

Chapter 10

The Theory of Games

In looking at the supply side of markets, we looked first at markets where there were a very large number of sellers (perfect competition) so sellers were able to make their decisions just on the basis of the market price; they didn't really have to think too much about what their competitors were doing. We then looked at monopoly, where the sole seller had no competition to worry about (although we did look briefly at situations where the monopolist did worry about the entry of *potential* competitors). Most real-world markets are characterized neither by perfect competition nor by monopoly, but by situations where the number of sellers may be quite small, at least in a local geographic area. In such situations, firms have to be very much aware of what their competitors are doing and how their competitors might respond to any decisions they make.

The Theory of Games, first proposed by John von Neumann and Oskar Morgenstern in 1944, is a theoretical approach to analyzing situations like these. Although the analysis of economic behavior in situations involving small numbers of competitors has been going on since at least the 19th century, the Theory of Games brought a new order to such analyses, incorporating within its domain many of the old approaches and concepts, but also providing new ways of thinking about these situations systematically. Particularly since the last quarter of the 20th century, game theory has become the primary approach taken by most economists to study a wide variety of strategic situations, situations where decision-makers try to take into account the behavior of other decision-makers whose actions interact with their own in producing the outcomes experienced by everyone.

10.1 Basics

A **game** is any strategic situation involving two or more decision-makers whose actions jointly determine the outcomes for all of them. Any game can be characterized by the following features and concepts:

- **Players:** The decision-makers in the game, following the practice in typ-

ical board and parlor games like chess and checkers, are called players.

- **Actions:** At any point in the game, if a player makes a decision, it would be to “play” an action. The action is his or her move at that point in the game.
- **Strategies:** A strategy is a complete description of all the possible actions that a player would take at all possible junctures of the game. Some games, like chess, may be so complex and have such a large number of possible combinations of actions that it may be impossible to fully specify a strategy. However, for many games it is possible to fully articulate all the possible situations that might arise and to specify all the actions that a player might make at each of those situations.
- **Payoffs:** At the end of every game, all the players receive payoffs. These may be monetary, such as profits or monetary winnings, non-monetary such as utility levels, or simply outcomes like winning and losing. The units in which the payoffs are measured are the units of whatever the players care about, such as profits, money, utility, etc. The goal of each player during the game is assumed to be maximization (or minimization) of her payoff. When the payoffs of all the players at every possible conclusion of a game add up to zero, the game is called a **zero-sum game**. Such games typically involve one player (the loser) paying the other player (the winner), so that the payoffs net to zero.
- **Information and Beliefs:** It is important to be clear about what each player does or does not know at each point of the game and what beliefs players hold. **Games of complete information** are games in which each player knows all the possible strategies, payoffs and objectives of each of the other players. **Games of incomplete information** are games in which players have different degrees of information, knowledge, and beliefs about the game.
- **Stages:** Games can have just one stage or multiple stages. A single-stage game is one in which players can make only one move and they do so simultaneously, such as the games “rock-paper-scissors” or “matching pennies.” Such games are called **static games**. A multi-stage game is one which will typically have players making moves in sequence, as in chess. Such games are called **dynamic games**.
- **Repetitions:** Static games can be played just once, in which case they are called **one-shot games**, or they may be played repeatedly, in which case they are called **repeated games**. Real life games of market competition may be modeled as repeated games or as dynamic games.
- **Equilibrium:** The game theorist’s notion of a “solution” to a game is to find its equilibrium, a situation where players are playing a combination of strategies such that no player has any incentive to change her strategy.

There are different notions of equilibrium in games and we will study some as we move forward.

- **Form:** Games are generally presented in one of two forms. **Normal form games** are presented in the form of a matrix (easier to do when there are only two players), while **extensive form games** are presented in the form of a decision tree. We will see examples of both forms as we look at some actual games.

10.2 Static Games of Complete Information

The simplest games are static one-shot games between two players who each are fully informed about all the details of the game. Such games are called static games of complete information. There are a number of classic examples of games of this kind that bring out various aspects of games and their solutions. We will look at a number of these classic example games.

10.2.1 The Prisoners' Dilemma

This is perhaps the most famous of all games, as it captures one of the most fundamental dilemmas of competition. If the two players in this game could cooperate, they would both do well. But they both have an incentive to cheat on any cooperative agreement, and, when they do cheat, they both end up worse off.

The traditional scenario of the game, the one that gives it its name, goes something like this. Two men committed a murder, after having broken into a house and stealing some valuables. The police have arrested them, separated them so they can no longer communicate with one another, and have told each the following: We have proof that you stole the valuables and therefore know that we can have you both sentenced to 2 years in prison. We do not yet have proof that you committed the murder. Why don't you confess? If you do, and your partner doesn't, we will reward you by reducing your sentence to 1 year, and punish your partner by sentencing him to 20 years. If you don't confess, and your partner does, you will end up in jail for 20 years and he will be let off with just 1 year. Finally, if you both confess, we'll arrange to have you both sentenced to 10 years in prison.

