

Finally, recall that throughout this chapter we assumed the existence of various minima and maxima. So the results of this chapter apply only to a specific class of strictly competitive and zero-sum games. This class includes finite games. We shall return to this matter in a later chapter.

Chapter 8

Mixed Extensions

We now study a special case of infinite strategic games that are obtained in a canonic way from the finite games, by allowing mixed strategies. Below $[0, 1]$ stands for the real interval $\{r \in \mathbb{R} \mid 0 \leq r \leq 1\}$. By a ***probability distribution*** over a finite non-empty set A we mean a function

$$\pi : A \rightarrow [0, 1]$$

such that $\sum_{a \in A} \pi(a) = 1$. We denote the set of probability distributions over A by ΔA.

8.1 Mixed strategies

Consider now a finite strategic game $G := (S_1, \ldots, S_n, p_1, \ldots, p_n)$. By a ***mixed strategy*** of player i in G we mean a probability distribution over S_i. So ΔS_i is the set of mixed strategies available to player i. In what follows, we denote a mixed strategy of player i by m_i and a joint mixed strategy of the players by m.

Given a mixed strategy m_i of player i we define

$$support(m_i) := \{a \in S_i \mid m_i(a) > 0\}$$

and call this set the ***support*** of m_i. In specific examples we write a mixed strategy m_i as the sum $\sum_{a \in A} m_i(a) \cdot a$, where A is the support of m_i.

Note that in contrast to S_i the set ΔS_i is infinite. When referring to the mixed strategies, as in the previous chapters, we use the '$_{-i}$' notation. So for $m \in \Delta S_1 \times \cdots \times \Delta S_n$ we have $m_{-i} = (m_j)_{j \neq i}$, etc.

We can identify each strategy $s_i \in S_i$ with the mixed strategy that puts 'all the weight' on the strategy s_i. In this context s_i will be called a ***pure strategy***. Consequently we can view S_i as a subset of ΔS_i and S_{-i} as a subset of $\times_{j \neq i} \Delta S_j$.

By a ***mixed extension*** of $(S_1, \ldots, S_n, p_1, \ldots, p_n)$ we mean the strategic game

$$(\Delta S_1, \ldots, \Delta S_n, p_1, \ldots, p_n),$$

where each function p_i is extended in a canonic way from $S := S_1 \times \cdots \times S_n$ to $M := \Delta S_1 \times \cdots \times \Delta S_n$ by first viewing each joint mixed strategy $m = (m_1, \ldots, m_n) \in M$ as a probability distribution over S, by putting for $s \in S$

$$m(s) := m_1(s_1) \cdot \ldots \cdot m_n(s_n),$$

and then by putting

$$p_i(m) := \sum_{s \in S} m(s) \cdot p_i(s).$$

Example 8.1 Reconsider the Battle of the Sexes game from Chapter 1. Suppose that player 1 (man) chooses the mixed strategy $\frac{1}{2}F + \frac{1}{2}B$, while player 2 (woman) chooses the mixed strategy $\frac{1}{4}F + \frac{3}{4}B$. This pair m of the mixed strategies determines a probability distribution over the set of joint strategies, that we list to the left of the bimatrix of the game:

	F	B
F	$\frac{1}{8}$	$\frac{3}{8}$
B	$\frac{1}{8}$	$\frac{3}{8}$

	F	B
F	$2, 1$	$0, 0$
B	$0, 0$	$1, 2$

To compute the payoff of player 1 for this mixed strategy m we multiply each of his payoffs for a joint strategy by its probability and sum it up:

$$p_1(m) = \frac{1}{8}2 + \frac{3}{8}0 + \frac{1}{8}0 + \frac{3}{8}1 = \frac{5}{8}.$$

Analogously

$$p_2(m) = \frac{1}{8}1 + \frac{3}{8}0 + \frac{1}{8}0 + \frac{3}{8}2 = \frac{7}{8}.$$

□

This example suggests the computation of the payoffs in two-player games using matrix multiplication. First, we view each bimatrix of such a game as a pair of matrices $(\mathbf{A}, \mathbf{B})$. The first matrix represents the payoffs to player 1 and the second one to player 1. Assume now that player 1 has k strategies and player 2 has ℓ strategies. Then both $\mathbf{A}$ and $\mathbf{B}$ are $k \times \ell$ matrices. Further, each mixed strategy of player 1 can be viewed as a row vector $\mathbf{p}$ of length k (i.e., a $1 \times k$ matrix) and each mixed strategy of player 2 as a row vector $\mathbf{q}$ of length ℓ (i.e., a $1 \times \ell$ matrix). Since $\mathbf{p}$ and $\mathbf{q}$ represent mixed strategies, we have $\mathbf{p} \in \Delta^{k-1}$ and $\mathbf{q} \in \Delta^{\ell-1}$, where for all $m \geq 0$

$$\Delta^{m-1} := \{(x_1, \ldots, x_m) \mid \sum_{i=1}^{m} x_i = 1 \text{ and } \forall i \in \{1, \ldots, m\}\ x_i \geq 0\}.$$

Δ^{m-1} is called the $(m-1)$-***dimensional unit simplex***.

