Chapter 2
Strategic Dominance

2.1 Prisoner's Dilemma

Let us start with perhaps the most famous example in Game Theory, the *Prisoner's Dilemma*.[1] This is a two-player normal-form (simultaneous move) game. Two suspects, A (a woman) and B (a man), are taken into custody. The district attorney is sure that they have committed a crime together but does not have enough evidence. They are interrogated in separate rooms and cannot communicate. Both A and B have the same options: defection from the accomplice to reveal the evidence or cooperation with the other to keep silent. What will they do?

To predict a strategic outcome of this situation, we need to understand the suspects' objectives. The suspects should want to minimize possible punishment or penalties. The DA tells the suspects separately as follows."If you defect from your accomplice and tell me where the evidence is, I will release you right away. But if the other suspect defects first, I will prosecute only you and with the evidence you will be sentenced for a long time."

A careful suspect might ask, "what happens if both of us defect?" In this case the plea bargain does not work, and both suspects get prosecuted, sharing the guilt. Yet another possibility is that no one defects, in which case we assume that the prosecutor still goes for the trial. But this trial is not as scary as the one with strong evidence after defection.

Let us formulate the payoffs for the suspects in terms of the lengths of sentence of the four possible outcomes. If player A defects and player B does not, A gets 0 years of sentence with the plea bargain and B gets 5 years. Similarly, if player B defects and player A does not, then A gets 5 years and B gets 0. Since a shorter sentence is better, we should assign a smaller number to 5 years of sentence than 0. As an example, let us define the payoff as -1 times the term of sentence. Then no prosecution means a payoff value of 0, and a 5 year sentence is a payoff of -5.

[1] Alternatively, it is written as "Prisoners' Dilemma". This book follows Luce and Raiffa [4]. Also notice that the players are not yet prisoners.

T. Fujiwara-Greve, *Non-Cooperative Game Theory*, Monographs in Mathematical Economics 1, DOI 10.1007/978-4-431-55645-9_2

Table 2.1 Player A's payoff matrix

A's strategy\B's strategy	Defect	Cooperate
Defect	-3	0
Cooperate	-5	-1

If both suspects defect, no plea bargain holds, and the prosecutor gets strong evidence of the crime. But the suspects share the guilt, so let's assume that both get a 3 year sentence, and the payoff is -3 each. If both cooperate with each other and keep silent, although they still will be prosecuted, the evidence is weak. Assume that, in this case, they both get a 1 year sentence, and a payoff of -1.

Now the game is specified. If you were player A (or B), what would you do to maximize your payoff (or minimize the term of sentence)?

2.2 Strict Dominance

Although the game is completely specified in the previous section by words, it is convenient to introduce a mathematical formulation for later generalizations. Denote the set of strategies of a player by adding a subscript for the name of the player to the letter S (which stands for Strategy). That is, player A's set of strategies is denoted as $S_A = \{C, D\}$ (where C stands for cooperation and D stands for defection), and player B's set of strategies is denoted as $S_B = \{C, D\}$ as well. Player A's payoff (or term of sentence) cannot be determined by her strategy alone. Rather, it is determined by a combination of both players' strategies. The set of all combinations of both players' strategies is the Cartesian product $S_A \times S_B$ of the two strategy sets.[2] Hence, player A's payoff is computed by a *payoff function* $u_A : S_A \times S_B \to \Re$. To be concrete, for the current Prisoner's Dilemma game, u_A is specified as follows:

$$u_A(D, D) = -3, u_A(C, D) = -5, u_A(D, C) = 0, u_A(C, C) = -1. \tag{2.1}$$

(Note that the first coordinate in the parentheses is player A's strategy, and the second coordinate is player B's.)

Let us consider player A's payoff maximization. Table 2.1 is a *matrix representation* of the payoff function (2.1), where the rows are her strategies and the columns are player B's strategies. It is easy to see from this table that, if B defects, it is better for A to also defect, because defection gives her a 3 year sentence while cooperation (with B) gives her 5 years. If player B cooperates with player A, still it is better for A to defect, because defection leads to release, but cooperation means a 1 year sentence.

[2] If there is a natural order among players in the description of the game, the product is usually taken that way. In this case the alphabetical order of the names suggests that A is the first coordinate. One can equivalently formulate the product and the game with player B as the first player.

