CHAPTER EIGHT

Game Theory

This chapter provides an introduction to noncooperative game theory, a tool used to understand the strategic interactions among two or more agents. The range of applications of game theory has been growing constantly, including all areas of economics (from labor economics to macroeconomics) and other fields such as political science and biology. Game theory is particularly useful in understanding the interaction between firms in an oligopoly, so the concepts learned here will be used extensively in Chapter 15. We begin with the central concept of Nash equilibrium and study its application in simple games. We then go on to study refinements of Nash equilibrium that are used in games with more complicated timing and information structures.

BASIC CONCEPTS

Thus far in Part 3 of this text, we have studied individual decisions made in isolation. In this chapter we study decision making in a more complicated, strategic setting. In a strategic setting, a person may no longer have an obvious choice that is best for him or her. What is best for one decision-maker may depend on what the other is doing and vice versa.

For example, consider the strategic interaction between drivers and the police. Whether drivers prefer to speed may depend on whether the police set up speed traps. Whether the police find speed traps valuable depends on how much drivers speed. This confusing circularity would seem to make it difficult to make much headway in analyzing strategic behavior. In fact, the tools of game theory will allow us to push the analysis nearly as far, for example, as our analysis of consumer utility maximization in Chapter 4.

There are two major tasks involved when using game theory to analyze an economic situation. The first is to distill the situation into a simple game. Because the analysis involved in strategic settings quickly grows more complicated than in simple decision problems, it is important to simplify the setting as much as possible by retaining only a few essential elements. There is a certain art to distilling games from situations that is hard to teach. The examples in the text and problems in this chapter can serve as models that may help in approaching new situations.

The second task is to "solve" the given game, which results in a prediction about what will happen. To solve a game, one takes an equilibrium concept (e.g., Nash equilibrium) and runs through the calculations required to apply it to the given game. Much of the chapter will be devoted to learning the most widely used equilibrium concepts and to practicing the calculations necessary to apply them to particular games.

A *game* is an abstract model of a strategic situation. Even the most basic games have three essential elements: players, strategies, and payoffs. In complicated settings, it is sometimes also necessary to specify additional elements such as the sequence of moves and the information that players have when they move (who knows what when) to describe the game fully.

Players

Each decision-maker in a game is called a *player*. These players may be individuals (as in poker games), firms (as in markets with few firms), or entire nations (as in military conflicts). A player is characterized as having the ability to choose from among a set of possible actions. Usually the number of players is fixed throughout the "play" of the game. Games are sometimes characterized by the number of players involved (two-player, three-player, or n-player games). As does much of the economic literature, this chapter often focuses on two-player games because this is the simplest strategic setting.

We will label the players with numbers; thus, in a two-player game we will have players 1 and 2. In an n-player game we will have players 1, 2,…, n, with the generic player labeled i.

Strategies

Each course of action open to a player during the game is called a *strategy*. Depending on the game being examined, a strategy may be a simple action (drive over the speed limit or not) or a complex plan of action that may be contingent on earlier play in the game (say, speeding only if the driver has observed speed traps less than a quarter of the time in past drives). Many aspects of game theory can be illustrated in games in which players choose between just two possible actions.

Let S_1 denote the set of strategies open to player 1, S_2 the set open to player 2, and (more generally) S_i the set open to player i. Let $s_1 \in S_1$ be a particular strategy chosen by player 1 from the set of possibilities, $s_2 \in S_2$ the particular strategy chosen by player 2, and $s_i \in S_i$ for player i. A strategy *profile* will refer to a listing of particular strategies chosen by each of a group of players.

Payoffs

The final return to each player at the conclusion of a game is called a *payoff*. Payoffs are measured in levels of utility obtained by the players. For simplicity, monetary payoffs (say, profits for firms) are often used. More generally, payoffs can incorporate nonmonetary factors such as prestige, emotion, risk preferences, and so forth.

In a two-player game, $u_1(s_1, s_2)$ denotes player 1's payoff given that he or she chooses s_1 and the other player chooses s_2 and similarly $u_2(s_2, s_1)$ denotes player 2's payoff.[1] The fact that player 1's payoff may depend on player 2's strategy (and vice versa) is where the strategic interdependence shows up. In an n-player game, we can write the payoff of a generic player i as $u_i(s_i, s_{-i})$, which depends on player i's own strategy s_i and the profile $s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$ of the strategies of all players other than i.

PRISONERS' DILEMMA

The Prisoners' Dilemma, introduced by A. W. Tucker in the 1940s, is one of the most famous games studied in game theory and will serve here as a nice example to illustrate all the notation just introduced. The title stems from the following situation. Two suspects are arrested for a crime. The district attorney has little evidence in the case and is eager to extract a confession. She separates the suspects and tells each: "If you fink on your companion but your companion doesn't fink on you, I can promise you a reduced

[1]Technically, these are the von Neumann–Morgenstern utility functions from the previous chapter.

(one-year) sentence, whereas your companion will get four years. If you both fink on each other, you will each get a three-year sentence." Each suspect also knows that if neither of them finks then the lack of evidence will result in being tried for a lesser crime for which the punishment is a two-year sentence.

Boiled down to its essence, the Prisoners' Dilemma has two strategic players: the suspects, labeled 1 and 2. (There is also a district attorney, but because her actions have already been fully specified, there is no reason to complicate the game and include her in the specification.) Each player has two possible strategies open to him: fink or remain silent. Therefore, we write their strategy sets as $S_1 = S_2 = \{\text{fink, silent}\}$. To avoid negative numbers we will specify payoffs as the years of freedom over the next four years. For example, if suspect 1 finks and suspect 2 does not, suspect 1 will enjoy three years of freedom and suspect 2 none, that is, $u_1(\text{fink, silent}) = 3$ and $u_2(\text{silent, fink}) = 0$.

Normal form

The Prisoners' Dilemma (and games like it) can be summarized by the matrix shown in Figure 8.1, called the *normal form* of the game. Each of the four boxes represents a different combination of strategies and shows the players' payoffs for that combination. The usual convention is to have player 1's strategies in the row headings and player 2's in the column headings and to list the payoffs in order of player 1, then player 2 in each box.

Thinking strategically about the Prisoners' Dilemma

Although we have not discussed how to solve games yet, it is worth thinking about what we might predict will happen in the Prisoners' Dilemma. Studying Figure 8.1, on first thought one might predict that both will be silent. This gives the most total years of freedom for both (four) compared with any other outcome. Thinking a bit deeper, this may not be the best prediction in the game. Imagine ourselves in player 1's position for a moment. We do not know what player 2 will do yet because we have not solved out the game, so let's investigate each possibility. Suppose player 2 chose to fink. By finking ourselves we would earn one year of freedom versus none if we remained silent, so finking is better for us. Suppose player 2 chose to remain silent. Finking is still better for us than remaining silent because we get three rather than two years of freedom. Regardless of what the other player does, finking is better for us than being silent because it results in an extra year of freedom. Because players are symmetric, the same reasoning holds if we

FIGURE 8.1

Normal Form for the Prisoners' Dilemma

		Suspect 2	
		Fink	Silent
Suspect 1	Fink	$u_1 = 1, u_2 = 1$	$u_1 = 3, u_2 = 0$
	Silent	$u_1 = 0, u_2 = 3$	$u_1 = 2, u_2 = 2$

imagine ourselves in player 2's position. Therefore, the best prediction in the Prisoners' Dilemma is that both will fink. When we formally introduce the main solution concept—Nash equilibrium—we will indeed find that both finking is a Nash equilibrium.

The prediction has a paradoxical property: By both finking, the suspects only enjoy one year of freedom, but if they were both silent they would both do better, enjoying two years of freedom. The paradox should not be taken to imply that players are stupid or that our prediction is wrong. Rather, it reveals a central insight from game theory that pitting players against each other in strategic situations sometimes leads to outcomes that are inefficient for the players.[2] The suspects might try to avoid the extra prison time by coming to an agreement beforehand to remain silent, perhaps reinforced by threats to retaliate afterward if one or the other finks. Introducing agreements and threats leads to a game that differs from the basic Prisoners' Dilemma, a game that should be analyzed on its own terms using the tools we will develop shortly.

Solving the Prisoners' Dilemma was easy because there were only two players and two strategies and because the strategic calculations involved were fairly straightforward. It would be useful to have a systematic way of solving this as well as more complicated games. Nash equilibrium provides us with such a systematic solution.

NASH EQUILIBRIUM

In the economic theory of markets, the concept of equilibrium is developed to indicate a situation in which both suppliers and demanders are content with the market outcome. Given the equilibrium price and quantity, no market participant has an incentive to change his or her behavior. In the strategic setting of game theory, we will adopt a related notion of equilibrium, formalized by John Nash in the 1950s, called *Nash equilibrium*.[3] Nash equilibrium involves strategic choices that, once made, provide no incentives for the players to alter their behavior further. A Nash equilibrium is a strategy for each player that is the best choice for each player given the others' equilibrium strategies.

The next several sections provide a formal definition of Nash equilibrium, apply the concept to the Prisoners' Dilemma, and then demonstrate a shortcut (involving underlining payoffs) for picking Nash equilibria out of the normal form. As at other points in the chapter, the reader who wants to avoid wading through a lot of math can skip over the notation and definitions and jump right to the applications without losing too much of the basic insight behind game theory.

A formal definition

Nash equilibrium can be defined simply in terms of *best responses*. In an n-player game, strategy s_i is a best response to rivals' strategies s_{-i} if player i cannot obtain a strictly higher payoff with any other possible strategy, $s_i' \in S_i$, given that rivals are playing s_{-i}.

DEFINITION

Best response. s_i is a best response for player i to rivals' strategies s_{-i}, denoted $s_i \in BR_i(s_{-i})$, if

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \quad \text{for all} \quad s_i' \in S_i. \qquad \textbf{(8.1)}$$

[2] When we say the outcome is inefficient, we are focusing just on the suspects' utilities; if the focus were shifted to society at large, then both finking might be a good outcome for the criminal justice system—presumably the motivation behind the district attorney's offer.

[3] John Nash, "Equilibrium Points in n-Person Games," *Proceedings of the National Academy of Sciences* 36 (1950): 48–49. Nash is the principal figure in the 2001 film *A Beautiful Mind* (see Problem 8.5 for a game-theory example from the film) and co-winner of the 1994 Nobel Prize in economics.

A technicality embedded in the definition is that there may be a set of best responses rather than a unique one; that is why we used the set inclusion notation $s_i \in BR_i(s_{-i})$. There may be a tie for the best response, in which case the set $BR_i(s_{-i})$ will contain more than one element. If there is not a tie, then there will be a single best response s_i and we can simply write $s_i = BR_i(s_{-i})$.

We can now define a Nash equilibrium in an n-player game as follows.

DEFINITION

Nash equilibrium. A Nash equilibrium is a strategy profile $(s_1^*, s_2^*, \ldots, s_n^*)$ such that, for each player $i = 1, 2, \ldots, n$, s_i^* is a best response to the other players' equilibrium strategies s_{-i}^*. That is, $s_i^* \in BR_i(s_{-i}^*)$.

These definitions involve a lot of notation. The notation is a bit simpler in a two-player game. In a two-player game, (s_1^*, s_2^*) is a Nash equilibrium if s_1^* and s_2^* are mutual best responses against each other:

$$u_1(s_1^*, s_2^*) \geq u_1(s_1, s_2^*) \quad \text{for all} \quad s_1 \in S_1 \tag{8.2}$$

and

$$u_2(s_1^*, s_2^*) \geq u_2(s_2, s_1^*) \quad \text{for all} \quad s_2 \in S_2. \tag{8.3}$$

A Nash equilibrium is stable in that, even if all players revealed their strategies to each other, no player would have an incentive to deviate from his or her equilibrium strategy and choose something else. Nonequilibrium strategies are not stable in this way. If an outcome is not a Nash equilibrium, then at least one player must benefit from deviating. Hyper-rational players could be expected to solve the inference problem and deduce that all would play a Nash equilibrium (especially if there is a unique Nash equilibrium). Even if players are not hyper-rational, over the long run we can expect their play to converge to a Nash equilibrium as they abandon strategies that are not mutual best responses.

Besides this stability property, another reason Nash equilibrium is used so widely in economics is that it is guaranteed to exist for all games we will study (allowing for mixed strategies, to be defined below; Nash equilibria in pure strategies do not have to exist). The mathematics behind this existence result are discussed at length in the Extensions to this chapter. Nash equilibrium has some drawbacks. There may be multiple Nash equilibria, making it hard to come up with a unique prediction. Also, the definition of Nash equilibrium leaves unclear how a player can choose a best-response strategy before knowing how rivals will play.

Nash equilibrium in the Prisoners' Dilemma

Let's apply the concepts of best response and Nash equilibrium to the example of the Prisoners' Dilemma. Our educated guess was that both players will end up finking. We will show that both finking is a Nash equilibrium of the game. To do this, we need to show that finking is a best response to the other players' finking. Refer to the payoff matrix in Figure 8.1. If player 2 finks, we are in the first column of the matrix. If player 1 also finks, his payoff is 1; if he is silent, his payoff is 0. Because he earns the most from finking given player 2 finks, finking is player 1's best response to player 2's finking. Because players are symmetric, the same logic implies that player 2's finking is a best response to player 1's finking. Therefore, both finking is indeed a Nash equilibrium.

We can show more: that both finking is the only Nash equilibrium. To do so, we need to rule out the other three outcomes. Consider the outcome in which player 1 finks and player 2 is silent, abbreviated (fink, silent), the upper right corner of the matrix. This is not a Nash equilibrium. Given that player 1 finks, as we have already said, player 2's best response is to fink, not to be silent. Symmetrically, the outcome in which player 1 is silent and player 2 finks in the lower left corner of the matrix is not a Nash equilibrium. That leaves the outcome in which both are silent. Given that player 2 is silent, we focus our attention on the second column of the matrix: The two rows in that column show that player 1's payoff is 2 from being silent and 3 from finking. Therefore, silent is not a best response to fink; thus, both being silent cannot be a Nash equilibrium.

To rule out a Nash equilibrium, it is enough to find just one player who is not playing a best response and thus would want to deviate to some other strategy. Considering the outcome (fink, silent), although player 1 would not deviate from this outcome (he earns 3, which is the most possible), player 2 would prefer to deviate from silent to fink. Symmetrically, considering the outcome (silent, fink), although player 2 does not want to deviate, player 1 prefers to deviate from silent to fink, so this is not a Nash equilibrium. Considering the outcome (silent, silent), both players prefer to deviate to another strategy, more than enough to rule out this outcome as a Nash equilibrium.

Underlining best-response payoffs

A quick way to find the Nash equilibria of a game is to underline best-response payoffs in the matrix. The underlining procedure is demonstrated for the Prisoners' Dilemma in Figure 8.2. The first step is to underline the payoffs corresponding to player 1's best responses. Player 1's best response is to fink if player 2 finks, so we underline $u_1 = 1$ in the upper left box, and to fink if player 2 is silent, so we underline $u_1 = 3$ in the upper left box. Next, we move to underlining the payoffs corresponding to player 2's best responses. Player 2's best response is to fink if player 1 finks, so we underline $u_2 = 1$ in the upper left box, and to fink if player 1 is silent, so we underline $u_2 = 3$ in the lower left box.

