

A Primer on Strategic Games

Krzysztof R. Apt*

Abstract

This is a short introduction to the subject of strategic games. We focus on the concepts of best response, Nash equilibrium, strict and weak dominance, and mixed strategies, and study the relation between these concepts in the context of the iterated elimination of strategies. Also, we discuss some variants of the original definition of a strategic game. Finally, we introduce the basics of mechanism design and use pre-Bayesian games to explain it.

1 Introduction

Mathematical game theory, as launched by Von Neumann and Morgenstern in their seminal book, von Neumann and Morgenstern [1944], followed by Nash's contributions Nash [1950, 1951], 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 successfully 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 maximising players. In strategic games each player has a payoff function that he aims to maximise 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 this primer is to provide a simple introduction to the most common concepts used in strategic games: best response, Nash equilibrium, dominated strategies and mixed strategies and to clarify the relation between these concepts.

In the first part we consider the case of games with *complete information*. In the second part we discuss strategic games with *incomplete information*, by introducing first the basics of the theory of *mechanism design* that deals with ways of preventing *strategic behaviour*, i.e., manipulations aiming at maximising one's profit. We focus on the concepts, examples and results, and leave simple proofs as exercises.

*Centrum for Mathematics and Computer Science (CWI), Science Park 123, 1098 XG Amsterdam, the Netherlands, and University of Amsterdam.

2 Basic concepts

Assume a set $\{1, \dots, n\}$ of players, where $n > 1$. A **strategic game** (or **non-cooperative game**) for n players, written as $(S_1, \dots, S_n, p_1, \dots, p_n)$, consists of

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

for each player i .

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 maximise his payoff,
- players have **common knowledge** of the game and of each others' rationality.¹

Here are three classic examples of strategic two-player games to which we shall return in a moment. We represent such games in the form of a bimatrix, the entries of which are the corresponding payoffs to the row and column players.

Prisoner's Dilemma

	<i>C</i>	<i>D</i>
<i>C</i>	2, 2	0, 3
<i>D</i>	3, 0	1, 1

Battle of the Sexes

	<i>F</i>	<i>B</i>
<i>F</i>	2, 1	0, 0
<i>B</i>	0, 0	1, 2

Matching Pennies

	<i>H</i>	<i>T</i>
<i>H</i>	1, -1	-1, 1
<i>T</i>	-1, 1	1, -1

We introduce now some basic notions that will allow us to discuss and analyse strategic games in a meaningful way. Fix a strategic game

$$(S_1, \dots, S_n, p_1, \dots, p_n).$$

We denote $S_1 \times \dots \times S_n$ by S , call each element $s \in S$ a **joint strategy**, or a **strategy profile**, denote the i th element of s by s_i , and abbreviate the sequence $(s_j)_{j \neq i}$ to s_{-i} .

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

Occasionally we write (s_i, s_{-i}) instead of s . Finally, we abbreviate $\times_{j \neq i} S_j$ to S_{-i} and use the ‘ $_{-i}$ ’ notation for other sequences and Cartesian products.

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

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

So a joint strategy is a Nash equilibrium if no player can achieve a higher payoff by *unilaterally* switching to another strategy.

Finally, we call a joint strategy s **Pareto efficient** if for no joint strategy s'

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

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

Some games, like the Prisoner’s Dilemma, have a unique Nash equilibrium, namely (D, D) , while some other ones, like the Matching Pennies, have no Nash equilibrium. Yet other games, like the Battle of the Sexes, have multiple Nash equilibria, namely (F, F) and (B, B) . One of the peculiarities of the Prisoner’s Dilemma game is that its Nash equilibrium is the only outcome that is not Pareto efficient.

Let us return now to our analysis of an arbitrary strategic game $(S_1, \dots, S_n, p_1, \dots, p_n)$. Let s_i, s'_i be strategies of player i . We say that s_i **strictly dominates** s'_i (or equivalently, that s'_i is **strictly dominated by** s_i) if

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

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

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

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

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

Further, we say that s_i is **strictly dominant** if it strictly dominates all other strategies of player i and define analogously a **weakly dominant** and a **dominant** strategy.

Clearly, a rational player will not choose a strictly dominated strategy. 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 Pareto efficient outcome (C, C) yields for each of them a strictly higher payoff.

The Prisoner’s Dilemma game can be easily generalised to n players as follows. Assume that each player has two strategies, C and D . Denote by C^n the joint strategy in which

each strategy equals C and similarly with D^n . Further, given a joint strategy s_{-i} of the opponents of player i denote by $|s_{-i}(C)|$ the number of C strategies in s_{-i} .

Assume now that k_i and l_i , where $i \in \{1, \dots, n\}$, are real numbers such that for all $i \in \{1, \dots, n\}$ we have $k_i(n-1) > l_i > 0$. We put

$$p_i(s) := \begin{cases} k_i |s_{-i}(C)| + l_i & \text{if } s_i = D \\ k_i |s_{-i}(C)| & \text{if } s_i = C. \end{cases}$$

Note that for $n = 2, k_i = 2$ and $l_i = 1$ we get the original Prisoner's Dilemma game.

Then for all players i we have $p_i(C^n) = k_i(n-1) > l_i = p_i(D^n)$, so for all players the strategy profile C^n yields a strictly higher payoff than D^n . Yet for all players i strategy C is strictly dominated by strategy D , since for all $s_{-i} \in S_{-i}$ we have $p_i(D, s_{-i}) - p_i(C, s_{-i}) = l_i > 0$.

Whether a rational player will never choose a weakly dominated strategy is a more subtle issue that we shall not pursue here.

By definition, no player achieves a higher payoff by switching from a dominant strategy to another strategy. This explains the following obvious observation.

Note 2.1 (Dominant Strategy). *Consider a strategic game G . Suppose that s is a joint strategy such that each s_i is a dominant strategy. Then it is a Nash equilibrium of G .*

In particular, the conclusion of the lemma holds if each s_i is a strictly or a weakly dominant strategy. In the former case, when the game is finite, we can additionally assert (see the IESDS Theorem 3.2 below) that s is a unique Nash equilibrium of G . This stronger claim does not hold if each s_i is a weakly dominant strategy. 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) .

The converse of the above Note of course is not true. Indeed, there are games in which no strategy is dominant, and yet they have a Nash equilibrium. An example is the Battle of the Sexes game that has two Nash equilibria, but no dominant strategy.

So to find a Nash equilibrium (or Nash equilibria) of a game it does not suffice to check whether a dominant strategy exists. In what follows we investigate whether iterated elimination of strategies can be of help.

3 Iterated elimination of strategies I

3.1 Elimination of strictly dominated strategies

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 $G := (S_1, \dots, S_n, p_1, \dots, p_n)$ and (possibly empty) sets of strategies R_1, \dots, R_n such that $R_i \subseteq S_i$ for $i \in \{1, \dots, n\}$ we say that $R := (R_1, \dots, R_n, p_1, \dots, p_n)$ is a **restriction** of G . Here of course we view each p_i as a function on the subset $R_1 \times \dots \times R_n$ of $S_1 \times \dots \times S_n$. In what follows, given a restriction R we denote by R_i the set of strategies of player i in R .

We now introduce the following notion of reduction between the restrictions R and R' of G :

$$R \rightarrow_S R'$$

when $R \neq R'$, $\forall i \in \{1, \dots, n\} R'_i \subseteq R_i$ and

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

That is, $R \rightarrow_S R'$ when R' results from R by removing from it some strictly dominated strategies.

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 3.1. 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. 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 each player is left in R with exactly one strategy, we say that G is **solved by IESDS**.

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

Theorem 3.2 (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.*

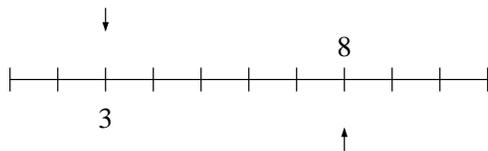
Exercise 1. Provide the proof. □

Example 3.3. A nice example of a game that is solved by IESDS is the **location game** due to Hotelling [1929]. Assume 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 attracted.

