CHAPTER 29

GAME APPLICATIONS

In the last chapter we described a number of important concepts in game theory and illustrated them using a few examples. In this chapter we examine four important issues in game theory—cooperation, competition, coexistence, and commitment—and see how they work in various strategic interactions.

In order to do this, we first develop an important analytic tool, **best response curves**, which can be used to solve for equilibria in games.

29.1 Best Response Curves

Consider a two-person game, and put yourself in the position of one of the players. For any choice the other player can make, your **best response** is the choice that maximizes your payoff. If there are several choices that maximize your payoff, then your best response will be the set of all such choices.

For example, consider the game depicted in Table 29.1, which we used to illustrate the concept of a Nash equilibrium. If the column player chooses left, row's best response is to choose top; if column chooses right, then

A simple game

		Column: Left	Column: Right
Row	Top	2, 1	0, 0
	Bottom	0, 0	1, 2

row's best response is to choose bottom. Similarly, the best responses for column are to play left in response to top and to play right in response to bottom.

We can write this out in a little table:

Column's choice:	Left	Right
Row's best response:	Top	Bottom
Row's choice:	Top	Bottom
Column's best response:	Left	Right

Notice that if column thinks that row will play top, then column will want to play left, *and* if row thinks that column will play left, row will want to play top. So the pair of choices (top, left) are mutually consistent in the sense that each player is making an optimal response to the other player's choice.

Consider a general two-person game in which row has choices $r_1, \ldots, r_R$ and column has choices $c_1, \ldots, c_C$. For each choice r that row makes, let $b_c(r)$ be a best response for column, and for each choice c that column makes, let $b_r(c)$ be a best response for row. Then a **Nash equilibrium** is a pair of strategies (r^*, c^*) such that

$$c^* = b_c(r^*)$$
$$r^* = b_r(c^*).$$

The concept of Nash equilibrium formalizes the idea of "mutual consistency." If row expects column to play left, then row will choose to play top, and if column expects row to play top, column will want to play left. So it is the *beliefs* and the *actions* of the players that are mutually consistent in a Nash equilibrium.

Note that in some cases one of the players may be indifferent among several best responses. This is why we only require that c^* be *one* of column's best responses, and r^* be *one* of row's best responses. If there is a unique best response for each choice then the best response *curves* can be represented as best response *functions.*

This way of looking at the concept of a Nash equilibrium makes it clear that it is simply a generalization of the Cournot equilibrium described in Chapter 27. In the Cournot case, the choice variable is the amount of output produced, which is a continuous variable. The Cournot equilibrium has the property that each firm is choosing its profit-maximizing output, given the choice of the other firm.

The Bertrand equilibrium, also described in Chapter 27, is a Nash equilibrium in pricing strategies. Each firm chooses the price that maximizes its profit, given the choice that it thinks the other firm will make.

These examples show how the best response curve generalizes the earlier models, and allows for a relatively simple way to solve for Nash equilibrium. These properties make best response curves a very helpful tool to solve for an equilibrium of a game.

29.2 Mixed Strategies

Let us use best response functions to analyze the game shown in Table 29.2.

Table 29.2

Solving for Nash equilibrium.

		Ms. Column	
		Left	Right
Mr. Row	Top	2, 1	0, 0
	Bottom	0, 0	1, 2

We are interested in looking for mixed strategy equilibria as well as pure strategy equilibria, so we let r be the probability that row plays top, and $(1 - r)$ the probability that he plays bottom. Similarly, let c be the probability that column plays left, and $(1-c)$ the probability that she plays right. The pure strategies occur when r and c equal 0 or 1.

Let us calculate row's expected payoff if he chooses probability r of playing top and column chooses probability c of playing left. Look at the following array

Combination	Probability	Payoff to Row
Top, Left	rc	2
Bottom, Left	$(1-r)c$	0
Top, Right	$r(1-c)$	0
Bottom, Right	$(1-r)(1-c)$	1

To calculate the expected payoff to row, we weight row's payoffs in the third column by the probability that they occur, given in the second column, and add these up. The answer is

$$\text{Row's payoff} = 2rc + (1-r)(1-c),$$

which we can multiply out to be

$$\text{Row's payoff} = 2rc + 1 - r - c + rc.$$

Now suppose that row contemplates increasing r by Δr. How will his payoff change?

$$\begin{aligned}\Delta\text{payoff to row} &= 2c\,\Delta r - \Delta r + c\,\Delta r\\ &= (3c-1)\Delta r.\end{aligned}$$

This expression will be positive when $3c > 1$ and negative when $3c < 1$. Hence, row will want to increase r whenever $c > 1/3$, decrease r when $c < 1/3$, and be happy with any value of $0 \leq r \leq 1$ when $c = 1/3$.