The game is presented in its normal form in Figure 10.1. Each player has just two strategies: "Confess" and "Don't Confess." In the normal form, the game is presented in the form of a matrix, with the strategies of one player, who we will call the "row player," presented in the rows, and the strategies of the other player, the "column player," presented in the columns. The cells of the matrix show the payoffs if that particular pair of strategies is played, with the row player's payoff listed first and the column player's payoff listed second. The payoffs in this example are the number of years to which each player gets sentenced, and to indicate that these are "bads" that the players would want to minimize, we have added minus signs to the payoffs. To take one cell as

		Column player's Strategies	
		Confess	Don't Confess
Row Player's Strategies	Confess	-10, -10	-1, -20
	Don't Confess	-20, -1	-2, -2

Figure 10.1: The Prisoners' Dilemma Game

an example, suppose the row player confesses (top row) and the column player doesn't confess (right column), the row player gets 1 year and the column player gets 20 years. Similarly for the other cells. In general, in a 2x2 matrix of this sort, the strategy of the row player presented in the top row (here, Confess) is called "Top" (T), and the strategy presented in the bottom row (here, Don't Confess) is called "Bottom" (B). Similarly, the column player's strategies are called "Left" (L, here Confess) and "Right" (R, here Don't Confess).

How will the players play this game? It turns out that, in all likelihood, they will both Confess. We can see why this is by considering the reasoning of the row player. He might think: I don't know what my partner is going to do, but let me think about what my best strategy would be depending upon what he does. Suppose he confesses. Then I'm better off confessing, since I get only 10 years rather than 20. So, if he confesses, I should confess too. But what if he doesn't confess? Well, in that case, I'm better off confessing again, since I'll get only 1 year rather than 2. So, no matter what he does, it's better for me to confess. The column player could go through the same process of reasoning, and, as the problem is perfectly symmetric, would arrive at the same conclusion. Thus both players would conclude that Confess is their best strategy, and both would end up with 10 years in jail.

Note that the payoffs under this outcome (-10,-10) are the only set of payoffs that the two players would unanimously agree is worse than another set of possible payoffs: (-2,-2) if neither confesses. If both players could somehow agree to not confess, and somehow enforce this agreement, they would both be better off. But any agreement to not confess is unlikely to stick. Suppose, for example, that, as the two are being bundled into the police van upon their arrest, Row says to Column: "Look, they're going to play Prisoners' Dilemma on us. Let's agree to not confess!" And suppose Column agrees, and they swear to abide by the agreement. Once the two are separated, they each go through their thought

process, realize they are unambiguously better off if they confess, especially if there is a cache of money and jewelry hidden away that they can access once they are out of jail, and therefore they each have a powerful incentive to renege on their agreement. That could be one reason why criminal organizations like the Mafia have the strong code of *omerta*, silence under all circumstances, the breaking of which is punishable by death. In effect, the organization changes the payoffs by making Confess a very expensive strategy and thereby helps enforce agreements to Don't Confess.

The Prisoners' Dilemma game is a rather special one in that both players have what are called **dominant strategies**, strategies that are the best choice no matter what the other player does. The strategy pair (Confess, Confess) is therefore a **dominant strategy equilibrium** and is considered a very robust outcome for the game were it to actually be played.

It is worth noting that Confess is the dominant strategy for each player, *even if they didn't actually commit the murder*. Thus this kind of payoff structure can elicit false confessions for the police.

		Column player's Strategies	
		L	R
Row Player's Strategies	T	5, 5	20, 0
	B	0, 20	10, 10

Figure 10.2: Payoff Matrix for the Generic Prisoners' Dilemma Game

A more generic version of the Prisoners' Dilemma game is presented in Figure 10.2. Here, the row player has the two strategies Top (T) and Bottom (B) and the column player has the two strategies Left (L) and Right (R). Payoffs are now good, so each player wants to maximize his or her payoff. Row's dominant strategy is T and Column's is L, so (T, L) is the dominant strategy equilibrium. The payoffs are (5, 5), even though this outcome is not Pareto-efficient, since the outcome (10, 10) is Pareto superior to it.

