

Strategic Games

2.1. Introduction to Strategic Games

A *strategic game* is a static model that describes interactive situations among several players.¹ According to this model, all the players make their decisions simultaneously and independently. From the mathematical perspective, strategic games are very simple objects. They are characterized by the strategies available to the players along with their payoff functions. Even though one may think of the payoffs of the players as money, we have already seen in Chapter 1 that payoff functions may be representations of more general preferences of the players over the set of possible outcomes. These general preferences may account for other sources of utility such as unselfishness, solidarity, or personal affinities.

Throughout this book it is implicitly assumed that each player is rational in the sense that he tries to maximize his own payoff. Moreover, for a rational player, there is no bound in the complexity of the computations he can make or in the sophistication of his strategies.²

We start this chapter by formally introducing the concept of strategic game and then we move to the most widely studied solution concept in game theory: Nash equilibrium. We discuss some important classes of games, with special emphasis on zero-sum games. Later we study other solution concepts different from Nash equilibrium (Section 2.9) and, towards

¹Strategic games are also known as *games in normal form*.

²This assumption is standard in classic game theory and in most of the fields in which it is applied, especially in economics. Rubinstein (1998) offers a deep treatment of different directions in which the rationality of the agents can be bounded along with the corresponding implications.

the end of the chapter, we try to establish a bridge between equilibrium behavior and rational behavior (Section 2.12).

We denote the set of players of a game by $N := \{1, \dots, n\}$.

Definition 2.1.1. An n -player strategic game with set of players N is a pair $G := (A, u)$ whose elements are the following:

Sets of strategies: For each $i \in N$, A_i is the nonempty set of strategies of player i and $A := \prod_{i=1}^n A_i$ is the set of strategy profiles.

Payoff functions: For each $i \in N$, $u_i: A \rightarrow \mathbb{R}$ is the payoff function of player i and $u := \prod_{i=1}^n u_i$; u_i assigns, to each strategy profile $a \in A$, the payoff that player i gets if a is played.

Remark 2.1.1. In a play of G , each player $i \in N$ chooses, simultaneously and independently, a strategy $a_i \in A_i$.³ Then, each player i gets payoff $u_i(a)$. One can imagine that, before playing the game, the players can communicate among themselves; if this is the case, during the course of this communication they can only make nonbinding agreements.

Remark 2.1.2. In the kind of interactive situations that strategic games model, the following elements are implicitly involved:

- $\{A_i\}_{i \in N}$, the strategy sets of the players.
- R , the set of possible outcomes.
- A function $f: A \rightarrow R$ that assigns, to each strategy profile a , its corresponding outcome.
- $\{\succeq_i\}_{i \in N}$, the preferences of the players over the outcomes in R . They are assumed to be complete, transitive, and representable through a utility function.⁴
- $\{U_i\}_{i \in N}$, the utility functions of the players, which represent their preferences on R .

Hence, a strategic game is a “simplification” in which, for each $i \in N$ and each $a \in A$, $u_i(a) = U_i(f(a))$.

Below, we show some examples of strategic games.

Example 2.1.1. (Prisoner’s dilemma). Two suspects in a severe crime and a small robbery are put into separate cells. They are known to be guilty in the robbery but the police have no evidence for the crime. Both of them are given the chance to confess. If both confess the crime, each of them will spend 10 years in jail. If only one confesses, he will act as a witness against

³In Section 2.11 we introduce some structure to the situation in which the strategic game takes place and, in that setting, players can correlate their choices.

⁴That is, there is a countable subset of R that is order dense in R (see Theorem 1.2.3).

the other, who will spend 15 years in jail, and will receive no punishment. Finally, if no one confesses, they will be judged for the small robbery and each of them will spend 1 year in jail. Following the common terminology for this game, we refer to the confession as “defect” (D) and to no confession as “not defect” ND . Then, the prisoner’s dilemma game is a strategic game (A, u) in which

- $A_1 = A_2 = \{ND, D\}$;
- $u_1(ND, ND) = -1$, $u_1(ND, D) = -15$, $u_1(D, ND) = 0$, and $u_1(D, D) = -10$; and
- $u_2(ND, ND) = -1$, $u_2(ND, D) = 0$, $u_2(D, ND) = -15$, and $u_2(D, D) = -10$.

Figure 2.1.1 shows a more convenient representation of this game, which is the standard way to represent strategic games with finite (and small) strategy sets.

	ND	D
ND	$-1, -1$	$-15, 0$
D	$0, -15$	$-10, -10$

Figure 2.1.1. The prisoner’s dilemma.

The prisoner’s dilemma game is a classic in game theory. It has been widely used not only within theoretical settings, but also for more applied purposes in sociology and behavioral economics. The “cooperative” outcome, $(-1, -1)$, is quite good for both players; it is almost all they can get in the game. Yet, for each player, the strategy D leads to a strictly higher payoff than ND , regardless of the strategy chosen by the other player. Hence, a rational decision maker should always play D . Thus, if both players behave rationally, they get payoffs $(-10, -10)$, which are much worse than the payoffs in the “cooperative” outcome. Many real life situations can be seen as prisoner’s dilemma games. For instance, during the nuclear race between the US and the USSR in the Cold War, both countries could decide whether or not to produce nuclear weapons; in this situation, the payoffs would have the structure of those in Figure 2.1.1. \diamond

Example 2.1.2. (Cournot oligopoly (Cournot 1838)). The set N corresponds with the set of producers of a certain commodity. Each producer $i \in N$ has to choose a strategy $a_i \in [0, \infty)$ that denotes the number of units of the commodity that he will produce and bring to the market; $c_i(a_i)$ denotes the total cost player i has to face when choosing strategy a_i . The price of one unit of the commodity in the market depends on $\sum_{i \in N} a_i$ and is denoted

by $\pi(\sum_{i \in N} a_i)$. This situation can be modeled by the strategic game $G = (A, u)$, where:

- $A_i = [0, \infty)$ and
- for each $i \in N$ and each $a \in A$, $u_i(a) = \pi(\sum_{j \in N} a_j) a_i - c_i(a_i)$. \diamond

Example 2.1.3. (A first-price auction). An object is to be auctioned. The set N corresponds with the potential bidders. Player i 's valuation of the object is v_i . Suppose that $v_1 > v_2 > \dots > v_n > 0$. The rules of the auction are as follows:

- The players submit bids simultaneously.
- The object is assigned to the player with the lowest index among those who have submitted the highest bid.⁵
- The player who gets the object pays his bid.