Now that the best-response payoffs have been underlined, we look for boxes in which every player's payoff is underlined. These boxes correspond to Nash equilibria. (There may be additional Nash equilibria involving mixed strategies, defined later in the chapter.) In Figure 8.2, only in the upper left box are both payoffs underlined, verifying that (fink, fink)—and none of the other outcomes—is a Nash equilibrium.

FIGURE 8.2

Underlining Procedure in the Prisoners' Dilemma

		Suspect 2	
		Fink	Silent
Suspect 1	Fink	$\underline{u_1 = 1}, \underline{u_2 = 1}$	$\underline{u_1 = 3}, u_2 = 0$
	Silent	$u_1 = 0, \underline{u_2 = 3}$	$u_1 = 2, u_2 = 2$

Dominant strategies

(Fink, fink) is a Nash equilibrium in the Prisoners' Dilemma because finking is a best response to the other player's finking. We can say more: Finking is the best response to all the other player's strategies, fink and silent. (This can be seen, among other ways, from the underlining procedure shown in Figure 8.2: All player 1's payoffs are underlined in the row in which he plays fink, and all player 2's payoffs are underlined in the column in which he plays fink.)

A strategy that is a best response to any strategy the other players might choose is called a *dominant strategy*. Players do not always have dominant strategies, but when they do there is strong reason to believe they will play that way. Complicated strategic considerations do not matter when a player has a dominant strategy because what is best for that player is independent of what others are doing.

DEFINITION

Dominant strategy. A dominant strategy is a strategy s_i^* for player i that is a best response to all strategy profiles of other players. That is, $s_i^* \in BR_i(s_{-i})$ for all s_{-i}.

Note the difference between a Nash equilibrium strategy and a dominant strategy. A strategy that is part of a Nash equilibrium need only be a best response to one strategy profile of other players—namely, their equilibrium strategies. A dominant strategy must be a best response not just to the Nash equilibrium strategies of other players but to all the strategies of those players.

If all players in a game have a dominant strategy, then we say the game has a *dominant strategy equilibrium*. As well as being the Nash equilibrium of the Prisoners' Dilemma, (fink, fink) is a dominant strategy equilibrium. It is generally true for all games that a dominant strategy equilibrium, if it exists, is also a Nash equilibrium and is the unique such equilibrium.

Battle of the Sexes

The famous Battle of the Sexes game is another example that illustrates the concepts of best response and Nash equilibrium. The story goes that a wife (player 1) and husband (player 2) would like to meet each other for an evening out. They can go either to the ballet or to a boxing match. Both prefer to spend time together than apart. Conditional on being together, the wife prefers to go to the ballet and the husband to the boxing match. The normal form of the game is presented in Figure 8.3. For brevity we dispense with the

FIGURE 8.3

Normal Form for the Battle of the Sexes

		Player 2 (Husband)	
		Ballet	Boxing
Player 1 (Wife)	Ballet	2, 1	0, 0
	Boxing	0, 0	1, 2

u_1 and u_2 labels on the payoffs and simply re-emphasize the convention that the first payoff is player 1's and the second is player 2's.

We will examine the four boxes in Figure 8.3 and determine which are Nash equilibria and which are not. Start with the outcome in which both players choose ballet, written (ballet, ballet), the upper left corner of the payoff matrix. Given that the husband plays ballet, the wife's best response is to play ballet (this gives her her highest payoff in the matrix of 2). Using notation, ballet $= BR_1$(ballet). [We do not need the fancy set-inclusion symbol as in "ballet $\in BR_1$(ballet)" because the husband has only one best response to the wife's choosing ballet.] Given that the wife plays ballet, the husband's best response is to play ballet. If he deviated to boxing, then he would earn 0 rather than 1 because they would end up not coordinating. Using notation, ballet $=$ BR_2(ballet). Thus, (ballet, ballet) is indeed a Nash equilibrium. Symmetrically, (boxing, boxing) is a Nash equilibrium.

Consider the outcome (ballet, boxing) in the upper left corner of the matrix. Given the husband chooses boxing, the wife earns 0 from choosing ballet but 1 from choosing boxing; therefore, ballet is not a best response for the wife to the husband's choosing boxing. In notation, ballet $\notin BR_1$(boxing). Hence (ballet, boxing) cannot be a Nash equilibrium. [The husband's strategy of boxing is not a best response to the wife's playing ballet either; thus, both players would prefer to deviate from (ballet, boxing), although we only need to find one player who would want to deviate to rule out an outcome as a Nash equilibrium.] Symmetrically, (boxing, ballet) is not a Nash equilibrium either.

The Battle of the Sexes is an example of a game with more than one Nash equilibrium (in fact, it has three—a third in mixed strategies, as we will see). It is hard to say which of the two we have found thus far is more plausible because they are symmetric. Therefore, it is difficult to make a firm prediction in this game. The Battle of the Sexes is also an example of a game with no dominant strategies. A player prefers to play ballet if the other plays ballet and boxing if the other plays boxing.

Figure 8.4 applies the underlining procedure, used to find Nash equilibria quickly, to the Battle of the Sexes. The procedure verifies that the two outcomes in which the players succeed in coordinating are Nash equilibria and the two outcomes in which they do not coordinate are not.

Examples 8.1 and 8.2 provide additional practice in finding Nash equilibria in more complicated settings (a game that has many ties for best responses in Example 8.1 and a game that has three strategies for each player in Example 8.2).

FIGURE 8.4

Underlining Procedure in the Battle of the Sexes

		Player 2 (Husband)	
		Ballet	Boxing
Player 1 (Wife)	Ballet	(2, 1)	0, 0
	Boxing	0, 0	(1, 2)

EXAMPLE 8.1 The Prisoners' Dilemma Redux

In this variation on the Prisoners' Dilemma, a suspect is convicted and receives a sentence of four years if he is finked on and goes free if not. The district attorney does not reward finking. Figure 8.5 presents the normal form for the game before and after applying the procedure for underlining best responses. Payoffs are again restated in terms of years of freedom.

FIGURE 8.5 The Prisoners' Dilemma Redux

(a) Normal form

Suspect 1 \ Suspect 2	Fink	Silent
Fink	0, 0	1, 0
Silent	0, 1	1, 1

(b) Underlining procedure

Suspect 1 \ Suspect 2	Fink	Silent
Fink	(<u>0</u>, <u>0</u>)	(<u>1</u>, <u>0</u>)
Silent	(<u>0</u>, <u>1</u>)	(<u>1</u>, <u>1</u>)

Ties for best responses are rife. For example, given player 2 finks, player 1's payoff is 0 whether he finks or is silent. Thus, there is a tie for player 1's best response to player 2's finking. This is an example of the set of best responses containing more than one element: BR_1 (fink) = {fink, silent}.

The underlining procedure shows that there is a Nash equilibrium in each of the four boxes. Given that suspects receive no personal reward or penalty for finking, they are both indifferent between finking and being silent; thus, any outcome can be a Nash equilibrium.

QUERY: Does any player have a dominant strategy?

EXAMPLE 8.2 Rock, Paper, Scissors

Rock, Paper, Scissors is a children's game in which the two players simultaneously display one of three hand symbols. Figure 8.6 presents the normal form. The zero payoffs along the diagonal show that if players adopt the same strategy then no payments are made. In other cases, the payoffs indicate a \$1 payment from loser to winner under the usual hierarchy (rock breaks scissors, scissors cut paper, paper covers rock).

As anyone who has played this game knows, and as the underlining procedure reveals, none of the nine boxes represents a Nash equilibrium. Any strategy pair is unstable because it offers at least one of the players an incentive to deviate. For example, (scissors, scissors) provides an incentive for either player 1 or 2 to choose rock; (paper, rock) provides an incentive for player 2 to choose scissors.

FIGURE 8.6 Rock, Paper, Scissors

(a) Normal form

Player 1 \ Player 2	Rock	Paper	Scissors
Rock	0, 0	−1, 1	1, −1
Paper	1, −1	0, 0	−1, 1
Scissors	−1, 1	1, −1	0, 0

(b) Underlining procedure

Player 1 \ Player 2	Rock	Paper	Scissors
Rock	0, 0	−1, <u>1</u>	<u>1</u>, −1
Paper	<u>1</u>, −1	0, 0	−1, <u>1</u>
Scissors	−1, <u>1</u>	<u>1</u>, −1	0, 0

The game does have a Nash equilibrium—not any of the nine boxes in the figure but in mixed strategies, defined in the next section.

QUERY: Does any player have a dominant strategy? Why is (paper, scissors) not a Nash equilibrium?

MIXED STRATEGIES

Players' strategies can be more complicated than simply choosing an action with certainty. In this section we study *mixed strategies*, which have the player randomly select from several possible actions. By contrast, the strategies considered in the examples thus far have a player choose one action or another with certainty; these are called *pure strategies*. For example, in the Battle of the Sexes, we have considered the pure strategies of choosing either ballet or boxing for sure. A possible mixed strategy in this game would be to flip a coin and then attend the ballet if and only if the coin comes up heads, yielding a 50–50 chance of showing up at either event.

Although at first glance it may seem bizarre to have players flipping coins to determine how they will play, there are good reasons for studying mixed strategies. First, some games (such as Rock, Paper, Scissors) have no Nash equilibria in pure strategies. As we will see in the section on existence, such games will always have a Nash equilibrium in mixed strategies; therefore, allowing for mixed strategies will enable us to make predictions in such games where it was impossible to do so otherwise. Second, strategies involving randomization are natural and familiar in certain settings. Students are familiar with the setting of class exams. Class time is usually too limited for the professor to examine students on every topic taught in class, but it may be sufficient to test students on a subset of topics to induce them to study all the material. If students knew which topics were on the test, then they might be inclined to study only those and not the others; therefore, the professor must choose the topics at random to get the students to study everything. Random strategies are also familiar in sports (the same soccer player sometimes shoots to the right of the net and sometimes to the left on penalty kicks) and in card games (the poker player sometimes folds and sometimes bluffs with a similarly poor hand at different times).[4]

Formal definitions

To be more formal, suppose that player i has a set of M possible actions $A_i = \{a_i^1, \ldots, a_i^m, \ldots, a_i^M\}$, where the subscript refers to the player and the superscript to the different choices. A mixed strategy is a probability distribution over the M actions, $s_i = (\sigma_i^1, \ldots, \sigma_i^m, \ldots, \sigma_i^M)$, where σ_i^m is a number between 0 and 1 that indicates the probability of player i playing action a_i^m. The probabilities in s_i must sum to unity: $\sigma_i^1 + \cdots + \sigma_i^m + \cdots + \sigma_i^M = 1$.

In the Battle of the Sexes, for example, both players have the same two actions of ballet and boxing, so we can write $A_1 = A_2 = \{\text{ballet, boxing}\}$. We can write a mixed strategy as a pair of probabilities $(\sigma, 1 - \sigma)$, where σ is the probability that the player chooses ballet. The probabilities must sum to unity, and so, with two actions, once the probability of one action is specified, the probability of the other is determined. Mixed strategy (1/3, 2/3) means that the player plays ballet with probability 1/3 and boxing with probability 2/3; (1/2, 1/2) means that the player is equally likely to play ballet or boxing; (1, 0) means that the player chooses ballet with certainty; and (0, 1) means that the player chooses boxing with certainty.

In our terminology, a mixed strategy is a general category that includes the special case of a pure strategy. A pure strategy is the special case in which only one action is played with positive probability. Mixed strategies that involve two or more actions being played with positive probability are called *strictly mixed strategies*. Returning to the examples from the previous paragraph of mixed strategies in the Battle of the Sexes, all four strategies (1/3, 2/3), (1/2, 1/2), (1, 0), and (0, 1) are mixed strategies. The first two are strictly mixed, and the second two are pure strategies.

With this notation for actions and mixed strategies behind us, we do not need new definitions for best response, Nash equilibrium, and dominant strategy. The definitions introduced when s_i was taken to be a pure strategy also apply to the case in which s_i is taken to be a mixed strategy. The only change is that the payoff function $u_i(s_i, s_{-i})$, rather

[4]A third reason is that mixed strategies can be "purified" by specifying a more complicated game in which one or the other action is better for the player for privately known reasons and where that action is played with certainty. For example, a history professor might decide to ask an exam question about World War I because, unbeknownst to the students, she recently read an interesting journal article about it. See John Harsanyi, "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points," *International Journal of Game Theory* 2 (1973): 1–23. Harsanyi was a co-winner (along with Nash) of the 1994 Nobel Prize in economics.

than being a certain payoff, must be reinterpreted as the expected value of a random payoff, with probabilities given by the strategies s_i and s_{-i}. Example 8.3 provides some practice in computing expected payoffs in the Battle of the Sexes.

EXAMPLE 8.3 Expected Payoffs in the Battle of the Sexes

Let's compute players' expected payoffs if the wife chooses the mixed strategy (1/9, 8/9) and the husband (4/5, 1/5) in the Battle of the Sexes. The wife's expected payoff is

$$\begin{aligned} U_1\left(\left(\frac{1}{9}, \frac{8}{9}\right), \left(\frac{4}{5}, \frac{1}{5}\right)\right) &= \left(\frac{1}{9}\right)\left(\frac{4}{5}\right)U_1(\text{ballet, ballet}) + \left(\frac{1}{9}\right)\left(\frac{1}{5}\right)U_1(\text{ballet, boxing}) \\ &\quad + \left(\frac{8}{9}\right)\left(\frac{4}{5}\right)U_1(\text{boxing, ballet}) + \left(\frac{8}{9}\right)\left(\frac{1}{5}\right)U_1(\text{boxing, boxing}) \\ &= \left(\frac{1}{9}\right)\left(\frac{4}{5}\right)(2) + \left(\frac{1}{9}\right)\left(\frac{1}{5}\right)(0) + \left(\frac{8}{9}\right)\left(\frac{4}{5}\right)(0) + \left(\frac{8}{9}\right)\left(\frac{1}{5}\right)(1) \\ &= \frac{16}{45}. \end{aligned} \tag{8.4}$$

To understand Equation 8.4, it is helpful to review the concept of expected value from Chapter 2. The expected value of a random variable equals the sum over all outcomes of the probability of the outcome multiplied by the value of the random variable in that outcome. In the Battle of the Sexes, there are four outcomes, corresponding to the four boxes in Figure 8.3. Because players randomize independently, the probability of reaching a particular box equals the product of the probabilities that each player plays the strategy leading to that box. Thus, for example, the probability (boxing, ballet)—that is, the wife plays boxing and the husband plays ballet—equals (8/9) × (4/5). The probabilities of the four outcomes are multiplied by the value of the relevant random variable (in this case, players 1's payoff) in each outcome.