The Prisoners' Dilemma game has been found to have many applications in many different fields, including biology, ecology, nuclear strategy, economics, and so on. In economics, it is seen as a simple way to describe many competitive situations in which cooperation between competitors would be beneficial to them but difficult to implement. For example, suppose the two players in

the game were rival firms trying to decide how much to spend on advertising. Suppose T and L represent high levels of advertising and B and R represent low levels of advertising. Suppose, with both firms engaging in low levels of advertising, each firm's profits would be 10, but with both firms engaging in high levels of advertising, each firm's profits would be 5, because advertising is expensive and each firm's advertising neutralises the other's. However, if one firm advertises a lot and the other advertises a little, suppose the payoffs are 20 and 0 respectively, because the high-advertiser can steal a lot of the low-advertiser's customers. Thus the payoff matrix in Figure 10.2 can represent this competitive situation. We can see that high advertising is the dominant strategy for each firm, resulting in the Pareto-inefficient outcome of (5, 5). High advertising becomes the dominant strategy because, if your rival advertises high, you must advertise high to protect your market share, while, if your rival advertises low, you want to advertise high in order to steal his customers and increase your market share. Thus we have a classic Prisoners' Dilemma type of game situation.

10.2.2 Battle of the Sexes

The next game we will look at is called by the colorful title "Battle of the Sexes." This game describes a situation in which the players might agree that cooperation is desirable, but they differ on the nature of the cooperation they would like to see. The outcome could be a deadlock.

		Colleen's Strategies	
		L Ballgame	R Theatre
Rob's Strategies	T Ballgame	3, 2	1, 1
	B Theatre	0, 0	2, 3

Figure 10.3: Payoff Matrix for the Battle of the Sexes Game

The story told to paint a picture for this game is that a couple are planning to go out for the evening. They love one another dearly and so they both agree they would like to go out together, but one, say Rob (he will be the Row player), wants to go to a baseball game, while the other, say Colleen (who will be the Column player) wants to go to the theatre. A typical payoff matrix for this game

might look like the one in Figure 10.3. It shows that Rob and Colleen would like to be together (the payoffs at (T, L) and (B, R) are strictly greater than the others), but Rob prefers they both go to the baseball game, while Colleen prefers they both go to the theatre. If they go their separate ways, Rob to the baseball and Colleen to the theatre (T, R), they are both worse off, while if they each go to the other's favorite event (B, L) they are even more worse off.

This game does not have a dominant strategy equilibrium. If Colleen absolutely insists on going to the theatre, Rob is better off going to the theatre as well (rather than his separate way), while if Colleen agrees to go to the theatre, Rob would clearly prefer that. So Rob's optimal strategy depends upon what Colleen does; he does not have a dominant strategy. Similarly, Colleen's optimal strategy depends upon what Rob does, so she does not have a dominant strategy either.

We therefore need a different equilibrium concept for this game than we did for the Prisoners' Dilemma. The most widely used equilibrium concept in game theory is that of Nash equilibrium. The **Nash Equilibrium** of a game is a set of strategies, one for each player, such that each player's payoff is optimized, conditional on the strategies of every other player. That is, conditional on what everyone else is doing, no player can improve his or her payoff by changing their strategy.

The way to check if any set of strategies is a Nash equilibrium is to look to see if any player can improve their payoff, holding all the other strategies fixed. In the present game, consider the strategy pair (T, L), i.e., where Rob plays Top (goes to the baseball) and Colleen plays Left (also goes to the baseball). The payoffs are (3, 2). Rob's only alternative is to play Bottom, and, if he did, with Colleen continuing to play L, his payoff would fall to 0. Thus Rob has no incentive to deviate from (T, L). Colleen's only alternative is to play Right, and her payoff falls to 1 if she does that, so she has no incentive to deviate either. Thus (T, L) is a Nash equilibrium.

Let's look at (T, R), where the payoffs are (1, 1). Here, if Rob changed his strategy to Bottom, his payoff would go up to 2. So Rob has an incentive to deviate from (T, R), and that is enough for us to determine that (T, R) is not a Nash equilibrium. In fact, even Colleen has an incentive to deviate from (T, R).

Let's look at (B, L), where the payoffs are (0, 0). Here, Colleen could raise her payoff to 3 by switching to R, and Rob could raise his to 3 also by switching to T. Thus (B, L) is not a Nash equilibrium.

Finally, let's look at (B, R), with payoffs (2, 3). Here, neither player can improve by changing strategy, so (B, R) is also a Nash equilibrium. Thus the Battle of the Sexes game has two Nash equilibria: (T, L) and (B, R).

There are two things that need to be said at this stage of the discussion. First, there is nothing in the structure of the game or in the method of solution that would help us determine which of these two Nash equilibria would actually be achieved. Each spouse could fight tooth and nail for "their" equilibrium and the couple could end up at an impasse (many real-world couples do!). Second, there is nothing that guarantees that the outcome *will* be one of the Nash

equilibria anyway. Although most authors presume that rationality will push players towards the Nash equilibria, there is actually no assurance that they will actually end up there.