This auction can be modeled by the following strategic game G . For each $i \in N$, $A_i = [0, \infty)$ and, for each $a \in A$,

$$u_i(a) = \begin{cases} v_i - a_i & i = \min\{j \in N : a_j = \max_{l \in N} a_l\} \\ 0 & \text{otherwise.} \end{cases} \quad \diamond$$

Example 2.1.4. (A second-price auction). In a second-price auction, the rules are the same as in a first-price auction, except that the player who gets the object pays the highest of the bids of the other players. The strategic game that corresponds with this situation is the same as in Example 2.1.3, but now

$$u_i(a) = \begin{cases} v_i - \max\{a_j : j \in N, j \neq i\} & i = \min\{j \in N : a_j = \max_{l \in N} a_l\} \\ 0 & \text{otherwise.} \end{cases} \quad \diamond$$

Note that the model of strategic games is a static one, in the sense that it assumes that players choose their strategies *simultaneously*. Nevertheless, multistage interactive situations can also be represented by strategic games (although some information is lost in the process). Why is this reasonable? Suppose that several players are involved in a multistage conflict and that each of them wants to make a strategic analysis of the situation. Before the play starts, each player must consider all the possible circumstances in which he may have to make a decision and the corresponding consequences of every possible decision. This can be made constructing the strategic form of the game. We illustrate this in the following example.

Example 2.1.5. (A multistage interactive situation). Consider the following situation with two players and three stages:

⁵This is just an example. Many other tie-breaking rules and, indeed, many other auction formats can be defined.

- Stage 1:** Player 1 chooses between left (L_1) and right (R_1).
- Stage 2:** Player 2, after observing the choice of player 1, chooses between left (L_2) and right (R_2) when player 1 has chosen left and between l_2 and r_2 when player 1 has chosen right.
- Stage 3:** If left has been chosen in the first two stages, then player 1 chooses again between left (l_1) and right (r_1) and the game finishes. Otherwise, no more choices are made and the game finishes.

The payoffs for all possible outcomes are shown in the tree depicted in Figure 2.1.2 (the first and second components of the payoff vectors are the payoffs to players 1 and 2, respectively).




Figure 2.1.2. Tree representation of the game in Example 2.1.5.

Player 1 has to make a decision when the game starts; besides, he may have to make a decision at stage 3. A strategy of player 1 is a plan for all the circumstances in which he may have to make a decision. Thus the strategy set of player 1 is:

$$A_1 = \{L_1l_1, L_1r_1, R_1l_1, R_1r_1\},$$

where, for instance, L_1r_1 means “I plan to play L_1 at stage 1 and r_1 at stage 3”. Player 2, after observing player 1’s choice at stage 1, has to make a decision at stage 2. Hence, his strategy set is:

$$A_2 = \{L_2l_2, L_2r_2, R_2l_2, R_2r_2\},$$

where, for instance, R_2l_2 means “I plan to play R_2 if player 1 has played L_1 , and to play l_2 if player 1 has played R_1 ”. Figure 2.1.3 represents the strategic game associated with this multistage situation.⁶ \diamond

⁶It may seem strange to distinguish between R_1l_1 and R_1r_1 . Indeed, both strategies are indistinguishable in the strategic form. However, as we will see in Chapter 3, the behavior in nodes that are not reached may make a difference when studying the equilibrium concepts of extensive games (through the tree/extensive representation).

	L_2l_2	L_2r_2	R_2l_2	R_2r_2
L_1l_1	8,0	8,0	1,2	1,2
L_1r_1	3,8	3,8	1,2	1,2
R_1l_1	2,1	0,4	2,1	0,4
R_1r_1	2,1	0,4	2,1	0,4

Figure 2.1.3. Representation of Example 2.1.5 as a strategic game.

2.2. Nash Equilibrium in Strategic Games

The next step is to propose solution concepts that aim to describe how rational agents should behave. The most important solution concept for strategic games is Nash equilibrium. It was introduced in Nash (1950b, 1951) and, because of its impact in both economic theory and game theory, John Nash was one of the laureates of the Nobel Prize in Economics in 1994. In Kohlberg (1990), the main idea underlying Nash equilibrium is said to be “to make a bold simplification and, instead of asking how the process of deductions might unfold, ask where its rest points may be”. In fact, a Nash equilibrium of a strategic game is simply a strategy profile such that no player gains when unilaterally deviating from it; *i.e.*, the Nash equilibrium concept searches for *rest points* of the interactive situation described by the strategic game.

Given a game $G = (A, u)$ and a strategy profile $a \in A$, let (a_{-i}, \hat{a}_i) denote the profile $(a_1, \dots, a_{i-1}, \hat{a}_i, a_{i+1}, \dots, a_n)$.

Definition 2.2.1. Let $G = (A, u)$ be a strategic game. A *Nash equilibrium* of G is a strategy profile $a^* \in A$ such that, for each $i \in N$ and each $\hat{a}_i \in A_i$,

$$u_i(a^*) \geq u_i(a_{-i}^*, \hat{a}_i).$$

Next, we study the Nash equilibria of the strategic games we presented in the previous section.

Example 2.2.1. The only Nash equilibrium of the prisoner’s dilemma is $a^* = (D, D)$. Moreover, as we have already argued, this is the rational behavior in a noncooperative environment. \diamond

Example 2.2.2. We now study the Nash equilibria in a Cournot model like the one in Example 2.1.2. We do it under the following assumptions:

- We deal with a duopoly, *i.e.*, $n = 2$.
- For each $i \in \{1, 2\}$, $c_i(a_i) = ca_i$, where $c > 0$.
- Let d be a fixed number, $d > c$. The price function is given by:

$$\pi(a_1 + a_2) = \begin{cases} d - (a_1 + a_2) & a_1 + a_2 < d \\ 0 & \text{otherwise.} \end{cases}$$

For each $i \in \{1, 2\}$ and each $a \in A$, the payoff functions of the associated strategic game are

$$u_i(a) = \begin{cases} a_i(d - a_1 - a_2 - c) & a_1 + a_2 < d \\ -a_i c & \text{otherwise.} \end{cases}$$