To be more specific we assume that the vendors choose a location from the set $\{1, \dots, n\}$ 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, they end up with a fractional payoff.

In general, we have the following game:

- each set of strategies consists of the set $\{1, \dots, n\}$,

- each payoff function p_i is defined by:

$$p_i(s_i, s_{3-i}) := \begin{cases} \frac{s_i + s_{3-i} - 1}{2} & \text{if } s_i < s_{3-i} \\ n - \frac{s_i + s_{3-i} - 1}{2} & \text{if } s_i > s_{3-i} \\ \frac{n}{2} & \text{if } s_i = s_{3-i}. \end{cases}$$

It is easy to see 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 , and so on. \square

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 game theory parlance: is strict dominance *order independent*? The answer is positive. The following result was established independently by Gilboa et al. [1990] and Stegeman [1990].

Theorem 3.4 (Order Independence I). *Given a finite strategic game all iterated eliminations of strictly dominated strategies yield the same outcome.*

As noted by Dufwenberg and Stegeman [2002] the above result does not hold for infinite strategic games.

Example 3.5. 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.

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. \square

The above example 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. An interested reader is referred to Dufwenberg and Stegeman [2002] and Apt [2007] where two different options are proposed and some limited order independence results are established.

The above example also shows that in the IESDS Theorem 3.2(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.

3.2 Elimination of weakly dominated strategies

Analogous considerations can be carried out for the elimination of weakly dominated strategies, by considering the appropriate reduction relation \rightarrow_W defined in the expected way. Below we abbreviate iterated elimination of weakly dominated strategies to **IEWDS**.

However, in the case of IEWDS 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 3.2 does not hold when reformulated for weak dominance.

On the other hand, some partial results are still valid here.

Theorem 3.6 (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 .*

Exercise 2. Provide the proof. □

Example 3.7. A nice example of a game that is solved by IEWDS is the **Beauty Contest game** due to Moulin [1986]. In this game there are $n > 2$ players, each with the set of strategies equal $\{1, \dots, 100\}$. 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 $\frac{1}{2}, 0, \frac{1}{2}$.

One can check that this game is solved by IEWDS and results in the joint strategy $(1, \dots, 1)$. Hence, by the IEWDS Theorem 3.6 this joint strategy is a (not necessarily unique; we shall return to this question in Section 5) Nash equilibrium. □

Exercise 3. Show that the Beauty Contest game is indeed solved by IEWDS. □

Note that in contrast to the IESDS Theorem 3.2 we do not claim in part (ii) of the IEWDS Theorem 3.6 that the resulting joint strategy is a *unique* Nash equilibrium. In fact, such a stronger claim does not hold. 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 3.8. Consider the following game:

	<i>L</i>	<i>M</i>	<i>R</i>
<i>T</i>	0, 1	1, 0	0, 0
<i>B</i>	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

	<i>L</i>	<i>M</i>
<i>T</i>	0, 1	1, 0
<i>B</i>	0, 0	0, 0

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

	<i>L</i>
<i>T</i>	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’. The reader may also check that the second Nash equilibrium, (B, L) , can be found using IEWDS, as well, but not the third one, (B, R) . □

To summarise, 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 3.6 does not hold for infinite games. Indeed, Example 3.5 applies here, as well.

3.3 Elimination of never best responses

Finally, we consider the process of eliminating strategies that are never best responses to a joint strategy of the opponents. To motivate this procedure 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. However, 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 $R \neq R'$, $\forall i \in \{1, \dots, n\} R'_i \subseteq R_i$ and

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

That is, $R \rightarrow_N R'$ when R' results from R by removing from it some strategies that are never best responses.

We then focus on the iterated elimination of never best responses, in short **IENBR**, obtained by using the \rightarrow_N^* relation. The following counterpart of the IESDS Theorem 3.2 then holds.

Theorem 3.9 (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.*

Exercise 4. Provide the proof. □

Further, as shown by Apt [2005], we have the following analogue of the Order Independence I Theorem 3.4.

Theorem 3.10 (Order Independence II). *Given a finite strategic game all iterated eliminations of never best responses yield the same outcome.*

In the case of infinite games we encounter the same problems as in the case of IESDS as Example 3.5 readily applies to IENBR, as well. In particular, 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 3.11. Consider the following infinite variant of the location game considered in Example 3.3. 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:

$$p_i(s_i, s_{3-i}) := \begin{cases} \frac{s_i + s_{3-i}}{2} & \text{if } s_i < s_{3-i} \\ 100 - \frac{s_i + s_{3-i}}{2} & \text{if } s_i > s_{3-i} \\ 50 & \text{if } s_i = s_{3-i}. \end{cases}$$

It is easy to check that in this game no strategy strictly or weakly dominates another one. On the other hand each strategy 50 is a best response to some strategy, namely to 50, 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 straightforward to check that this is the case. Moreover, by the IENBR Theorem 3.9(i) we know that this is a unique Nash equilibrium. \square

4 Mixed extension

We now study a special case of infinite strategic games that are obtained in a canonical 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 .

Consider now a finite strategic game $G := (S_1, \dots, S_n, p_1, \dots, p_n)$. 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

$$\text{support}(m_i) := \{a \in S_i \mid m_i(a) > 0\}$$

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 sections, we use the ‘ $_{-i}$ ’ notation. So for $m \in \Delta S_1 \times \dots \times \Delta S_n$ 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 $(S_1, \dots, S_n, p_1, \dots, p_n)$ we mean the strategic game

$$(\Delta S_1, \dots, \Delta S_n, p_1, \dots, p_n),$$

where each function p_i is extended in a canonical way from $S := S_1 \times \dots \times S_n$ to $M := \Delta S_1 \times \dots \times \Delta S_n$ by first viewing each joint mixed strategy $m = (m_1, \dots, m_n) \in M$ as a probability distribution over S , by putting for $s \in S$

$$m(s) := m_1(s_1) \cdot \dots \cdot m_n(s_n),$$

and then by putting

$$p_i(m) := \sum_{s \in S} m(s) \cdot p_i(s).$$

The notion of a Nash equilibrium readily applies to mixed extensions. In this context 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.

Lemma 4.1 (Characterisation). *Consider a finite strategic game $(S_1, \dots, S_n, p_1, \dots, p_n)$. 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 $i \in \{1, \dots, n\}$ and all $m'_i \in \Delta S_i$,

(ii) for all $i \in \{1, \dots, n\}$ and all $s_i \in S_i$

$$p_i(m) \geq p_i(s_i, m_{-i}),$$

(iii) for all $i \in \{1, \dots, n\}$ and all $s_i \in \text{support}(m_i)$

$$p_i(m) = p_i(s_i, m_{-i})$$

and for all $i \in \{1, \dots, n\}$ and all $s_i \notin \text{support}(m_i)$

$$p_i(m) \geq p_i(s_i, m_{-i}).$$

Exercise 5. Provide the proof. □

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.

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 Section 3. Take

$$\begin{aligned} m_1 &:= r_1 \cdot F + (1 - r_1) \cdot B, \\ m_2 &:= r_2 \cdot F + (1 - r_2) \cdot B, \end{aligned}$$

where $0 < r_1, r_2 < 1$. By definition

$$\begin{aligned} p_1(m_1, m_2) &= 2 \cdot r_1 \cdot r_2 + (1 - r_1) \cdot (1 - r_2), \\ p_2(m_1, m_2) &= r_1 \cdot r_2 + 2 \cdot (1 - r_1) \cdot (1 - r_2). \end{aligned}$$

Suppose now that (m_1, m_2) is a Nash equilibrium in mixed strategies. By the equivalence between (i) and (iii) of the Characterisation Lemma 4.1 $p_1(F, m_2) = p_1(B, m_2)$, 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

$$p_1(m_1, m_2) = p_2(m_1, m_2) = \frac{2}{3}.$$

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 due to Nash [1950].