The Battle of the Sexes game is illustrative of many problems. Where two agents attempt to negotiate anything, reaching some sort of agreement would be better for both, but each has a vested interest in a different equilibrium than the other. For example, in negotiations between the management of a firm and its labor union over a new labor contract, both might agree that they would want to avoid a strike or a lock-out. Yet each may press for a co-operative equilibrium more favorable to their interests, the union for higher wages and shorter hours, and management for the opposite. Thus this situation is much like the Battle of the Sexes. Another example is where makers of products might try to agree on standards that make their products mutually compatible. Here experience seems to suggest that the impasse tends to win. IBM and Apple were unable to agree on a common operating system, Sony and the VHS consortium were unable to settle on a standard format for videotape, and in recent years manufacturers of mobile phone handsets have failed to agree on a common system either.

10.2.3 Matching Pennies and Mixed Strategies

In the game Matching Pennies, each of the two players, say Rowan and Columbine, puts a covered coin on the palm of their hand and then they simultaneously reveal the coins to one another. If the coins match (i.e., both are Heads or both are Tails), Columbine pays Rowan. If the coins don't match (one is Heads and the other is Tails), Rowan pays Columbine. The payoff matrix is shown in Figure 10.4. Note that the payoffs in each cell of the payoff matrix of this game sum to zero; this is therefore an example of a **zero-sum game**.

		Columbine's Strategies	
		Heads	Tails
Rowan's Strategies	Heads	1, -1	-1, 1
	Tails	-1, 1	1, -1

Figure 10.4: Payoff Matrix for the Matching Pennies Game

An examination of the four strategy pairs in this game reveals that none of them is a Nash equilibrium. Consider

- (H, H): Columbine would want to deviate.
- (H, T): Rowan would want to deviate.
- (T, H): Rowan would want to deviate.
- (T, T): Columbine would want to deviate.

Is there no equilibrium in this game?

A little thought reveals a rather obvious solution to this game. Each player could simply flip the coin, or place it randomly on their hand without looking. In that case, each player's expected payoff is zero and neither can do any better by doing anything differently. So this would be an equilibrium. What we have here is a **mixed strategy Nash equilibrium**. There is no Nash equilibrium in pure strategies, i.e., in either of the available strategies played with certainty, but there *is* an equilibrium in mixed strategies, where the players play the pure strategies with certain probabilities. In this mixed strategy equilibrium, each player is playing the pure strategies with probability $\frac{1}{2}$.

The idea of just flipping the coins seemed intuitive, but it would be useful to have a systematic way of finding mixed strategy equilibria. The process goes as follows. Suppose Rowan is thinking of playing H with probability p and T with probability $(1-p)$, and Columbine is thinking of playing H with probability q and T with probability $(1-q)$. Then Rowan's expected payoff is

$$\begin{aligned} E\pi_R &= p\{q \cdot 1 + (1-q)(-1)\} + (1-p)\{q(-1) + (1-q) \cdot 1\} \\ &= p(q-1+q) + (1-p)(-q+1-q) \\ &= p(2q-1) + (1-p)(1-2q). \end{aligned}$$

Then, to choose p so as to maximize $E\pi_R$, Rowan would set

$$\frac{\partial E\pi_R}{\partial p} = 2q - 1 - 1 + 2q = 0,$$

which yields the "solution"

$$q = \frac{1}{2}.$$

Note that p has dropped out of this first-order condition! There is no particular value of p at which $E\pi_R$ is maximized. However, we have found that, if $q = \frac{1}{2}$, $\frac{\partial E\pi_R}{\partial p} = 0$, that is, $E\pi_R$ is invariant to changes in p .

A similar analysis for Columbine will show that there is no particular value of q where $E\pi_C$ is maximized, but, if $p = \frac{1}{2}$, $E\pi_C$ will be invariant to changes in q . Thus, if Rowan plays (H, T) with probabilities $(\frac{1}{2}, \frac{1}{2})$ and Columbine plays (H, T) with probabilities $(\frac{1}{2}, \frac{1}{2})$ neither player will have any incentive to change their decision. Thus we have a Nash equilibrium in mixed strategies. It turns out that we can always find a Nash equilibrium of a game, if not in pure strategies, then in mixed strategies.