In summary, strategy D gives a strictly greater payoff than that of strategy C, regardless of the opponent's strategy choice. Hence, the payoff-maximizing player A would not choose strategy C.

This logic can be generalized. Consider an arbitrary n-player normal-form game. Let $\{1, 2, \ldots, n\}$ be the set of players with the generic element i, S_i be the set of strategies of player i, and $u_i : S_1 \times S_2 \times \cdots \times S_n \to \Re$ be the payoff function of player i. For convenience we also introduce the following notation:

$$S := S_1 \times S_2 \times \cdots \times S_n,$$
$$S_{-i} := S_1 \times \cdots \times S_{i-1} \times S_{i+1} \times \cdots \times S_n.$$

S is the set of all *strategy combinations* (*strategy profiles*) of all players. S_{-i} is the set of strategy combinations of players other than i.

Definition 2.1 Given a player i, a strategy $s_i \in S_i$ is *strictly dominated*[3] by another strategy $s_i' \in S_i$ if, for any $s_{-i} \in S_{-i}$,

$$u_i(s_i, s_{-i}) < u_i(s_i', s_{-i}).$$

That is, for any possible strategy combination s_{-i} by the other players, s_i gives strictly less payoff value than s_i' does. Hence, at least s_i' is better than s_i in terms of payoff maximization. In general, we postulate that **a payoff-maximizing (rational) player would not use a strategy that is strictly dominated by some other strategy**.

For the above Prisoner's Dilemma game, player A would not use strategy C, and hence she must choose strategy D. This indirect way of predicting the choice of strategy D is important. In general, a player may have more than two strategies, and all we can say right now is that choosing a strategy that is strictly dominated by some other strategy is in contradiction with payoff maximization. With more than two strategies, it is often insufficient to "predict" that a player would choose a strategy that is not strictly dominated.[4]

Next, we formulate the payoff function of player B. Recall that the first coordinate of a strategy combination is player A's strategy. Hence, we should write:

$$u_B(D, D) = -3, u_B(C, D) = 0, u_B(D, C) = -5, u_B(C, C) = -1. \tag{2.2}$$

[3]For the extended definition of strict dominance under "mixed strategies", see Sect. 3.6.

[4]However, if there is a strategy that strictly dominates all other strategies (of the relevant player), then this strategy (called the *dominant strategy*) should be chosen by a rational player. In the Prisoner's Dilemma, the strategy that is not strictly dominated by some other strategy coincides with the dominant strategy, and thus it is easy to predict an outcome by the dominant strategy. In general, there are few games with dominant strategies. Hence we do not emphasize the prediction that a player chooses the dominant strategy. Note also that, if there is a strategy combination such that all players are using a dominant strategy, the combination is called a *dominant-strategy equilibrium*.

We can see that this payoff function is symmetric[5] (we can permutate the players to obtain the same function) to player A's payoff function (2.1). Therefore, by the same logic for player A, strategy C is strictly dominated by strategy D for player B, and we predict that B does not choose strategy C.

To summarize, our forecast for the Prisoner's Dilemma game is that both players defect, and the social outcome would be the strategy combination (D, D).

What is the dilemma here? In order to choose a strategy, players can use simple logic. They do not even need to know the other's payoff function. But the dilemma occurs when they realize the resulting payoff combination. With the rational outcome of (D, D), both players get 3 year sentences, but if they choose the outcome (C, C), then they could have reduced the sentence to 1 year each! That is, the strategy combination (D, D) is socially not desirable. This social valuation can be mathematically defined as follows.

Definition 2.2 A strategy combination $(s_1, s_2, \ldots, s_n) \in S$ is *efficient*[6] if there is no strategy combination $(s'_1, s'_2, \ldots, s'_n) \in S$ such that
(1) for any player $i \in \{1, 2, \ldots, n\}$, $u_i(s_1, s_2, \ldots, s_n) \leqq u_i(s'_1, s'_2, \ldots, s'_n)$, and
(2) for some player $j \in \{1, 2, \ldots, n\}$, $u_j(s_1, s_2, \ldots, s_n) < u_j(s'_1, s'_2, \ldots, s'_n)$.

Using this definition, the rational players' predicted strategy choice (D, D) in the Prisoner's Dilemma is not efficient, because there exists (C, C) which satisfies (1) and (2).