In the case of our example we have

$$\mathbf{p} = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \end{pmatrix}, \mathbf{q} = \begin{pmatrix} \frac{1}{4} & \frac{3}{4} \end{pmatrix}, \mathbf{A} = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}, \mathbf{B} = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}.$$

Now, the payoff functions can be defined as follows:

$$p_1(\mathbf{p}, \mathbf{q}) = \sum_{i=1}^{k} \sum_{j=1}^{\ell} \mathbf{p}_i \mathbf{q}_j \mathbf{A}_{ij} = \mathbf{p}\mathbf{A}\mathbf{q}^T$$

and

$$p_2(\mathbf{p}, \mathbf{q}) = \sum_{i=1}^{k} \sum_{j=1}^{\ell} \mathbf{p}_i \mathbf{q}_j \mathbf{B}_{ij} = \mathbf{p}\mathbf{B}\mathbf{q}^T.$$

8.2 Nash equilibria in mixed strategies

In the context of a mixed extension we talk about a ***pure Nash equilibrium***, when each of the constituent strategies is pure, and refer to an arbitrary Nash equilibrium of the mixed extension as a ***Nash equilibrium in mixed strategies*** of the initial finite game. In what follows, when we use the letter m we implicitly refer to the latter Nash equilibrium.

Below we shall need the following notion. Given a probability distribution π over a finite non-empty multiset[1] A of reals, we call

$$\sum_{r \in A} \pi(r) \cdot r$$

[1]This reference to a multiset is relevant.

a ***convex combination of the elements of*** A. For instance, given the multiset $A := \{\!\{4, 2, 2\}\!\}$, $\frac{1}{3}4 + \frac{1}{3}2 + \frac{1}{3}2$, so $\frac{8}{3}$, is a convex combination of the elements of A.

To see the use of this notion when discussing mixed strategies note that for every joint mixed strategy m we have

$$p_i(m) = \sum_{s_i \in support(m_i)} m_i(s_i) \cdot p_i(s_i, m_{-i}).$$

That is, $p_i(m)$ is a convex combination of the elements of the multiset

$$\{\!\{p_i(s_i, m_{-i}) \mid s_i \in support(m_i)\}\!\}.$$

We shall employ the following simple observations on convex combinations.

Note 8.2 (Convex Combination) *Consider a convex combination*

$$cc := \sum_{r \in A} \pi(r) \cdot r$$

of the elements of a finite multiset A of reals. Then

(i) $\max A \geq cc$,

(ii) $cc \geq \max A$ *iff*

- $cc = r$ *for all* $r \in A$ *such that* $\pi(r) > 0$,
- $cc \geq r$ *for all* $r \in A$ *such that* $\pi(r) = 0$.

□

Lemma 8.3 (Characterization) *Consider a finite strategic game*

$$(S_1, \ldots, S_n, p_1, \ldots, p_n).$$

The following statements are equivalent:

(i) m *is a Nash equilibrium in mixed strategies, i.e.,*

$$p_i(m) \geq p_i(m'_i, m_{-i})$$

for all $i \in \{1, \ldots, n\}$ *and all* $m'_i \in \Delta S_i$, *(ii) for all* $i \in \{1, \ldots, n\}$ *and all* $s_i \in S_i$

$$p_i(m) \geq p_i(s_i, m_{-i}),$$

(iii) for all $i \in \{1, \ldots, n\}$ *and all* $s_i \in support(m_i)$

$$p_i(m) = p_i(s_i, m_{-i})$$

and for all $i \in \{1, \ldots, n\}$ *and all* $s_i \notin support(m_i)$

$$p_i(m) \geq p_i(s_i, m_{-i}).$$

Note that the equivalence between (i) and (ii) implies that each Nash equilibrium of the initial game is a pure Nash equilibrium of the mixed extension. In turn, the equivalence between (i) and (iii) provides us with a straightforward way of testing whether a joint mixed strategy is a Nash equilibrium.