10.2.4 Mixed Strategies in the Battle of the Sexes game

As another example of mixed strategies equilibria, we could look at the possibility of such an equilibrium in the Battle of the Sexes game. Recall the payoff matrix from Figure 10.3. Suppose Rob contemplates playing T and B with probabilities p and $(1-p)$, and Colleen contemplates playing L and R with probabilities q and $(1-q)$. Then Rob's expected payoff would be

$$\begin{aligned} E\pi_R &= p\{3q + 1 \cdot (1 - q)\} + (1 - p)\{0 \cdot q + 2 \cdot (1 - q)\} \\ &= p(1 + 2q) + (1 - p)(2 - 2q). \end{aligned}$$

To maximize $E\pi_R$, Rob would set

$$\frac{\partial E\pi_R}{\partial p} = 1 + 2q - 2 + 2q = 0,$$

which yields the "solution"

$$q = \frac{1}{4}.$$

So if Colleen plays L and R with probabilities $\frac{1}{4}$ and $\frac{3}{4}$, Rob's expected payoff will be invariant with changes in p . It can be shown that then $E\pi_R = \frac{3}{2}$.

We can similarly write down Colleen's expected payoff:

$$\begin{aligned} E\pi_C &= q\{2p + 0 \cdot (1 - p)\} + (1 - q)\{1 \cdot p + 3 \cdot (1 - p)\} \\ &= q(2p) + (1 - q)(3 - 2p). \end{aligned}$$

For maximization,

$$\frac{\partial E\pi_C}{\partial q} = 2p - 3 + 2p = 0,$$

which simplifies to

$$p = \frac{3}{4}.$$

So if Rob plays T and B with probabilities $\frac{3}{4}$ and $\frac{1}{4}$, Colleen's expected payoff will be invariant to changes in q . Her expected payoff will also be $\frac{3}{2}$.

We therefore have a Nash equilibrium in mixed strategies for the Battle of the Sexes game, in which Rob plays T and B with probabilities $\frac{3}{4}$ and $\frac{1}{4}$ and Colleen plays L and R with probabilities $\frac{1}{4}$ and $\frac{3}{4}$.

10.3 Sequential Games

Most real-world games do not involve just one set of simultaneous moves by all the players, but rather a sequence of moves, where one player moves first, the other responds, and so on. Such games with multiple stages are called sequential or dynamic games.

We can illustrate these types of games by looking at a sequential version of the Battle of the Sexes game. Suppose the baseball game and the theatre

are located in completely different parts of the city and Rob and Colleen are each at work and plan to meet at the location of their evening date. No matter where they decide to go, Rob will have the longer drive and so he would have to leave work earlier than Colleen, and will have first choice in deciding where to actually go. In this way, he decides where to go first and informs Colleen where he is headed. At that point, she can choose to join Rob where he is already going, or choose to go her separate way. However, she can announce to Rob before he leaves work what her intentions are.

In this two-stage version of the game, since he moves first, Rob still has just two strategies available to him: to go to the baseball game or to the theatre. But Colleen is going to move later, so she is going to find herself at one of two different decision points. Rob could have gone to the baseball game or to the theatre, and, for each of these possibilities, Colleen has two possible actions she can take, to go to the baseball game or to the theatre. So Colleen has four possible strategies:

1. If Rob plays T, she plays L, and, if Rob plays B, she plays L. This strategy will be denoted LL, the first L denoting her response to his playing T and the second her response to his playing B.
2. If Rob plays T, she plays L, and, if Rob plays B, she plays R. This strategy will be denoted LR.
3. If Rob plays B, she plays R, and, if Rob plays B, she plays L. This strategy will be denoted RL.
4. If Rob plays B, she plays R, and, if Rob plays B, she plays R. This strategy will be denoted RR.

		Colleen's Strategies			
		LL	LR	RL	RR
Rob's Strategies	T	3, 2	3, 2	1, 1	1, 1
	B	0, 0	2, 3	0, 0	2, 3

Figure 10.5: Normal form of the Sequential Battle of the Sexes Game

The payoff matrix for this game can then be written as in Figure 10.5. The payoffs are obtained from the static version of the game whose payoff matrix was

in Figure 10.3. Here, for example, (T, LL) results in payoffs of (3, 2) because LL means that Colleen will play L if Rob plays T, so we end up with the payoff that (T, L) would have yielded in the static game. (T, RL) results in payoffs of (1, 1) since RL means that Colleen plays R if Rob plays T, resulting in the outcome from (T, R) in the static game.

If we examine the payoff matrix in Figure 10.5, we see that this game has three Nash equilibria in pure strategies: (T, LL), (T, LR) and (B, RR). Of these three equilibria, the first two are somewhat similar in that they both involve Colleen playing L if Rob plays T. Since that gives Rob his most preferred option, he has no incentive to deviate from playing T. And since L is *ex post* Colleen's best option if Rob plays T, she has no incentive to deviate from either of her strategies either.