This is one of the most fundamental problems in game theory: **the outcome that results from payoff-maximizing players' decisions is not always socially efficient**. However, it is in a way plausible, because non-cooperative game theory assumes that each player independently chooses decisions and does not consider the effects on others. Making self-interested players act in consideration of others is a completely different problem from finding a strategic outcome of arbitrary games. Game theory's main target is the latter, and the former should be formulated as a problem of designing games or finding a class of games with predicted outcomes (*equilibria*) that possess particular characteristics. For most of this book, we describe the main problem of the game theory, to provide a coherent prediction to arbitrary games. The Prisoner's Dilemma is "solved" by constructing cooperative equilibria for the infinitely repeated version of it (Chap. 5), or by adding incomplete information to the finitely repeated version (Chap. 8).

[5] Its matrix representation is symmetrical to the one in Table 2.1.

[6] Strictly speaking, this definition is *Strong Efficiency*. There is also *Weak Efficiency*, which requires that there is no $(s'_1, s'_2, \ldots, s'_n) \in S$ such that $u_i(s_1, s_2, \ldots, s_n) < u_i(s'_1, s'_2, \ldots, s'_n)$ for all $i \in \{1, 2, \ldots, n\}$.

Table 2.2 Player A's new payoff matrix

A's strategy\B's strategy	Defect	Cooperate
Defect	−3	−1
Cooperate	−5	0

2.3 Common Knowledge of a Game

So far, each player has only needed to know her/his own payoff function in order to make a rational decision. However, such method does not work with even a small change in the game structure. For example, let us modify the above Prisoner's Dilemma a little. Keeping player B's payoff function intact, assume that player A wants to defect if B defects, but she wants to cooperate with him if he does so. That is, we switch the last two payoff values in (2.1):

$$u_A(D, D) = -3, u_A(C, D) = -5, u_A(D, C) = -1, u_A(C, C) = 0. \tag{2.3}$$

The new matrix representation is Table 2.2.

Now, no strategy of player A is strictly dominated by some other strategy. What should player A do? Suppose that A knows the structure of the game, specifically, that A knows B's payoff function (2.2) and that he is rational. Then A can predict what B will do. By the logic in Sect. 2.2, player A can predict that B would choose strategy D. Then, if A defects, she would get −3, while if she cooperates, she would receive −5. Therefore it is rational for A to choose strategy D.

In this way, if players know all others' payoff functions and their rationality, it may be possible to "read" others' strategic decisions and find out their own rational choices. In this book, until the end of Chap. 5, we restrict our attention to games with complete information, which means that all players not only know the structure of the game (the set of players, the set of strategies of each player, and the payoff function of each player) and rationality of all players, but also they know that "all players know the structure of the game and rationality of all players", and they know that they know that "all players know the structure of the game and rationality of all players", and so on, *ad infinitum*. This infinitely deep knowledge is called *common knowledge* and is essential for the next analysis.[7]

2.4 Iterative Elimination of Strictly Dominated Strategies

The point of the analysis of the modified Prisoner's Dilemma in Sect. 2.3 was that, when the game and rationality are common knowledge, it may be possible to **eliminate** strictly dominated strategies of some player(s) from consideration and to solve for an outcome. Since a player takes into account the other player's decision-making,

[7]For a more detailed explanation of common knowledge, see, for example, Aumann [1], Chap. 5 of Osborne and Rubinstein [7], and Perea [8].

Table 2.3 Modified Prisoner's Dilemma

A's strategy\B's strategy	Defect	Cooperate
Defect	$-3, -3$	$-1, -5$
Cooperate	$-5, \quad 0$	$0, -1$

let us combine the two players' payoff matrices. For example, the modified Prisoner's Dilemma can be expressed by the (double) matrix representation of Table 2.3, where the first coordinate of the payoff combinations is player A's payoff (based on Eq. (2.3)) and the second coordinate is B's (based on Eq. (2.2)).

Table 2.3 in fact completely specifies a game. The set of players is shown at the top left cell, and player A's strategies are represented by rows (hence she is the *row player*), while player B's strategies are represented by columns (hence he is the *column player*). Payoff functions are defined by two matrices of first coordinates (for A) and second coordinates (for B). This double matrix representation of a game allows us to encompass the problem of both players' optimization in the same table and is quite convenient. Note, though, that the row player A can only choose the top or bottom row, and the column player B can only choose the left or right column. This means that, if players try to move from (D, D) to (C, C), they cannot make that happen unilaterally.