Similarly, the payoff to column is given by

$$\text{Column's payoff} = cr + 2(1-c)(1-r).$$

Column's payoff will change when c changes by Δc according to

$$\begin{aligned}\Delta\,\text{payoff to column} &= r\,\Delta c + 2r\,\Delta c - 2\Delta c\\ &= (3r-2)\Delta c.\end{aligned}$$

Hence column will want to increase c whenever $r > 2/3$, decrease c when $r < 2/3$, and be happy with any value of $0 \leq c \leq 1$ when $r = 2/3$.

We can use this information to plot the best response curves. Start with row. If column choses $c = 0$, row will want to make r as small as possible, so $r = 0$ is the best response to $c = 0$. This choice will continue to be the best response up until $c = 1/3$, at which point *any* value of r between 0 and 1 is a best response. For all $c > 1/3$, the best response row can make is $r = 1$.

These curves are depicted in Figure 29.1. It is easy to see that they cross in three places: $(0,0)$, $(2/3,1/3)$, and $(1,1)$, which correspond to the three Nash equilibria of this game. Two of these strategies are pure strategies, and one is a mixed strategy.

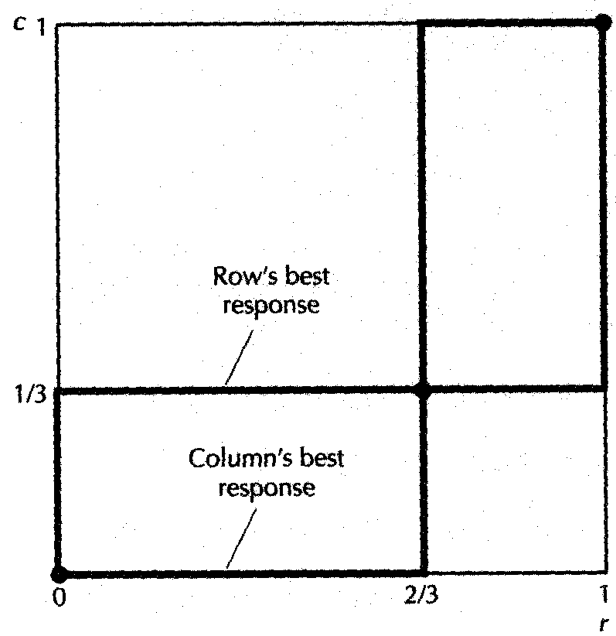

Best response curves. The two curves depict the best response of row and column to each other's choices. The intersections of the curves are Nash equilibria. In this case there are three equilibria, two with pure strategies and one with mixed strategies.

29.3 Games of Coordination

Armed with the tools of the last section we can examine our first class of games, **coordination games**. These are games where the payoffs to the players are highest when they can coordinate their strategies. The problem, in practice, is to develop mechanisms that enable this coordination.

Battle of the Sexes

The classic example of a coordination game is the so-called battle of the sexes. In this game, a boy and a girl want to meet at a movie but haven't had a chance to arrange which one. Alas, they forgot their cell phones, so they have no way to coordinate their meeting and have to guess which movie the other will want to attend.

The boy wants to see the latest action flick, while the girl would rather go to an art film, but they would both rather go to the same movie than not meet up at all Payoffs consistent with these preferences are shown in

The battle of the sexes.

		Girl	
		Action	Art
Boy	Action	2, 1	0, 0
	Art	0, 0	1, 2

Table 29.3. Note the defining feature of coordination games: the payoffs are higher when the players coordinate their actions than when they don't.

What are the Nash equilibria of this game? Luckily, this is just the game we used in the last section to illustrate best response curves. We saw there that there are three equilibria: both choose action, both choose art, or each chooses his or her preferred choice with probability 2/3.

Since all of these are possible equilibria, it is hard to say what will happen from this description alone. Generally, we would look to considerations outside the formal description of the game to resolve the problem. For example, suppose that the art film was a closer destination for one of the two players. Then both players might reasonably suppose that would be the equilibrium choice.

When players have good reasons to believe that one of the equilibria is more "natural" than the others, it is called a **focal point** of the game.

Prisoner's Dilemma

The prisoner's dilemma, which we discussed extensively in the last chapter, is also a coordination game. Recall the story: two prisoners can either confess, thereby implicating the other, or deny committing a crime. The payoffs are shown in Table 29.4.

The striking feature of the prisoner's dilemma is that confessing is a dominant strategy, even though coordination (both choose deny) is far superior in terms of the total payoff. Coordination would allow the prisoners to choose the best payoff, but the problem is that there is no easy way to make it happen in a single-shot game.