The interesting Nash equilibrium is (B, RR). If Colleen manages to convince Rob that she is going to play R no matter what he does (that's what RR implies, that she plays R whether he plays T or B), he can be induced to play B, since that is his best response to her playing R. That is why (B, RR) is a Nash equilibrium. The trouble with this equilibrium is that Colleen's threat to play R no matter what Rob does is not credible, because it is not in her best interest to play R if he plays T. In other words, she could threaten to go to the theatre (play R) no matter what he does, but, once he has left his office and has been driving for 15 minutes towards the baseball game (by which time he would not be able to reach the theatre in time if he changes his destination), and if he now calls her and tells her he is already part of the way to the baseball game, it would no longer be rational for her to follow through on her threat. Thus Rob could call her bluff and simply play T, in which case Colleen's best response would be to abandon her threat and play L anyway.

The preceding analysis calls into question the type of Nash equilibrium embodied in the strategy pair (B, RR) and suggests a new notion of equilibrium for multi-stage games. This is called **subgame perfect equilibrium** (SPE) and is defined as a Nash equilibrium in which each player plays rationally at all stages of the game. In other words, it is an equilibrium that does not involve any non-credible threats. The Nash equilibrium (B, RR) is not a SPE because it involves Colleen's threat to play R if Rob plays T, a threat that is not credible, because it would not be rational for Colleen to do so. Similarly, the Nash equilibrium (T, LL) is not a SPE because it involves a threat by Colleen to play L if Rob plays B, another threat that is not credible. The only SPE of this game is (T, LR) since LR involves Colleen playing L if Rob plays T (which is indeed her best-response to T) and playing R if Rob plays B (which is her best-response again). Thus (T, LR) is a SPE, and we have already shown that the other two Nash equilibria are not SPEs.

Having the first move gives Rob the advantage in this game and the SPE gives him the outcome that he would have preferred from among the two previous Nash equilibria of the static version of the game. Colleen's threat to play R no matter what Rob does is really a bluff, and Rob can call her bluff and simply play T in the knowledge that it would then be rational for her to abandon her threat and play L. But what if Colleen could somehow pre-commit to carry out

her “irrational” threat? Perhaps she could then turn the game in her favour. For example, suppose they already have tickets to both events and new tickets are impossible to buy at this point. What if, before Rob has left his office, Colleen calls him on Skype and, in full view of him on the video feed, burns her ticket to the baseball game? She has now *pre-committed* to going to the theatre because she no longer has a ticket to the baseball game! In this case, her strategy RR does not involve a non-credible threat, because she has effectively eliminated the possibility of playing L. Thus the strategy pair (B, RR) is now a SPE and the outcome will be *Colleen's* preferred one. A credible pre-commitment to an otherwise non-credible threat has turned the game to Colleen's favour.

The foregoing analysis has analyzed the game in its normal form, and this has yielded interesting insights into the way threats can be used in games.

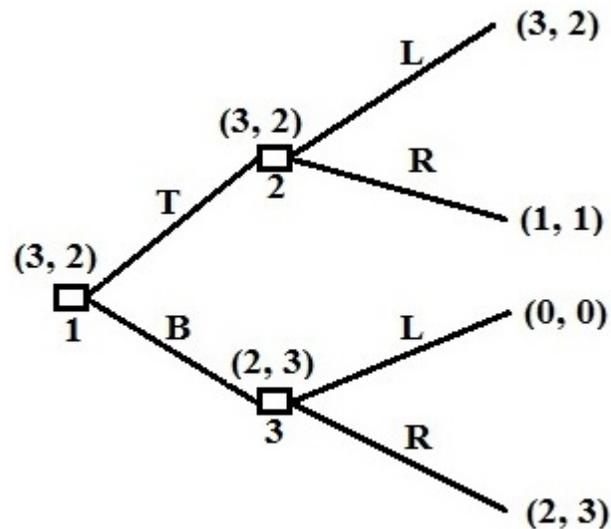


Figure 10.6: Extensive Form of the Sequential Battle of the Sexes Game

Another way of analyzing this game is by looking at it in **extensive form**, that is, in the form of a decision tree, and then solving it by backward induction. The extensive form of this game is presented in Figure 10.6. Rob moves first and can select T or B, which are the two branches emanating from decision node 1. Colleen moves second and can find herself either at decision node number 2 (if Rob chose T) or decision node 3 (if Rob chose B). At each of these nodes, Colleen could choose L or R, so there are four possible outcomes, and the payoffs for each are written at the ends of the respective branches of the tree.

The extensive form of the game is solved by backward induction. We ask (or, properly, Rob asks) what Colleen would do at each of the two decision nodes where she might find herself. At node 2, her best response to Rob playing T would be to play L, since she would get a payoff of 2 rather than 1 by playing L.