Theorem 4.2 (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's Theorem follows directly from the following result due to Kakutani [1941].²

Theorem 4.3 (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 $x^* \in \Phi(x^*)$.* □

²Recall that a subset A of \mathbb{R}^n is called **compact** if it is closed and bounded, and is called **convex** if for any $\mathbf{x}, \mathbf{y} \in A$ and $\alpha \in [0, 1]$ we have $\alpha \mathbf{x} + (1 - \alpha) \mathbf{y} \in A$.

Proof of Nash's Theorem. Fix a finite strategic game $(S_1, \dots, S_n, p_1, \dots, p_n)$. 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 $best : \Delta S_1 \times \dots \times \Delta S_n \rightarrow \mathcal{P}(\Delta S_1 \times \dots \times \Delta S_n)$ by

$$best(m) := best_1(m_{-1}) \times \dots \times best_n(m_{-n}).$$

It is now straightforward to check that m is a Nash equilibrium iff $m \in best(m)$. Moreover, one can easily check that the function $best(\cdot)$ satisfies the conditions of Kakutani's Theorem. The fact that for every joint mixed strategy m , $best(m)$ is non-empty is a direct consequence of the Extreme Value Theorem stating that every real-valued continuous function on a compact subset of \mathcal{R}^ℓ attains a maximum. \square

5 Iterated elimination of strategies II

The notions of dominance apply in particular to mixed extensions of finite strategic games. But we can also consider dominance of a *pure* strategy by a *mixed* strategy. Given a finite strategic game $G := (S_1, \dots, S_n, p_1, \dots, p_n)$, we say that a (pure) strategy s_i of player i is ***strictly dominated by*** a mixed strategy m_i if

$$\forall s_{-i} \in S_{-i} \quad p_i(m_i, s_{-i}) > p_i(s_i, s_{-i}),$$

and that s_i is ***weakly dominated by*** a mixed strategy m_i if

$$\forall s_{-i} \in S_{-i} \quad p_i(m_i, s_{-i}) \geq p_i(s_i, s_{-i}) \text{ and } \exists s_{-i} \in S_{-i} \quad p_i(m_i, s_{-i}) > p_i(s_i, s_{-i}).$$

In what follows we discuss for these two forms of dominance the counterparts of the results presented in Section 3.

5.1 Elimination of strictly dominated strategies

Strict dominance by a mixed strategy leads to a stronger notion of strategy elimination. For example, in the game

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

the strategy B is strictly dominated neither by T nor M but is strictly dominated by $\frac{1}{2} \cdot T + \frac{1}{2} \cdot M$.

We now focus on iterated elimination of pure strategies that are strictly dominated by a mixed strategy. As in Section 3 we would like to clarify whether it affects the Nash equilibria, in this case equilibria in mixed strategies.

Instead of the lengthy wording ‘the iterated elimination of strategies strictly dominated by a mixed strategy’ we write **IESDMS**. We have then the following counterpart of the IESDS Theorem 3.2, where we refer to Nash equilibria in mixed strategies. Given a restriction G' of G and a joint mixed strategy m of G , when we say that m is a Nash equilibrium of G' we implicitly stipulate that each strategy used (with positive probability) in m is a strategy in G' .

Theorem 5.1 (IESDMS). *Suppose that G is a finite strategic game.*

- (i) *If G' is an outcome of IESDMS from G , then m is a Nash equilibrium of G iff it is a Nash equilibrium of G' .*
- (ii) *If G is solved by IESDMS, then the resulting joint strategy is a unique Nash equilibrium of G (in, possibly, mixed strategies).*

Exercise 6. Provide the proof. □

To illustrate the use of this result let us return to the Beauty Contest game discussed in Example 3.7. We explained there why $(1, \dots, 1)$ is a Nash equilibrium. Now we can draw a stronger conclusion.

Example 5.2. One can show that the Beauty Contest game is solved by IESDMS in 99 rounds. In each round the highest strategy of each player is removed and eventually each player is left with the strategy 1. On account of the above theorem we now conclude that $(1, \dots, 1)$ is a *unique* Nash equilibrium. □

Exercise 7. Show that the Beauty Contest game is indeed solved by IESDMS in 99 rounds. □

As in the case of strict dominance by a pure strategy we now address the question of whether the outcome of IESDMS is unique. The answer, as before, is positive. The following result was established by Osborne and Rubinstein [1994].

Theorem 5.3 (Order independence III). *All iterated eliminations of strategies strictly dominated by a mixed strategy yield the same outcome.*

5.2 Elimination of weakly dominated strategies

Next, we consider iterated elimination of pure strategies that are weakly dominated by a mixed strategy.

As already noticed in Subsection 3.2 an elimination by means of weakly dominated strategies can result in a loss of Nash equilibria. Clearly, the same observation applies here. We also have the following counterpart of the IEWDS Theorem 3.6, where we refer to Nash equilibria in mixed strategies. Instead of ‘the iterated elimination of strategies weakly dominated by a mixed strategy’ we write **IEWDMS**.

Theorem 5.4 (IEWDMS). *Suppose that G is a finite strategic game.*

- (i) If G' is an outcome of IEWDMS from G and m is a Nash equilibrium of G' , then m is a Nash equilibrium of G .
- (ii) If G is solved by IEWDMS, then the resulting joint strategy is a Nash equilibrium of G .

Here is a simple application of this theorem.

Corollary 5.5. *Every mixed extension of a finite strategic game has a Nash equilibrium such that no strategy used in it is weakly dominated by a mixed strategy.*

Proof. It suffices to apply Nash's Theorem 4.2 to an outcome of IEWDMS and use item (i) of the above theorem. \square

Finally, observe that the outcome of IEWMDS does not need to be unique. In fact, Example 3.8 applies here, as well.

5.3 Rationalizability

Finally, we consider iterated elimination of strategies that are never best responses to a joint mixed strategy of the opponents. Following Bernheim [1984] and Pearce [1984], strategies that survive such an elimination process are called rationalizable strategies.³

Formally, we define rationalizable strategies as follows. Consider a restriction R of a finite strategic game G . Let

$$\mathcal{RAT}(R) := (S'_1, \dots, S'_n),$$

where for all $i \in \{1, \dots, n\}$

$$S'_i := \{s_i \in R_i \mid \exists m_{-i} \in \times_{j \neq i} \Delta R_j \text{ } s_i \text{ is a best response to } m_{-i} \text{ in } G\}.$$

Note the use of G instead of R in the definition of S'_i . We shall comment on it below.

Consider now the outcome $G_{\mathcal{RAT}}$ of iterating \mathcal{RAT} starting with G . We call then the strategies present in the restriction $G_{\mathcal{RAT}}$ **rationalizable**.

We have the following counterpart of the IESDMS Theorem 5.1, due to Bernheim [1984].

Theorem 5.6. *Assume a finite strategic game G .*

- (i) *Then m is a Nash equilibrium of G iff it is a Nash equilibrium of $G_{\mathcal{RAT}}$.*
- (ii) *If each player has in $G_{\mathcal{RAT}}$ exactly one strategy, then the resulting joint strategy is a unique Nash equilibrium of G .*

Exercise 8. Provide the proof. \square

³More precisely, in each of these papers a different definition is used; see Apt [2007] for an analysis of the conditions for which these definitions coincide.

In the context of rationalizability a joint mixed strategy of the opponents is referred to as a **belief**. The definition of rationalizability is generic in the class of beliefs w.r.t. which best responses are collected. For example, we could use here joint pure strategies of the opponents, or probability distributions over the Cartesian product of the opponents' strategy sets, so the elements of the set ΔS_{-i} (extending in an expected way the payoff functions). In the first case we talk about **point beliefs** and in the second case about **correlated beliefs**.