Proof.
$(i) \Rightarrow (ii)$ Immediate.
$(ii) \Rightarrow (iii)$ We noticed already that $p_i(m)$ is a convex combination of the elements of the multiset

$$A := \{\!\{p_i(s_i, m_{-i}) \mid s_i \in support(m_i)\}\!\}.$$

So this implication is a consequence of part (ii) of the Convex Combination Note 8.2.
$(iii) \Rightarrow (i)$ Consider the multiset

$$A := \{\!\{p_i(s_i, m_{-i}) \mid s_i \in S_i\}\!\}.$$

But for all $m'_i \in \Delta S_i$, in particular m_i, the payoff $p_i(m'_i, m_{-i})$ is a convex combination of the elements of the multiset A.

So by the assumptions and part (ii) of the Convex Combination Note 8.2

$$p_i(m) \geq max\ A,$$

and by part (i) of the above Note

$$max\ A \geq p_i(m'_i, m_{-i}).$$

Hence $p_i(m) \geq p_i(m'_i, m_{-i})$. □

We now illustrate the use of the above theorem by finding in the Battle of the Sexes game a Nash equilibrium in mixed strategies, in addition to the two pure ones exhibited in Chapter 2. Take

$$m_1 := r_1 \cdot F + (1 - r_1) \cdot B,$$
$$m_2 := r_2 \cdot F + (1 - r_2) \cdot B,$$

where $0 < r_1, r_2 < 1$. By definition

$$p_1(m_1, m_2) = 2 \cdot r_1 \cdot r_2 + (1 - r_1) \cdot (1 - r_2),$$
$$p_2(m_1, m_2) = r_1 \cdot r_2 + 2 \cdot (1 - r_1) \cdot (1 - r_2).$$

Suppose now that (m_1, m_2) is a Nash equilibrium in mixed strategies. By the equivalence between (i) and (iii) of the Characterization Lemma 8.3 $p_1(F, m_2) = p_1(B, m_2)$, i.e., (using $r_1 = 1$ and $r_1 = 0$ in the above formula for $p_1(\cdot)$) $2 \cdot r_2 = 1 - r_2$, and $p_2(m_1, F) = p_2(m_1, B)$, i.e., (using $r_2 = 1$ and $r_2 = 0$ in the above formula for $p_2(\cdot)$) $r_1 = 2 \cdot (1 - r_1)$. So $r_2 = \frac{1}{3}$ and $r_1 = \frac{2}{3}$.

This implies that for these values of r_1 and r_2, (m_1, m_2) is a Nash equilibrium in mixed strategies and we have

$$p_1(m_1, m_2) = p_2(m_1, m_2) = \tfrac{2}{3}.$$

8.3 Nash theorem

We now establish a fundamental result about games that are mixed extensions. In what follows we shall use the following result from the calculus.

Theorem 8.4 (Extreme Value Theorem) *Suppose that A is a non-empty compact subset of $\mathbb{R}^n$ and*

$$f : A \rightarrow \mathbb{R}$$

is a continuous function. Then f attains a minimum and a maximum. □

The example of the Matching Pennies game illustrated that some strategic games do not have a Nash equilibrium. In the case of mixed extensions the situation changes and we have the following fundamental result established by J. Nash in 1950.

Theorem 8.5 (Nash) *Every mixed extension of a finite strategic game has a Nash equilibrium.*

In other words, every finite strategic game has a Nash equilibrium in mixed strategies. In the case of the Matching Pennies game it is straightforward to check that $(\frac{1}{2} \cdot H + \frac{1}{2} \cdot T, \frac{1}{2} \cdot H + \frac{1}{2} \cdot T)$ is such a Nash equilibrium. In this equilibrium the payoffs to each player are 0.

Nash Theorem follows directly from the following result.[2]

Theorem 8.6 (Kakutani) *Suppose that A is a non-empty compact and convex subset of $\mathbb{R}^n$ and*

$$\Phi : A \rightarrow \mathcal{P}(A)$$

such that

- *$\Phi(x)$ is non-empty and convex for all $x \in A$,*
- *the* ***graph*** *of Φ, so the set $\{(x, y) \mid y \in \Phi(x)\}$, is closed.*

Then $x^* \in A$ exists such that $x^* \in \Phi(x^*)$. □

Proof of Nash Theorem. Fix a finite strategic game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$. Define the function $best_i : \times_{j \neq i} \Delta S_j \rightarrow \mathcal{P}(\Delta S_i)$ by

$$best_i(m_{-i}) := \{m_i \in \Delta S_i \mid m_i \text{ is a best response to } m_{-i}\}.$$

Then define the function $best : \Delta S_1 \times \cdots \times \Delta S_n \rightarrow \mathcal{P}(\Delta S_1 \times \cdots \times \Delta S_n)$ by

$$best(m) := best_1(m_{-1}) \times \cdots \times best_n(m_{-n}).$$

It is now straightforward to check that m is a Nash equilibrium iff $m \in best(m)$. Moreover, one easily can check that the function $best(\cdot)$ satisfies the conditions of Kakutani Theorem. The fact that for every joint mixed strategy m, $best(m)$ is non-empty is a direct consequence of the Extreme Value Theorem 8.4. □

Ever since Nash established his celebrated Theorem, a search has continued to generalize his result to a larger class of games. A motivation for this endevour has been existence of natural infinite games that are not mixed extensions of finite games. As an example of such an early result let us mention the following theorem stablished independently in 1952 by Debreu, Fan and Glickstein.