Thus Rob can reason that, if he plays T, Colleen will play L and the outcome will be $(3, 2)$. Accordingly, we write $(3, 2)$ as the payoff above decision node 2. Similarly, we ask what Colleen would do at node 3. Here, her best response to Rob playing B is to play R, and so the outcome would be $(2, 3)$. Therefore, Rob can deduce that, if he plays B, Colleen will play R and the outcome will be $(2, 3)$. Rob can now look at his options at decision node 1 and can see that, if he plays T, the outcome will be $(3, 2)$ and, if he plays B, the outcome will be $(2, 3)$. Rob can therefore see that his better option is to choose T. Colleen will then choose L and the payoffs will be $(3, 2)$. We therefore find the same outcome as we did with the SPE of the normal form game.

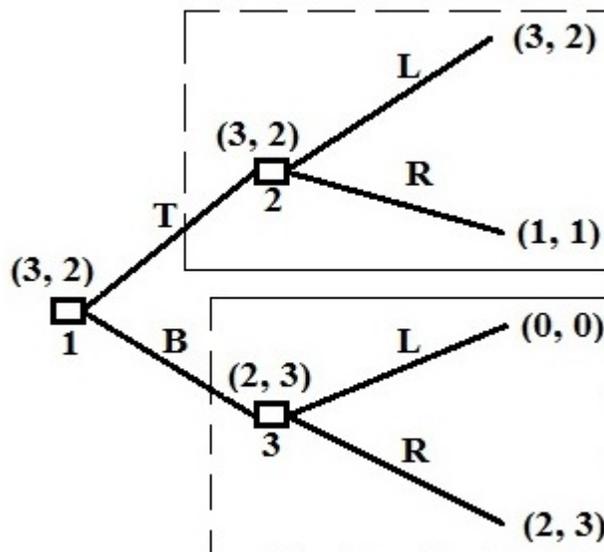


Figure 10.7: Sub-games of the Extensive Form Battle of the Sexes Game

We could actually derive the SPE specifically from the extensive form of the game. Think of Colleen's two decision nodes being the starting points of two "sub-games," marked by the dotted line outlines in Figure 10.7. We can then ask what her rational choice would be in each of these sub-games. We have already seen that her rational choice in the upper sub-game is to choose L in response to Rob's choice of T, and in the lower sub-game is to choose R in response to Rob's choice of B. Thus her rational strategy is LR. We have also already seen that Rob's choice in the "sub-game" around decision node 1 is to choose T. Putting the two players' rational strategies together, we have the SPE as (T, LR) .

10.4 Repeated Games

We have looked at single-stage and multi-stage games, but have had the players play them just once. In most real-world situations, games are played not just once but repeatedly. For example, married couples don't have just one evening to plan, they have thousands. Firms that are competing with one another may find themselves in a Prisoners' Dilemma type of situation, but they know they will face the same situation again and again. What happens to the equilibrium outcome of the game when it is played repeatedly? Do the players keep playing the same strategies again and again, or do they change their behavior once they realize they will face the same opponent again tomorrow? These are the questions economists try to answer when they look at repeated games.

In thinking of the repeated Battle of the Sexes, it is quite clear that a threat such as RR on the part of Colleen, while being a non-credible threat in the one-shot game, may be a valid strategy in the repeated game. Imagine that Colleen threatened to go to the theatre no matter what Rob did, and suppose Rob decided to "call her bluff" and go to the baseball game. Colleen could well want to follow through with her threat in order to establish a "reputation," to convince Rob that she doesn't make empty threats. Her hope would be that in future Rob would be more likely to believe her threats and would therefore not call her bluffs. Of course, Rob could also issue his own threats and could keep calling her bluffs, and the couple could end up warring with each other and finish in divorce court! But it is not clear what the equilibrium would be.

Let us focus on the repeated Prisoners' Dilemma, since that is the game that has been studied the most, both theoretically and experimentally. It turns out that the outcome is likely to depend upon whether the game is played a known finite number of times, or an infinite number of times (which can be closely modeled by looking at a game which is played an unknown but finite number of times). The question is, might players cooperate with each other more when the game is played repeatedly? Think of "Confess" as being the non-cooperative, or "defect," strategy in the game and think of "Don't Confess" as being the cooperative strategy since it is mutually beneficial if both players play this strategy. Might repetition encourage players to not confess more often?

If the game is played a finite number of times, theory tells us that hope for cooperation unravels. Suppose the game is to be played five times. Then, before the fifth round, the situation would have reverted to the one-shot game. Both players will realize that cooperation is no longer beneficial and so both would defect. Looking forward before the fourth round, rational players would realize that defection is the rational strategy on the fifth round and would therefore see no benefit in cooperating in the fourth round. By this process, cooperation unravels as a viable strategy at any stage of the game, and so both players would defect, or play the non-cooperative strategy, at all times. In fact, this result has been shown to be generalizable to any game which has a unique Nash equilibrium in its static version. If the game is repeated a known finite number of times, the only subgame perfect equilibrium of the repeated game is the static Nash equilibrium at each stage of the game.