In the case of point beliefs we can apply the elimination procedure entailed by \mathcal{RAT} to arbitrary games. To avoid discussion of the outcomes reached in the case of infinite iterations we focus on a result for a limited case. We refer here to Nash equilibria in pure strategies.

Theorem 5.7. *Assume a strategic game G . Consider the definition of the \mathcal{RAT} operator for the case of point beliefs and suppose that the outcome $G_{\mathcal{RAT}}$ is reached in finitely many steps.*

- (i) *Then s is a Nash equilibrium of G iff it is a Nash equilibrium of $G_{\mathcal{RAT}}$.*
- (ii) *If each player is left in $G_{\mathcal{RAT}}$ with exactly one strategy, then the resulting joint strategy is a unique Nash equilibrium of G .*

Exercise 9. Provide the proof. □

A subtle point is that when G is infinite, the restriction $G_{\mathcal{RAT}}$ may have empty strategy sets (and hence no joint strategy).

Example 5.8. Bertrand competition, originally proposed by Bertrand [1883], is a game concerned with a simultaneous selection of prices for the same product by two firms. The product is then sold by the firm that chose a lower price. In the case of a tie the product is sold by both firms and the profits are split.

Consider a version in which the range of possible prices is the left-open real interval $(0, 100]$ and the demand equals $100 - p$, where p is the lower price. So in this game G there are two players, each with the set $(0, 100]$ of strategies and the payoff functions are defined by:

$$p_1(s_1, s_2) := \begin{cases} s_1(100 - s_1) & \text{if } s_1 < s_2 \\ \frac{s_1(100 - s_1)}{2} & \text{if } s_1 = s_2 \\ 0 & \text{if } s_1 > s_2 \end{cases}$$

$$p_2(s_1, s_2) := \begin{cases} s_2(100 - s_2) & \text{if } s_2 < s_1 \\ \frac{s_2(100 - s_2)}{2} & \text{if } s_1 = s_2 \\ 0 & \text{if } s_2 > s_1. \end{cases}$$

Consider now each player's best responses to the strategies of the opponent. Since $s_1 = 50$ maximises the value of $s_1(100 - s_1)$ in the interval $(0, 100]$, the strategy 50 is

the unique best response of the first player to any strategy $s_2 > 50$ of the second player. Further, no strategy is a best response to a strategy $s_2 \leq 50$. By symmetry the same holds for the strategies of the second player.

So the elimination of never best responses leaves each player with a single strategy, 50. In the second round we need to consider the best responses to these two strategies in the *original* game G . In G the strategy $s_1 = 49$ is a better response to $s_2 = 50$ than $s_1 = 50$ and symmetrically for the second player. So in the second round of elimination both strategies 50 are eliminated and we reach the restriction with the empty strategy sets. By Theorem 5.7 we conclude that the original game G has no Nash equilibrium. \square

Note that if we defined S'_i in the definition of the operator \mathcal{RAT} using the restriction R instead of the original game G , the iteration would stop in the above example after the first round. Such a modified definition of the \mathcal{RAT} operator is actually an instance of the IENBR (iterated elimination of never best responses) in which at each stage all never best responses are eliminated. So for the above game G we can then conclude by the IENBR Theorem 3.9(i) that it has at most one equilibrium, namely $(50, 50)$, and then check separately that in fact it is not a Nash equilibrium.

5.4 A comparison between the introduced notions

We introduced so far the notions of strict dominance, weak dominance, and a best response, and related them to the notion of a Nash equilibrium. To conclude this section we clarify the connections between the notions of dominance and of best response.

Clearly, if a strategy is strictly dominated, then it is a never best response. However, the converse fails. Further, there is no relation between the notions of weak dominance and never best response. Indeed, in the game considered in Subsection 3.3 strategy C is a never best response, yet it is neither strictly nor weakly dominated. Further, in the game given in Example 3.8 strategy M is weakly dominated and is also a best response to B .

The situation changes in the case of mixed extensions of two-player finite games. Below, by a ***totally mixed strategy*** we mean a mixed strategy with full support, i.e., one in which each strategy is used with a strictly positive probability. The following results were established by Pearce [1984].

Theorem 5.9. *Consider a finite two-player strategic game.*

- (i) *A pure strategy is strictly dominated by a mixed strategy iff it is not a best response to a mixed strategy.*
- (ii) *A pure strategy is weakly dominated by a mixed strategy iff it is not a best response to a totally mixed strategy.*

We only prove here part (i). Pearce [1984] provides a short, but a bit tricky proof based on Nash's Theorem 4.2. The proof we provide, due to Fudenberg and Tirole [1991], is a bit more intuitive.

We shall use the following result, see, e.g., Rockafellar [1996].

Theorem 5.10 (Separating Hyperplane). *Let A and B be disjoint convex subsets of \mathbb{R}^k . Then there exists a non-zero $c \in \mathbb{R}^k$ and $d \in \mathbb{R}$ such that*

$$c \cdot x \geq d \text{ for all } x \in A,$$

$$c \cdot y \leq d \text{ for all } y \in B.$$

Proof of Theorem 5.9(i).

Clearly, if a pure strategy is strictly dominated by a mixed strategy, then it is not a best response to a mixed strategy. To prove the converse, fix a two-player strategic game (S_1, S_2, p_1, p_2) . Also fix $i \in \{1, 2\}$ and abbreviate $3 - i$ to $-i$.

Suppose that a strategy $s_i \in S_i$ is not strictly dominated by a mixed strategy. Let

$$A := \{x \in \mathbb{R}^{|S_{-i}|} \mid \forall s_{-i} \in S_{-i} \ x_{s_{-i}} > 0\}$$

and

$$B := \{(p_i(m_i, s_{-i}) - p_i(s_i, s_{-i}))_{s_{-i} \in S_{-i}} \mid m_i \in \Delta S_i\}.$$

By the choice of s_i the sets A and B are disjoint. Moreover, both sets are convex subsets of $\mathbb{R}^{|S_{-i}|}$.

By the Separating Hyperplane Theorem 5.10 for some non-zero $c \in \mathbb{R}^{|S_{-i}|}$ and $d \in \mathbb{R}$

$$c \cdot x \geq d \text{ for all } x \in A, \tag{1}$$

$$c \cdot y \leq d \text{ for all } y \in B. \tag{2}$$

But $\mathbf{0} \in B$, so by (2) $d \geq 0$. Hence by (1) and the definition of A for all $s_{-i} \in S_{-i}$ we have $c_{s_{-i}} \geq 0$. Again by (1) and the definition of A this excludes the contingency that $d > 0$, i.e., $d = 0$. Hence by (2)

$$\sum_{s_{-i} \in S_{-i}} c_{s_{-i}} p_i(m_i, s_{-i}) \leq \sum_{s_{-i} \in S_{-i}} c_{s_{-i}} p_i(s_i, s_{-i}) \text{ for all } m_i \in \Delta S_i. \tag{3}$$

Let $\bar{c} := \sum_{s_{-i} \in S_{-i}} c_{s_{-i}}$. By the assumption $\bar{c} \neq 0$. Take

$$m_{-i} := \sum_{s_{-i} \in S_{-i}} \frac{c_{s_{-i}}}{\bar{c}} s_{-i}.$$

Then (3) can be rewritten as

$$p_i(m_i, m_{-i}) \leq p_i(s_i, m_{-i}) \text{ for all } m_i \in \Delta S_i,$$

i.e., s_i is a best response to m_{-i} . □

6 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 order on A .

In Osborne and Rubinstein [1994] a strategic game is defined as a sequence

$$(S_1, \dots, S_n, \succeq_1, \dots, \succeq_n),$$

where each \succeq_i is player's i **preference relation** defined on the set $S_1 \times \dots \times S_n$ of joint strategies.