One way out of the prisoner's dilemma is to enlarge the game by adding new choices. We saw in the last chapter that an indefinitely repeated prisoner's dilemma game could achieve the cooperative outcome via strategies like tit for tat, in which players rewarded cooperation and punished lack of cooperation through their future actions. The extra strategic consideration

The prisoner's dilemma.

		Player B	
		Confess	Deny
Player A	**Confess**	−3, −3	0, −6
	Deny	−6, 0	−1, −1

here is that refusing to cooperate today may result in extended punishment later on.

Another way to "solve" the prisoner's dilemma is to add the possibility of contracting. For example, both players could sign a contract saying that they will stick with the cooperative strategy. If either of them reneges on the contract, he or she will have to pay a fine or be punished in some way. Contracts are very helpful in achieving all sorts of outcomes, but they rely on the existence of a legal system that will enforce such contracts. This makes sense for business negotiations but is not an appropriate assumption in other contexts, such as military games or international negotiations.

Assurance Games

Consider the U.S.-U.S.S.R. arms race of the 1950s in which each country could build nuclear missiles or refrain from building them. The payoffs to these strategies might look like those shown in Table 29.5. The best outcome for both parties is to refrain from building the missiles, giving a payoff of (4, 4). But if one refrains while the other builds, the payoff will be 3 to the builder and 1 to the refrainer. The payoff if they both build missile sites is (2, 2).

It is not hard to see that there are two pure strategy Nash equilibria, (refrain, refrain) and (build, build). However, (refrain, refrain) is better for both parties. The trouble is, neither party knows which choice the other will make. Before committing to refrain, each party wants some *assurance* that the other will refrain.

One way to achieve this assurance is for one of the players to move first, by opening itself to inspection, say. Note that this can be unilateral, at least as long as one believes the payoffs in the game. If one player announces that it is refraining from deploying nuclear missiles and gives the other player sufficient evidence of its choice, it can rest assured that the other player will also refrain.

An arms race.

		U.S.S.R. Refrain	U.S.S.R. Build
U.S.	Refrain	4, 4	1, 3
U.S.	Build	3, 1	2, 2

Chicken

Our last coordination game is based on an automobile game popularized in the movies. Two teenagers start at opposite ends of the street and drive in a straight line toward each other. The first to swerve loses face; if neither swerves, they both crash into each other. Some possible payoffs are shown in Table 29.6.

There are two pure strategy Nash equilibria, (row swerves, column doesn't) and (column swerves, row doesn't). Column prefers the first equilibrium and row the second, but each equilibrium is better than a crash. Note the difference between this and the assurance game; there, both players were better off doing the same thing (building or refraining) than doing different things. Here, both players are worse off doing the same thing (driving straight or swerving) than if they did different things.

Each player knows that if he can commit himself to driving straight, the other will chicken out. But of course, each player also knows that it would be crazy to crash into each other. So how can one of the players enforce his preferred equilibrium?

One important strategy is commitment. Suppose that row ostentatiously fastened a steering wheel lock on his car before starting out. Column, recognizing that row now has no choice but to go straight, would choose to swerve. Of course if both players put on a lock, the outcome would be disastrous!

How to Coordinate

If you are a player in a coordination game, you may want to get the other player to cooperate at an equilibrium that you both like (the assurance game), cooperate an an equilibrium one of you likes (battle of the sexes), play something other than the equilibrium strategy (the prisoner's dilemma), or make a choice leading to your preferred outcome (chicken).

In the assurance game, the battle of the sexes, and chicken, this can be accomplished by one player's moving first, and committing herself to a particular choice. The other player can then observe the choice and respond accordingly. In the prisoner's dilemma, this strategy doesn't work: if one player chooses not to confess, it is in the other's interest to do so. Instead of sequential moves, repetition and contracting are major ways to "solve" the prisoner's dilemma.

Chicken.

		Column: Swerve	Column: Straight
Row	Swerve	0, 0	-1, 1
	Straight	1, -1	-2, -2

29.4 Games of Competition

The opposite pole from cooperation is competition. This is the famous case of **zero-sum games**, so called because the payoff to one player is equal to the losses of the other.

Most sports are effectively zero-sum games: a point awarded to one team is equivalent to a point subtracted from the other team. Competition is fierce in such games because the players' interests are diametrically opposed.

Let us illustrate a zero-sum game by looking at soccer, known as football in most of the world. Row is kicking a penalty shot and column is defending. Row can kick to the left or kick to the right; column can favor one side and defend to the left or defend to the right in order to deflect the kick.

We will express the payoffs to these strategies in terms of expected points. Obviously row will be more successful if column jumps the wrong way. On the other hand, the game may not be perfectly symmetric since row may be better at kicking in one direction than another and column may be better at defending one direction or the other.

Let us assume that row will score 80 percent of the time if he kicks to the left and column jumps to the right but only 50 percent of the time if column jumps to the left. If row kicks to the right, we will assume that he succeeds 90 percent of the time if column jumps to the left but 20 percent

Penalty point in soccer.