If the game is played an infinite number of times, or at least if there is no known end point for the game, it is possible that more cooperative behavior might emerge. This is borne out by results of experiments in which players play the game for real money. It is found that cooperative behavior frequently emerges in actual practice. And in a tournament of strategies playing against one another, the strategy that generated the highest cumulative payoff and won the tournament was the “Tit for Tat” strategy. This strategy simply plays whatever the opponent played on the previous round. If the opponent defected on the previous round, Tit for Tat defects, and if the opponent cooperated on the previous round, Tit for Tat cooperates.

10.5 Summary

In this chapter, we have looked at the Theory of Games, an approach to thinking about strategic behavior in situations where multiple decision-makers interact and whose decisions mutually affect everyone. The basic points were as follows:

- The fundamental solution concept in games is of equilibrium. We studied several types of equilibria.
- A Dominant Strategy Equilibrium is one where every player in the game is playing a strategy that is their optimal one regardless of what other players play.
- A Nash Equilibrium is one where every player is playing their best response to what all other players are playing, i.e., every player is playing their optimal strategy conditional on what every other player is playing.
- A Mixed Strategy is one where players play their (pure) strategies according to probabilities they choose. It is possible for Nash equilibrium to involve mixed strategies and in fact every static game has a mixed strategy Nash equilibrium.
- A Subgame Perfect Equilibrium is a solution concept for multi-stage games and is one where every player plays their optimal strategy at every stage of the game. In other words, a SPE involves no non-credible threats.
- When games are played repeatedly for a known finite number of times, there is a tendency for the static game Nash equilibrium to be played at each repetition.
- However, when games are repeated an infinite (or an unknown finite) number of times, departures from static Nash equilibrium can occur and more cooperative behavior tends to emerge.

10.6 Exercises

1. AIR, Inc. and RMS, Inc. are in competition. Each can follow one of two strategies, it can either co-operate, or compete. The payoff matrix, showing profits in millions of dollars, looks like this:

		RMS Strategies	
		Cooperate	Compete
AIR Strategies	Cooperate	40, 30	10, 60
	Compete	80, 10	30, 20

- (a) Identify any Nash equilibria in pure strategies to this game. Explain your reasons.
- (b) Prove that there does not exist a Nash equilibrium in mixed strategies. In one sentence, provide an intuitive reason for this result.
- (c) If AIR acquired RMS, it could co-ordinate the two firms' strategies. (Assume there are no other advantages or disadvantages to the merger.) What is the maximum amount AIR would be willing to pay for RMS? What is the minimum amount RMS would require to be willing to be acquired?
2. Consider the Battle of the Sexes game with the payoff matrix above when played as a static game of full information (the husband is the row player and the wife is the column player).

	Left	Right
Top	2, 1	0, 0
Bottom	0, 0	1, 2

Suppose that the game changes to a dynamic game in which the husband moves first.

- (a) Write down the normal form of the dynamic game and find any Nash equilibria of this game. Explain your reasons.
- (b) Now write down the extensive form of this dynamic game and find any subgame perfect equilibria.
3. Grand Corp. and Little Corp. are the only two producers of blodgets. Each can produce blodgets at a per-unit cost of \$4. Presently, Grand has an 80 per cent market share. The two firms agree to form a cartel, and the

agreement is that they will preserve market shares at 80 and 20 per cent respectively, while trying to maximize joint profits. The demand curve for blodgets is

$$P = 44 - 4Q.$$

- (a) How much should each firm produce for the cartel agreement to be realized?
 - (b) Suppose each of the firms is thinking about cheating on their cartel agreement by producing one unit more than the agreement had specified. They also know that their “partner” may cheat in this way. Each firm therefore has two possible output levels: the output agreed upon [your answer to (a)], or an output level one unit higher than that. By suitable calculations, construct a payoff matrix for the game the two firms find themselves playing.
 - (c) By analyzing the game, what would you predict the two firms to do? Comment briefly on this outcome as compared to the cartel outcome.
4. Consider a static game with the following payoff matrix:

	Left	Right
Top	10, 5	5, 3
Bottom	5, 0	20, 2

- (a) Find all Nash equilibria in pure strategies for this game.
- (b) Find all Nash equilibria in mixed strategies for this game.
- (c) Write down the normal form for the dynamic version of this game in which the row player moves first.
- (d) Find all Nash equilibria in pure strategies for the dynamic game of part (c) and also find the subgame perfect equilibrium of that game.