Next we compute the wife's expected payoff if she plays the pure strategy of going to ballet [the same as the mixed strategy (1, 0)] and the husband continues to play the mixed strategy (4/5, 1/5). Now there are only two relevant outcomes, given by the two boxes in the row in which the wife plays ballet. The probabilities of the two outcomes are given by the probabilities in the husband's mixed strategy. Therefore,

$$\begin{aligned} U_1\left(\text{ballet}, \left(\frac{4}{5}, \frac{1}{5}\right)\right) &= \left(\frac{4}{5}\right)U_1(\text{ballet, ballet}) + \left(\frac{1}{5}\right)U_1(\text{ballet, boxing}) \\ &= \left(\frac{4}{5}\right)(2) + \left(\frac{1}{5}\right)(0) = \frac{8}{5}. \end{aligned} \tag{8.5}$$

Finally, we will compute the general expression for the wife's expected payoff when she plays mixed strategy $(w, 1 - w)$ and the husband plays $(h, 1 - h)$: If the wife plays ballet with probability w and the husband with probability h, then

$$\begin{aligned} U_1((w, 1-w), (h, 1-h)) &= (w)(h)U_1(\text{ballet, ballet}) + (w)(1-h)U_1(\text{ballet, boxing}) \\ &\quad + (1-w)(h)U_1(\text{boxing, ballet}) \\ &\quad + (1-w)(1-h)U_1(\text{boxing, boxing}) \\ &= (w)(h)(2) + (w)(1-h)(0) + (1-w)(h)(0) \\ &\quad + (1-w)(1-h)(1) \\ &= 1 - h - w + 3hw. \end{aligned} \tag{8.6}$$

QUERY: What is the husband's expected payoff in each case? Show that his expected payoff is $2 - 2h - 2w + 3hw$ in the general case. Given the husband plays the mixed strategy (4/5, 1/5), what strategy provides the wife with the highest payoff?

Computing mixed-strategy equilibria

Computing Nash equilibria of a game when strictly mixed strategies are involved is a bit more complicated than when pure strategies are involved. Before wading in, we can save a lot of work by asking whether the game even has a Nash equilibrium in strictly mixed strategies. If it does not, having found all the pure-strategy Nash equilibria, then one has finished analyzing the game. The key to guessing whether a game has a Nash equilibrium in strictly mixed strategies is the surprising result that almost all games have an odd number of Nash equilibria.[5]

Let's apply this insight to some of the examples considered thus far. We found an odd number (one) of pure-strategy Nash equilibria in the Prisoners' Dilemma, suggesting we need not look further for one in strictly mixed strategies. In the Battle of the Sexes, we found an even number (two) of pure-strategy Nash equilibria, suggesting the existence of a third one in strictly mixed strategies. Example 8.2—Rock, Paper, Scissors—has no pure-strategy Nash equilibria. To arrive at an odd number of Nash equilibria, we would expect to find one Nash equilibrium in strictly mixed strategies.

EXAMPLE 8.4 Mixed-Strategy Nash Equilibrium in the Battle of the Sexes

A general mixed strategy for the wife in the Battle of the Sexes is $(w, 1 - w)$ and for the husband is $(h, 1 - h)$, where w and h are the probabilities of playing ballet for the wife and husband, respectively. We will compute values of w and h that make up Nash equilibria. Both players have a continuum of possible strategies between 0 and 1. Therefore, we cannot write these strategies in the rows and columns of a matrix and underline best-response payoffs to find the Nash equilibria. Instead, we will use graphical methods to solve for the Nash equilibria.

Given players' general mixed strategies, we saw in Example 8.3 that the wife's expected payoff is

$$U_1((w, 1 - w), (h, 1 - h)) = 1 - h - w + 3hw. \qquad \textbf{(8.7)}$$

As Equation 8.7 shows, the wife's best response depends on h. If $h < 1/3$, she wants to set w as low as possible: $w = 0$. If $h > 1/3$, her best response is to set w as high as possible: $w = 1$. When $h = 1/3$, her expected payoff equals 2/3 regardless of what w she chooses. In this case there is a tie for the best response, including any w from 0 to 1.

In Example 8.3, we stated that the husband's expected payoff is

$$U_2((h, 1 - h), (w, 1 - w)) = 2 - 2h - 2w + 3hw. \qquad \textbf{(8.8)}$$

When $w < 2/3$, his expected payoff is maximized by $h = 0$; when $w > 2/3$, his expected payoff is maximized by $h = 1$; and when $w = 2/3$, he is indifferent among all values of h, obtaining an expected payoff of 2/3 regardless.

The best responses are graphed in Figure 8.7. The Nash equilibria are given by the intersection points between the best responses. At these intersection points, both players are best responding to each other, which is what is required for the outcome to be a Nash equilibrium. There are three Nash equilibria. The points E_1 and E_2 are the pure-strategy Nash equilibria we found before, with E_1 corresponding to the pure-strategy Nash equilibrium in which both play boxing and E_2 to that in which both play ballet. Point E_3 is the strictly mixed-strategy Nash equilibrium, which can be spelled out as "the wife plays ballet with probability 2/3 and boxing with probability 1/3 and the husband plays ballet with probability 1/3 and boxing with probability 2/3." More succinctly, having defined w and h, we may write the equilibrium as "$w^* = 2/3$ and $h^* = 1/3$."

[5] John Harsanyi, "Oddness of the Number of Equilibrium Points: A New Proof," *International Journal of Game Theory* 2 (1973): 235–50. Games in which there are ties between payoffs may have an even or infinite number of Nash equilibria. Example 8.1, the Prisoners' Dilemma Redux, has several payoff ties. The game has four pure-strategy Nash equilibria and an infinite number of different mixed-strategy equilibria.

FIGURE 8.7 Nash Equilibria in Mixed Strategies in the Battle of the Sexes

Ballet is chosen by the wife with probability w and by the husband with probability h. Players' best responses are graphed on the same set of axes. The three intersection points E_1, E_2, and E_3 are Nash equilibria. The Nash equilibrium in strictly mixed strategies, E_3, is $w^* = 2/3$ and $h^* = 1/3$.

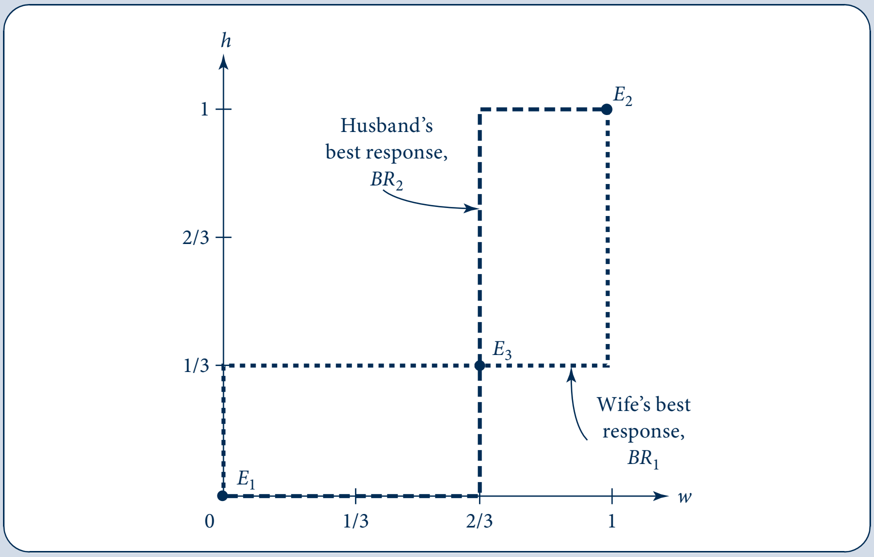

QUERY: What is a player's expected payoff in the Nash equilibrium in strictly mixed strategies? How does this payoff compare with those in the pure-strategy Nash equilibria? What arguments might be offered that one or another of the three Nash equilibria might be the best prediction in this game?

Example 8.4 runs through the lengthy calculations involved in finding all the Nash equilibria of the Battle of the Sexes, those in pure strategies and those in strictly mixed strategies. A shortcut to finding the Nash equilibrium in strictly mixed strategies is based on the insight that a player will be willing to randomize between two actions in equilibrium only if he or she gets the same expected payoff from playing either action or, in other words, is indifferent between the two actions in equilibrium. Otherwise, one of the two actions would provide a higher expected payoff, and the player would prefer to play that action with certainty.

Suppose the husband is playing mixed strategy $(h, 1-h)$, that is, playing ballet with probability h and boxing with probability $1-h$. The wife's expected payoff from playing ballet is

$$U_1(\text{ballet}, (h, 1-h)) = (h)(2) + (1-h)(0) = 2h. \quad \textbf{(8.9)}$$

Her expected payoff from playing boxing is

$$U_1(\text{boxing}, (h, 1-h)) = (h)(0) + (1-h)(1) = 1-h. \quad \textbf{(8.10)}$$

For the wife to be indifferent between ballet and boxing in equilibrium, Equations 8.9 and 8.10 must be equal: $2h = 1-h$, implying $h^* = 1/3$. Similar calculations based on the husband's indifference between playing ballet and boxing in equilibrium show that the wife's probability of playing ballet in the strictly mixed strategy Nash equilibrium is $w^* = 2/3$. (Work through these calculations as an exercise.)

Notice that the wife's indifference condition does not "pin down" her equilibrium mixed strategy. The wife's indifference condition cannot pin down her own equilibrium mixed strategy because, given that she is indifferent between the two actions in equilibrium, her overall expected payoff is the same no matter what probability distribution she plays over the two actions. Rather, the wife's indifference condition pins down the other player's—the husband's—mixed strategy. There is a unique probability distribution he can use to play ballet and boxing that makes her indifferent between the two actions and thus makes her willing to randomize. Given any probability of his playing ballet and boxing other than (1/3, 2/3), it would not be a stable outcome for her to randomize.

Thus, two principles should be kept in mind when seeking Nash equilibria in strictly mixed strategies. One is that a player randomizes over only those actions among which he or she is indifferent, given other players' equilibrium mixed strategies. The second is that one player's indifference condition pins down the *other* player's mixed strategy.

EXISTENCE OF EQUILIBRIUM

One of the reasons Nash equilibrium is so widely used is that a Nash equilibrium is guaranteed to exist in a wide class of games. This is not true for some other equilibrium concepts. Consider the dominant strategy equilibrium concept. The Prisoners' Dilemma has a dominant strategy equilibrium (both suspects fink), but most games do not. Indeed, there are many games—including, for example, the Battle of the Sexes—in which no player has a dominant strategy, let alone all the players. In such games, we cannot make predictions using dominant strategy equilibrium, but we can using Nash equilibrium.

The Extensions section at the end of this chapter will provide the technical details behind John Nash's proof of the existence of his equilibrium in all finite games (games with a finite number of players and a finite number of actions). The existence theorem does not guarantee the existence of a pure-strategy Nash equilibrium. We already saw an example: Rock, Paper, Scissors in Example 8.2. However, if a finite game does not have a pure-strategy Nash equilibrium, the theorem guarantees that it will have a mixed-strategy Nash equilibrium. The proof of Nash's theorem is similar to the proof in Chapter 13 of the existence of prices leading to a general competitive equilibrium. The Extensions section includes an existence theorem for games with a continuum of actions, as studied in the next section.

CONTINUUM OF ACTIONS

Most of the insight from economic situations can often be gained by distilling the situation down to a few or even two actions, as with all the games studied thus far. Other times, additional insight can be gained by allowing a continuum of actions. To be clear, we have already encountered a continuum of *strategies*—in our discussion of mixed strategies—but still the probability distributions in mixed strategies were over a finite number of actions. In this section we focus on continuum of *actions*.

Some settings are more realistically modeled via a continuous range of actions. In Chapter 15, for example, we will study competition between strategic firms. In one model (Bertrand), firms set prices; in another (Cournot), firms set quantities. It is natural to allow firms to choose any non-negative price or quantity rather than artificially restricting them to just two prices (say, \$2 or \$5) or two quantities (say, 100 or 1,000 units). Continuous actions have several other advantages. The familiar methods from calculus can often be used to solve for Nash equilibria. It is also possible to analyze how the equilibrium actions vary with changes in underlying parameters. With the Cournot model, for example, we might want to know how equilibrium quantities change with a small increase in a firm's marginal costs or a demand parameter.

Tragedy of the Commons

Example 8.5 illustrates how to solve for the Nash equilibrium when the game (in this case, the Tragedy of the Commons) involves a continuum of actions. The first step is to write down the payoff for each player as a function of all players' actions. The next step is to compute the first-order condition associated with each player's payoff maximum. This will give an equation that can be rearranged into the best response of each player as a function of all other players' actions. There will be one equation for each player. With n players, the system of n equations for the n unknown equilibrium actions can be solved simultaneously by either algebraic or graphical methods.

EXAMPLE 8.5 Tragedy of the Commons

The term *Tragedy of the Commons* has come to signify environmental problems of overuse that arise when scarce resources are treated as common property.[6] A game-theoretic illustration of this issue can be developed by assuming that two herders decide how many sheep to graze on the village commons. The problem is that the commons is small and can rapidly succumb to overgrazing.

To add some mathematical structure to the problem, let q_i be the number of sheep that herder $i = 1, 2$ grazes on the commons, and suppose that the per-sheep value of grazing on the commons (in terms of wool and sheep-milk cheese) is

$$v(q_1, q_2) = 120 - (q_1 + q_2). \quad \textbf{(8.11)}$$

This function implies that the value of grazing a given number of sheep is lower the more sheep are around competing for grass. We cannot use a matrix to represent the normal form of this game of continuous actions. Instead, the normal form is simply a listing of the herders' payoff functions

$$\begin{aligned} u_1(q_1, q_2) &= q_1 v(q_1, q_2) = q_1(120 - q_1 - q_2), \\ u_2(q_1, q_2) &= q_2 v(q_1, q_2) = q_2(120 - q_1 - q_2). \end{aligned} \quad \textbf{(8.12)}$$

To find the Nash equilibrium, we solve herder 1's value-maximization problem:

$$\max_{q_1}\{q_1(120 - q_1 - q_2)\}. \quad \textbf{(8.13)}$$

The first-order condition for a maximum is

$$120 - 2q_1 - q_2 = 0 \quad \textbf{(8.14)}$$

or, rearranging,

$$q_1 = 60 - \frac{q_2}{2} = BR_1(q_2). \quad \textbf{(8.15)}$$

Similar steps show that herder 2's best response is

$$q_2 = 60 - \frac{q_1}{2} = BR_2(q_1). \quad \textbf{(8.16)}$$

The Nash equilibrium is given by the pair (q_1^*, q_2^*) that satisfies Equations 8.15 and 8.16 simultaneously. Taking an algebraic approach to the simultaneous solution, Equation 8.16 can be substituted into Equation 8.15, which yields

[6]This term was popularized by G. Hardin, "The Tragedy of the Commons," *Science* 162 (1968): 1243–48.

$$q_1 = 60 - \frac{1}{2}\left(60 - \frac{q_1}{2}\right); \tag{8.17}$$

on rearranging, this implies $q_1^* = 40$. Substituting $q_1^* = 40$ into Equation 8.17 implies $q_2^* = 40$ as well. Thus, each herder will graze 40 sheep on the common. Each earns a payoff of 1,600, as can be seen by substituting $q_1^* = q_2^* = 40$ into the payoff function in Equation 8.13.