Let us repeat the logic of Sect. 2.3 using Table 2.3. Assume that the structure of the game, i.e., Table 2.3, as well as the rationality of both players are common knowledge. Player A does not have a strictly dominated strategy, but player B's strategy C is strictly dominated by strategy D, and thus rational player B would not choose strategy C. Moreover, player A knows this. Therefore, A can consider a reduced Table 2.4, eliminating strategy C from B's choices.

Focusing on Table 2.4, player A can compare the first coordinates of the rows to conclude that strategy C is strictly dominated by strategy D. Hence, the outcome becomes (D, D).

The above example does not require common knowledge (that is, the infinitely many layers of knowledge) of the game and rationality. It suffices that player A knows the game and rationality of player B. For more complex games, we need deeper knowledge. Consider the game represented by Table 2.5, where the two players are called 1 and 2, player 1 has three strategies called x, y, and z, and player 2 has three strategies called X, Y, and Z. (This game, with three strategies for each of the two players, has 3×3 payoff matrix for each player and belongs to the class called 3×3 *games*. In this sense, the Prisoner's Dilemma is a 2×2 *game*.)

To predict a strategic outcome for the game represented by Table 2.5, consider player 1, who compares the first coordinates of payoff pairs across rows. There is no

Table 2.4 Reduced matrix representation

A's strategy\B's strategy	Defect
Defect	$-3, -3$
Cooperate	$-5, \quad 0$

Table 2.5 A 3 × 3 game

1\2	X	Y	Z
x	3, 5	2, 2	2, 3
y	2, 2	0, 4	4, 1
z	1, 1	1, 2	1, 5

strict dominance relationship between strategies x and y, because if player 2 chooses strategy X or Y, x is better than y, but if player 2 chooses strategy Z, y is better than x. However, strategy z is strictly dominated by strategy x. Therefore, the rational player 1 would not choose strategy z. Similarly, player 2 chooses among the columns X, Y, and Z by comparing the second coordinates of the payoff values. There is no strict dominance among the three strategies of player 2 in Table 2.5. If player 1 chooses strategy x, then strategy X is better for player 2 than the other two strategies. If player 1 chooses y, then Y is the best, and if player 1 uses z, then strategy Z is the best for player 2.

So far, the players can reason using only the knowledge of their own payoff functions. Next, player 2, knowing the payoff function and rationality of player 1, can predict that player 1 would not use strategy z. Hence player 2 can focus on a reduced Table 2.6 after eliminating strategy z.

In the reduced Table 2.6, strategy Z is strictly dominated by strategy X for player 2. However, for player 1, strategies x and y are not strictly dominated by another.

By the common knowledge of the game, player 1 also knows the payoff function and rationality of player 2. In addition, player 1 knows that player 2 knows the payoff function and rationality of player 1. Therefore, player 1 can deduce that player 2 would eliminate the possibility of strategy z by player 1 and focus on Table 2.6. From that, player 1 can predict that player 2 would not choose strategy Z. This reasoning gives us a further reduced game, Table 2.7.

In Table 2.7, strategy y is strictly dominated by strategy x. Thus, we can predict strategy x as player 1's rational choice. This can be also predicted by player 2, who will choose strategy X. (Note that in this last step, we use the assumption that player 2 knows that player 1 knows that player 2 knows player 1's payoff function.) Therefore, the predicted outcome of the game represented by Table 2.5 is unique and is (x, X).

Table 2.6 Reduced 3 × 3 game

1\2	X	Y	Z
x	3, 5	2, 2	2, 3
y	2, 2	0, 4	4, 1

Table 2.7 Further reduced 3 × 3 game

1\2	X	Y
x	3, 5	2, 2
y	2, 2	0, 4

Table 2.8 Weak domination

1\2	L	R
U	11, 0	10, 0
D	10, 0	10, 0

This is an example where we can iteratively eliminate strictly dominated strategies (for all players) to arrive at a prediction of the strategic outcome of a game. That is, the iterative elimination of strictly dominated strategies is an equilibrium concept. Two remarks are in order. First, iterative elimination of an arbitrary size game is possible only under the common knowledge of the game and rationality. Second, in general, the remaining outcome after iterative elimination may not be unique. See, for example, the game of Table 2.8.