By definition, a Nash equilibrium of this game, sometimes called a Cournot equilibrium, is a pair $(a_1^*, a_2^*) \in A_1 \times A_2$ such that i) for each $\hat{a}_1 \in A_1$, $u_1(a_1^*, a_2^*) \geq u_1(\hat{a}_1, a_2^*)$ and ii) for each $\hat{a}_2 \in A_2$, $u_2(a_1^*, a_2^*) \geq u_2(a_1^*, \hat{a}_2)$. Now, we compute a Nash equilibrium of this game. For each $i \in \{1, 2\}$ and each $a \in A$, let $f_i(a) := a_i(d - a_1 - a_2 - c)$. Then,

$$\frac{\partial f_1}{\partial a_1}(a) = -2a_1 + d - a_2 - c \quad \text{and} \quad \frac{\partial f_2}{\partial a_2}(a) = -2a_2 + d - a_1 - c.$$

Hence,

$$\frac{\partial f_1}{\partial a_1}(a) = 0 \Leftrightarrow a_1 = \frac{d - a_2 - c}{2} \quad \text{and} \quad \frac{\partial f_2}{\partial a_2}(a) = 0 \Leftrightarrow a_2 = \frac{d - a_1 - c}{2}.$$

Note that, for each $a \in A$, $\frac{\partial^2 f_1}{\partial a_1^2}(a) = \frac{\partial^2 f_2}{\partial a_2^2}(a) = -2$.

For each $i \in N$ and each $a_{-i} \in A_{-i}$, define the set $\text{BR}_i(a_{-i}) := \{\hat{a}_i : \text{for each } \tilde{a}_i \in A_i, u_i(a_{-i}, \hat{a}_i) \geq u_i(a_{-i}, \tilde{a}_i)\}$, where BR stands for "best reply". For each $a_2 \in A_2$, if $a_2 < d - c$, $\text{BR}_1(a_2) = \frac{d - a_2 - c}{2}$ and $\text{BR}_1(a_2) = 0$ otherwise. Analogously, for each $a_1 \in A_1$, if $a_1 < d - c$, $\text{BR}_2(a_1) = \frac{d - a_1 - c}{2}$ and $\text{BR}_2(a_1) = 0$ otherwise. Hence, since $a^* = (\frac{d-c}{3}, \frac{d-c}{3})$ is the unique solution of the system

$$\begin{cases} a_1 = \frac{d - a_2 - c}{2} \\ a_2 = \frac{d - a_1 - c}{2} \end{cases}$$

a^* is the unique Nash equilibrium of this game. Note that, for each $i \in \{1, 2\}$, $u_i(a^*) = \frac{(d-c)^2}{9}$.

Observe that, if instead of two duopolists, the market consists of a single monopolist, then his payoff function would be

$$u(a) = \begin{cases} a(d - a - c) & a < d \\ -ac & \text{otherwise.} \end{cases}$$

Let $f(a) := a(d - a - c)$. Now, $f'(a) = 0$ if and only if $a = (d - c)/2$. Since, for each $a \in \mathbb{R}$, $f''(a) = -2$, then the optimal production level and cost for a monopolist are $\bar{a} = \frac{d-c}{2}$ and $u(\bar{a}) = \frac{(d-c)^2}{4}$. Therefore, the profit of the monopolist is more than the sum of the profits of the two producers in the equilibrium a^* . Moreover, since $\bar{a} < a_1^* + a_2^*$, the market price is smaller in the duopoly case. \diamond

Example 2.2.3. It is easy to check that in the first-price auction described in Example 2.1.3, the set of Nash equilibria consists of the strategy profiles $a^* \in [0, \infty)^n$ satisfying the following three conditions:

- $a_1^* \in [v_2, v_1]$,
- for each $j \in N \setminus \{1\}$, $a_j^* \leq a_1^*$, and
- there is $j \in N \setminus \{1\}$ such that $a_j^* = a_1^*$.

Note that, in a Nash equilibrium, player 1 always gets the object. \diamond

Example 2.2.4. It is easy to check that in the second-price auction described in Example 2.1.3, the strategy profile $a^* = (v_1, \dots, v_n)$ satisfies the following condition: for each $i \in N$, each $\hat{a}_i \in A_i$, and each $a_{-i} \in A_{-i}$, $u_i(a_{-i}, v_i) \geq u_i(a_{-i}, \hat{a}_i)$. This condition implies that a^* is a Nash equilibrium of this game. Note that, if a^* is played, player 1 gets the object. However, there are Nash equilibria of the second-price auction in which player 1 is not the winner. Take, for instance, $a = (0, v_1 + 1, 0, \dots, 0)$. \diamond

Example 2.2.5. Consider the strategic game in Example 2.1.5. The strategy profile (L_1l_1, R_2r_2) is the unique Nash equilibrium of that game. Observe that, in order to find a Nash equilibrium in a game in which players have finite strategy sets, given the “tabular” representation, we only have to find a cell for which the payoff of player 1 is a maximum of its column and the payoff of player 2 is the maximum of its row. \diamond

Note that, in all the strategic games that we have presented above, there is at least one Nash equilibrium. However, not every strategic game has a Nash equilibrium.

Example 2.2.6. (Matching pennies) Players 1 and 2 have to choose, simultaneously and independently, a natural number. If the sum of the chosen numbers is even, then player 1 wins; if the sum is odd, then player 2 wins. Strategically, all that matters in this game is whether a player chooses an even number (E) or an odd one (O). The strategic game which models this situation is given in Figure 2.2.1. This game does not have any Nash equilibria. \diamond

	E	O
E	$1, -1$	$-1, 1$
O	$-1, 1$	$1, -1$

Figure 2.2.1. The matching pennies game.

Next, we present Nash theorem, which provides a sufficient condition for the existence of a Nash equilibrium in a strategic game. To state and

prove Nash theorem we need some previous concepts and a classic result on correspondences: Kakutani fixed-point theorem.⁷