In Apt et al. [2008] another modification of strategic games is considered, called a **strategic game with parametrised 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 parametrised by a joint strategy s_{-i} of his opponents. In Apt et al. [2008] only strict preferences were considered and so defined finite games with parametrised preferences were compared with the concept of **CP-nets** (Conditional Preference nets), a formalism used for representing conditional and qualitative preferences, see, e.g., Boutilier et al. [2004].

Next, in Roux et al. [2008] **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 formalise gene regulation networks and some aspects of security.

Finally, let us mention another generalisation of strategic games, called **graphical games**, introduced by Kearns et al. [2001]. These games stress the locality in taking a 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, \dots, 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 $s' \not\succeq_i s$. However, other introduced notions can be defined only for some variants. In particular, Pareto efficiency cannot be defined for strategic games with parametrised 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 3.2, carry over directly to the strategic games as defined in Osborne and Rubinstein [1994] or in Apt et al. [2008]. 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.

7 Mechanism design

Mechanism design is one of the important areas of economics. The 2007 Nobel Prize in Economics went to three economists who laid its foundations. To quote from The Economist 2007, mechanism design deals with the problem of ‘how to arrange our economic interactions so that, when everyone behaves in a self-interested manner, the result is something we all like’. So these interactions are supposed to yield desired social decisions when each agent is interested in maximising only his own utility.

In mechanism design one is interested in the ways of inducing the players to submit true information. This subject is closely related to game theory, though it focuses on other issues. In the next section we shall clarify this connection. To discuss mechanism design in more detail we need to introduce some basic concepts.

Assume a set $\{1, \dots, n\}$ of players with $n > 1$, a non-empty set of *decisions* D , and for each player i

- a non-empty set of *types* Θ_i , and
- an *initial utility function* $v_i : D \times \Theta_i \rightarrow \mathbb{R}$.

In this context a type is some private information known only to the player, for example, in the case of an auction, the player’s valuation of the items for sale.

When discussing types and sets of types we use then the same abbreviations as in Section 2. In particular, we define $\Theta := \Theta_1 \times \dots \times \Theta_n$ and for $\theta \in \Theta$ we have $(\theta_i, \theta_{-i}) = \theta$.

A *decision rule* is a function $f : \Theta \rightarrow D$. We call the tuple

$$(D, \Theta_1, \dots, \Theta_n, v_1, \dots, v_n, f)$$

a *decision problem*.

Decision problems are considered in the presence of a *central authority* who takes decisions on the basis of the information provided by the players. Given a decision problem the desired decision is obtained through the following sequence of events, where f is a given, publicly known, decision rule:

- each player i receives (becomes aware of) his type $\theta_i \in \Theta_i$,
- each player i announces to the central authority a type $\theta'_i \in \Theta_i$; this yields a joint type $\theta' := (\theta'_1, \dots, \theta'_n)$,
- the central authority then takes the decision $d := f(\theta')$ and communicates it to each player,
- the resulting initial utility for player i is then $v_i(d, \theta_i)$.

The difficulty in taking decisions through the above described sequence of events is that players are assumed to be rational, that is they want to maximise their utility. As a result they may submit false information to manipulate the outcome (decision). To better understand the notion of a decision problem consider the following two natural examples.

Example 7.1. [Sealed-bid Auction]

We consider a *sealed-bid auction* in which there is a single object for sale. Each player (bidder) simultaneously submits to the central authority his type (bid) in a sealed envelope and the object is allocated to the highest bidder.

Given a sequence $a := (a_1, \dots, a_j)$ of reals denote the least l such that $a_l = \max_{k \in \{1, \dots, j\}} a_k$ by $\operatorname{argmax} a$. Then we can model a sealed-bid auction as the following decision problem $(D, \Theta_1, \dots, \Theta_n, v_1, \dots, v_n, f)$:

- $D = \{1, \dots, n\}$,
- for all $i \in \{1, \dots, n\}$, $\Theta_i = \mathbb{R}_+$; $\theta_i \in \Theta_i$ is player's i valuation of the object,
- for all $i \in \{1, \dots, n\}$, $v_i(d, \theta_i) := (d = i)\theta$, where $d = i$ is a Boolean expression with the value 0 or 1,
- $f(\theta) := \operatorname{argmax} \theta$.

Here decision $d \in D$ indicates to which player the object is sold. Further, $f(\theta) = i$, where

$$\theta_i = \max_{j \in \{1, \dots, n\}} \theta_j \text{ and } \forall j \in \{1, \dots, i-1\} \theta_j < \theta_i.$$

So we assume that in the case of a tie the object is allocated to the highest bidder with the lowest index. □

Example 7.2. [Public project problem]

This problem deals with the task of taking a joint decision concerning construction of a *public good*,⁴ for example a bridge. Each player reports to the central authority his appreciation of the gain from the project when it takes place. If the sum of the appreciations exceeds the cost of the project, the project takes place and each player has to pay the same fraction of the cost. Otherwise the project is cancelled.

This problem corresponds to the following decision problem, where c , with $c > 0$, is the cost of the project:

- $D = \{0, 1\}$ (reflecting whether a project is cancelled or takes place),
- for all $i \in \{1, \dots, n\}$, $\Theta_i = \mathbb{R}_+$,
- for all $i \in \{1, \dots, n\}$, $v_i(d, \theta_i) := d(\theta_i - \frac{c}{n})$,
- $f(\theta) := \begin{cases} 1 & \text{if } \sum_{i=1}^n \theta_i \geq c \\ 0 & \text{otherwise.} \end{cases}$

If the project takes place ($d = 1$), $\frac{c}{n}$ is the cost share of the project for each player. □

⁴In Economics public goods are so-called not excludable and non-rival goods. To quote from Mankiw [2001]: ‘People cannot be prevented from using a public good, and one person’s enjoyment of a public good does not reduce another person’s enjoyment of it.’

Let us return now to the decision rules. We call a decision rule f **efficient** if for all $\theta \in \Theta$ and $d' \in D$

$$\sum_{i=1}^n v_i(f(\theta), \theta_i) \geq \sum_{i=1}^n v_i(d', \theta_i).$$

Intuitively, this means that for all $\theta \in \Theta$, $f(\theta)$ is a decision that maximises the **initial social welfare** from a decision d , defined by $\sum_{i=1}^n v_i(d, \theta_i)$. It is easy to check that the decision rules used in Examples 7.1 and 7.2 are efficient.

Let us return now to the subject of manipulations. As an example, consider the case of the public project problem. A player whose type (that is, appreciation of the gain from the project) exceeds the cost share $\frac{c}{n}$ should manipulate the outcome and announce the type c . This will guarantee that the project will take place, irrespective of the types announced by the other players. Analogously, a player whose type is lower than $\frac{c}{n}$ should submit the type 0 to minimise the chance that the project will take place.

To prevent such manipulations we use **taxes**, which are transfer payments between the players and central authority. This leads to a modification of the initial decision problem $(D, \Theta_1, \dots, \Theta_n, v_1, \dots, v_n, f)$ to the following one:

- the set of decisions is $D \times \mathbb{R}^n$,
- the decision rule is a function $(f, t) : \Theta \rightarrow D \times \mathbb{R}^n$, where $t : \Theta \rightarrow \mathbb{R}^n$ and $(f, t)(\theta) := (f(\theta), t(\theta))$,
- the **final utility function** of player i is the function $u_i : D \times \mathbb{R}^n \times \Theta_i \rightarrow \mathbb{R}$ defined by

$$u_i(d, t_1, \dots, t_n, \theta_i) := v_i(d, \theta_i) + t_i.$$

We call then $(D \times \mathbb{R}^n, \Theta_1, \dots, \Theta_n, u_1, \dots, u_n, (f, t))$ a **direct mechanism** and refer to t as the **tax function**.