A convenient property of the iterative elimination process of strictly dominated strategies is that the order of elimination does not matter, i.e., when there are multiple strategies that are strictly dominated by some other strategy, the resulting outcome(s) of the elimination process do not change by the order of deletion among them. This is because, even if some strictly dominated strategies are not eliminated at one round, they will be still strictly dominated in the later rounds and therefore will be eliminated in the end.[8] (The elimination process stops when there are no more strictly dominated strategies for any player.)

2.5 Weak Dominance

Recall that the strict dominance relationship between strategies requires that the strict inequality among payoffs must hold, for any strategy combination by other players. To extend this idea, one may think it irrational to choose a strategy that never has a greater payoff than some other strategy and in at least one case has a strictly lower payoff. Consider the game represented by Table 2.8.

In this game, no player has a strictly dominated strategy. If player 2 chooses strategy R, strategy D gives player 1 the same payoff as strategy U does. For player 2, both strategies give the same payoff for any strategy by player 1.

However, choosing strategy D is not so rational for player 1, because strategy U guarantees the same payoff as that of strategy D for any strategy by player 2, and in the case where player 2 chooses strategy L, it gives a greater payoff. Let us formalize this idea.

Definition 2.3 Given a player i, a strategy $s_i \in S_i$ is *weakly dominated* by another strategy $s'_i \in S_i$ if,
(1) for any $s_{-i} \in S_{-i},\ u_i(s_i, s_{-i}) \leqq u_i(s'_i, s_{-i})$, and
(2) for some $s'_{-i} \in S_{-i},\ u_i(s_i, s'_{-i}) < u_i(s'_i, s'_{-i})$.

[8]For a formal proof, see Ritzberger [9], Theorem 5.1.

Table 2.9 Weak domination 2

1\2	L	R
U	11, 0	10, 10
D	10, 0	10, 0

Let us consider the iterative elimination of **weakly** dominated strategies. This is a stronger equilibrium concept than the iterative elimination of strictly dominated strategies, because the process may eliminate more strategies. For example, in the game represented by Table 2.8, strict dominance does not eliminate any strategy (i.e., the "equilibria" by iterative elimination of strictly dominated strategies is the set of all strategy combinations), while weak dominance can eliminate strategy D for player 1. Nonetheless, player 2's strategies do not have a weak dominance relationship, hence the process stops there with the "equilibria" of (U, L) and (U, R). Thus the prediction by iterative elimination of weakly dominated strategies may not be unique either.

Moreover, a fundamental weakness of the iterative elimination of weakly dominated strategies is that, unlike with strict domination, the order of elimination affects the resulting outcomes. Consider a modified game of Table 2.8, represented by Table 2.9.

Suppose that we eliminate player 2's weakly dominated strategy L first. Then the reduced game does not have even a weakly dominated strategy. Therefore the "prediction" by this process is $\{(U, R), (D, R)\}$. However, if we eliminate player 1's weakly dominated strategy D first, then strategy L for player 2 is strictly dominated by R in the reduced game and can be eliminated. Hence under this order of elimination, the resulting outcome is (U, R) only. Therefore, the iterative elimination of weakly dominated strategy is not a well-defined equilibrium concept.

In addition, if we consider Nash equilibrium (Chap. 3) as the standard equilibrium concept, iterative elimination of weakly dominated strategies may delete a strategy that is a part of a Nash equilibrium (see Problem 3.9). Therefore, it is safer to apply iterative elimination for strictly dominated strategies only. Iterative elimination of strictly dominated strategies does not eliminate a strategy that is part of a Nash equilibrium (see Sect. 3.5). However, if there are many equilibria in a game, it is better to focus on the ones that do not involve weakly dominated strategies. (This is advocated by Luce and Raiffa [4] and Kohlberg and Mertens [3].)

2.6 Maximin Criterion

Most games do not have dominance among any player's strategies. Can players deduce some strategy combination by some reasoning, at least based on their objective of payoff-maximization? Let us focus on two-player, *zero-sum* games.

Definition 2.4 A two-player, *zero-sum game* is a normal-form game $G = (\{1, 2\}, S_1, S_2, u_1, u_2)$ such that the sum of the two players' payoff values is zero for any strategy combination. That is, for any $(s_1, s_2) \in S$,

$$u_1(s_1, s_2) = -u_2(s_1, s_2).$$

For example, games in which one player wins and the other loses can be formulated as zero-sum games.