[2] Recall that a subset A of $\mathbb{R}^n$ is called ***compact*** if it is closed and bounded.

Theorem 8.7 *Consider a strategic game such that*

- *each strategy set is a non-empty compact convex subset of* $\mathbb{R}^n$*,*
- *each payoff function* p_i *is continuous and quasi-concave in the ith argument.*[3]

Then a Nash equilibrium exists.

More recent work in this area focused on existence of Nash equilibria in games with non-continuous payoff functions.

8.4 Minimax theorem

Let us return now to strictly competitive games that we studied in Chapter 7. First note the following lemma.

Lemma 8.8 *Consider a strategic game* $(S_1, \ldots, S_n, p_1, \ldots, p_n)$ *that is a mixed extension. Then*

(i) For all $s_i \in S_i$*,* $\min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i})$ *exists.*

(ii) $\max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i})$ *exists.*

(iii) For all $s_{-i} \in S_{-i}$*,* $\max_{s_i \in S_i} p_i(s_i, s_{-i})$ *exists.*

(iv) $\min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} p_i(s_i, s_{-i})$ *exists.*

Proof. It is a direct consequence of the Extreme Value Theorem 8.4. □

This lemma implies that we can apply the results of Chapter 7 to each strictly competitive game that is a mixed extension. Indeed, it ensures that the minima and maxima the existence of which we assumed in the proofs given there always exist. However, equipped with the knowledge that each such game has a Nash equilibrium we can now draw additional conclusions.

Theorem 8.9 *Consider a strictly competitive game that is a mixed extension. For* $i = 1, 2$ *we have* $maxmin_i = minmax_i$*.*

[3] Recall that the function $p_i : S \rightarrow \mathbb{R}$ is ***quasi-concave in the ith argument*** if the set $\{s'_i \in S_i \mid p_i(s'_i, s_{-i}) \geq p_i(s)\}$ is convex for all $s \in S$.

Proof. By the Nash Theorem 8.5 and the Strictly Competitive Games Theorem 7.2(*ii*). □

The formulation 'a strictly competitive game that is a mixed extension' is rather awkward and it is tempting to write instead 'the mixed extension of a strictly competitive game'. However, one can show that the mixed extension of a strictly competitive game does not need to be a strictly competitive game, see Exercise 8.1.

On the other hand we have the following simple observation.

Note 8.10 (Mixed Extension) *The mixed extension of a zero-sum game is a zero-sum game.*

Proof. Fix a finite zero-sum game (S_1, S_2, p_1, p_2). For each joint strategy m we have

$$p_1(m)+p_2(m) = \sum_{s \in S} m(s)p_1(s)+\sum_{s \in S} m(s)p_2(s) = \sum_{s \in S} m(s)(p_1(s)+p_2(s)) = 0.$$

□

This means that for finite zero-sum games we have the following result, originally established by von Neumann in 1928.

Theorem 8.11 (Minimax) *Consider a finite zero-sum game* $G := (S_1, S_2, p_1, p_2)$. *Then for* $i = 1, 2$

$$\max_{m_i \in M_i} \min_{m_{-i} \in M_{-i}} p_i(m_i, m_{-i}) = \min_{m_{-i} \in M_{-i}} \max_{m_i \in M_i} p_i(m_i, m_{-i}).$$

Proof. By the Mixed Extension Note 8.10 the mixed extension of G is zero-sum, so strictly competitive. It suffices to use Theorem 8.9 and expand the definitions of $minmax_i$ and $maxmin_i$. □

Finally, note that using the matrix notation we can rewrite the above equalities as follows, where $\mathbf{A}$ is an arbitrary $k \times \ell$ matrix (that is the reward matrix of a zero-sum game):

$$\max_{\mathbf{p} \in \Delta^{k-1}} \min_{\mathbf{q} \in \Delta^{\ell-1}} \mathbf{pAq}^T = \min_{\mathbf{q} \in \Delta^{\ell-1}} \max_{\mathbf{p} \in \Delta^{k-1}} \mathbf{pAq}^T.$$

So the Minimax Theorem can be alternatively viewed as a theorem about matrices and unit simplices. This formulation of the Minimax Theorem has been generalized in many ways to a statement

$$\max_{x \in X} \min_{y \in Y} f(x, y) = \min_{y \in Y} \max_{x \in X} f(x, y),$$

where X and Y are appropriate sets replacing the unit simplices and $f : X \times Y \to \mathbb{R}$ is an appropriate function replacing the payoff function. Such theorems are called Minimax theorems.