Equations 8.15 and 8.16 can also be solved simultaneously using graphical methods. Figure 8.8 plots the two best responses on a graph with player 1's action on the horizontal axis and

FIGURE 8.8 Best-Response Diagram for the Tragedy of the Commons

The intersection, E_1, between the two herders' best responses is the Nash equilibrium. An increase in the per-sheep value of grazing in the Tragedy of the Commons shifts out herder 1's best response, resulting in a Nash equilibrium E_2 in which herder 1 grazes more sheep (and herder 2, fewer sheep) than in the original Nash equilibrium.

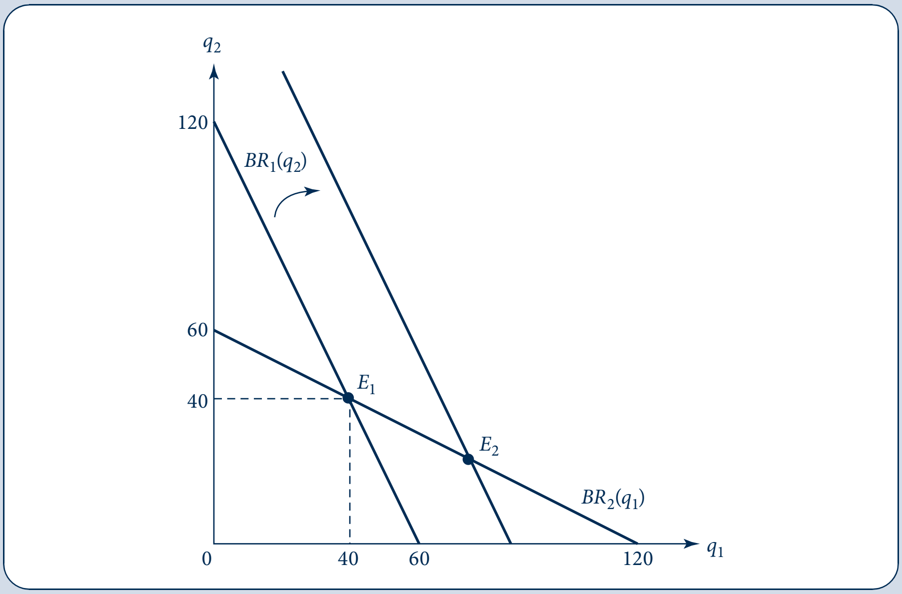

player 2's on the vertical axis. These best responses are simply lines and thus are easy to graph in this example. (To be consistent with the axis labels, the inverse of Equation 8.15 is actually what is graphed.) The two best responses intersect at the Nash equilibrium E_1.

The graphical method is useful for showing how the Nash equilibrium shifts with changes in the parameters of the problem. Suppose the per-sheep value of grazing increases for the first herder while the second remains as in Equation 8.11, perhaps because the first herder starts raising merino sheep with more valuable wool. This change would shift the best response out for herder 1 while leaving herder 2's the same. The new intersection point (E_2 in Figure 8.8), which is the new Nash equilibrium, involves more sheep for 1 and fewer for 2.

The Nash equilibrium is not the best use of the commons. In the original problem, both herders' per-sheep value of grazing is given by Equation 8.11. If both grazed only 30 sheep, then each would earn a payoff of 1,800, as can be seen by substituting $q_1 = q_2 = 30$ into Equation 8.13. Indeed, the "joint payoff maximization" problem

$$\max_{q_1, q_2}\{(q_1 + q_2)v(q_1, q_2)\} = \max_{q_1, q_2}\{(q_1 + q_2)(120 - q_1 - q_2)\} \tag{8.18}$$

is solved by $q_1 = q_2 = 30$ or, more generally, by any q_1 and q_2 that sum to 60.

QUERY: How would the Nash equilibrium shift if both herders' benefits increased by the same amount? What about a decrease in (only) herder 2's benefit from grazing?

As Example 8.5 shows, graphical methods are particularly convenient for quickly determining how the equilibrium shifts with changes in the underlying parameters. The example shifted the benefit of grazing to one of herders. This exercise nicely illustrates the nature of strategic interaction. Herder 2's payoff function has not changed (only herder 1's has), yet his equilibrium action changes. The second herder observes the first's higher benefit, anticipates that the first will increase the number of sheep he grazes, and reduces his own grazing in response.

The Tragedy of the Commons shares with the Prisoners' Dilemma the feature that the Nash equilibrium is less efficient for all players than some other outcome. In the Prisoners' Dilemma, both fink in equilibrium when it would be more efficient for both to be silent. In the Tragedy of the Commons, the herders graze more sheep in equilibrium than is efficient. This insight may explain why ocean fishing grounds and other common resources can end up being overused even to the point of exhaustion if their use is left unregulated. More detail on such problems—involving what we will call *negative externalities*—is provided in Chapter 19.

SEQUENTIAL GAMES

In some games, the order of moves matters. For example, in a bicycle race with a staggered start, it may help to go last and thus know the time to beat. On the other hand, competition to establish a new high-definition video format may be won by the first firm to market its technology, thereby capturing an installed base of consumers.

Sequential games differ from the simultaneous games we have considered thus far in that a player who moves later in the game can observe how others have played up to that moment. The player can use this information to form more sophisticated strategies than simply choosing an action; the player's strategy can be a contingent plan with the action played depending on what the other players have done.

To illustrate the new concepts raised by sequential games—and, in particular, to make a stark contrast between sequential and simultaneous games—we take a simultaneous game we have discussed already, the Battle of the Sexes, and turn it into a sequential game.

Sequential Battle of the Sexes

Consider the Battle of the Sexes game analyzed previously with all the same actions and payoffs, but now change the timing of moves. Rather than the wife and husband making a simultaneous choice, the wife moves first, choosing ballet or boxing; the husband observes this choice (say, the wife calls him from her chosen location), and then the husband makes his choice. The wife's possible strategies have not changed: She can choose the simple actions ballet or boxing (or perhaps a mixed strategy involving both actions, although this will not be a relevant consideration in the sequential game). The husband's set of possible strategies has expanded. For each of the wife's two actions, he can choose one of two actions; therefore, he has four possible strategies, which are listed in Table 8.1.

TABLE 8.1 HUSBAND'S CONTINGENT STRATEGIES

Contingent Strategy	Written in Conditional Format
Always go to the ballet	(ballet | ballet, ballet | boxing)
Follow his wife	(ballet | ballet, boxing | boxing)
Do the opposite	(boxing | ballet, ballet | boxing)
Always go to boxing	(boxing | ballet, boxing | boxing)

The vertical bar in the husband's strategies means "conditional on" and thus, for example, "boxing | ballet" should be read as "the husband chooses boxing conditional on the wife's choosing ballet."

Given that the husband has four pure strategies rather than just two, the normal form (given in Figure 8.9) must now be expanded to eight boxes. Roughly speaking, the normal form is twice as complicated as that for the simultaneous version of the game in Figure 8.2. This motivates a new way to represent games, called the *extensive form*, which is especially convenient for sequential games.

Extensive form

The *extensive form* of a game shows the order of moves as branches of a tree rather than collapsing everything down into a matrix. The extensive form for the sequential Battle of the Sexes is shown in Figure 8.10a. The action proceeds from left to right. Each node (shown as a dot on the tree) represents a decision point for the player indicated there. The first move belongs to the wife. After any action she might take, the husband gets to move. Payoffs are listed at the end of the tree in the same order (player 1's, player 2's) as in the normal form.

Contrast Figure 8.10a with Figure 8.10b, which shows the extensive form for the simultaneous version of the game. It is hard to harmonize an extensive form, in which moves happen in progression, with a simultaneous game, in which everything happens at the same time. The trick is to pick one of the two players to occupy the role of the second mover but then highlight that he or she is not really the second mover by connecting his or her decision nodes in the same *information set*, the dotted oval around the nodes. The dotted oval in Figure 8.10b indicates that the husband does not know his wife's move when he chooses his action. It does not matter which player is picked for first and second mover in a simultaneous game; we picked the husband in the figure to make the extensive form in Figure 8.10b look as much like Figure 8.10a as possible.

The similarity between the two extensive forms illustrates the point that that form does not grow in complexity for sequential games the way the normal form does. We

FIGURE 8.9

Normal Form for the Sequential Battle of the Sexes

		Husband			
		(Ballet | Ballet Ballet | Boxing)	(Ballet | Ballet Boxing | Boxing)	(Boxing | Ballet Ballet | Boxing)	(Boxing | Ballet Boxing | Boxing)
Wife	Ballet	2, 1	2, 1	0, 0	0, 0
	Boxing	0, 0	1, 2	0, 0	1, 2

FIGURE 8.10

Extensive Form for the Battle of the Sexes

In the sequential version (a), the husband moves second, after observing his wife's move. In the simultaneous version (b), he does not know her choice when he moves, so his decision nodes must be connected in one information set.

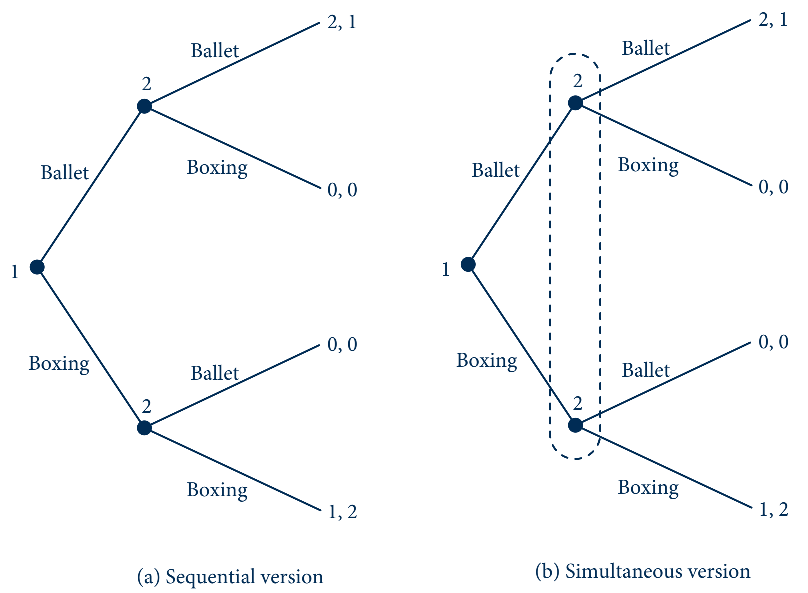

next will draw on both normal and extensive forms in our analysis of the sequential Battle of the Sexes.

Nash equilibria

To solve for the Nash equilibria, return to the normal form in Figure 8.9. Applying the method of underlining best-response payoffs—being careful to underline both payoffs in cases of ties for the best response—reveals three pure-strategy Nash equilibria:

1. wife plays ballet, husband plays (ballet | ballet, ballet | boxing);
2. wife plays ballet, husband plays (ballet | ballet, boxing | boxing);
3. wife plays boxing, husband plays (boxing | ballet, boxing | boxing).

As with the simultaneous version of the Battle of the Sexes, here again we have multiple equilibria. Yet now game theory offers a good way to select among the equilibria. Consider the third Nash equilibrium. The husband's strategy (boxing | ballet, boxing | boxing) involves the implicit threat that he will choose boxing even if his wife chooses ballet. This threat is sufficient to deter her from choosing ballet. Given that she chooses boxing in equilibrium, his strategy earns him 2, which is the best he can do in any outcome. Thus, the outcome is a Nash equilibrium. But the husband's threat is not credible—that is, it is an empty threat. If the wife really were to choose ballet first, then he would give up a payoff of 1 by choosing boxing rather than ballet. It is clear why he would want to threaten to choose boxing, but it is not clear that such a threat should be believed. Similarly, the husband's strategy (ballet | ballet, ballet | boxing) in the first Nash equilibrium also involves an empty threat: that he will choose ballet if his wife chooses boxing. (This is an odd threat to make because he does not gain from making it, but it is an empty threat nonetheless.)

Another way to understand empty versus credible threats is by using the concept of the *equilibrium path*, the connected path through the extensive form implied by equilibrium strategies. In Figure 8.11, which reproduces the extensive form of the sequential Battle of the Sexes from Figure 8.10, a dotted line is used to identify the equilibrium path for the third of the listed Nash equilibria. The third outcome is a Nash equilibrium because the strategies are rational along the equilibrium path. However, following the wife's choosing ballet—an event that is off the equilibrium path—the husband's strategy is irrational. The concept of subgame-perfect equilibrium in the next section will rule out irrational play both on and off the equilibrium path.

Subgame-perfect equilibrium

Game theory offers a formal way of selecting the reasonable Nash equilibria in sequential games using the concept of subgame-perfect equilibrium. Subgame-perfect equilibrium is a refinement that rules out empty threats by requiring strategies to be rational even for contingencies that do not arise in equilibrium.

Before defining subgame-perfect equilibrium formally, we need a few preliminary definitions. A *subgame* is a part of the extensive form beginning with a decision node and including everything that branches out to the right of it. A *proper subgame* is a subgame

FIGURE 8.11

Equilibrium Path

In the third of the Nash equilibria listed for the sequential Battle of the Sexes, the wife plays boxing and the husband plays (boxing | ballet, boxing | boxing), tracing out the branches indicated with thick lines (both solid and dashed). The dashed line is the equilibrium path; the rest of the tree is referred to as being "off the equilibrium path."

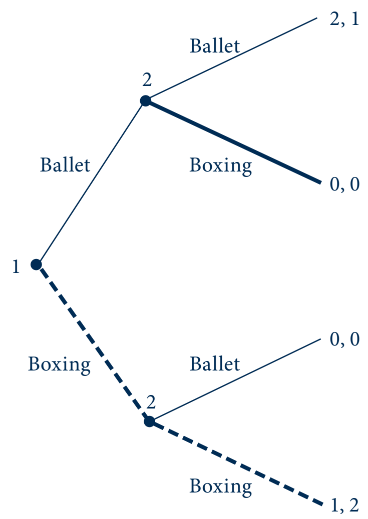

that starts at a decision node not connected to another in an information set. Conceptually, this means that the player who moves first in a proper subgame knows the actions played by others that have led up to that point. It is easier to see what a proper subgame is than to define it in words. Figure 8.12 shows the extensive forms from the simultaneous and sequential versions of the Battle of the Sexes with boxes drawn around the proper subgames in each. The sequential version (a) has three proper subgames: the game itself and two lower subgames starting with decision nodes where the husband gets to move. The simultaneous version (b) has only one decision node—the topmost node—not connected to another in an information set. Hence this verion has only one subgame: the whole game itself.

DEFINITION

Subgame-perfect equilibrium. A subgame-perfect equilibrium is a strategy profile $(s_1^*, s_2^*, \ldots, s_n^*)$ that is a Nash equilibrium on every proper subgame.

A subgame-perfect equilibrium is always a Nash equilibrium. This is true because the whole game is a proper subgame of itself; thus, a subgame-perfect equilibrium must be a Nash equilibrium for the whole game. In the simultaneous version of the Battle of the Sexes, there is nothing more to say because there are no subgames other than the whole game itself.