In a two-player, zero-sum game, maximizing one's own payoff is equivalent to minimizing the opponent's payoff. Therefore, when the structure of the game and player rationality are common knowledge, a player can reason that, if her chosen strategy is to be a part of a stable outcome, then it must be that the opponent's strategy against it is the one that minimizes her payoff, given her strategy. We can examine the payoff matrices of the players to find out which strategy combination is consistent with this reasoning by both players. For the game represented by Table 2.10, player 1 can deduce that if she chooses strategy x, then player 2 would choose strategy Y, which gives her payoff of −4. This "payoff that you receive from a strategy when your opponent minimizes your payoff, given the strategy" is called the *reservation payoff*. In a two-player, zero-sum game, each strategy is associated with a reservation payoff, as Table 2.11 shows for the game of Table 2.10. Note that the row player's reservation payoffs are shown in the rows of the rightmost column, while the column player's are shown at the bottom of the columns.

One way to proceed with this reasoning is that, because player 2 will try to minimize player 1's payoff, player 1 must choose a strategy with the highest reservation payoff. That is, player 1 should solve the following maximization problem:

$$\max_{s_1 \in S_1} \min_{s_2 \in S_2} u_1(s_1, s_2).$$

Table 2.10 A zero-sum game

1\2	X	Y
x	3, −3	−4, 4
y	−2, 2	−1, 1
z	2, −2	1, −1

Table 2.11 Reservation payoffs

1\2	X	Y	1's reservation payoff
x	3, −3	−4, 4	−4
y	−2, 2	−1, 1	−2
z	2, −2	1, −1	1
2's reservation payoff	−3	−1	

Table 2.12 A non-zero-sum game

1\2	A	B	1's reservation payoff
a	2, −2	0, 0	0
b	1, −1	3, 3	1
2's reservation payoff	−2	0	

Strategically choosing one's action by solving the above problem is called *maximin criterion* (Luce and Raiffa [4]), and the above value is called the *maxmin value* of player 1.

Similarly, player 2 solves

$$\max_{s_2 \in S_2} \min_{s_1 \in S_1} u_2(s_1, s_2).$$

By the definition of a zero-sum game, it holds that $u_2(s_1, s_2) = -u_1(s_1, s_2)$. Hence the latter problem is equivalent to

$$\min_{s_2 \in S_2} \max_{s_1 \in S_1} u_1(s_1, s_2).$$

This is the value that player 2 wants player 1 to receive. The value of this minimization is called the *minmax value* of player 1.[9]

For the example of Table 2.11, the strategy that maximizes player 1's reservation payoff is strategy z, and the one for player 2 is strategy Y. Moreover, when the two players play this strategy combination, their maximal reservation payoffs realize in this game. Therefore, we can predict that the outcome of this game is (z, Y). (In general, however, we need to expand the set of strategies to *mixed strategies*, defined in Sect. 3.6, in order to warrant the consistency of maxmin value and minmax value for zero-sum games. This is the cornerstone result called the Minimax Theorem by von Neumann [5] and von Neumann and Morgenstern [6]. For general Minimax Theorems, see Fan [2])

It is, however, easy to see that the maximin criterion does not work in general games. Consider the game represented by Table 2.12. The strategy that maximizes player 1's reservation payoff is b, while the one for player 2 is B. However, for the strategy combination (b, B), their payoffs are not the maximal reservation payoffs. In other words, it is not consistent for the players to choose a strategy that maximizes their own payoff under the assumption that "the opponent minimizes my payoff".

Nonetheless, it is still a valid reasoning to maximize your payoff under the assumption that **your opponent maximizes her/his own payoff**. This idea leads to the notion of Nash equilibrium later. In two-player, zero-sum games, your opponent's minimization of your payoff and her maximization of her own payoff are equivalent, but in general games, they are not equivalent. The latter concerns rationality of the opponent.

[9]The definitions of maxmin value and minmax value in this section are within the "pure" strategies. A general definition of the minmax value is given in Sect. 5.7.