So when the received (true) type of player i is θ_i and his announced type is θ'_i , his final utility is

$$u_i((f, t)(\theta'_i, \theta_{-i}), \theta_i) = v_i(f(\theta'_i, \theta_{-i}), \theta_i) + t_i(\theta'_i, \theta_{-i}),$$

where θ_{-i} are the types announced by the other players.

In each direct mechanism, given the vector θ of announced types, $t(\theta) := (t_1(\theta), \dots, t_n(\theta))$ is the vector of the resulting payments. If $t_i(\theta) \geq 0$, player i *receives* from the central authority $t_i(\theta)$, and if $t_i(\theta) < 0$, he *pays* to the central authority $|t_i(\theta)|$.

The following definition then captures the idea that taxes prevent manipulations. We say that a direct mechanism with tax function t is **incentive compatible** if for all $\theta \in \Theta$, $i \in \{1, \dots, n\}$ and $\theta'_i \in \Theta_i$

$$u_i((f, t)(\theta_i, \theta_{-i}), \theta_i) \geq u_i((f, t)(\theta'_i, \theta_{-i}), \theta_i).$$

Intuitively, this means that for each player i announcing one's true type (θ_i) is better than announcing another type (θ'_i). That is, false announcements, i.e., manipulations, do not pay off.

From now on we focus on specific incentive compatible direct mechanisms. Each **Groves mechanism** is a direct mechanism obtained by using a tax function $t(\cdot) := (t_1(\cdot), \dots, t_n(\cdot))$, where for all $i \in \{1, \dots, n\}$

- $t_i : \Theta \rightarrow \mathbb{R}$ is defined by $t_i(\theta) := g_i(\theta) + h_i(\theta_{-i})$, where
- $g_i(\theta) := \sum_{j \neq i} v_j(f(\theta), \theta_j)$,
- $h_i : \Theta_{-i} \rightarrow \mathbb{R}$ is an arbitrary function.

Note that, not accidentally, $v_i(f(\theta), \theta_i) + g_i(\theta)$ is simply the initial social welfare from the decision $f(\theta)$.

The importance of Groves mechanisms is then revealed by the following crucial result due to Groves [1973].

Theorem 7.3 (Groves). *Consider a decision problem $(D, \Theta_1, \dots, \Theta_n, v_1, \dots, v_n, f)$ with an efficient decision rule f . Then each Groves mechanism is incentive compatible.*

Proof. The proof is remarkably straightforward. Since f is efficient, for all $\theta \in \Theta$, $i \in \{1, \dots, n\}$ and $\theta'_i \in \Theta_i$ we have

$$\begin{aligned} u_i((f, t)(\theta_i, \theta_{-i}), \theta_i) &= \sum_{j=1}^n v_j(f(\theta_i, \theta_{-i}), \theta_j) + h_i(\theta_{-i}) \\ &\geq \sum_{j=1}^n v_j(f(\theta'_i, \theta_{-i}), \theta_j) + h_i(\theta_{-i}) \\ &= u_i((f, t)(\theta'_i, \theta_{-i}), \theta_i). \end{aligned}$$

□

When for a given direct mechanism for all $\theta \in \Theta$ we have $\sum_{i=1}^n t_i(\theta) \leq 0$, the mechanism is called **feasible**, which means that it can be realised without external financing.

Each Groves mechanism is uniquely determined by the functions h_1, \dots, h_n . A special case, called the **pivotal mechanism**, is obtained by using

$$h_i(\theta_{-i}) := - \max_{d \in D} \sum_{j \neq i} v_j(d, \theta_j).$$

So then

$$t_i(\theta) = \sum_{j \neq i} v_j(f(\theta), \theta_j) - \max_{d \in D} \sum_{j \neq i} v_j(d, \theta_j).$$

Hence for all θ and $i \in \{1, \dots, n\}$ we have $t_i(\theta) \leq 0$, which means that the pivotal mechanism is feasible and that each player needs to make the payment $|t_i(\theta)|$ to the central authority.

We noted already that the decision rules used in Examples 7.1 and 7.2 are efficient. So in each example Groves' Theorem 7.3 applies and in particular the pivotal mechanism is incentive compatible. Let us see now the details.

Re: Example 7.1 Given a sequence θ of reals we denote by θ^* its reordering from the largest to the smallest element. So for example, for $\theta = (1, 5, 4, 3, 2)$ we have $(\theta_{-2})_2^* = 3$ since $\theta_{-2} = (1, 4, 3, 2)$.

To compute the taxes in the sealed-bid auction in the case of the pivotal mechanism we use the following observation.

Note 7.4. *In the sealed-bid auction we have for the pivotal mechanism*

$$t_i(\theta) = \begin{cases} -\theta_2^* & \text{if } i = \operatorname{argsmax} \theta \\ 0 & \text{otherwise.} \end{cases}$$

Exercise 10. Provide the proof. □

So the highest bidder wins the object and pays for it the amount $\max_{j \neq i} \theta_j$.

The resulting sealed-bid auction was introduced by Vickrey [1961] and is called a **Vickrey auction**. To illustrate it suppose there are three players, A, B, and C whose true types (bids) are respectively 18, 21, and 24. When they bid truthfully the object is allocated to player C whose tax (payment) according to Note 7.4 is 21, so the second price offered. Table 1 summarises the situation.

player	type	tax	u_i
A	18	0	0
B	21	0	0
C	24	-21	3

Table 1: The pivotal mechanism for the sealed-bid auction

This explains why this auction is alternatively called a **second-price auction**. By Groves' Theorem 7.3 this auction is incentive compatible. In contrast, the **first-price auction**, in which the winner pays the price he offered (so the first, or the highest price), is not incentive compatible. Indeed, reconsider the above example. If player C submits 22 instead of his true type 24, he then wins the object but needs to pay 22 instead of 24. More formally, in the direct mechanism corresponding to the first-price auction we have

$$u_C((f, t)(18, 21, 22), 24) = 24 - 22 = 2 > 0 = u_C((f, t)(18, 21, 24), 24),$$

which contradicts incentive compatibility for the joint type (18, 21, 24). □

Re: Example 7.2 To compute the taxes in the public project problem in the case of the pivotal mechanism we use the following observation.

Note 7.5. *In the public project problem we have for the pivotal mechanism*

$$t_i(\theta) = \begin{cases} 0 & \text{if } \sum_{j \neq i} \theta_j \geq \frac{n-1}{n}c \text{ and } \sum_{j=1}^n \theta_j \geq c \\ \sum_{j \neq i} \theta_j - \frac{n-1}{n}c & \text{if } \sum_{j \neq i} \theta_j < \frac{n-1}{n}c \text{ and } \sum_{j=1}^n \theta_j \geq c \\ 0 & \text{if } \sum_{j \neq i} \theta_j \leq \frac{n-1}{n}c \text{ and } \sum_{j=1}^n \theta_j < c \\ \frac{n-1}{n}c - \sum_{j \neq i} \theta_j & \text{if } \sum_{j \neq i} \theta_j > \frac{n-1}{n}c \text{ and } \sum_{j=1}^n \theta_j < c. \end{cases}$$

Exercise 11. Provide the proof. □

To illustrate the pivotal mechanism suppose that $c = 30$ and that there are three players, A, B, and C whose true types are respectively 6, 7, and 25. When these types are announced the project takes place and Table 2 summarises the taxes that players need to pay and their final utilities. The taxes were computed using Note 7.5.

player	type	tax	u_i
A	6	0	-4
B	7	0	-3
C	25	-7	8

Table 2: The pivotal mechanism for the public project problem

Suppose now that the true types of players are respectively 4, 3 and 22 and, as before, $c = 30$. When these types are also the announced types, the project does not take place. Still, some players need to pay a tax, as Table 3 illustrates. One can show that this deficiency is shared by all feasible incentive compatible direct mechanisms for the public project, see [Mas-Collel et al., 1995, page 861-862].

player	type	tax	u_i
A	4	-5	-5
B	3	-6	-6
C	22	0	0

Table 3: The pivotal mechanism for the public project problem

8 Pre-Bayesian games

Mechanism design, as introduced in the previous section, can be explained in game-theoretic terms using pre-Bayesian games, introduced by Ashlagi et al. [2006] (see also Hyafil and Boutilier [2004] and Aghassi and Bertsimas [2006]). In strategic games, after each player selected his strategy, each player knows the payoff of *every other player*. This is not the case in pre-Bayesian games in which each player has a private type on which he can condition his strategy. This distinguishing feature of pre-Bayesian games explains why they form a class of games with *incomplete information*. Formally, they are defined as follows.