Let $X \subset \mathbb{R}^n$ and $Y \subset \mathbb{R}^m$. A *correspondence* F from X to Y is a map $F: X \rightarrow 2^Y$. A correspondence F is *nonempty-valued*, *closed-valued*, or *convex-valued* if, for each $x \in X$, $F(x)$ is, respectively, a nonempty, closed, or convex subset of Y . Now, we move to the definition of continuity of a correspondence. First, recall the standard definition of continuity of a function. A function $f: X \rightarrow Y$ is *continuous* if, for each sequence $\{x^k\} \subset X$ converging to $\bar{x} \in X$ and each open set $Y^* \subset Y$ such that $f(\bar{x}) \in Y^*$, there is $k_0 \in \mathbb{N}$ such that, for each $k \geq k_0$, $f(x^k) \in Y^*$.⁸ We present two generalizations of the above definition. A correspondence F is *upper hemicontinuous* if, for each sequence $\{x^k\} \subset X$ converging to $\bar{x} \in X$ and each open set $Y^* \subset Y$ such that $F(\bar{x}) \subset Y^*$, there is $k_0 \in \mathbb{N}$ such that, for each $k \geq k_0$, $F(x^k) \subset Y^*$. A correspondence F is *lower hemicontinuous* if, for each sequence $\{x^k\} \subset X$ converging to $\bar{x} \in X$ and each open set $Y^* \subset Y$ such that $F(\bar{x}) \cap Y^* \neq \emptyset$, there is $k_0 \in \mathbb{N}$ such that, for each $k \geq k_0$, $F(x^k) \cap Y^* \neq \emptyset$. Note that, if the correspondence F is a function (*i.e.*, only selects singletons in 2^Y), both upper and lower hemicontinuity properties reduce to the standard continuity of functions. Moreover, Figure 2.2.2 gives some intuition for the concepts of upper and lower hemicontinuity. Given a convergent sequence in the domain of the correspondence, upper hemicontinuity allows for “explosions” in the limit, whereas lower hemicontinuity allows for “implosions”; see $F(\bar{x})$ in Figures 2.2.2 (a) and 2.2.2 (b), respectively.

Theorem 2.2.1 (Kakutani fixed-point theorem). *Let $X \subset \mathbb{R}^n$ be a nonempty, convex, and compact set. Let $F: X \rightarrow X$ be an upper hemicontinuous, nonempty-valued, closed-valued, and convex-valued correspondence. Then, there is $\bar{x} \in X$ such that $\bar{x} \in F(\bar{x})$, *i.e.*, F has a fixed-point.*

Proof. Refer to Section 2.13. □

Definition 2.2.2. Let $G = (A, u)$ be a strategic game such that, for each $i \in N$, i) there is $m_i \in \mathbb{N}$ such that A_i is a nonempty and compact subset of \mathbb{R}^{m_i} and ii) u_i is continuous. Then, for each $i \in N$, i 's *best reply correspondence*, $\text{BR}_i: A_{-i} \rightarrow A_i$, is defined, for each $a_{-i} \in A_{-i}$, by

$$\text{BR}_i(a_{-i}) := \{a_i \in A_i : u_i(a_{-i}, a_i) = \max_{\tilde{a}_i \in A_i} u_i(a_{-i}, \tilde{a}_i)\}.$$
⁹

⁷The proof of the Kakutani fixed-point theorem is quite technical and requires some auxiliary notations and results. Since these auxiliary elements do not have much to do with game theory, we have relegated all of them, along with the proof of the Kakutani theorem, to an independent section (Section 2.13).

⁸We have defined sequential continuity instead of continuity but recall that, for metric spaces, they are equivalent.

⁹The assumptions on the strategy spaces and on the payoff functions guarantee that the best reply correspondences are well defined.

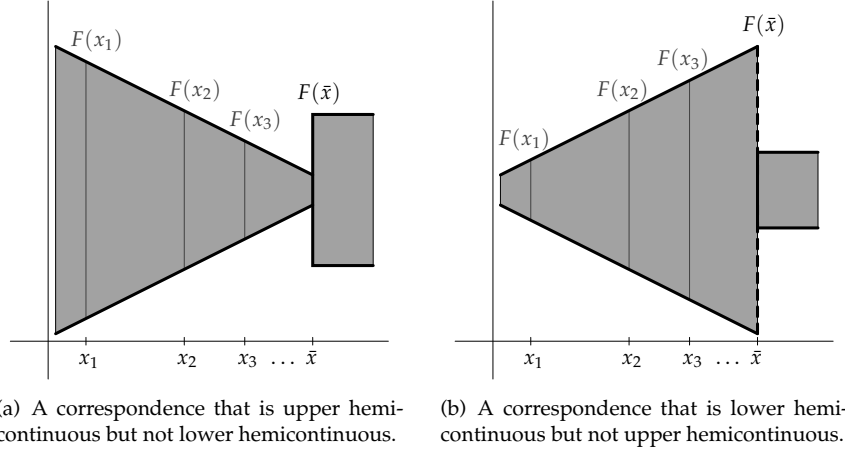


Figure 2.2.2. Two correspondences from \mathbb{R} to $2^{\mathbb{R}}$.

Let $BR: A \rightarrow A$ be defined, for each $a \in A$, as $BR(a) := \prod_{i \in N} BR_i(a_{-i})$.

Let $m \in \mathbb{N}$ and $A \subset \mathbb{R}^m$ be a convex set. A function $f: A \rightarrow \mathbb{R}$ is *quasi-concave* if, for each $r \in \mathbb{R}$, the set $\{a \in A : f(a) \geq r\}$ is convex or, equivalently, if, for each $a, \tilde{a} \in A$ and each $\alpha \in [0, 1]$, $f(\alpha a + (1 - \alpha)\tilde{a}) \geq \min\{f(a), f(\tilde{a})\}$. Quasi-concavity implies concavity, which requires $f(\alpha a + (1 - \alpha)\tilde{a}) \geq \alpha f(a) + (1 - \alpha)f(\tilde{a})$, but it is substantially more general. For instance, the convex function $f(x) = e^x$ is quasi-concave and so it is any monotone function from \mathbb{R} to \mathbb{R} .

Proposition 2.2.2. *Let $G = (A, u)$ be a strategic game such that, for each $i \in N$,*

- i) A_i is a nonempty and compact subset of \mathbb{R}^{m_i} .
- ii) u_i is continuous.
- iii) For each a_{-i} , $u_i(a_{-i}, \cdot)$ is quasi-concave on A_i .

Then, for each $i \in N$, BR_i is an upper hemicontinuous, nonempty-valued, closed-valued, and convex-valued correspondence. Therefore, BR also satisfies the previous properties.

Proof. Let $i \in N$.

Nonempty-valuedness: Obvious, since every continuous function defined on a compact set reaches a maximum.

Closed-valuedness: Also straightforward by the continuity of the payoff functions and the compactness of the sets of strategies.

Convex-valuedness: Let $a_{-i} \in A_{-i}$ and $\tilde{a}_i \in \text{BR}_i(a_{-i})$. Let $r := u_i(a_{-i}, \tilde{a}_i)$. Then, $\text{BR}_i(a_{-i}) = \{a_i \in A_i : u_i(a_{-i}, a_i) \geq r\}$. Convex-valuedness is implied by the quasi-concavity of u .