In the sequential version, subgame-perfect equilibrium has more bite. Strategies must not only form a Nash equilibrium on the whole game itself; they must also form Nash

FIGURE 8.12

Proper Subgames in the Battle of the Sexes

The sequential version in (a) has three proper subgames, labeled *A*, *B*, and *C*. The simultaneous version in (b) has only one proper subgame: the whole game itself, labeled *D*.

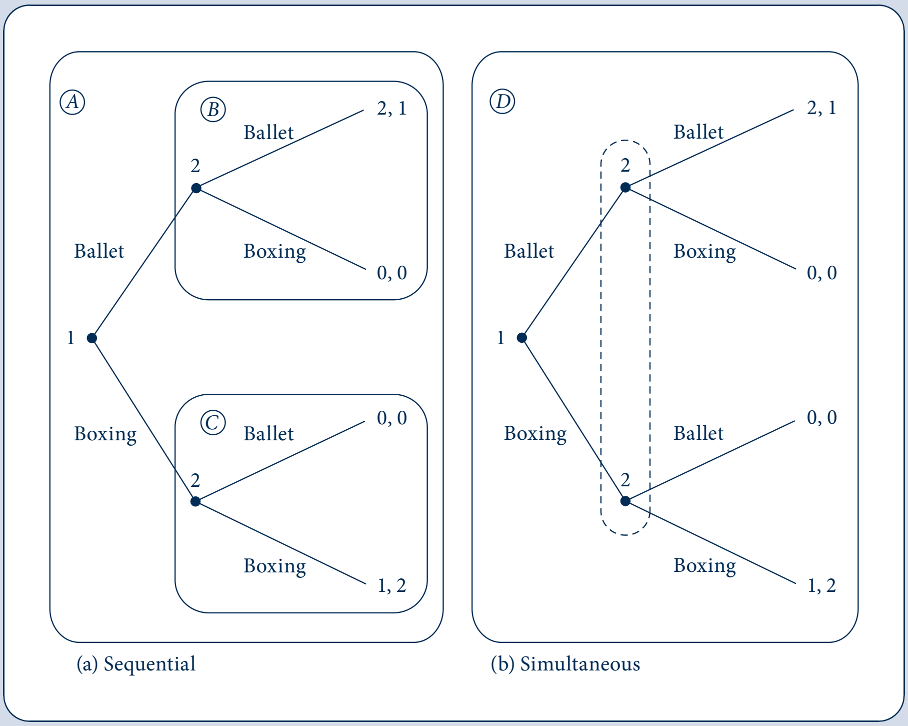

equilibria on the two proper subgames starting with the decision points at which the husband moves. These subgames are simple decision problems, so it is easy to compute the corresponding Nash equilibria. For subgame *B*, beginning with the husband's decision node following his wife's choosing ballet, he has a simple decision between ballet (which earns him a payoff of 1) and boxing (which earns him a payoff of 0). The Nash equilibrium in this simple decision subgame is for the husband to choose ballet. For the other subgame, *C*, he has a simple decision between ballet, which earns him 0, and boxing, which earns him 2. The Nash equilibrium in this simple decision subgame is for him to choose boxing. Therefore, the husband has only one strategy that can be part of a subgame-perfect equilibrium: (ballet | ballet, boxing | boxing). Any other strategy has him playing something that is not a Nash equilibrium for some proper subgame. Returning to the three enumerated Nash equilibria, only the second is subgame perfect; the first and the third are not. For example, the third equilibrium, in which the husband always goes to boxing, is ruled out as a subgame-perfect equilibrium because the husband's strategy (boxing | boxing) is not a Nash equilibrium in proper subgame *B*. Thus, subgame-perfect equilibrium rules out the empty threat (of always going to boxing) that we were uncomfortable with earlier.

More generally, subgame-perfect equilibrium rules out any sort of empty threat in a sequential game. In effect, Nash equilibrium requires behavior to be rational only on the equilibrium path. Players can choose potentially irrational actions on other parts of the extensive form. In particular, one player can threaten to damage both to scare the other from choosing certain actions. Subgame-perfect equilibrium requires rational behavior both on and off the equilibrium path. Threats to play irrationally—that is, threats to choose something other than one's best response—are ruled out as being empty.

Backward induction

Our approach to solving for the equilibrium in the sequential Battle of the Sexes was to find all the Nash equilibria using the normal form and then to seek among those for the subgame-perfect equilibrium. A shortcut for finding the subgame-perfect equilibrium directly is to use *backward induction*, the process of solving for equilibrium by working backward from the end of the game to the beginning. Backward induction works as follows. Identify all the subgames at the bottom of the extensive form. Find the Nash equilibria on these subgames. Replace the (potentially complicated) subgames with the actions and payoffs resulting from Nash equilibrium play on these subgames. Then move up to the next level of subgames and repeat the procedure.

Figure 8.13 illustrates the use of backward induction in the sequential Battle of the Sexes. First, we compute the Nash equilibria of the bottom-most subgames at the husband's decision nodes. In the subgame following his wife's choosing ballet, he would choose ballet, giving payoffs 2 for her and 1 for him. In the subgame following his wife's choosing boxing, he would choose boxing, giving payoffs 1 for her and 2 for him. Next, substitute the husband's equilibrium strategies for the subgames themselves. The resulting game is a simple decision problem for the wife (drawn in the lower panel of the figure): a choice between ballet, which would give her a payoff of 2, and boxing, which would give her a payoff of 1. The Nash equilibrium of this game is for her to choose the action with the higher payoff, ballet. In sum, backward induction allows us to jump straight to the subgame-perfect equilibrium in which the wife chooses ballet and the husband chooses (ballet | ballet, boxing | boxing), bypassing the other Nash equilibria.

Backward induction is particularly useful in games that feature many rounds of sequential play. As rounds are added, it quickly becomes too hard to solve for all the Nash

FIGURE 8.13

Applying Backward Induction

The last subgames (where player 2 moves) are replaced by the Nash equilibria on these subgames. The simple game that results at right can be solved for player 1's equilibrium action.

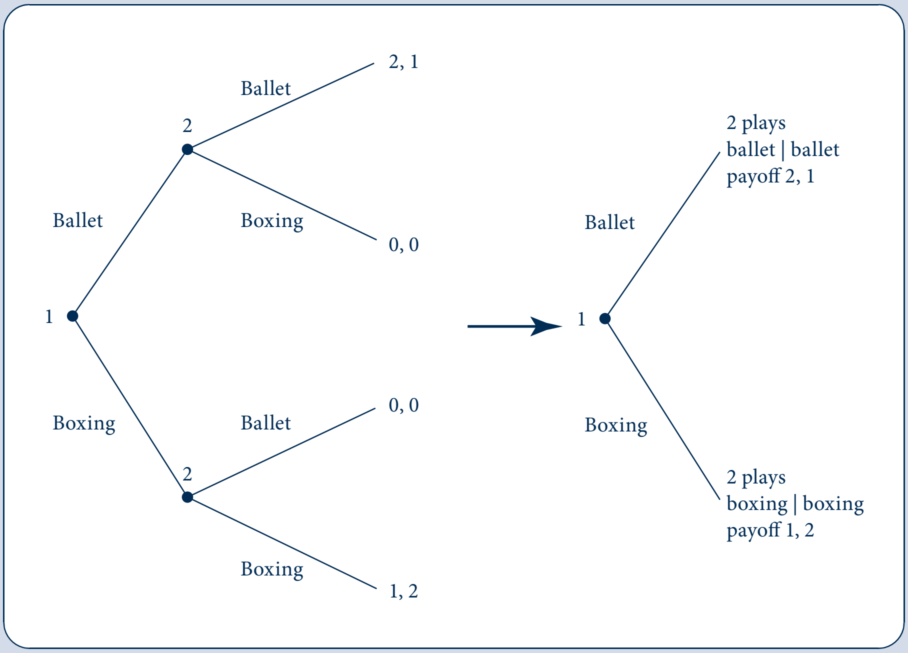

equilibria and then to sort through which are subgame-perfect. With backward induction, an additional round is simply accommodated by adding another iteration of the procedure.

REPEATED GAMES

In the games examined thus far, each player makes one choice and the game ends. In many real-world settings, players play the same game over and over again. For example, the players in the Prisoners' Dilemma may anticipate committing future crimes and thus playing future Prisoners' Dilemmas together. Gasoline stations located across the street from each other, when they set their prices each morning, effectively play a new pricing game every day. The simple constituent game (e.g., the Prisoners' Dilemma or the gasoline-pricing game) that is played repeatedly is called the *stage game*. As we saw with the Prisoners' Dilemma, the equilibrium in one play of the stage game may be worse for all players than some other, more cooperative, outcome. Repeated play of the stage game opens up the possibility of cooperation in equilibrium. Players can adopt *trigger strategies*, whereby they continue to cooperate as long as all have cooperated up to that point but revert to playing the Nash equilibrium if anyone deviates from cooperation. We will investigate the conditions under which trigger strategies work to increase players' payoffs. As is standard in game theory, we will focus on subgame-perfect equilibria of the repeated games.

Finitely repeated games

For many stage games, repeating them a known, finite number of times does not increase the possibility for cooperation. To see this point concretely, suppose the Prisoners'

Dilemma were played repeatedly for T periods. Use backward induction to solve for the subgame-perfect equilibrium. The lowest subgame is the Prisoners' Dilemma stage game played in period T. Regardless of what happened before, the Nash equilibrium on this subgame is for both to fink. Folding the game back to period $T-1$, trigger strategies that condition period T play on what happens in period $T-1$ are ruled out. Although a player might like to promise to play cooperatively in period T and thus reward the other for playing cooperatively in period $T-1$, we have just seen that nothing that happens in period $T-1$ affects what happens subsequently because players both fink in period T regardless. It is as though period $T-1$ were the last, and the Nash equilibrium of this subgame is again for both to fink. Working backward in this way, we see that players will fink each period; that is, players will simply repeat the Nash equilibrium of the stage game T times.

Reinhard Selten, winner of the Nobel Prize in economics for his contributions to game theory, showed that this logic is general: For any stage game with a unique Nash equilibrium, the unique subgame-perfect equilibrium of the finitely repeated game involves playing the Nash equilibrium every period.[7]

If the stage game has multiple Nash equilibria, it may be possible to achieve some cooperation in a finitely repeated game. Players can use trigger strategies, sustaining cooperation in early periods on an outcome that is not an equilibrium of the stage game, by threatening to play in later periods the Nash equilibrium that yields a worse outcome for the player who deviates from cooperation.[8] Rather than delving into the details of finitely repeated games, we will instead turn to infinitely repeated games, which greatly expand the possibility of cooperation.

Infinitely repeated games

With finitely repeated games, the folk theorem applies only if the stage game has multiple equilibria. If, like the Prisoners' Dilemma, the stage game has only one Nash equilibrium, then Selten's result tells us that the finitely repeated game has only one subgame-perfect equilibrium: repeating the stage-game Nash equilibrium each period. Backward induction starting from the last period T unravels any other outcomes.

With infinitely repeated games, however, there is no definite ending period T from which to start backward induction. Outcomes involving cooperation do not necessarily end up unraveling. Under some conditions the opposite may be the case, with essentially anything being possible in equilibrium of the infinitely repeated game. This result is sometimes called the folk theorem because it was part of the "folk wisdom" of game theory before anyone bothered to prove it formally.

One difficulty with infinitely repeated games involves adding up payoffs across periods. An infinite stream of low payoffs sums to infinity just as an infinite stream of high payoffs. How can the two streams be ranked? We will circumvent this problem with the aid of discounting. Let δ be the discount factor (discussed in the Chapter 17 Appendix) measuring how much a payoff unit is worth if received one period in the future rather than today. In Chapter 17 we show that δ is inversely related to the interest rate.[9] If the interest rate is high, then a person would much rather receive payment today than next period because investing

---

[7] R. Selten, "A Simple Model of Imperfect Competition, Where 4 Are Few and 6 Are Many," *International Journal of Game Theory* 2 (1973): 141–201.

[8] J. P. Benoit and V. Krishna, "Finitely Repeated Games," *Econometrica* 53 (1985): 890–940.

[9] Beware of the subtle difference between the formulas for the present value of an annuity stream used here and in Chapter 17 Appendix. There the payments came at the end of the period rather than at the beginning as assumed here. So here the present value of \$1 payment per period from now on is

$$\$1 + \$1 \cdot \delta + \$1 \cdot \delta^2 + \$1 \cdot \delta^3 + \ldots = \frac{\$1}{1-\delta}.$$

today's payment would provide a return of principal plus a large interest payment next period. Besides the interest rate, δ can also incorporate uncertainty about whether the game continues in future periods. The higher the probability that the game ends after the current period, the lower the expected return from stage games that might not actually be played.

Factoring in a probability that the repeated game ends after each period makes the setting of an infinitely repeated game more believable. The crucial issue with an infinitely repeated game is not that it goes on forever but that its end is indeterminate. Interpreted in this way, there is a sense in which infinitely repeated games are more realistic than finitely repeated games with large T. Suppose we expect two neighboring gasoline stations to play a pricing game each day until electric cars replace gasoline-powered ones. It is unlikely the gasoline stations would know that electric cars were coming in exactly $T = 2{,}000$ days. More realistically, the gasoline stations will be uncertain about the end of gasoline-powered cars; thus, the end of their pricing game is indeterminate.

Players can try to sustain cooperation using trigger strategies. Trigger strategies have them continuing to cooperate as long as no one has deviated; deviation triggers some sort of punishment. The key question in determining whether trigger strategies "work" is whether the punishment can be severe enough to deter the deviation in the first place.

Suppose both players use the following specific trigger strategy in the Prisoners' Dilemma: Continue being silent if no one has deviated; fink forever afterward if anyone has deviated to fink in the past. To show that this trigger strategy forms a subgame-perfect equilibrium, we need to check that a player cannot gain from deviating. Along the equilibrium path, both players are silent every period; this provides each with a payoff of 2 every period for a present discounted value of

$$\begin{aligned} V^{eq} &= 2 + 2\delta + 2\delta^2 + 2\delta^3 + \cdots \\ &= 2(1 + \delta + \delta^2 + \delta^3 + \cdots) \\ &= \frac{2}{1-\delta}. \end{aligned} \tag{8.19}$$

A player who deviates by finking earns 3 in that period, but then both players fink every period from then on—each earning 1 per period for a total presented discounted payoff of

$$\begin{aligned} V^{dev} &= 3 + (1)(\delta) + (1)(\delta^2) + (1)(\delta^3) + \cdots \\ &= 3 + \delta(1 + \delta + \delta^2 + \cdots) \\ &= 3 + \frac{\delta}{1-\delta}. \end{aligned} \tag{8.20}$$

The trigger strategies form a subgame-perfect equilibrium if $V^{eq} \geq V^{dev}$, implying that

$$\frac{2}{1-\delta} \geq 3 + \frac{\delta}{1-\delta}. \tag{8.21}$$

After multiplying through by $1 - \delta$ and rearranging, we obtain $\delta \geq 1/2$. In other words, players will find continued cooperative play desirable provided they do not discount future gains from such cooperation too highly. If $\delta < 1/2$, then no cooperation is possible in the infinitely repeated Prisoners' Dilemma; the only subgame-perfect equilibrium involves finking every period.