Exercise 8.1 Find a 2×2 strictly competitive game such that its mixed extension is not a strictly competitive game.

Exercise 8.2 Prove that the Matching Pennies game has exactly one Nash equilibrium in mixed strategies. □

Exercise 8.3 Find all Nash equilibria in mixed strategies of the Rock, Paper, Scissors game. □

Chapter 9

Alternative Concepts

In the presentation until now we heavily relied on the definition of a strategic game and focused several times on the crucial notion of a Nash equilibrium. However, both the concept of an equilibrium and of a strategic game can be defined in alternative ways. Here we discuss some alternative definitions and explain their consequences.

9.1 Other equilibria notions

Nash equilibrium is a most popular and most widely used notion of an equilibrium. However, there are many other natural alternatives. In this section we briefly discuss three alternative equilibria notions. To define them fix a strategic game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$.

Strict Nash equilibrium We call a joint strategy s a ***strict Nash equilibrium*** if

$$\forall i \in \{1, \ldots, n\} \; \forall s'_i \in S_i \setminus \{s_i\} \; p_i(s_i, s_{-i}) > p_i(s'_i, s_{-i}).$$

So a joint strategy is a strict Nash equilibrium if each player achieves a *strictly lower* payoff by unilaterally switching to another strategy.

Obviously every strict Nash equilibrium is a Nash equilibrium and the converse does not need to hold.

Consider now the Battle of the Sexes game. Its pure Nash equilibria that we identified in Chapter 1 are clearly strict. However, its Nash equilibrium in mixed strategy we identified in Example 8.1 of Section 8.1 is not strict.

Indeed, the following simple observation holds.

Note 9.1 *Consider a mixed extension of a finite strategic game. Every strict Nash equilibrium is a Nash equilibrium in pure strategies.*

Proof. It is a direct consequence of the Characterization Lemma 8.3. □

Consequently each finite game with no Nash equilibrium in pure strategies, for instance the Matching Pennies game, has no strict Nash equilibrium in mixed strategies. So the analogue of Nash theorem does not hold for strict Nash equilibria, which makes this equilibrium notion less useful.

ϵ-Nash equilibrium The idea of an ϵ-Nash equilibrium formalizes the intuition that a joint strategy can be also be satisfactory for the players when each of them can gain only very little from deviating from his strategy.

Let $\epsilon > 0$ be a small positive real. We call a joint strategy s an ***ϵ-Nash equilibrium*** if

$$\forall i \in \{1, \ldots, n\}\ \forall s'_i \in S_i\ p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i}) - \epsilon.$$

So a joint strategy is an ϵ-Nash equilibrium if no player can gain more than ϵ by unilaterally switching to another strategy. In this context ϵ can be interpreted either as the amount of uncertainty about the payoffs or as the gain from switching to another strategy.

Clearly, a joint strategy is a Nash equilibrium iff it is an ϵ-Nash equilibrium for every $\epsilon > 0$. However, the payoffs in an ϵ-Nash equilibrium can be substantially lower than in a Nash equilibrium. Consider for example the following game:

	L	R
T	$1, \quad 1$	$0, \quad 0$
B	$1+\epsilon, \quad 1$	$100, 100$

This game has a unique Nash equilibrium (B, R), which obviously is also an ϵ-Nash equilibrium. However, (T, L) is also an ϵ-Nash equilibrium.

Strong Nash equilibrium Another variation of the notion of a Nash equilibrium focusses on the concept of a coalition, by which we mean a non-empty subset of all players.

Given a subset $K := \{k_1, \ldots, k_m\}$ of $N := \{1, \ldots, n\}$ we abbreviate the sequence $(s_{k_1}, \ldots, s_{k_m})$ of strategies to s_K and $S_{k_1} \times \cdots \times S_{k_m}$ to S_K.

We call a joint strategy s a ***strong Nash equilibrium*** if for all coalitions K there does not exist $s'_K \in S_K$ such that

$$p_i(s'_K, s_{N \setminus K}) > p_i(s_K, s_{N \setminus K}) \text{ for all } i \in K.$$

So a joint strategy is a strong Nash equilibrium if no coalition can profit from deviating from it, where by "profit from" we mean that each member of the coalition gets a strictly higher payoff. The notion of a strong Nash equilibrium generalizes the notion of a Nash equilibrium by considering possible deviations of coalitions instead of individual players.