Upper hemicontinuity: Suppose BR_i is not upper hemicontinuous. Then, there are a sequence $\{a^k\} \subset A_{-i}$ converging to $\bar{a} \in A_{-i}$ and an open set $B^* \subset A_i$ with $\text{BR}_i(\bar{a}) \subset B^*$, satisfying that, for each $k_0 \in \mathbb{N}$, there is $k \geq k_0$ such that $\text{BR}_i(a^k) \not\subset B^*$. The latter implies that there is a sequence $\{\tilde{a}^m\} \subset A_i$ such that, for each $m \in \mathbb{N}$, $\tilde{a}^m \in \text{BR}_i(a^m) \setminus B^*$. Since A_i is compact, $\{\tilde{a}^m\}$ has a convergent subsequence. Assume, without loss of generality, that $\{\tilde{a}^m\}$ itself converges and let $\hat{a} \in A_i$ be its limit. Since B^* is an open set, then $A_i \setminus B^*$ is closed. Hence, $\hat{a} \in A_i \setminus B^*$ and, therefore, $\hat{a} \notin \text{BR}_i(\bar{a})$. For each $m \in \mathbb{N}$ and each $a \in A_i$, $u_i(a^m, \tilde{a}^m) \geq u_i(a^m, a)$. Letting $m \rightarrow \infty$ and using the continuity of u_i , we have that, for each $a \in A_i$, $u_i(\bar{a}, \hat{a}) \geq u_i(\bar{a}, a)$. Hence, $\hat{a} \in \text{BR}_i(\bar{a})$ and we have a contradiction.

The result for BR is now straightforward. \square

Exercise 2.6 asks you to show that, under the assumptions of Proposition 2.2.2, the BR_i functions need not be lower hemicontinuous.

We present below the main result of this section, whose proof is an immediate consequence of Kakutani theorem and Proposition 2.2.2.

Theorem 2.2.3 (Nash theorem¹⁰). *Let $G = (A, u)$ be a strategic game such that, for each $i \in N$,*

- i) A_i is a nonempty, convex, and compact subset of \mathbb{R}^{m_i} .
- ii) u_i is continuous.
- iii) For each a_{-i} , $u_i(a_{-i}, \cdot)$ is quasi-concave on A_i .¹¹

Then, the game G has, at least, one Nash equilibrium.

Proof. If a is a fixed point of the correspondence $\text{BR}: A \rightarrow A$, then a is a Nash equilibrium of G . By Proposition 2.2.2, BR satisfies the conditions of Kakutani theorem (recall that now we have further assumed that the A_i sets are convex). Hence, BR has a fixed point. \square

It is now an easy exercise to show that the set of Nash equilibria is closed (Exercise 2.5 asks the reader to prove this). Moreover, the role of

¹⁰In reality, what we present here is a generalization of the original Nash theorem. This version was proved in Rosen (1965). We present the Nash theorem in its original form in Theorem 2.4.1 (Section 2.4).

¹¹In this context, the quasi-concavity of the function $u_i(a_{-i}, \cdot)$ can be phrased as follows: provided that all players different from player i are playing according to a_{-i} , we have that, for each $K \in \mathbb{R}$, if two strategies give at least payoff K to player i , then so does any convex combination of them.

quasi-concavity in the above theorem is to ensure that any convex combination of best responses is a best response, *i.e.*, BR is convex-valued.

The reader can easily verify that we cannot apply Nash theorem to any of the examples treated above. However, as we will see during this chapter, the sufficient condition provided by Nash theorem applies to a wide class of games.

2.3. Two-Player Zero-Sum Games

Many of the early works in game theory concentrated on a very special class of strategic games, namely *two-player zero-sum games*. They are characterized by the fact that the two players involved in the game have completely opposite interests.

Definition 2.3.1. A *two-player zero-sum game* is a strategic game given by $G = (\{A_1, A_2\}, \{u_1, u_2\})$, where, for each strategy profile $(a_1, a_2) \in A_1 \times A_2$, $u_1(a_1, a_2) + u_2(a_1, a_2) = 0$.

Note that in the previous definition we have imposed no restriction on the sets A_1 and A_2 . We discuss the special case in which A_1 and A_2 are finite in Section 2.6; there we present and prove one of the classic results in game theory, the minimax theorem.

To characterize a two-player zero-sum game it is sufficient to give the payoff function of one of the players. Usually, player 1's payoff function is provided. Thus, when dealing with two-player zero-sum games we say that G is the triple (A_1, A_2, u_1) . We assume that u_1 is bounded on $A_1 \times A_2$.

As mentioned above, the class of two-player zero-sum games was the first to be studied by game theorists and they represent situations in which players have totally opposite interests: whenever one player prefers (a_1, a_2) to (\hat{a}_1, \hat{a}_2) , then the other player prefers (\hat{a}_1, \hat{a}_2) to (a_1, a_2) . These games were initially analyzed by John von Neumann, who introduced the following concepts.

Definition 2.3.2. Let $G = (A_1, A_2, u_1)$ be a two-player zero-sum game.

Lower value: Let $\underline{\Delta}: A_1 \rightarrow \mathbb{R}$ be defined, for each $a_1 \in A_1$, by $\underline{\Delta}(a_1) := \inf_{a_2 \in A_2} u_1(a_1, a_2)$, *i.e.*, the worst payoff that player 1 can get if he plays a_1 .

The *lower value* of G , denoted by $\underline{\lambda}$, is given by

$$\underline{\lambda} := \sup_{a_1 \in A_1} \underline{\Delta}(a_1).$$

It provides the payoff that player 1 can guarantee for himself in G (or, at least, get arbitrarily close to it).

Upper value: Let $\bar{\Lambda}: A_2 \rightarrow \mathbb{R}$ be defined, for each $a_2 \in A_2$, by $\bar{\Lambda}(a_2) := \sup_{a_1 \in A_1} u_1(a_1, a_2)$, i.e., the maximum loss that player 2 may suffer when he plays a_2 .

The *upper value* of G , denoted by $\bar{\lambda}$, is given by

$$\bar{\lambda} := \inf_{a_2 \in A_2} \bar{\Lambda}(a_2).$$

It provides the supremum of losses that player 2 may suffer in G , i.e., the maximum payoff that player 1 may get in G (or, at least, get arbitrarily close to it).