The trigger strategy we considered has players revert to the stage-game Nash equilibrium of finking each period forever. This strategy, which involves the harshest possible punishment for deviation, is called the *grim strategy*. Less harsh punishments include the so-called tit-for-tat strategy, which involves only one round of punishment for cheating. Because the grim strategy involves the harshest punishment possible, it elicits cooperation for the largest range of cases (the lowest value of δ) of any strategy. Harsh punishments work well because, if players succeed in cooperating, they never experience the losses from the punishment in equilibrium.[10]

The discount factor δ is crucial in determining whether trigger strategies can sustain cooperation in the Prisoners' Dilemma or, indeed, in any stage game. As δ approaches 1, grim-strategy punishments become infinitely harsh because they involve an unending stream of undiscounted losses. Infinite punishments can be used to sustain a wide range of possible outcomes. This is the logic behind the *folk theorem for infinitely repeated games*. Take any stage-game payoff for a player between Nash equilibrium one and the highest one that appears anywhere in the payoff matrix. Let V be the present discounted value of the infinite stream of this payoff. The folk theorem says that the player can earn V in some subgame-perfect equilibrium for δ close enough to 1.[11]

INCOMPLETE INFORMATION

In the games studied thus far, players knew everything there was to know about the setup of the game, including each others' strategy sets and payoffs. Matters become more complicated, and potentially more interesting, if some players have information about the game that others do not. Poker would be different if all hands were played face up. The fun of playing poker comes from knowing what is in your hand but not others'. Incomplete information arises in many other real-world contexts besides parlor games. A sports team may try to hide the injury of a star player from future opponents to prevent them from exploiting this weakness. Firms' production technologies may be trade secrets, and thus firms may not know whether they face efficient or weak competitors. This section (and the next two) will introduce the tools needed to analyze games of incomplete information. The analysis integrates the material on game theory developed thus far in this chapter with the material on uncertainty and information from the previous chapter.

Games of incomplete information can quickly become complicated. Players who lack full information about the game will try to use what they do know to make inferences about what they do not. The inference process can be involved. In poker, for example, knowing what is in your hand can tell you something about what is in others'. A player who holds two aces knows that others are less likely to hold aces because two of the four aces are not available. Information on others' hands can also come from the size of their bets or from their facial expressions (of course, a big bet may be a bluff and a facial expression may be faked). Probability theory provides a formula, called *Bayes' rule*, for making inferences about hidden information. We will encounter Bayes' rule in a later section. The relevance of Bayes' rule in games of incomplete information has led them to be called *Bayesian games.*

To limit the complexity of the analysis, we will focus on the simplest possible setting throughout. We will focus on two-player games in which one of the players (player 1) has private information and the other (player 2) does not. The analysis of games of incomplete information is divided into two sections. The next section begins with the simple case in which the players move simultaneously. The subsequent section then

---

[10]Nobel Prize–winning economist Gary Becker introduced a related point, the maximal punishment principle for crime. The principle says that even minor crimes should receive draconian punishments, which can deter crime with minimal expenditure on policing. The punishments are costless to society because no crimes are committed in equilibrum, so punishments never have to be carried out. See G. Becker, "Crime and Punishment: An Economic Approach," *Journal of Political Economy* 76 (1968): 169–217. Less harsh punishments may be suitable in settings involving uncertainty. For example, citizens may not be certain about the penal code; police may not be certain they have arrested the guilty party.

[11]A more powerful version of the folk theorem was proved by D. Fudenberg and E. Maskin ("The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica* 54 (1986) 533–56). Payoffs below even the Nash equilibrium ones can be generated by some subgame-perfect equilibrium, payoffs all the way down to players' minmax level (the lowest level a player can be reduced to by all other players working against him or her).

analyzes games in which the informed player 1 moves first. Such games, called *signaling* games, are more complicated than simultaneous games because player 1's action may signal something about his or her private information to the uninformed player 2. We will introduce Bayes' rule at that point to help analyze player 2's inference about player 1's hidden information based on observations of player 1's action.

SIMULTANEOUS BAYESIAN GAMES

In this section we study a two-player, simultaneous-move game in which player 1 has private information but player 2 does not. (We will use "he" for player 1 and "she" for player 2 to facilitate the exposition.) We begin by studying how to model private information.

Player types and beliefs

John Harsanyi, who received the Nobel Prize in economics for his work on games with incomplete information, provided a simple way to model private information by introducing player characteristics or *types*.[12] Player 1 can be one of a number of possible such types, denoted t. Player 1 knows his own type. Player 2 is uncertain about t and must decide on her strategy based on beliefs about t.

Formally, the game begins at an initial node, called a *chance node*, at which a particular value t_k is randomly drawn for player 1's type t from a set of possible types $T = \{t_1, \ldots, t_k, \ldots, t_K\}$. Let $\Pr(t_k)$ be the probability of drawing the particular type t_k. Player 1 sees which type is drawn. Player 2 does not see the draw and only knows the probabilities, using them to form her beliefs about player 1's type. Thus, the probability that player 2 places on player 1's being of type t_k is $\Pr(t_k)$.

Because player 1 observes his type t before moving, his strategy can be conditioned on t. Conditioning on this information may be a big benefit to a player. In poker, for example, the stronger a player's hand, the more likely the player is to win the pot and the more aggressively the player may want to bid. Let $s_1(t)$ be player 1's strategy contingent on his type. Because player 2 does not observe t, her strategy is simply the unconditional one, s_2. As with games of complete information, players' payoffs depend on strategies. In Bayesian games, payoffs may also depend on types. Therefore, we write player 1's payoff as $u_1(s_1(t), s_2, t)$ and player 2's as $u_2(s_2, s_1(t), t)$. Note that t appears in two places in player 2's payoff function. Player 1's type may have a direct effect on player 2's payoffs. Player 1's type also has an indirect effect through its effect on player 1's strategy $s_1(t)$, which in turn affects player 2's payoffs. Because player 2's payoffs depend on t in these two ways, her beliefs about t will be crucial in the calculation of her optimal strategy.

Figure 8.14 provides a simple example of a simultaneous Bayesian game. Each player chooses one of two actions. All payoffs are known except for player 1's payoff when 1 chooses U and 2 chooses L. Player 1's payoff in outcome (U, L) is identified as his type, t. There are two possible values for player 1's type, $t = 6$ and $t = 0$, each occurring with equal probability. Player 1 knows his type before moving. Player 2's beliefs are that each type has probability 1/2. The extensive form is drawn in Figure 8.15.

Bayesian–Nash equilibrium

Extending Nash equilibrium to Bayesian games requires two small matters of interpretation. First, recall that player 1 may play a different action for each of his types. Equilibrium requires that player 1's strategy be a best response for each and every one of his types. Second, recall that player 2 is uncertain about player 1's type. Equilibrium requires

[12] J. Harsanyi, "Games with Incomplete Information Played by Bayesian Players," *Management Science* 14 (1967–68): 159–82, 320–34, 486–502.

FIGURE 8.14

Simple Game of Incomplete Information

$t = 6$ with probability 1/2 and $t = 0$ with probability 1/2.

Player 1 \ Player 2	L	R
U	t, 2	0, 0
D	2, 0	2, 4

FIGURE 8.15

Extensive Form for Simple Game of Incomplete Information

This figure translates Figure 8.14 into an extensive-form game. The initial chance node is indicated by an open circle. Player 2's decision nodes are in the same information set because she does not observe player 1's type or action before moving.

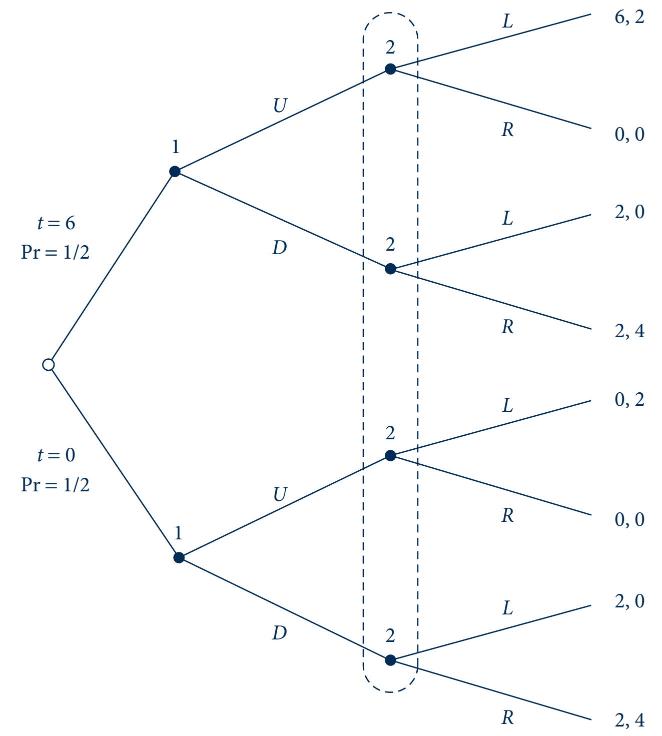

that player 2's strategy maximize an expected payoff, where the expectation is taken with respect to her beliefs about player 1's type. We encountered expected payoffs in our discussion of mixed strategies. The calculations involved in computing the best response to the pure strategies of different types of rivals in a game of incomplete information are similar to the calculations involved in computing the best response to a rival's mixed strategy in a game of complete information.

Interpreted in this way, Nash equilibrium in the setting of a Bayesian game is called *Bayesian–Nash equilibrium*. Next we provide a formal definition of the concept for reference. Given that the notation is fairly dense, it may be easier to first skip to Examples 8.6 and 8.7, which provide a blueprint on how to solve for equilibria in Bayesian games you might come across.

DEFINITION

Bayesian–Nash equilibrium. In a two-player, simultaneous-move game in which player 1 has private information, a Bayesian–Nash equilibrium is a strategy profile $(s_1^*(t), s_2^*)$ such that $s_1^*(t)$ is a best response to s_2^* for each type $t \in T$ of player 1,

$$U_1(s_1^*(t), s_2^*, t) \geq U_1(s_1', s_2^*, t) \quad \text{for all} \quad s_1' \in S_1, \tag{8.22}$$

and such that s_2^* is a best response to $s_1^*(t)$ given player 2's beliefs $\Pr(t_k)$ about player 1's types:

$$\sum_{t_k \in T} \Pr(t_k) U_2(s_2^*, s_1^*(t_k), t_k) \geq \sum_{t_k \in T} \Pr(t_k) U_2(s_2', s_1^*(t_k), t_k) \quad \text{for all} \quad s_2' \in S_2. \tag{8.23}$$

Because the difference between Nash equilibrium and Bayesian–Nash equilibrium is only a matter of interpretation, all our previous results for Nash equilibrium (including the existence proof) apply to Bayesian–Nash equilibrium as well.

EXAMPLE 8.6 Bayesian–Nash Equilibrium of Game in Figure 8.15

To solve for the Bayesian–Nash equilibrium of the game in Figure 8.15, first solve for the informed player's (player 1's) best responses for each of his types. If player 1 is of type $t = 0$, then he would choose D rather than U because he earns 0 by playing U and 2 by playing D regardless of what player 2 does. If player 1 is of type $t = 6$, then his best response is U to player 2's playing L and D to her playing R. This leaves only two possible candidates for an equilibrium in pure strategies:

$$1 \text{ plays } (U|t = 6, D|t = 0) \text{ and 2 plays } L;$$
$$1 \text{ plays } (D|t = 6, D|t = 0) \text{ and 2 plays } R.$$

The first candidate cannot be an equilibrium because, given that player 1 plays $(U|t = 6, D|t = 0)$, player 2 earns an expected payoff of 1 from playing L. Player 2 would gain by deviating to R, earning an expected payoff of 2.

The second candidate is a Bayesian–Nash equilibrium. Given that player 2 plays R, player 1's best response is to play D, providing a payoff of 2 rather than 0 regardless of his type. Given that both types of player 1 play D, player 2's best response is to play R, providing a payoff of 4 rather than 0.

QUERY: If the probability that player 1 is of type $t = 6$ is high enough, can the first candidate be a Bayesian–Nash equilibrium? If so, compute the threshold probability.

EXAMPLE 8.7 Tragedy of the Commons as a Bayesian Game

For an example of a Bayesian game with continuous actions, consider the Tragedy of the Commons in Example 8.5 but now suppose that herder 1 has private information regarding his value of grazing per sheep:

$$v_1(q_1, q_2, t) = t - (q_1 + q_2), \tag{8.24}$$

where herder 1's type is $t = 130$ (the "high" type) with probability 2/3 and $t = 100$ (the "low" type) with probability 1/3. Herder 2's value remains the same as in Equation 8.11.

To solve for the Bayesian–Nash equilibrium, we first solve for the informed player's (herder 1's) best responses for each of his types. For any type t and rival's strategy q_2, herder 1's value-maximization problem is

$$\max_{q_1}\{q_1 v_1(q_1, q_2, t)\} = \max_{q_1}\{q_1(t - q_1 - q_2)\}. \tag{8.25}$$

The first-order condition for a maximum is

$$t - 2q_1 - q_2 = 0. \tag{8.26}$$

Rearranging and then substituting the values $t = 130$ and $t = 100$, we obtain

$$q_{1H} = 65 - \frac{q_2}{2} \quad \text{and} \quad q_{1L} = 50 - \frac{q_2}{2}, \tag{8.27}$$

where q_{1H} is the quantity for the "high" type of herder 1 (i.e., the $t = 100$ type) and q_{1L} for the "low" type (the $t = 100$ type).

Next we solve for herder 2's best response. Herder 2's expected payoff is

$$\frac{2}{3}[q_2(120 - q_{1H} - q_2)] + \frac{1}{3}[q_2(120 - q_{1L} - q_2)] = q_2(120 - \overline{q}_1 - q_2), \tag{8.28}$$

where

$$\overline{q}_1 = \frac{2}{3}q_{1H} + \frac{1}{3}q_{1L}. \tag{8.29}$$

Rearranging the first-order condition from the maximization of Equation 8.28 with respect to q_2 gives

$$q_2 = 60 - \frac{\overline{q}_1}{2}. \tag{8.30}$$

Substituting for q_{1H} and q_{1L} from Equation 8.27 into Equation 8.29 and then substituting the resulting expression for $\overline{q}_1$ into Equation 8.30 yields

$$q_2 = 30 + \frac{q_2}{4}, \tag{8.31}$$

implying that $q_2^* = 40$. Substituting $q_2^* = 40$ back into Equation 8.27 implies $q_{1H}^* = 45$ and $q_{1L}^* = 30$.

Figure 8.16 depicts the Bayesian–Nash equilibrium graphically. Herder 2 imagines playing against an average type of herder 1, whose average best response is given by the thick dashed line. The intersection of this best response and herder 2's at point B determines herder 2's equilibrium quantity, $q_2^* = 40$. The best response of the low (resp. high) type of herder 1 to $q_2^* = 40$ is given by point A (resp. point C). For comparison, the full-information Nash equilibria are drawn when herder 1 is known to be the low type (point A') or the high type (point C').

