Lesson 14

Game Theory

Roberto Serrano and Allan M. Feldman

Version C

1. Introduction

In the last lesson we discussed duopoly markets in which two firms compete to sell a product. In such markets, the firms behave strategically; each firm must think about what the *other firm* is doing in order to decide what it should do itself. The theory of duopoly was originally developed in the 19th century, but it led to the theory of games in the 20th century. The first major book in game theory, published in 1944, was *Theory of Games and Economic Behavior*, by John von Neumann (1903-1957) and Oskar Morgenstern (1902-1977). We will return to the contributions of Von Neumann and Morgenstern in Lesson 19, on uncertainty and expected utility.

A group of people (or teams, firms, armies, countries) are in a *game* if their decision problems are interdependent, in the sense that the actions that all of them take influence the outcomes for everyone. *Game theory* is the study of games; it can also be called *interactive decision theory.* Many real-life interactions can be viewed as games. Obviously football, soccer, and baseball games are *games.* But so are the interactions of duopolists, the political campaigns between parties before an election, and the interactions of armed forces and countries. Even some interactions between animal or plant species in nature can be modeled as games. In fact, game theory has been used in many different fields in recent decades, including economics, political science, psychology, sociology, computer science, and biology.

This brief lesson is not meant to replace a formal course in game theory; it is only an introduction. The general emphasis is on how *strategic behavior* affects the interactions among rational players in a game. We will provide some basic definitions, and we will discuss a number of well-known simple examples. We will start with a description of the *prisoners' dilemma*, where we will introduce the idea of a *dominant strategy equilibrium.* We will briefly discuss *repeated games* in the prisoners' dilemma context, and *tit for tat strategies.* Then we will describe the *battle of the sexes*, and introduce the concept of *Nash equilibrium*. We will discuss the possibilities of there being multiple Nash equilibria, or no (pure strategy) Nash equilibria, and we discuss the idea of mixed strategy equilibria. We will then present an *expanded battle of the sexes*, and we will see that in game theory, an expansion of choices may make players worse off instead of better off. At the end of the lesson, we will describe *sequential move games*, and we will briefly discuss *threats*.

2. The Prisoners' Dilemma, and the Idea of Dominant Strategy Equilibrium

The most well-known example in game theory is the *prisoners' dilemma*. (It was developed around 1950 by Merrill M. Flood (1908-1991) and Melvin Dresher (1911-1992) of the RAND Corporation. It was so-named by Albert W. Tucker (1905-1995), a Princeton University mathematics professor.)

Consider the following. A crime is committed. The police arrive at the scene and arrest two suspects. Each of the suspects is taken to the police station for interrogation, and they are placed in separate cells. The cells are cold and nasty. The police interrogate them separately, and without any lawyers present. A police officer tells each one: "You can keep your mouth shut and refuse to testify. Or, you can confess and testify at trial."

We use some special and potentially confusing terminology to describe this choice. If a suspect refuses to testify, we say that he has chosen to *cooperate* with his fellow suspect. If a suspect confesses and testifies at trial, we say that he has chosen to *defect* from his fellow suspect. The reader will need to remember that to "cooperate" means to cooperate *with the other suspect, not with the police*, and also to remember that to "defect" means to defect *from the other suspect*.

The officer goes on: "If both of you refuse to testify, we will only be able to convict you on a minor charge, which will result in a sentence of 6 months in prison for each of you. If both of you confess and testify, you will each get 5 years in prison. If one of you refuses to testify (i.e., "cooperates") while the other confesses and testifies (i.e., "defects"), the one who testifies will go free, and the one who refuses to testify will get a full 10 years in prison."

The officer concludes: "That's what we're offering you, you lowlife hooligan. Think it over. We'll be back tomorrow to hear what you have to say."

We now consider this question: given this information, how should a rational suspect behave?

Should the suspects "cooperate" with each other (and tell the police nothing) or should they "defect" from each other (and confess)?

Table 14.1 below shows the prisoners' dilemma game. In game theory, the people playing the game are called players, so we now refer to our suspects as players. Player 1 chooses the rows in the table, while player 2 chooses the columns. Each of them has two possible actions to choose: "Cooperate" or "Defect." Each of the four action combinations results in payoffs to each player, in the form of prison time to be served. The outcomes are shown as the vectors in the cells of Table 14.1. The first entry is always the outcome for player 1, and the second is always the outcome for player 2. For instance, if player 1 defects while player 2 cooperates (bottom row, left column of the table), prison time for player 1 is None, and prison time for player 2 is 10 years. Note that these outcomes are "bads" rather than "goods"; each player wants to *minimize* his outcome.

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	6 months, 6 months	10 years, None
	Defect	None, 10 years	5 years, 5 years

Table 14.1: The prisoners' dilemma.

Each suspect wants to minimize his own jail time. But each must think about what the other suspect will do.

Let us now analyze the problem carefully. Here's how player 1 thinks about the game. He considers what player 2 might do. If player 2 cooperates, they are in the first column of the table. In this case, player 1 gets 6 months if he cooperates (first row), and no prison time if he defects (second row). Therefore, if player 2 cooperates, player 1 will defect. On the other hand, if player 2 defects, they are in the second column of the table. In this case, player 1 gets 10 years if he cooperates (first row), and 5 years if he defects (second row). Therefore if player 2 defects, player 1 will defect.

We now realize that whatever action player 2 chooses, player 1 will want to defect. We leave it to the reader to do the same type of analysis for player 2, whose payoffs are the second entries in each of the payoff vectors. When you do this, you will conclude that player 2 will want to defect, whatever action player 1 chooses.