Note that $\underline{\lambda} \leq \bar{\lambda}$ since, for each $a_1 \in A_1$ and each $a_2 \in A_2$,

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

Definition 2.3.3. A two-player zero-sum game $G = (A_1, A_2, u_1)$ is said to be *strictly determined* or to have a *value* if its lower value and its upper value coincide, i.e., if $\underline{\lambda} = \bar{\lambda}$. In such a case, $V := \underline{\lambda} = \bar{\lambda}$ is the *value* of the game.

Definition 2.3.4. Let $G = (A_1, A_2, u_1)$ be a two-player zero-sum game with value V .

- i) A strategy $a_1 \in A_1$ is *optimal for player 1* if $V = \underline{\Delta}(a_1)$.
- ii) A strategy $a_2 \in A_2$ is *optimal for player 2* if $V = \bar{\Lambda}(a_2)$.

Under the existence of optimal strategies, the value of a zero-sum game is the payoff that player 1 can guarantee for himself; similarly, it is the opposite of the payoff that player 2 can guarantee for himself. In those two-player zero-sum games that do not have a value, the situation is not strictly determined in the sense that it is not clear how the payoff $\bar{\lambda} - \underline{\lambda}$ is going to be allocated.

The following examples illustrate the different possibilities that may arise regarding the value and the optimal strategies of a two-player zero-sum game.

Example 2.3.1. (A finite two-player zero-sum game that is strictly determined). Consider the two-player zero-sum game in Figure 2.3.1. Clearly,

	L_2	R_2
L_1	2	2
R_1	1	3

Figure 2.3.1. A strictly determined game.

$\underline{\Delta}(L_1) = 2$ and $\underline{\Delta}(R_1) = 1$, so $\underline{\lambda} = 2$. Besides, $\bar{\Lambda}(L_2) = 2$ and $\bar{\Lambda}(R_2) = 3$, so

$\bar{\lambda} = 2$. Hence, the value of this game is 2 and L_1 and L_2 are optimal strategies for players 1 and 2, respectively. The latter is a general feature of finite two-player zero-sum games (those in which the strategy sets of the players are finite): if they have a value, then both players have optimal strategies. \diamond

Example 2.3.2. (The matching pennies revisited). In this game, introduced in Example 2.2.6, $\underline{\lambda} = -1$ and $\bar{\lambda} = 1$. Hence, it is not strictly determined. However, consider the following extension of the game. Each player, instead of selecting either E or O , can randomize over the two choices. Denote by a_1 and a_2 the probability that players 1 and 2 (respectively) choose E . Consider now the infinite two-player zero-sum game $([0, 1], [0, 1], u_1)$ where, for each pair $(a_1, a_2) \in [0, 1] \times [0, 1]$, $u_1(a_1, a_2) = 4a_1a_2 - 2a_1 - 2a_2 + 1$. This new game we have defined is just the *mixed extension* of the matching pennies game; we formally define this concept in the next section. We now show that the game $(\{[0, 1], [0, 1]\}, u_1)$ defined above is strictly determined. For each $a_1 \in [0, 1]$,

$$\underline{\Delta}(a_1) = \inf_{a_2 \in [0, 1]} ((4a_1 - 2)a_2 - 2a_1 + 1) = \begin{cases} 2a_1 - 1 & a_1 < 1/2 \\ 0 & a_1 = 1/2 \\ 1 - 2a_1 & a_1 > 1/2 \end{cases}$$

and, for each $a_2 \in [0, 1]$,

$$\bar{\Delta}(a_2) = \sup_{a_1 \in [0, 1]} ((4a_2 - 2)a_1 - 2a_2 + 1) = \begin{cases} 1 - 2a_2 & a_2 < 1/2 \\ 0 & a_2 = 1/2 \\ 2a_2 - 1 & a_2 > 1/2. \end{cases}$$

Hence, $\underline{\lambda} = \bar{\lambda} = 0$, that is, the value of the game is 0. Moreover, $a_1 = a_2 = 1/2$ is the only optimal strategy of either player; the best thing a player can do in this game is to be completely unpredictable. \diamond

Example 2.3.3. (An infinite two-player zero-sum game that is not strictly determined). Take the two-player zero-sum game $([0, 1], [0, 1], u_1)$, where, for each $(a_1, a_2) \in [0, 1] \times [0, 1]$, $u_1(a_1, a_2) = \frac{1}{1+(a_1-a_2)^2}$. For each $a_1 \in [0, 1]$,

$$\underline{\Delta}(a_1) = \inf_{a_2 \in [0, 1]} \frac{1}{1+(a_1-a_2)^2} = \begin{cases} \frac{1}{1+(a_1-1)^2} & a_1 \leq 1/2 \\ \frac{1}{1+a_1^2} & a_1 \geq 1/2 \end{cases}$$

and $\underline{\lambda} = \underline{\Delta}(1/2) = 4/5$. For each $a_2 \in [0, 1]$,

$$\bar{\Delta}(a_2) = \sup_{a_1 \in [0, 1]} \frac{1}{1+(a_1-a_2)^2} = 1.$$

Hence $\underline{\lambda} < \bar{\lambda}$ and the game is not strictly determined. \diamond

Example 2.3.4. (An infinite two-player zero-sum game with a value and with optimal strategies only for one player). Consider the two-player zero-sum game $((0,1), (0,1), u_1)$, where, for each pair $(a_1, a_2) \in (0,1) \times (0,1)$, $u_1(a_1, a_2) = a_1 a_2$. It is easy to check that, for each $a_1 \in (0,1)$, $\Delta(a_1) = 0$. Hence, $\lambda = 0$ and, for each $a_2 \in (0,1)$, $\bar{\Lambda}(a_2) = a_2$. Thus, $\bar{\lambda} = 0$. Therefore, the game is strictly determined, its value is zero, and the set of optimal strategies of player 1 is the whole interval $(0,1)$, but player 2 does not have optimal strategies. \diamond

Example 2.3.5. (An infinite two-player zero-sum game with a value but without optimal strategies for any of the players). Consider the two-player zero-sum game $((0,1), (1,2), u_1)$, where, for each pair $(a_1, a_2) \in (0,1) \times (1,2)$, $u_1(a_1, a_2) = a_1 a_2$. It is easy to check that, for each $a_1 \in (0,1)$, $\Delta(a_1) = a_1$. Hence, $\lambda = 1$ and, for each $a_2 \in (1,2)$, $\bar{\Lambda}(a_2) = a_2$. Thus, $\bar{\lambda} = 1$. Therefore, the game is strictly determined with value one, but players do not have optimal strategies. \diamond

So far, we have not discussed Nash equilibria of two-player zero-sum games, even though they are strategic games. We may wonder what is the connection between von Neumann's theory and Nash's theory, *i.e.*, what is the relation between the Nash equilibrium concept and the profiles of optimal strategies in two-player zero-sum games. We address this point below.