Assume a set $\{1, \dots, n\}$ of players, where $n > 1$. A *pre-Bayesian game* for n players consists of

- a non-empty set A_i of **actions**,
- a non-empty set Θ_i of **types**,
- a **payoff function** $p_i : A_1 \times \dots \times A_n \times \Theta_i \rightarrow \mathbb{R}$,

for each player i .

Let $A := A_1 \times \dots \times A_n$. In a pre-Bayesian game Nature (an external agent) moves first and provides each player i with a type $\theta_i \in \Theta_i$. Each player knows only his type. Subsequently the players simultaneously select their actions. The payoff function of each player now depends on his type, so after all players selected their actions, each player knows his payoff but does not know the payoffs of the other players. Note that given a pre-Bayesian game, every joint type $\theta \in \Theta$ uniquely determines a strategic game, to which we refer below as a θ -game.

A **strategy** for player i in a pre-Bayesian game is a function $s_i : \Theta_i \rightarrow A_i$. The previously introduced notions can be naturally adjusted to pre-Bayesian games. In particular, a joint strategy $s(\cdot) := (s_1(\cdot), \dots, s_n(\cdot))$ is called an **ex-post equilibrium** if

$$\forall \theta \in \Theta \forall i \in \{1, \dots, n\} \forall a_i \in A_i p_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i) \geq p_i(a_i, s_{-i}(\theta_{-i}), \theta_i),$$

where $s_{-i}(\theta_{-i})$ is an abbreviation for the sequence of actions $(s_j(\theta_j))_{j \neq i}$.

In turn, a strategy $s_i(\cdot)$ for player i is called **dominant** if

$$\forall \theta_i \in \Theta_i \forall a \in A p_i(s_i(\theta_i), a_{-i}, \theta_i) \geq p_i(a_i, a_{-i}, \theta_i).$$

So $s(\cdot)$ is an ex-post equilibrium iff for every joint type $\theta \in \Theta$ the sequence of actions $(s_1(\theta_1), \dots, s_n(\theta_n))$ is a Nash equilibrium in the corresponding θ -game. Further, $s_i(\cdot)$ is a dominant strategy of player i iff for every type $\theta_i \in \Theta_i$, $s_i(\theta_i)$ is a dominant strategy of player i in every (θ_i, θ_{-i}) -game.

We also have the following immediate counterpart of the Dominant Strategy Note 2.1.

Note 8.1 (Dominant Strategy). *Consider a pre-Bayesian game G . Suppose that $s(\cdot)$ is a joint strategy such that each $s_i(\cdot)$ is a dominant strategy. Then it is an ex-post equilibrium of G .*

Example 8.2. As an example of a pre-Bayesian game, suppose that

- $\Theta_1 = \{U, D\}$, $\Theta_2 = \{L, R\}$,
- $A_1 = A_2 = \{F, B\}$,

and consider the pre-Bayesian game uniquely determined by the following four θ -games. Here and below we marked the payoffs in Nash equilibria in these θ -games in bold.

		L		R	
U	F	$\mathbf{2,1}$	$2,0$	F	$\mathbf{2,1}$
	B	$0,1$	$2,1$	B	$0,0$
D	F	$3,1$	$2,0$	F	$3,0$
	B	$\mathbf{5,1}$	$4,1$	B	$\mathbf{4,1}$

This shows that the strategies $s_1(\cdot)$ and $s_2(\cdot)$ such that

$$s_1(U) := F, s_1(D) := B, s_2(L) = F, s_2(R) = B$$

form here an ex-post equilibrium. □

However, there is a crucial difference between strategic games and pre-Bayesian games.

Example 8.3. Consider the following pre-Bayesian game:

- $\Theta_1 = \{U, B\}, \Theta_2 = \{L, R\},$
- $A_1 = A_2 = \{C, D\}.$

		L		R	
U	C	$2,2$	$0,0$	C	$2,1$
	D	$3,0$	$\mathbf{1,1}$	D	$3,0$
B	C	$\mathbf{1,2}$	$3,0$	C	$\mathbf{1,1}$
	D	$0,0$	$2,1$	D	$0,0$

Even though each θ -game has a Nash equilibrium, they are so ‘positioned’ that the pre-Bayesian game has no ex-post equilibrium. Even more, if we consider a mixed extension of this game, then the situation does not change. The reason is that no new Nash equilibria are then added to the ‘constituent’ θ -games. (Indeed, each of them is solved by IESDS and hence by the IESDMS Theorem 5.1(ii) has a unique Nash equilibrium.) This shows that a mixed extension of a finite pre-Bayesian game does not need to have an ex-post equilibrium, which contrasts with the existence of Nash equilibria in mixed extensions of finite strategic games. □

To relate pre-Bayesian games to mechanism design we need one more notion. We say that a pre-Bayesian game is of a **revelation-type** if $A_i = \Theta_i$ for all $i \in \{1, \dots, n\}$. So in a revelation-type pre-Bayesian game the strategies of a player are the functions on his set of types. A strategy for player i is called then **truth-telling** if it is the identity function $\pi_i(\cdot)$ on Θ_i .

Now, as explained in Ashlagi et al. [2006] mechanism design can be viewed as an instance of the revelation-type pre-Bayesian games. Indeed, we have the following immediate, yet revealing observation.

Theorem 8.4. *Given a direct mechanism*

$$(D \times \mathbb{R}^n, \Theta_1, \dots, \Theta_n, u_1, \dots, u_n, (f, t))$$

associate with it a revelation-type pre-Bayesian game, in which each payoff function p_i is defined by

$$p_i((\theta'_i, \theta_{-i}), \theta_i) := u_i((f, t)(\theta'_i, \theta_{-i}), \theta_i).$$

Then the mechanism is incentive compatible iff in the associated pre-Bayesian game for each player truth-telling is a dominant strategy.

By Groves's Theorem 7.3 we conclude that in the pre-Bayesian game associated with a Groves mechanism, $(\pi_1(\cdot), \dots, \pi_n(\cdot))$ is a dominant strategy ex-post equilibrium.

9 Conclusions

9.1 Bibliographic remarks

Historically, the notion of an equilibrium in a strategic game occurred first in Cournot [1838] in his study of production levels of a homogeneous product in a duopoly competition. The celebrated von Neumann's Minimax Theorem proved by von Neumann [1928] establishes an existence of a Nash equilibrium in mixed strategies in two-player zero-sum games. An alternative proof of Nash's Theorem, given in Nash [1951], uses Brouwer's Fixed Point Theorem.

Ever since Nash established his celebrated theorem, a search has continued to generalise his result to a larger class of games. A motivation for this endeavour has been the 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 due to Debreu [1952], Fan [1952] and Glicksberg [1952].

Theorem 9.1. *Consider a strategic game such that*

- *each strategy set is a non-empty compact convex subset of a complete metric space,*
- *each payoff function p_i is continuous and quasi-concave in the i th argument.⁵*

Then a Nash equilibrium exists.

More recent work in this area focused on the existence of Nash equilibria in games with non-continuous payoff functions, see in particular Reny [1999] and Bich [2006].