Table 2.13 A three-player game

1 \ 2	X	Y
x	3, 5, 4	0, 2, 1
y	2, 2, 5	2, 4, 3

3: *L*

1 \ 2	X	Y
x	3, 5, 2	2, 2, 0
y	2, 2, 0	0, 4, 1

R

2.7 Matrix Representation of 3-Player Games

So far, we have shown examples of two-player games only. In this section, we describe how three-player games can be represented by matrices. For example, suppose that there are three players, 1, 2, and 3, and player 1's set of strategies is $\{x, y\}$, player 2's is $\{X, Y\}$, and player 3's is $\{L, R\}$. We can represent this three-player game by two tables, where in each table, the rows are player 1's strategies and the columns are player 2's strategies. Player 3 is assumed to choose the matrices (or the tables) corresponding to L or R. In this case, player 3 is called the *matrix player*.

The players' payoffs are usually listed in order of the row player's, then the column player's, and lastly the matrix player's. To compare player 3's payoffs among strategy combinations, we must jump from one table to the other. For example, when player 1 chooses strategy x and player 2 chooses Y, player 3 must compare the third coordinate's value 1 at the top-right corner of Table L, which corresponds to the strategy combination (x, Y, L), and the third coordinate's value 0 at the top-right of Table R, which corresponds to the strategy combination (x, Y, R). By comparing this way, our reader should be able to conclude that in the game represented by Table 2.13 (consisting of two tables), strategy R is strictly dominated by strategy L for player 3. In this way, three-person games can be nicely represented by matrices.

Games with four or more players cannot be represented easily with matrices, and thus these games are usually specified by words and equations.

Problems

2.1 Construct matrix representations for the following games.

(a) There are two sales clerks, Ms. A and Mr. B. They have the same abilities, and their sales performance depends only on their effort. They both have the same feasible strategies, to make an effort (strategy E), or not to make an effort (strategy N). They choose one of these strategies simultaneously. Their payoffs are the following "points", which are based on their relative performances.

If both players pick the same strategy, then the sales figures are the same, and therefore each player gets 1 point. If one player makes an effort while the other does not, then the player who makes an effort has a higher sales figure, so that (s)he gets 3 points and the other gets 0 points.

(b) Consider the same strategic situation, except for the payoff structure. If a player makes an effort, (s)he incurs a cost of 2.5 points. The sales points are awarded in the same way as in (a).

(c) Suppose that there are two electric appliance shops, P1 and P2. Every morning, both shops put prices on their sale items. Today's main sale item is a laptop computer. The wholesale price is \$500 per computer, which a shop must pay to the wholeseller if it sells the computer. For simplicity, assume that each shop charges either \$580 or \$550, and there are only two potential buyers. If the two shops charge the same price, each shop sells one unit, while if one shop charges a lower price than the other, it sells two units and the other shop sells nothing. A shop's payoff is its profit, namely, the total sales (charged price times the number of units sold) minus the total cost (wholesale price times the number of units sold).

2.2 Find the outcomes for the iterative elimination of strictly dominated strategies of the games in Problem 2.1(a), (b), and (c).

2.3 Consider the following two-player normal-form game. Analyze how the set of strategy combinations ("equilibria") that survive iterative elimination of weakly dominated strategies may be affected by the order of elimination.

P1\P2	Left	Center	Right
Up	2, 3	2, 2	1, 2
Middle	2, 1	2, 2	2, 2
Down	3, 1	2, 2	2, 2

2.4 Let us define the *second-price, sealed-bid auction* of n bidders (players). There is a good to be auctioned. The strategy of a player $i \in \{1, 2, \ldots, n\}$ is a *bid*, which is a non-negative real number. Players submit one bid each, simultaneously, in sealed envelopes. The player who submits the highest bid wins the auction (gets to buy the good) at the price of the second-highest bid. If there are multiple bidders who submitted the highest bid, one of them is selected as a winner with equal probability.

The payoff of a player $i \in \{1, 2, \ldots, n\}$ is defined as follows. Let a non-negative real number v_i be the *valuation* of the good for player i, which is the benefit i gets if i obtains the object. Let $\mathbf{b} = (b_1, b_2, \ldots, b_n)$ be the bid combination of all players. If i wins the auction, her payoff is $v_i - f(\mathbf{b})$ (where $f(\mathbf{b})$ is the second-highest number among $\{b_1, b_2, \ldots, b_n\}$). Otherwise her payoff is 0.

Prove that, for any player $i \in \{1, 2, \ldots, n\}$, the "honest" strategy such that $b_i^* = v_i$ (where player i bids her true valuation of the good) weakly dominates all other strategies.