In a game like this, actions that players might take are called *strategies.* A *dominant strategy* is a strategy which is optimal for a player, no matter what strategy the other player is choosing. In the prisoners' dilemma, the best thing for player 1 to do is to defect, no matter what player 2 might do. Therefore "Defect" is a dominant strategy for player 1. Similarly, "Defect" is a dominant strategy for player 2. When a pair of strategies are each dominant for the two players, the pair is called a *dominant strategy equilibrium* or a *solution in dominant strategies.* We now know that (Defect, Defect) is a *dominant strategy equilibrium* in the prisoners' dilemma. Rational players should choose dominant strategies if they exist; they clearly make sense, since a dominant strategy is the best for a player no matter what the other player is doing.

We conclude that the two suspects should both confess to the police, or defect from each other. Therefore they will each end up with a prison sentence of 5 years. Between the two of them, the total will be 10 years of prison. But this outcome is *very peculiar*, because if they had both chosen to keep their mouths shut, or cooperate with each other, they would have ended up with prison sentences of *only 6 months each, and a total of 1 year between the two of them.*

Back in Lesson 11 on perfectly competitive markets, we introduced the reader to Adam Smith's free market philosophy—his invisible hand theory. In brief, this is the theory that if the market is allowed to operate freely, with each consumer seeking to maximize his own utility and each firm seeking to maximize its own profits, with each of the players in the grand market game ignoring the welfare of all the others and doing the best it can for itself, the outcome will actually be best for society. That is, self-interested consumers and firms in a competitive market will end up maximizing social surplus, the sum of consumers' and producers' surplus.

But now note the dramatically different conclusion in the prisoners' dilemma. In this game, where we are focusing on the outcomes for the two suspects and ignoring the welfare of the police officers, the victims of the original crime, and the rest of society, the obvious and simple measure of social welfare for our two suspects is -1 times the sum of the two prison sentences. (We need the -1 to convert a *cost*—prison time—into a *benefit.*) But our analysis above indicates that each player, pursuing his own self-interest, maximizing his own welfare by minimizing his years here is that it may sometimes be in the interest of people to have reputations as being "crazy" or "tough," in order to induce beneficial changes in the behavior of others.

The moral of this story is that game theory can sometimes improve its predictions in explaining real-world phenomena by expanding its models.

4. The Battle of the Sexes, and the Idea of Nash Equilibrium

Most games are not as simple to solve as the prisoners' dilemma. That is, in most strategic situations, players do not have dominant strategies. In general, what each player will want to do will depend on what the other players are doing. Consequently, each player's conjectures about the behavior of the other players are crucial for determining his own behavior. For example, remember the first duopoly game of the last lesson, and its solution, the Cournot equilibrium (y_1^*, y_2^*). (Here y_1^* is firm 1's output, and y_2^* is firm 2's.) It is obvious that the Cournot equilibrium is not a dominant strategy equilibrium. If firm 2 decided to flood the market with product and drive the price down to zero, for example, firm 1 would not choose y_1^*. Rather, firm 1 would produce zero and save its production costs. This shows that producing y_1^* is not a dominant strategy for firm 1. The same argument applies to firm 2.

We will now analyze a new game, the *battle of the sexes.* This was first studied by R. Duncan Luce (1925-) and Howard Raiffa (1924-), in their 1957 book *Games and Decisions: Introduction and Critical Survey.*

A young woman (player 1) and her boyfriend (player 2) are out on Saturday night, driving in their own cars, on their way to meet each other for an evening together. Since this game was invented long before cellphones were around, they cannot communicate with each other. There are two options that they had talked about previously: a football game and an opera performance. But neither one of them can recall which option they had decided on. They like each other very much, and both would hate to spend the evening without the other. The young woman likes opera much better than football, but her boyfriend likes football better than opera. If the woman ends up at the opera with her boyfriend, her payoff is 3. But her payoff is 0 if she ends up at the opera without him. If the woman ends up at the football game with her boyfriend, her payoff is 1. But her payoff is 0 if she ends up at the football game without him. Similarly for the young man, if he ends up at the football game with her, his payoff is 3; if he ends up at the opera with her, his payoff is 1; and if he ends up at either place without her, his payoff is 0.

Table 14.2 shows the game. The rows of the table are the woman's possible strategies, and the columns are the man's. In other words, the woman chooses the row, and the man chooses the column. Each vector in each cell of the table shows the payoffs to the two players. For instance, if both of them choose football, they are in the first row, first column cell of the table. The payoff to the woman is then 1, and the payoff to the man is 3. Note that these payoffs, unlike the payoffs in the prisoners' dilemma game, are "goods" rather than "bads"; each player want to *maximize* rather than *minimize* her/his outcome.

		Man	
		Football	Opera
Woman	Football	1, 3	0, 0
	Opera	0, 0	3, 1

Table 14:2 The battle of the sexes.

What predictions can we make about this game? First of all, note that there are no dominant strategies. For either player, "Football" is better if she/he expects the other to choose "Football," but "Opera" is better if she/he expects the other to choose "Opera."

The standard equilibrium concept in the battle of the sexes is the Nash equilibrium, named for the famous 20th century economist, mathematician, and game theorist John Nash (1928-). A *Nash equilibrium* is a pair of strategies, one for each player, such that player 1's strategy is the best for her given player 2's strategy, and such that player 2's strategy is the best for him given player 1's strategy. Each player's strategy is a best response to the other's.

The reader should note that *a Cournot equilibrium in a duopoly model is a Nash equilibrium, and a Bertrand equilibrium in a duopoly model is also a Nash equilibrium in the corresponding duopoly game.* Moreover, *any dominant strategy equilibrium is a Nash equilibrium.* For example, (Defect, Defect) in the prisoners' dilemma is also a Nash equilibrium. This is because a dominant strategy for a player is always a best response for that player; therefore it is the best response when his opponent is playing his dominant strategy. But the reverse doesn't hold; and there will