Note that the unique Nash equilibrium of the Prisoner's Dilemma game is strict but not strong. For example, if both players deviate from D to C, then each of them gets a strictly higher payoff.

Correlated equilibrium The final concept of an equilibrium that we introduce is a generalization of Nash equilibrium in mixed strategies. Recall from Chapter 8 that given a finite strategic game $G := (S_1, \ldots, S_n, p_1, \ldots, p_n)$ each joint mixed strategy $m = (m_1, \ldots, m_n)$ induces a probability distribution over S, defined by

$$m(s) := m_1(s_1) \cdot \ldots \cdot m_n(s_n),$$

where $s \in S$.

We have then the following observation.

Note 9.2 (Nash Equilibrium in Mixed Strategies) *Consider a finite strategic game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$.*

Then m is a Nash equilibrium in mixed strategies iff for all $i \in \{1, \ldots, n\}$ and all $s'_i \in S_i$

$$\sum_{s \in S} m(s) \cdot p_i(s_i, s_{-i}) \geq \sum_{s \in S} m(s) \cdot p_i(s'_i, s_{-i}).$$

Proof. Fix $i \in \{1, \ldots, n\}$ and choose some $s'_i \in S_i$. Let

$$m'_i(s_i) := \begin{cases} 1 & \text{if } s_i = s'_i \\ 0 & \text{otherwise} \end{cases}$$

So m'_i is the mixed strategy that represents the pure strategy s'_i.

Let now $m' := (m_1, \ldots, m_{i-1}, m'_i, m_{i+1}, \ldots, m_n)$. We have

$$p_i(m) = \sum_{s \in S} m(s) \cdot p_i(s_i, s_{-i})$$

and

$$p_i(s'_i, m_{-i}) = \sum_{s \in S} m'(s) \cdot p_i(s_i, s_{-i}).$$

Further, one can check that

$$\sum_{s \in S} m'(s) \cdot p_i(s_i, s_{-i}) = \sum_{s \in S} m(s) \cdot p_i(s'_i, s_{-i}).$$

So the claim is a direct consequence of the equivalence between items (i) and (ii) of the Characterization Lemma 8.3. □

We now generalize the above inequality to an arbitrary probability distribution over S. This yields the following equilibrium notion. We call a probability distribution π over S a ***correlated equilibrium*** if for all $i \in \{1, \ldots, n\}$ and all $s'_i \in S_i$

$$\sum_{s \in S} \pi(s) \cdot p_i(s_i, s_{-i}) \geq \sum_{s \in S} \pi(s) \cdot p_i(s'_i, s_{-i}).$$

By the above Note every Nash equilibrium in mixed strategies is a correlated equilibrium. To see that the converse is not true consider the Battle of the Sexes game:

	F	B
F	$2,1$	$0,0$
B	$0,0$	$1,2$

It is easy to check that the following probability distribution forms a correlated equilibrium in this game:

	F	B
F	$\frac{1}{2}$	0
B	0	$\frac{1}{2}$

Intuitively, this equilibrium corresponds to a situation when an external observes flips a fair coin and gives each player a recommendation which strategy to choose.

Exercise 9.1 Check the above claim. □

9.2 Variations on the definition of strategic games

The notion of a strategic game is quantitative in the sense that it refers through payoffs to real numbers. A natural question to ask is: do the payoff values matter? The answer depends on which concepts we want to study. We mention here three qualitative variants of the definition of a strategic game in which the payoffs are replaced by preferences. By a ***preference relation*** on a set A we mean here a linear ordering on A.

In [48] a strategic game is defined as a sequence

$$(S_1, \ldots, S_n, \succeq_1, \ldots, \succeq_n),$$

where each $\succeq_i$ is player's i *preference relation* defined on the set $S_1 \times \cdots \times S_n$ of joint strategies.

In [5] another modification of strategic games is considered, called a ***strategic game with parametrized preferences***. In this approach each player i has a non-empty set of strategies S_i and a *preference relation* $\succeq_{s_{-i}}$ on S_i *parametrized* by a joint strategy s_{-i} of his opponents. In [5] only strict preferences were considered and so defined finite games with parametrized preferences were compared with the concept of ***CP-nets*** (Conditional Preference nets), a formalism used for representing conditional and qualitative preferences, see, e.g., [13].

Next, in [55] ***conversion/preference games*** are introduced. Such a game for n players consists of a set S of *situations* and for each player i a *preference relation* $\succeq_i$ on S and a *conversion relation* $\rightarrow_i$ on S. The definition is very general and no conditions are placed on the preference and conversion relations. These games are used to formalize gene regulation networks and some aspects of security.