QUERY: Suppose herder 1 is the high type. How does the number of sheep each herder grazes change as the game moves from incomplete to full information (moving from point C' to C)? What if herder 1 is the low type? Which type prefers full information and thus would like to signal its type? Which type prefers incomplete information and thus would like to hide its type? We will study the possibility player 1 can signal his type in the next section.

FIGURE 8.16 Equilibrium of the Bayesian Tragedy of the Commons

Best responses for herder 2 and both types of herder 1 are drawn as thick solid lines; the expected best response as perceived by 2 is drawn as the thick dashed line. The Bayesian–Nash equilibrium of the incomplete-information game is given by points *A* and *C*; Nash equilibria of the corresponding full-information games are given by points *A*′ and *C*′.

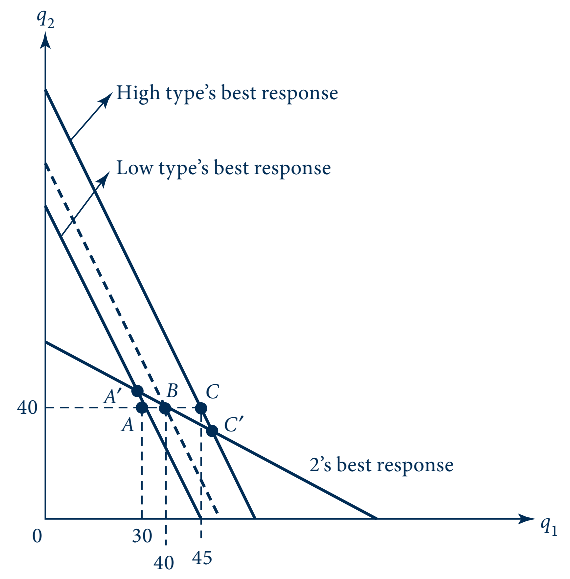

SIGNALING GAMES

In this section we move from simultaneous-move games of private information to sequential games in which the informed player, player 1, takes an action that is observable to player 2 before player 2 moves. Player 1's action provides information, a signal, that player 2 can use to update her beliefs about player 1's type, perhaps altering the way player 2 would play in the absence of such information. In poker, for example, player 2 may take a big raise by player 1 as a signal that he has a good hand, perhaps leading player 2 to fold. A firm considering whether to enter a market may take the incumbent firm's low price as a signal that the incumbent is a low-cost producer and thus a tough competitor, perhaps keeping the entrant out of the market. A prestigious college degree may signal that a job applicant is highly skilled.

The analysis of signaling games is more complicated than simultaneous games because we need to model how player 2 processes the information in player 1's signal and then updates her beliefs about player 1's type. To fix ideas, we will focus on a concrete application: a version of Michael Spence's model of job-market signaling, for which he won the Nobel Prize in economics.[13]

[13]M. Spence, "Job-Market Signaling," *Quarterly Journal of Economics* 87 (1973): 355–74.

Job-market signaling

Player 1 is a worker who can be one of two types, high-skilled ($t = H$) or low-skilled ($t = L$). Player 2 is a firm that considers hiring the applicant. A low-skilled worker is completely unproductive and generates no revenue for the firm; a high-skilled worker generates revenue π. If the applicant is hired, the firm must pay the worker w (think of this wage as being fixed by government regulation). Assume $\pi > w > 0$. Therefore, the firm wishes to hire the applicant if and only if he or she is high-skilled. But the firm cannot observe the applicant's skill; it can observe only the applicant's prior education. Let c_H be the high type's cost of obtaining an education and c_L the low type's cost. Assume $c_H < c_L$, implying that education requires less effort for the high-skilled applicant than the low-skilled one. We make the extreme assumption that education does not increase the worker's productivity directly. The applicant may still decide to obtain an education because of its value as a signal of ability to future employers.

Figure 8.17 shows the extensive form. Player 1 observes his or her type at the start; player 2 observes only player 1's action (education signal) before moving. Let Pr(H) and Pr(L) be player 2's beliefs before observing player 1's education signal that player 1 is high- or low-skilled. These are called player 1's *prior beliefs*. Observing player 1's action will lead player 2 to revise his or her beliefs to form what are called *posterior beliefs*. For

FIGURE 8.17

Job-Market Signaling

Player 1 (worker) observes his or her own type. Then player 1 chooses to become educated (E) or not (NE). After observing player 1's action, player 2 (firm) decides to make him or her a job offer (J) or not (NJ). The nodes in player 2's information sets are labeled $n_1, \ldots, n_4$ for reference.

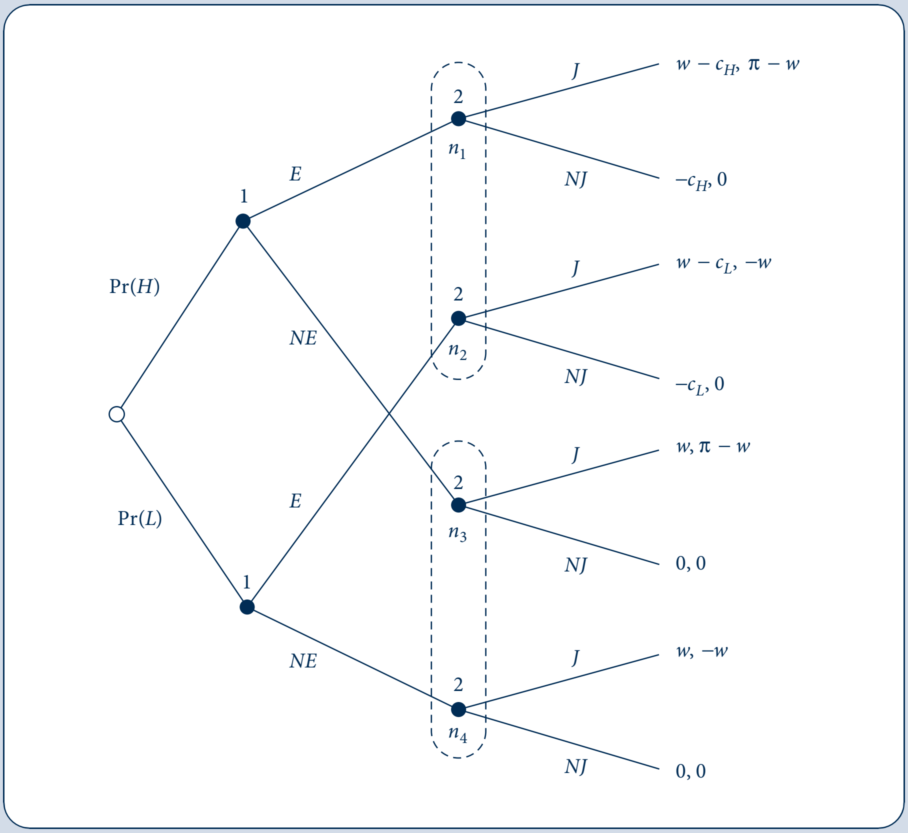

example, the probability that the worker is high-skilled is conditional on the worker's having obtained an education, $\Pr(H|E)$, and conditional on no education, $\Pr(H|NE)$.

Player 2's posterior beliefs are used to compute his or her best response to player 1's education decision. Suppose player 2 sees player 1 choose E. Then player 2's expected payoff from playing J is

$$\Pr(H|E)(\pi - w) + \Pr(L|E)(-w) = \Pr(H|E)\pi - w, \tag{8.32}$$

where the left side of this equation follows from the fact that because L and H are the only types, $\Pr(L|E) = 1 - \Pr(H|E)$. Player 2's payoff from playing NJ is 0. To determine the best response to E, player 2 compares the expected payoff in Equation 8.32 to 0. Player 2's best response is J if and only if $\Pr(H|E) > w/\pi$.

The question remains of how to compute posterior beliefs such as $\Pr(H|E)$. Rational players use a statistical formula, called *Bayes' rule*, to revise their prior beliefs to form posterior beliefs based on the observation of a signal.

Bayes' rule

Bayes' rule gives the following formula for computing player 2's posterior belief $\Pr(H|E)$[14]:

$$\Pr(H|E) = \frac{\Pr(E|H)\ \Pr(H)}{\Pr(E|H)\ \Pr(H) + \Pr(E|L)\ \Pr(L)}. \tag{8.33}$$

Similarly, $\Pr(H|E)$ is given by

$$\Pr(H|NE) = \frac{\Pr(NE|H)\ \Pr(H)}{\Pr(NE|H)\ \Pr(H) + \Pr(NE|L)\ \Pr(L)}. \tag{8.34}$$

Two sorts of probabilities appear on the left side of Equations 8.33 and 8.34:

- the prior beliefs $\Pr(H)$ and $\Pr(L)$;
- the conditional probabilities $\Pr(E|H)$, $\Pr(NE|L)$, and so forth.

The prior beliefs are given in the specification of the game by the probabilities of the different branches from the initial chance node. The conditional probabilities $\Pr(E|H)$, $\Pr(NE|L)$, and so forth are given by player 1's equilibrium strategy. For example, $\Pr(E|H)$ is the probability that player 1 plays E if he or she is of type H; $\Pr(NE|L)$ is the probability that player 1 plays NE if he or she is of type L; and so forth. As the schematic diagram in Figure 8.18 summarizes, Bayes' rule can be thought of as a "black box" that takes prior beliefs and strategies as inputs and gives as outputs the beliefs we must know to solve for an equilibrium of the game: player 2's posterior beliefs.

---

[14]Equation 8.33 can be derived from the definition of conditional probability in footnote 25 of Chapter 2. (Equation 8.34 can be derived similarly.) By definition,

$$\Pr(H|E) = \frac{\Pr(H \text{ and } E)}{\Pr(E)}.$$

Reversing the order of the two events in the conditional probability yields

$$\Pr(E|H) = \frac{\Pr(H \text{ and } E)}{\Pr(H)}$$

or, after rearranging,

$$\Pr(H \text{ and } E) = \Pr(E|H)\ \Pr(H).$$

Substituting the preceding equation into the first displayed equation of this footnote gives the numerator of Equation 8.33. The denominator follows because the events of player 1's being of type H or L are mutually exclusive and jointly exhaustive, so

$$\begin{aligned}\Pr(E) &= \Pr(E \text{ and } H) + \Pr(E \text{ and } L)\\ &= \Pr(E|H)\ \Pr(H) + \Pr(E|L)\ \Pr(L).\end{aligned}$$

FIGURE 8.18

Bayes' Rule as a Black Box

Bayes' rule is a formula for computing player 2's posterior beliefs from other pieces of information in the game.

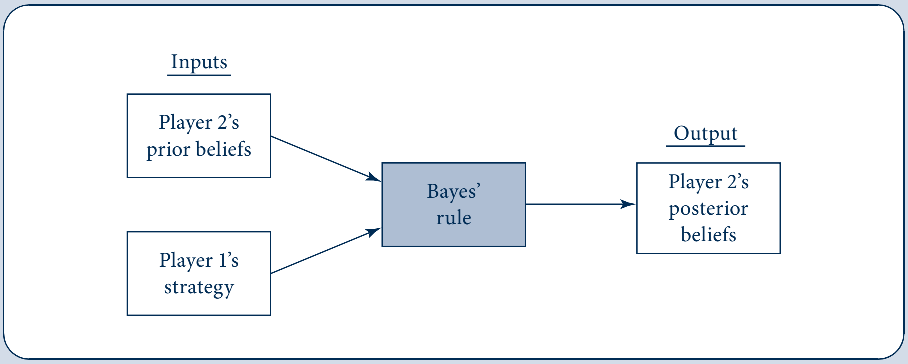

When player 1 plays a pure strategy, Bayes' rule often gives a simple result. Suppose, for example, that $\Pr(E|H) = 1$ and $\Pr(E|L) = 0$ or, in other words, that player 1 obtains an education if and only if he or she is high-skilled. Then Equation 8.33 implies

$$\Pr(H|E) = \frac{1 \cdot \Pr(H)}{1 \cdot \Pr(H) + 0 \cdot \Pr(L)} = 1. \quad \textbf{(8.35)}$$

That is, player 2 believes that player 1 must be high-skilled if it sees player 1 choose E. On the other hand, suppose that $\Pr(E|H) = \Pr(E|L) = 1$—that is, suppose player 1 obtains an education regardless of his or her type. Then Equation 8.33 implies

$$\Pr(H|E) = \frac{1 \cdot \Pr(H)}{1 \cdot \Pr(H) + 1 \cdot \Pr(L)} = \Pr(H), \quad \textbf{(8.36)}$$

because $\Pr(H) + \Pr(L) = 1$. That is, seeing player 1 play E provides no information about player 1's type, so player 2's posterior belief is the same as his or her prior one. More generally, if player 2 plays the mixed strategy $\Pr(E|H) = p$ and $\Pr(E|L) = q$, then Bayes' rule implies that

$$\Pr(H|E) = \frac{p\ \Pr(H)}{p\ \Pr(H) + q\ \Pr(L)}. \quad \textbf{(8.37)}$$

Perfect Bayesian equilibrium

With games of complete information, we moved from Nash equilibrium to the refinement of subgame-perfect equilibrium to rule out noncredible threats in sequential games. For the same reason, with games of incomplete information we move from Bayesian-Nash equilibrium to the refinement of perfect Bayesian equilibrium.

DEFINITION

Perfect Bayesian equilibrium. A perfect Bayesian equilibrium consists of a strategy profile and a set of beliefs such that

- at each information set, the strategy of the player moving there maximizes his or her expected payoff, where the expectation is taken with respect to his or her beliefs; and
- at each information set, where possible, the beliefs of the player moving there are formed using Bayes' rule (based on prior beliefs and other players' strategies).

The requirement that players play rationally at each information set is similar to the requirement from subgame-perfect equilibrium that play on every subgame form a Nash equilibrium. The requirement that players use Bayes' rule to update beliefs ensures that players incorporate the information from observing others' play in a rational way.

The remaining wrinkle in the definition of perfect Bayesian equilibrium is that Bayes' rule need only be used "where possible." Bayes' rule is useless following a completely unexpected event—in the context of a signaling model, an action that is not played in equilibrium by any type of player 1. For example, if neither H nor L type chooses E in the job-market signaling game, then the denominators of Equations 8.33 and 8.34 equal zero and the fraction is undefined. If Bayes' rule gives an undefined answer, then perfect Bayesian equilibrium puts no restrictions on player 2's posterior beliefs and thus we can assume any beliefs we like.

As we saw with games of complete information, signaling games may have multiple equilibria. The freedom to specify any beliefs when Bayes' rule gives an undefined answer may support additional perfect Bayesian equilibria. A systematic analysis of multiple equilibria starts by dividing the equilibria into three classes—separating, pooling, and hybrid. Then we look for perfect Bayesian equilibria within each class.

In a *separating equilibrium*, each type of player 1 chooses a different action. Therefore, player 2 learns player 1's type with certainty after observing player 1's action. The posterior beliefs that come from Bayes' rule are all zeros and ones. In a *pooling equilibrium*, different types of player 1 choose the same action. Observing player 1's action provides player 2 with no information about player 1's type. Pooling equilibria arise when one of player 1's types chooses an action that would otherwise be suboptimal to hide his or her private information. In a *hybrid equilibrium*, one type of player 1 plays a strictly mixed strategy; it is called a hybrid equilibrium because the mixed strategy sometimes results in the types being separated and sometimes pooled. Player 2 learns a little about player 1's type (Bayes' rule refines player 2's beliefs a bit) but does not learn player 1's type with certainty. Player 2 may respond to the uncertainty by playing a mixed strategy itself. The next three examples solve for the three different classes of equilibrium in the job-market signaling game.