2.5 Consider the following two-player normal-form game.

1\2	a	b	c
A	4, 1	5, 2	−2, 1
B	3, 3	4, 4	1, 2
C	2, 5	1, 1	−1, 6
D	5, 0	2, 4	0, 5

(a) Let K_0 be the fact that "player 1 does not take a strictly dominated strategy and she knows her own payoff function. Likewise, player 2 does not take a strictly dominated strategy and knows his own payoff function". What is the set of strategyf combinations deduced only from K_0? (You can use notation such as $\{A, B\} \times \{a, b\}$.)

(b) Let K_1 be the fact that "player 1 knows K_0 and player 2's payoff function. Player 2 also knows K_0 and player 1's payoff function". What is the set of strategy combinations deduced only from K_1?

(c) Let K_2 be the fact that "player 1 and player 2 both know K_0 and K_1". What is the set of strategy combinations deduced only from K_2?

(d) What is the set of strategy combinations after iterative elimination of strictly dominated strategies?

2.6 Consider the following three-player normal-form game

1 \ 2	L	R
U	2, 2, 2	3, 1, 1
D	1, 1, 3	2, 3, 4

3: A

1 \ 2	L	R
U	1, 2, 3	31, 1, 0
D	0, 3, 2	30, 30, 30

B

Player 1 chooses between rows U and D, player 2 chooses between columns L and R, and player 3 chooses between matrices A and B, simultaneously. The ith coordinate in each payoff combination is player i's payoff. Find the set of strategy combinations that survive iterative elimination of strictly dominated strategies.

2.7 Construct the matrix representation of the following game.

There is a duopoly market, in which firm 1 and firm 2 are the only producers. Each firm chooses a price in $\{2, 3, 4\}$ simultaneously. The payoff of a firm $i \in \{1, 2\}$ is its sales, which is its price multiplied by its demand. The demand of a firm $i \in \{1, 2\}$ is determined as follows. Let p_i be firm i's price and p_j be the opponent's. The demand of firm i is

$$D_i(p_i, p_j) = \begin{cases} (4.6 - p_i) & \text{if} p_i < p_j \\ \frac{1}{2}(4.6 - p_i) & \text{if} p_i = p_j \\ 0 & \text{if} p_i > p_j. \end{cases}$$

For example, if firm 1 chooses $p_1 = 3$ and firm 2 also chooses $p_2 = 3$, then each firm's sales (payoff) is $3 \cdot \frac{1}{2}(4.6 - 3) = 2.4$.

2.8 Continue to analyze the game in Problem 2.7.

(a) Find the set of strategy combinations that survive iterative elimination of strictly dominated strategies.

(b) Assume that firms can choose a price from $\{1, 2, 3, 4\}$. The payoff is determined in the same way as in Problem 2.7. Construct a matrix representation.

(c) Find the set of strategy combinations that survive iterative elimination of strictly dominated strategies for the game (b).

(d) Find the set of strategy combinations that survive iterative elimination of both strictly and weakly dominated strategies for the game (b).

References

1. Aumann R (1976) Agreeing to disagree. Ann Stat 4(6):1236–1239
2. Fan K (1953) Minimax theorems. Proc Nat Acad Sci 39:42–47
3. Kohlberg E, Mertens J-F (1986) On the strategic stability of equilibria. Econometrica 54(5):1003–1037
4. Luce D, Raiffa H (1957) Games and decisions. Wiley, New York
5. von Neumann J (1928) Zur Theorie der Gesellschaftsspiele. Mathematische Annalen 100:295–320 (English translation (1959): On the theory of games of strategy. In: Tucker A, Luce R (eds) Contributions to the theory of games. Princeton University Press, Princeton, pp 13–42
6. von Neumann J, Morgenstern O (1944) Theory of games and economic behavior. Princeton University Press, Princeton, NJ
7. Osborne M, Rubinstein A (1994) A course in game theory. MIT Press, Cambridge, MA
8. Perea A (2001) Rationality in extensive form games. Klouwer Academic Publishers, Dordrecht, the Netherlands
9. Ritzberger K (2002) Foundations of non-cooperative game theory. Oxford University Press, Oxford, UK

http://www.springer.com/978-4-431-55644-2

Non-Cooperative Game Theory
Fujiwara-Greve, T.
2015, XI, 260 p. 64 illus., Hardcover
ISBN: 978-4-431-55644-2