Let (A_1, A_2, u_1) be a two-player zero-sum game. A Nash equilibrium of G is a strategy profile $(a_1^*, a_2^*) \in A_1 \times A_2$ such that, for each $\hat{a}_1 \in A_1$ and each $\hat{a}_2 \in A_2$, $u_1(a_1^*, a_2^*) \geq u_1(\hat{a}_1, a_2^*)$ and $-u_1(a_1^*, a_2^*) \geq -u_1(a_1^*, \hat{a}_2)$; equivalently,

$$(2.3.1) \quad u_1(\hat{a}_1, a_2^*) \leq u_1(a_1^*, a_2^*) \leq u_1(a_1^*, \hat{a}_2).$$

A strategy profile $(a_1^*, a_2^*) \in A_1 \times A_2$ satisfying Eq. (2.3.1) is said to be a *saddle point* of u_1 . Hence, in two-player zero-sum games, Nash equilibria and saddle points of the payoff function of player 1 represent the same concept. We now present two propositions illustrating that Nash's theory for strategic games is a generalization of von Neumann's theory for two-player zero-sum games.

Proposition 2.3.1. *Let $G = (A_1, A_2, u_1)$ be a two-player zero-sum game and let $(a_1^*, a_2^*) \in A_1 \times A_2$ be a Nash equilibrium of G . Then:*

- i) G is strictly determined.
- ii) a_1^* is optimal for player 1 and a_2^* is optimal for player 2.
- iii) $V = u_1(a_1^*, a_2^*)$.

Proof. By Eq. (2.3.1) we have:

- $\underline{\lambda} = \sup_{\hat{a}_1 \in A_1} \underline{\Lambda}(\hat{a}_1) \geq \underline{\Lambda}(a_1^*) = \inf_{\hat{a}_2 \in A_2} u_1(a_1^*, \hat{a}_2) \geq u_1(a_1^*, a_2^*),$
- $u_1(a_1^*, a_2^*) \geq \sup_{\hat{a}_1 \in A_1} u_1(\hat{a}_1, a_2^*) = \bar{\Lambda}(a_2^*) \geq \inf_{\hat{a}_2 \in A_2} \bar{\Lambda}(\hat{a}_2) = \bar{\lambda}.$

Since we know that $\bar{\lambda} \geq \underline{\lambda}$, all the inequalities have to be equalities, which implies i), ii), and iii). \square

Proposition 2.3.2. *Let $G = (A_1, A_2, u_1)$ be a two-player zero-sum game. Let G be strictly determined and let $a_1 \in A_1$ and $a_2 \in A_2$ be optimal strategies of players 1 and 2, respectively. Then (a_1, a_2) is a Nash equilibrium of G and $V = u_1(a_1, a_2)$.*

Proof. Since a_1 and a_2 are optimal strategies we have that, for each $\hat{a}_1 \in A_1$ and each $\hat{a}_2 \in A_2$,

$$u_1(\hat{a}_1, a_2) \leq \bar{\Lambda}(a_2) = V = \underline{\Lambda}(a_1) \leq u_1(a_1, \hat{a}_2).$$

Taking $\hat{a}_1 = a_1$ and $\hat{a}_2 = a_2$, we have that $V = u_1(a_1, a_2)$. \square

Remark 2.3.1. In view of the propositions above, if (a_1^*, a_2^*) and (a_1, a_2) are Nash equilibria of a two-player zero-sum game G , then (a_1^*, a_2) and (a_1, a_2^*) are also Nash equilibria of G and, moreover,

$$u_1(a_1^*, a_2^*) = u_1(a_1, a_2) = u_1(a_1, a_2^*) = u_1(a_1^*, a_2).$$

The above observation is not true for every two-player game; see, for instance, the battle of the sexes in Example 2.5.1 (Section 2.5).

Remark 2.3.2. A strategic game $G = (\{A_1, A_2\}, u_1, u_2)$ is a two-player *constant-sum game* if there is $K \in \mathbb{R}$ such that, for each strategy profile $(a_1, a_2) \in A_1 \times A_2$, $u_1(a_1, a_2) + u_2(a_1, a_2) = K$. A zero-sum game is a constant-sum game ($K = 0$). However, from the strategic point of view, studying G is the same as studying the zero-sum game $\bar{G} := (A_1, A_2, u_1)$ because $(a_1, a_2) \in A_1 \times A_2$ is a Nash equilibrium of G if and only if it is a Nash equilibrium of \bar{G} . Therefore, the approach developed in this section can be readily extended to account for two-player constant-sum games.

2.4. Mixed Strategies in Finite Games

The main focus of the rest of the chapter is on finite games.

Definition 2.4.1. Let $G = (A, u)$ be a strategic game. We say that G is a *finite game* if, for each $i \in N$, $|A_i| < \infty$.

Since the sets of strategies in a finite game are not convex sets, Nash theorem cannot be applied to them. Moreover, we have already seen a finite game without Nash equilibria: the matching pennies (Example 2.2.6). However, there is a “theoretical trick” that allows us to extend the game and to guarantee the existence of Nash equilibria of the extended version of

every finite game: this trick consists of enlarging the strategic possibilities of the players and allowing them to choose not only the strategies they initially had (henceforth called *pure strategies*), but also the lotteries over their (finite) sets of pure strategies. This extension of the original game is called its *mixed extension*, and the strategies of the players in the mixed extension are called *mixed strategies*. Actually, we have already studied the mixed extension of the matching pennies game (see Example 2.3.2).

Although we have referred above to the mixed extension of a game as a “theoretical trick”, mixed strategies are natural in many practical situations. We briefly discuss this point in the following example (and also later in this book), in which we informally introduce the mixed extension of a strategic game. After the example, we provide the formal definition. For a much more detailed discussion on the importance, interpretations, and appropriateness of mixed strategies, Osborne and Rubinstein (1994, Section 3.2) can be consulted.