EXAMPLE 8.8 Separating Equilibrium in the Job-Market Signaling Game

A good guess for a separating equilibrium is that the high-skilled worker signals his or her type by getting an education and the low-skilled worker does not. Given these strategies, player 2's beliefs must be $\Pr(H|E) = \Pr(L|NE) = 1$ and $\Pr(H|NE) = \Pr(L|E) = 0$ according to Bayes' rule. Conditional on these beliefs, if player 2 observes that player 1 obtains an education, then player 2 knows it must be at node n_1 rather than n_2 in Figure 8.17. Its best response is to offer a job (J), given the payoff of $\pi - w > 0$. If player 2 observes that player 1 does not obtain an eduation, then player 2 knows it must be at node n_4 rather than n_3, and its best response is not to offer a job (NJ) because $0 > -w$.

The last step is to go back and check that player 1 would not want to deviate from the separating strategy ($E|H$, $NE|L$) given that player 2 plays ($J|E$, $NJ|NE$). Type H of player 1 earns $w - c_H$ by obtaining an education in equilibrium. If type H deviates and does not obtain an education, then he or she earns 0 because player 2 believes that player 1 is type L and does not offer a job. For type H not to prefer to deviate, it must be that $w - c_H > 0$. Next, turn to type L of player 1. Type L earns 0 by not obtaining an education in equilibrium. If type L deviates and obtains an education, then he or she earns $w - c_L$ because player 2 believes that player 1 is type H and offers a job. For type L not to prefer to deviate, we must have $w - c_L < 0$. Putting these conditions together, there is separating equilibrium in which the worker obtains an education if and only if he or she is high-skilled and in which the firm offers a job only to applicants with an education if and only if $c_H < w < c_L$.

Another possible separating equilibrium is for player 1 to obtain an education if and only if he or she is low-skilled. This is a bizarre outcome—because we expect education to be a signal of high rather than low skill—and fortunately we can rule it out as a perfect Bayesian equilibrium.

Player 2's best response would be to offer a job if and only if player 1 did not obtain an education. Type L would earn $-c_L$ from playing E and w from playing NE, so it would deviate to NE.

QUERY: Why does the worker sometimes obtain an education even though it does not raise his or her skill level? Would the separating equilibrium exist if a low-skilled worker could obtain an education more easily than a high-skilled one?

EXAMPLE 8.9 Pooling Equilibria in the Job-Market Signaling Game

Let's investigate a possible pooling equilibrium in which both types of player 1 choose E. For player 1 not to deviate from choosing E, player 2's strategy must be to offer a job if and only if the worker is educated—that is, $(J|E, NJ|NE)$. If player 2 does not offer jobs to educated workers, then player 1 might as well save the cost of obtaining an education and choose NE. If player 2 offers jobs to uneducated workers, then player 1 will again choose NE because he or she saves the cost of obtaining an education and still earns the wage from the job offer.

Next, we investigate when $(J|E, NJ|NE)$ is a best response for player 2. Player 2's posterior beliefs after seeing E are the same as his or her prior beliefs in this pooling equilibrium. Player 2's expected payoff from choosing J is

$$\begin{aligned} \Pr(H|E)(\pi - w) + \Pr(L|E)(-w) &= \Pr(H)(\pi - w) + \Pr(L)(-w) \\ &= \Pr(H)\pi - w. \end{aligned} \tag{8.38}$$

For J to be a best response to E, Equation 8.38 must exceed player 2's zero payoff from choosing NJ, which on rearranging implies that $\Pr(H) \geq w/\pi$. Player 2's posterior beliefs at nodes n_3 and n_4 are not pinned down by Bayes' rule because NE is never played in equilibrium and so seeing player 1 play NE is a completely unexpected event. Perfect Bayesian equilibrium allows us to specify any probability distribution we like for the posterior beliefs $\Pr(H|NE)$ at node n_3 and $\Pr(L|NE)$ at node n_4. Player 2's payoff from choosing NJ is 0. For NJ to be a best response to NE, 0 must exceed player 2's expected payoff from playing J:

$$0 > \Pr(H|NE)(\pi - w) + \Pr(L|NE)(-w) = \Pr(H|NE)\pi - w, \tag{8.39}$$

where the right side follows because $\Pr(H|NE) + \Pr(L|NE) = 1$. Rearranging yields $\Pr(H|NE) \leq w/\pi$.

In sum, for there to be a pooling equilibrium in which both types of player 1 obtain an education, we need $\Pr(H|NE) \leq w/\pi \leq \Pr(H)$. The firm has to be optimistic about the proportion of skilled workers in the population—$\Pr(H)$ must be sufficiently high—and pessimistic about the skill level of uneducated workers—$\Pr(H|NE)$ must be sufficiently low. In this equilibrium, type L pools with type H to prevent player 2 from learning anything about the worker's skill from the education signal.

The other possibility for a pooling equilibrium is for both types of player 1 to choose NE. There are a number of such equilibria depending on what is assumed about player 2's posterior beliefs out of equilibrium (i.e., player 2's beliefs after he or she observes player 1 choosing E). Perfect Bayesian equilibrium does not place any restrictions on these posterior beliefs. Problem 8.12 asks you to search for various of these equilibria and introduces a further refinement of perfect Bayesian equilibrium (the *intuitive criterion*) that helps rule out unreasonable out-of-equilibrium beliefs and thus implausible equilibria.

QUERY: Return to the pooling outcome in which both types of player 1 obtain an education. Consider player 2's posterior beliefs following the unexpected event that a worker shows up with no education. Perfect Bayesian equilibrium leaves us free to assume anything we want about these posterior beliefs. Suppose we assume that the firm obtains no information from the "no education" signal and so maintains its prior beliefs. Is the proposed pooling outcome an equilibrium? What if we assume that the firm takes "no education" as a bad signal of skill, believing that player 1's type is L for certain?

EXAMPLE 8.10 Hybrid Equilibria in the Job-Market Signaling Game

One possible hybrid equilibrium is for type H always to obtain an education and for type L to randomize, sometimes pretending to be a high type by obtaining an education. Type L randomizes between playing E and NE with probabilities e and $1 - e$. Player 2's strategy is to offer a job to an educated applicant with probability j and not to offer a job to an uneducated applicant.

We need to solve for the equilibrium values of the mixed strategies e^* and j^* and the posterior beliefs $\Pr(H|E)$ and $\Pr(H|NE)$ that are consistent with perfect Bayesian equilibrium. The posterior beliefs are computed using Bayes' rule:

$$\Pr(H|E) = \frac{\Pr(H)}{\Pr(H) + e\Pr(L)} = \frac{\Pr(H)}{\Pr(H) + e[1 - \Pr(H)]} \quad \textbf{(8.40)}$$

and $\Pr(H|NE) = 0$.

For type L of player 1 to be willing to play a strictly mixed strategy, he or she must get the same expected payoff from playing E—which equals $jw - c_L$, given player 2's mixed strategy—as from playing NE—which equals 0 given that player 2 does not offer a job to uneducated applicants. Hence $jw - c_L = 0$ or, solving for j, $j^* = c_L/w$.

Player 2 will play a strictly mixed strategy (conditional on observing E) only if he or she gets the same expected payoff from playing J, which equals

$$\Pr(H|E)(\pi - w) + \Pr(L|E)(-w) = \Pr(H|E)\pi - w, \quad \textbf{(8.41)}$$

as from playing NJ, which equals 0. Setting Equation 8.41 equal to 0, substituting for $\Pr(H|E)$ from Equation 8.40, and then solving for e gives

$$e^* = \frac{(\pi - w)\Pr(H)}{w[1 - \Pr(H)]}. \quad \textbf{(8.42)}$$

QUERY: To complete our analysis: In this equilibrium, type H of player 1 cannot prefer to deviate from E. Is this true? If so, can you show it? How does the probability of type L trying to "pool" with the high type by obtaining an education vary with player 2's prior belief that player 1 is the high type?

EXPERIMENTAL GAMES

Experimental economics is a recent branch of research that explores how well economic theory matches the behavior of experimental subjects in laboratory settings. The methods are similar to those used in experimental psychology—often conducted on campus using undergraduates as subjects—although experiments in economics tend to involve incentives in the form of explicit monetary payments paid to subjects. The importance of experimental economics was highlighted in 2002, when Vernon Smith received the Nobel Prize in economics for his pioneering work in the field. An important area in this field is the use of experimental methods to test game theory.

Experiments with the Prisoners' Dilemma

There have been hundreds of tests of whether players fink in the Prisoners' Dilemma as predicted by Nash equilibrium or whether they play the cooperative outcome of Silent. In one experiment, subjects played the game 20 times with each player being matched with a different, anonymous opponent to avoid repeated-game effects. Play converged to the Nash equilibrium as subjects gained experience with the game. Players played the cooperative action 43 percent of the time in the first five rounds, falling to only 20 percent of the time in the last five rounds.[15]

As is typical with experiments, subjects' behavior tended to be noisy. Although 80 percent of the decisions were consistent with Nash equilibrium play by the end of the experiment, 20 percent of them still were anomalous. Even when experimental play is roughly consistent with the predictions of theory, it is rarely entirely consistent.

Experiments with the Ultimatum Game

Experimental economics has also tested to see whether subgame-perfect equilibrium is a good predictor of behavior in sequential games. In one widely studied sequential game, the Ultimatum Game, the experimenter provides a pot of money to two players. The first mover (Proposer) proposes a split of this pot to the second mover. The second mover (Responder) then decides whether to accept the offer, in which case players are given the amount of money indicated, or reject the offer, in which case both players get nothing. In the subgame-perfect equilibrium, the Proposer offers a minimal share of the pot, and this is accepted by the Responder. One can see this by applying backward induction: The Responder should accept any positive division no matter how small; knowing this, the Proposer should offer the Responder only a minimal share.

In experiments, the division tends to be much more even than in the subgame-perfect equilibrium.[16] The most common offer is a 50–50 split. Responders tend to reject offers giving them less than 30 percent of the pot. This result is observed even when the pot is as high as $100, so that rejecting a 30 percent offer means turning down $30. Some economists have suggested that the money players receive may not be a true measure of their payoffs. They may care about other factors such as fairness and thus obtain a benefit from a more equal division of the pot. Even if a Proposer does not care directly about fairness, the fear that the Responder may care about fairness and thus might reject an uneven offer out of spite may lead the Proposer to propose an even split.

The departure of experimental behavior from the predictions of game theory was too systematic in the Ultimatum Game to be attributed to noisy play, leading some game theorists to rethink the theory and add an explicit consideration for fairness.[17]

Experiments with the Dictator Game

To test whether players care directly about fairness or act out of fear of the other player's spite, researchers experimented with a related game, the Dictator Game. In the Dictator Game, the Proposer chooses a split of the pot, and this split is implemented without input from the Responder. Proposers tend to offer a less-even split than in the Ultimatum Game but still offer the Responder some of the pot, suggesting that Proposers have some residual concern for fairness. The details of the experimental design are crucial, however, as one ingenious experiment showed.[18] The experiment was designed so that the experimenter would never learn which Proposers had made which offers. With this element of anonymity, Proposers almost never gave an equal split to Responders and indeed took the whole pot for themselves two thirds of the time. Proposers seem to care more about appearing fair to the experimenter than truly being fair.

---

[15]R. Cooper, D. V. DeJong, R. Forsythe, and T. W. Ross, "Cooperation Without Reputation: Experimental Evidence from Prisoner's Dilemma Games," *Games and Economic Behavior* (February 1996): 187–218.

[16]For a review of Ultimatum Game experiments and a textbook treatment of experimental economics more generally, see D. D. Davis and C. A. Holt, *Experimental Economics* (Princeton, NJ: Princeton University Press, 1993).

[17]See, for example, E. Fehr and K.M. Schmidt, "A Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics* (August 1999): 817–868.

[18]E. Hoffman, K. McCabe, K. Shachat, and V. Smith, "Preferences, Property Rights, and Anonymity in Bargaining Games," *Games and Economic Behavior* (November 1994): 346–80.

EVOLUTIONARY GAMES AND LEARNING

The frontier of game-theory research regards whether and how players come to play a Nash equilibrium. Hyper-rational players may deduce each others' strategies and instantly settle on the Nash equilibrium. How can they instantly coordinate on a single outcome when there are multiple Nash equilibria? What outcome would real-world players, for whom hyper-rational deductions may be too complex, settle on?

Game theorists have tried to model the dynamic process by which an equilibrium emerges over the long run from the play of a large population of agents who meet others at random and play a pairwise game. Game theorists analyze whether play converges to Nash equilibrium or some other outcome, which Nash equilibrium (if any) is converged to if there are multiple equilibria, and how long such convergence takes. Two models, which make varying assumptions about the level of players' rationality, have been most widely studied: an evolutionary model and a learning model.

In the evolutionary model, players do not make rational decisions; instead, they play the way they are genetically programmed. The more successful a player's strategy in the population, the more fit is the player and the more likely will the player survive to pass his or her genes on to future generations and thus the more likely the strategy spreads in the population.

Evolutionary models were initially developed by John Maynard Smith and other biologists to explain the evolution of such animal behavior as how hard a lion fights to win a mate or an ant fights to defend its colony. Although it may be more of a stretch to apply evolutionary models to humans, evolutionary models provide a convenient way of analyzing population dynamics and may have some direct bearing on how social conventions are passed down, perhaps through culture.

In a learning model, players are again matched at random with others from a large population. Players use their experiences of payoffs from past play to teach them how others are playing and how they themselves can best respond. Players usually are assumed to have a degree of rationality in that they can choose a static best response given their beliefs, may do some experimenting, and will update their beliefs according to some reasonable rule. Players are not fully rational in that they do not distort their strategies to affect others' learning and thus future play.

Game theorists have investigated whether more- or less-sophisticated learning strategies converge more or less quickly to a Nash equilibrium. Current research seeks to integrate theory with experimental study, trying to identify the specific algorithms that real-world subjects use when they learn to play games.

Summary

This chapter provided a structured way to think about strategic situations. We focused on the most important solution concept used in game theory, Nash equilibrium. We then progressed to several more refined solution concepts that are in standard use in game theory in more complicated settings (with sequential moves and incomplete information). Some of the principal results are as follows.

- All games have the same basic components: players, strategies, payoffs, and an information structure.
- Games can be written down in normal form (providing a payoff matrix or payoff functions) or extensive form (providing a game tree).
- Strategies can be simple actions, more complicated plans contingent on others' actions, or even probability distributions over simple actions (mixed strategies).
- A Nash equilibrium is a set of strategies, one for each player, that are mutual best responses. In other words, a player's strategy in a Nash equilibrium is optimal given that all others play their equilibrium strategies.
- A Nash equilibrium always exists in finite games (in mixed if not pure strategies).