Another generalization of strategic games, called ***graphical games***, introduced in [28]. These games stress the locality in taking decision. In a graphical game the payoff of each player depends only on the strategies of its neighbours in a given in advance graph structure over the set of players. Formally, such a game for n players with the corresponding strategy sets $S_1, \ldots, S_n$ is defined by assuming a neighbour function N that given a player i yields its set of neighbours $N(i)$. The payoff for player i is then a function p_i from $\times_{j \in N(i) \cup \{i\}} S_j$ to $\mathbb{R}$.

In all mentioned variants it is straightforward to define the notion of a Nash equilibrium. For example, in the conversion/preferences games it is defined as a situation s such that for all players i, if $s \rightarrow_i s'$, then $s' \not\succ_i s$. However, other introduced notions can be defined only for some variants. In particular, Pareto efficiency cannot be defined for strategic games with parametrized preferences since it requires a comparison of two arbitrary joint strategies. In turn, the notions of dominance cannot be defined for the conversion/preferences games, since they require the concept of a strategy for a player.

Various results concerning finite strategic games, for instance the IESDS Theorem 2.5, carry over directly to the the strategic games as defined in [48] or in [5]. On the other hand, in the variants of strategic games that rely on the notion of a preference we cannot consider mixed strategies, since the outcomes of playing different strategies by a player cannot be aggregated.

Bibliography

[1] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pages 295–304. IEEE Computer Society, 2004.

[2] E. Anshelevich, A. Dasgupta, J. M. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proc. of the 45th Symposium on Foundations of Computer Science (FOCS 2004)*, pages 295–304. IEEE Computer Society, 2004.

[3] K. R. Apt. The many faces of rationalizability. *The B.E. Journal of Theoretical Economics, 7(1)*, 2007. (Topics), Article 18, 39 pages. Available from `http://arxiv.org/abs/cs.GT/0608011`.

[4] K. R. Apt. Relative strength of strategy elimination procedures. *Economics Bulletin*, 3(21):1–9, 2007. Available from `http://www.economicsbulletin.com/`.

[5] K. R. Apt, F. Rossi, and K. B. Venable. Comparing the notions of optimality in CP-nets, strategic games and soft constraints. *Annals of Mathematics and Artificial Intelligence*, 52(1):25–54, 2008.

[6] K. R. Apt and S. Simon. A classification of weakly acyclic games. *Theory and Decision*, (78):501–524, 2015.

[7] K. R. Apt, S. Simon, and D. Wojtczak. Coordination games on directed graphs. In *Proc. of the 15th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2015)*, volume 215 of *EPTCS*, pages 67–80, 2016.

[8] R. Aumann. Game theory. In J. Eatwell, M. Milgate, and P. Newman, editors, *The New Palgrave: A Dictionary of Economics*. Stockton Press, New York, 1987.

[9] K. Basu. The traveler's dilemma: paradoxes of rationality in game theory. *American Economic Review*, 84(2):391–395, 1994.

[10] K. Basu. The traveler's dilemma. *Scientific American*, June:90–95, 2007.

[11] K. Binmore. *Fun and Games: A Text on Game Theory*. D.C. Heath, 1991.

[12] K. Binmore. *Playing for Real: A Text on Game Theory*. Oxford University Press, Oxford, 2007.

[13] C. Boutilier, R. I. Brafman, C. Domshlak, H. H. Hoos, and D. Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artif. Intell. Res. (JAIR)*, 21:135–191, 2004.

[14] D. Braess. Über ein paradoxon aus der verkehrsplanung. *Unternehmensforschung*, 12(1):258–268, 1968.

[15] I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, and L. Moscardelli. *Tight Bounds for Selfish and Greedy Load Balancing*, pages 311–322. Springer, 2006.

[16] G. Christodoulou and E. Koutsoupias. *On the Price of Anarchy and Stability of Correlated Equilibria of Linear Congestion Games*, pages 59–70. Springer, 2005.

[17] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 67–73. ACM, 2005.

[18] A. Cournot. *Recherches sur les Principes Mathématiques de la Théorie des Richesses*. Hachette, 1838. Republished in English as *Researches Into the Mathematical Principles of the Theory of Wealth.*

[19] M. Dufwenberg and M. Stegeman. Existence and uniqueness of maximal reductions under iterated strict dominance. *Econometrica*, 70(5):2007–2023, 2002.

[20] R. Engelberg and M. Schapira. Weakly-acyclic (internet) routing games. In *Proc. 4th International Symposium on Algorithmic Game Theory (SAGT11)*, volume 6982 of *Lecture Notes in Computer Science*, pages 290–301. Springer, 2011.

[21] A. Fabrikant, A. Jaggard, and M. Schapira. On the structure of weakly acyclic games. In *Proceedings of the Third International Symposium on Algorithmic Game Theory (SAGT 2010)*, volume 6386 of *Lecture Notes in Computer Science*, pages 126–137. Springer, 2010.