The issue of complexity of finding a Nash equilibrium has been a long standing open problem, clarified only recently, see Daskalakis et al. [2009] for an account of these developments. Iterated elimination of strictly dominated strategies and of weakly dominated strategies was introduced by Gale [1953] and Luce and Raiffa [1957]. The corresponding results summarised in Theorems 3.2, 3.6, 5.1 and 5.4 are folklore results.

⁵Recall that the function $p_i : S \rightarrow \mathbb{R}$ is *quasi-concave in the i th argument* if the set $\{s'_i \in S_i \mid p_i(s'_i, s_{-i}) \geq p_i(s)\}$ is convex for all $s \in S$.

Apt [2004] provides uniform proofs of various order independence results, including the Order Independence Theorems 3.4 and 5.3. The computational complexity of iterated elimination of strategies has been studied starting with Knuth et al. [1988], and with Brandt et al. [2009] as a recent contribution.

There is a lot of work on formal aspects of common knowledge and of its consequences for game theory. see, e.g., Aumann [1999] and Battigalli and Bonanno [1999].

9.2 Suggestions for further reading

Strategic games form a large research area and we have barely scratched its surface. There are several other equilibria notions and various other types of 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 1990s the leading textbooks were Myerson [1991], Binmore [1991], Fudenberg and Tirole [1991] and Osborne and Rubinstein [1994].

Moving to the next decade, Osborne [2005] is an excellent, broad in its scope, undergraduate level textbook, while Peters [2008] is probably the best book on the market on the graduate level. Undeservedly less known is the short and lucid Tijs [2003]. An elementary, short introduction, focusing on the concepts, is Shoham and Leyton-Brown [2008]. In turn, Ritzberger [2002] 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, Binmore [2007] is a thoroughly revised version of Binmore [1991].

Several textbooks on microeconomics include introductory chapters on game theory, including strategic games. Two good examples are Mas-Collel et al. [1995] and Jehle and Reny [2000]. Finally, Nisan et al. [2007] 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.

References

- M. Aghassi and D. Bertsimas. Robust game theory. *Math. Program.*, 107(1-2):231–273, 2006.
- K. R. Apt. Uniform proofs of order independence for various strategy elimination procedures. *The B.E. Journal of Theoretical Economics*, 4(1), 2004. (Contributions), Article 5, 48 pages. Available from <http://xxx.lanl.gov/abs/cs.GT/0403024>.
- K. R. Apt. Order independence and rationalizability. In *Proceedings 10th Conference on Theoretical Aspects of Reasoning about Knowledge (TARK '05)*, pages 22–38. The ACM Digital Library, 2005. Available from <http://portal.acm.org>.
- 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>.

- 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.
- I. Ashlagi, D. Monderer, and M. Tennenholtz. Resource selection games with unknown number of players. In *AAMAS '06: Proceedings 5th Int. Joint Conf. on Autonomous Agents and Multiagent Systems*, pages 819–825. ACM Press, 2006.
- R. Aumann. Interactive epistemology I: Knowledge. *International Journal of Game Theory*, 28(3):263–300, 1999.
- P. Battigalli and G. Bonanno. Recent results on belief, knowledge and the epistemic foundations of game theory. *Research in Economics*, 53(2):149–225, June 1999.
- B. D. Bernheim. Rationalizable strategic behavior. *Econometrica*, 52(4):1007–1028, 1984.
- J. Bertrand. Théorie mathématique de la richesse sociale. *Journal des Savants*, 67:499–508, 1883.
- P. Bich. A constructive and elementary proof of Reny’s theorem. Cahiers de la MSE b06001, Maison des Sciences Economiques, Université Paris Panthéon-Sorbonne, Jan. 2006. Available from <http://ideas.repec.org/p/mse/wpsorb/b06001.html>.
- K. Binmore. *Playing for Real: A Text on Game Theory*. Oxford University Press, Oxford, 2007.
- K. Binmore. *Fun and Games: A Text on Game Theory*. D.C. Heath, 1991.
- 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.
- F. Brandt, M. Brill, F. A. Fischer, and P. Harrenstein. On the complexity of iterated weak dominance in constant-sum games. In *Proceedings of the 2nd Symposium on Algorithmic Game Theory*, pages 287–298, 2009.
- 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*.
- C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. *Commun. ACM*, 52(2):89–97, 2009.
- G. Debreu. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences*, 38:886–893, 1952.
- M. Dufwenberg and M. Stegeman. Existence and uniqueness of maximal reductions under iterated strict dominance. *Econometrica*, 70(5):2007–2023, 2002.

- K. Fan. Fixed point and minimax theorems in locally convex topological linear spaces. *Proceedings of the National Academy of Sciences*, 38:121–126, 1952.
- D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, Massachusetts, 1991.
- D. Gale. Theory of n-person games with perfect information. *Proceedings of the National Academy of Sciences of the United States of America*, 39:496–501, 1953.
- I. Gilboa, E. Kalai, and E. Zemel. On the order of eliminating dominated strategies. *Operation Research Letters*, 9:85–89, 1990.
- I. L. Glicksberg. A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points. *Proceedings of the American Mathematical Society*, 3:170–174, 1952.
- T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- H. Hotelling. Stability in competition. *The Economic Journal*, 39:41–57, 1929.
- N. Hyafil and C. Boutilier. Regret minimizing equilibria and mechanisms for games with strict type uncertainty. In *Proceedings of the 20th Annual Conference on Uncertainty in Artificial Intelligence (UAI-04)*, pages 268–27, Arlington, Virginia, 2004. AUAI Press.
- G. Jehle and P. Reny. *Advanced Microeconomic Theory*. Addison Wesley, Reading, Massachusetts, second edition, 2000.
- S. Kakutani. A generalization of Brouwer’s fixed point theorem. *Duke Journal of Mathematics*, 8:457–459, 1941.
- 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.
- D. E. Knuth, C. H. Papadimitriou, and J. N. Tsitsiklis. A note on strategy elimination in bimatrix games. *Operations Research Letters*, 7(3):103–107, 1988.
- R. D. Luce and H. Raiffa. *Games and Decisions*. John Wiley and Sons, New York, 1957.
- N. G. Mankiw. *Principles of Economics*. Harcourt College Publishers, Orlando, Florida, second edition, 2001.
- A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, Oxford, 1995.
- H. Moulin. *Game Theory for the Social Sciences*. NYU Press, New York, second, revised edition, 1986.
- R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, Massachusetts, 1991.

- J. F. Nash. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences, USA*, 36:48–49, 1950.
- J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.
- N. Nisan, T. Roughgarden, É. Tardos, and V. J. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- M. J. Osborne. *An Introduction to Game Theory*. Oxford University Press, Oxford, 2005.
- M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, Cambridge, Massachusetts, 1994.
- D. G. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050, 1984.
- H. Peters. *Game Theory: A Multi-Leveled Approach*. Springer, Berlin, 2008.
- P. Reny. On the existence of pure and mixed strategy Nash equilibria in discontinuous games. *Econometrica*, 67(5):1029–1056, 1999.
- K. Ritzberger. *Foundations of Non-cooperative Game Theory*. Oxford University Press, Oxford, 2002.
- R. T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, 1996.
- S. L. Roux, P. Lescanne, and R. Vestergaard. Conversion/preference games. *CoRR*, abs/0811.0071, 2008.
- Y. Shoham and K. Leyton-Brown. *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Morgan and Claypool Publishers, Princeton, 2008.
- M. Stegeman. Deleting strictly eliminating dominated strategies. Working Paper 1990/6, Department of Economics, University of North Carolina, 1990.
- The Economist 2007. Intelligent design. *The Economist*, October 18th, 2007, 2008.
- S. Tijs. *Introduction to Game Theory*. Hindustan Book Agency, Gurgaon, India, 2003.
- W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–27, 1961.
- J. von Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100:295–320, 1928.
- J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.