Example 2.4.1. Consider the matching pennies game (see Example 2.2.6). Suppose that the players, besides choosing E or O , can choose a lottery L that selects E with probability $1/2$ and O with probability $1/2$ (think, for instance, of a coin toss). The players have von Neumann and Morgenstern utility functions, *i.e.*, their payoff functions can be extended to the set of mixed strategy profiles computing the mathematical expectation (Definition 1.3.4). Figure 2.4.1 represents the new game we are considering. Note

	E	O	L
E	1, -1	-1, 1	0, 0
O	-1, 1	1, -1	0, 0
L	0, 0	0, 0	0, 0

Figure 2.4.1. Matching pennies allowing for a coin toss.

that the payoff functions have been extended, taking into account that players choose their lotteries independently (we are in a strategic game). For instance:

$$u_1(L, L) = \frac{1}{4}u_1(E, E) + \frac{1}{4}u_1(E, O) + \frac{1}{4}u_1(O, E) + \frac{1}{4}u_1(O, O) = 0.$$

Observe that this game has a Nash equilibrium: (L, L) . The mixed extension of the matching pennies is a new strategic game in which players can choose not only L , but also any other lottery over $\{E, O\}$. It is easy to check that the only Nash equilibrium of the mixed extension of the matching pennies is (L, L) . One interpretation of this can be the following. In a matching pennies situation it is very important for both players that each one does not have any information of what will be his final choice (E or O). In order

to get this, it would be optimal for every player if the player himself does not know what his final choice will be: this reasoning justifies lottery L . \diamond

Definition 2.4.2. Let $G = (A, u)$ be a finite game. The *mixed extension* of G is the strategic game $E(G) := (S, u)$, whose elements are the following:

Sets of (mixed) strategies: For each $i \in N$, $S_i := \Delta A_i$ and $S := \prod_{i \in N} S_i$. For each $s \in S$ and each $a \in A$, let $s(a) := s_1(a_1) \cdot \dots \cdot s_n(a_n)$.¹²

Payoff functions: For each $s \in S$, $u_i(s) := \sum_{a \in A} u_i(a)s(a)$ and $u := \prod_{i=1}^n u_i$.

Remark 2.4.1. The mixed extension of a finite game only makes sense if the players have preferences over the set of lotteries on R (the set of possible outcomes) and their utility functions are von Neumann and Morgenstern utility functions (see Remark 2.1.2 for more details).

Remark 2.4.2. $E(G)$ is indeed an extension of G , in the sense that, for each player $i \in N$, each element of A_i (pure strategy) can be uniquely identified with an element of S_i (mixed strategy). In this sense, we can write $A_i \subset S_i$. Also the payoff functions in $E(G)$ are extensions of the payoff functions in G .

Remark 2.4.3. Let $m_i := |A_i|$. Then, for each $i \in N$, S_i can be identified with the simplex of \mathbb{R}^{m_i} given by:

$$\{s_i \in \mathbb{R}^{m_i} : \sum_{k=1}^{m_i} s_{i,k} = 1 \text{ and, for each } k \in \{1, \dots, m_i\}, s_{i,k} \geq 0\}.$$

Depending on the context, this vector notation for mixed strategies may be more convenient than the representation as functions and, because of this, both notations are used in the book.

Note that the mixed extension of a finite game satisfies the conditions of the Nash theorem. Hence, the mixed extension of a finite game always has, at least, one Nash equilibrium. Actually, this was the statement proved by Nash in his original paper. We formally write this result for the sake of completeness.

Theorem 2.4.1. *Let $G = (A, u)$ be a finite strategic game. Then, the mixed extension of G , $E(G)$, has, at least, one Nash equilibrium.*

Proof. Easily follows from Nash theorem. \square

¹²Since $s_i(a_i)$ is the probability that player i chooses strategy a_i , $s(a)$ is the probability that the profile a is played.

To finish this section, we give some definitions and basic results regarding finite games and their mixed extensions.

Definition 2.4.3. Let G be a finite game and $E(G)$ its mixed extension. Let $s_i \in S_i$ be a (mixed) strategy for player i and $s \in S$ a (mixed) strategy profile.

- i) The *support* of s_i is the set $\mathcal{S}(s_i) := \{a_i \in A_i : s_i(a_i) > 0\}$. Analogously, the support of s is the set $\mathcal{S}(s) := \prod_{i \in N} \mathcal{S}(s_i) = \{a \in A : s(a) > 0\}$.
- ii) We say that s_i is *completely mixed* if $\mathcal{S}(s_i) = A_i$. Analogously, we say that s is completely mixed if $\mathcal{S}(s) = A$ or, equivalently, if, for each $i \in N$, s_i is completely mixed.
- iii) The set of *pure best replies* of player i to s_{-i} is given by $\text{PBR}_i(s_{-i}) := \{a_i \in A_i : \text{for each } \hat{a}_i \in A_i, u_i(s_{-i}, a_i) \geq u_i(s_{-i}, \hat{a}_i)\}$. Let $\text{PBR}(s) := \prod_{i \in N} \text{PBR}_i(s_{-i})$.

Proposition 2.4.2. Let G be a finite game and $E(G)$ its mixed extension. Then, for each $i \in N$, each $s_i \in S_i$, and each $s \in S$, the following properties hold:

- i) $s_i \in \text{BR}_i(s_{-i})$ if and only if $\mathcal{S}(s_i) \subset \text{PBR}_i(s_{-i})$.
- ii) s is a Nash equilibrium of $E(G)$ if and only if $\mathcal{S}(s) \subset \text{PBR}(s)$.
- iii) s is a Nash equilibrium of $E(G)$ if and only if for each $\hat{a}_i \in A_i$ and each $i \in N$, $u_i(s) \geq u_i(s_{-i}, \hat{a}_i)$.

Proof. It is clear that $u_i(s) = \sum_{a_i \in A_i} u_i(s_{-i}, a_i) s_i(a_i)$. From this fact, the proposition immediately follows. \square

2.5. Bimatrix Games

In this section we discuss finite two-player games, which can be easily represented using matrix notation. Because of this, the notations in this section and in the next two are somewhat independent from the notations used in the rest of the book. The finite sets of strategies of the players are given by $L := \{1, \dots, l\}$ for player 1 and $M := \{1, \dots, m\}$ for player 2. Since $N = \{1, 2\}$, there is no need to use letters i and j to index the elements of N and, hence, in these sections we use them to index the strategy sets of player 1 and player 2, respectively. We denote matrices by capital calligraphic letters such as \mathcal{A} and \mathcal{B} with entries a_{ij} and b_{ij} . By $\mathcal{M}_{l \times m}$ we denote the set of all $l \times m$ matrices with real entries. The transpose of a matrix \mathcal{A} is denoted by \mathcal{A}^t . Finally, we use the notations a_i and a_j for the i -th row and the j -th column of \mathcal{A} , respectively.