[22] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, Massachusetts, 1991.

[23] R. Gardner. *Games for Business and Economics*. J. Wiley & Sons, Inc., New York, N.Y, 1995.

[24] I. Gilboa, E. Kalai, and E. Zemel. On the order of eliminating dominated strategies. *Operation Research Letters*, 9:85–89, 1990.

[25] G. Hardin. The tragedy of the commons. *Science*, 162:1243–1248, 1968.

[26] H. Hotelling. Stability in competition. *The Economic Journal*, 39:41–57, 1929.

[27] G. Jehle and P. Reny. *Advanced Microeconomic Theory*. Addison Wesley, Reading, Massachusetts, second edition, 2000.

[28] M. Kearns, M. Littman, and S. Singh. Graphical models for game theory. In *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence (UAI '01)*, pages 253–260. Morgan Kaufmann, 2001.

[29] D. König. Über eine Schlußweise aus dem Endlichen ins Unendliche. *Acta Litt. Ac. Sci.*, 3:121–130, 1927.

[30] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science*, pages 404–413. Springer, 1999.

[31] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009.

[32] R. D. Luce and H. Raiffa. *Games and Decisions.* John Wiley and Sons, New York, 1957.

[33] N. G. Mankiw. *Principles of Economics.* Harcourt College Publishers, Orlando, Florida, second edition, 2001.

[34] J. Marden, G. Arslan, and J. Shamma. Regret based dynamics: convergence in weakly acyclic games. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 194–201. IFAAMAS, 2007.

[35] A. Mas-Collel, M. D. Whinston, and J. R. Green. *Microeconomic Theory.* Oxford University Press, Oxford, 1995.

[36] M. Maschler, E. Solan, and S. Zamir. *Game Theory.* Cambridge University Press, Cambridge, 2013.

[37] I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behaviour*, 13:111–124, 1996.

[38] I. Milchtaich. Schedulers, potentials and weak potentials in weakly acyclic games. Working paper 2013-03, Bar-Ilan University, Department of Economics, 2013.

[39] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behaviour*, 14:124–143, 1996.

[40] H. Moulin. *Game Theory for the Social Sciences.* NYU Press, New York, second, revised edition, 1986.

[41] R. B. Myerson. *Game Theory: Analysis of Conflict.* Harvard University Press, Cambridge, Massachusetts, 1991.

[42] J. F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences, USA*, 36:48–49, 1950.

[43] J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.

[44] M. H. A. Newman. On theories with a combinatorial definition of "equivalence". *Annals of Math.*, 43(2):223–243, 1942.

[45] N. Nisan, T. Roughgarden, É. Tardos, and V. J. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[46] N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[47] M. J. Osborne. *An Introduction to Game Theory*. Oxford University Press, Oxford, 2005.

[48] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, Cambridge, Massachusetts, 1994.

[49] H. Peters. *Game Theory: A Multi-Leveled Approach*. Springer, Berlin, 2008.

[50] K. Ritzberger. *Foundations of Non-cooperative Game Theory*. Oxford University Press, Oxford, 2002.

[51] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, (2):65–67, 1973.

[52] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 513–522. ACM, 2009.

[53] T. Roughgarden. Intrinsic robustness of the price of anarchy. *Journal of the ACM*, 62(5):32:1–32:42, 2015.

[54] T. Roughgarden and Éva Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.

[55] S. L. Roux, P. Lescanne, and R. Vestergaard. Conversion/preference games. *CoRR*, abs/0811.0071, 2008.

[56] A. S. Schulz and N. S. Moses. On the performance of user equilibria in traffic networks. In *Proceedings of the 14th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 86–87. Society for Industrial and Applied Mathematics (SIAM), 2003.

[57] Y. Shoham and K. Leyton-Brown. *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Morgan and Claypool Publishers, Princeton, 2008.

[58] M. Stegeman. Deleting strictly eliminating dominated strategies. Working Paper 1990/6, Department of Economics, University of North Carolina, 1990.

[59] S. Tadelis. *Game Theory: an Introduction*. Princeton University Press, Princeton, 2013.

[60] É. Tardos and V. J. Vazirani. Basic solution concepts and computational issues. In N. Nisan, T. Roughgarden, É. Tardos, and V. J. Vazirani, editors, *Algorithmic Game Theory*, chapter 1, pages 3–28. Cambridge University Press, 2007.

[61] S. Tijs. *Introduction to Game Theory.* Hindustan Book Agency, Gurgaon, India, 2003.

[62] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–27, 1961.

[63] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

[64] H. P. Young. The evolution of conventions. *Econometrica*, 61(1):57–84, 1993.