		Column	
		Defend left	Defend right
Row	Kick left	50, −50	80, −80
	Kick right	90, −90	20, −20

of the time if column jumps to the right. These payoffs are illustrated in Table 29.7.

Note that the payoffs in each entry sum to zero, indicating that the players have diametrically opposed goals. Row wants to maximize his expected payoff, and column wants to maximize her expected payoff—which means she wants to *minimize* row's payoff.

Obviously, if column knows which way row will kick she will have a tremendous advantage. Row, recognizing this, will therefore try to keep column guessing. In particular, he will kick sometimes to his strong side and sometimes to his weak side. That is, he will pursue a **mixed strategy**.

If row kicks left with probability p, he will get an expected payoff of $50p + 90(1 - p)$ when column jumps left and $80p + 20(1 - p)$ when column jumps right. Row wants to make this expected payoff as big as possible, and column wants to make it as small as possible.

For example, suppose that row chooses to kick left half the time. If column jumps left, row will have an expected payoff of $50 \times 1/2 + 90 \times 1/2 = 70$, and if column jumps right, row will have an expected payoff of $80 \times 1/2 + 20 \times 1/2 = 50$.

Column, of course, can carry through this same reasoning. If column believes that row will kick to the left half the time, then column will want to jump to the right, since this is the choice that minimizes row's expected payoff (thereby maximizing column's expected payoff).

Figure 29.2 shows row's expected payoffs for different choices of p. This simply involves graphing the two functions $50p{+}90(1{-}p)$ and $80p{+}20(1{-}p)$. Since these two expressions are linear functions of p, the graphs are straight lines.

Row recognizes that column will always try to minimize his expected payoff. Thus, for any p, the best payoff he can hope for is the *minimum* of the payoffs given by the two strategies. We've illustrated this by the colored line in Figure 29.2.

Where does the maximum of these minimum payoffs occur? Obviously, it occurs at the peak of the colored line, or, equivalently, where the two

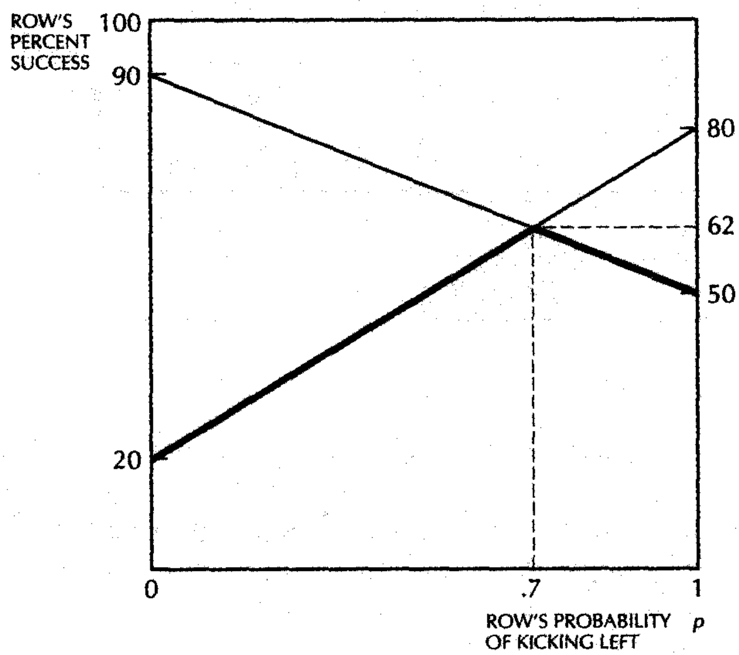

Row's strategy. The two curves show row's expected payoff as a function of p, the probability that he kicks to the left. Whatever p he chooses, column will try to minimize row's payoff.

lines intersect. We can calculate this value algebraically by solving

$$50p + 90(1-p) = 80p + 20(1-p)$$

for p. You should verify that the solution is $p = .7$.

Hence, if row kicks to the left 70 percent of the time and column responds optimally, row will have an expected payoff of $50 \times .7 + 90 \times .3 = 62$.

What about column? We can perform a similar analysis for her choices. Suppose column decides to jump to the left with probability q and jump to the right with probability $(1-q)$. Then row's expected payoff will be $50q + 80(1-q)$ if column jumps to the left and $90q + 20(1-q)$ if column jumps to the right. For each q, column will want to *minimize* row's payoff. But column recognizes that row wants to *maximize* this same payoff.

Hence, if column chooses to jump to the left with probability 1/2, she recognizes that row will get an expected payoff of $50 \times 1/2 + 80 \times 1/2 = 65$ if row kicks left and $90 \times 1/2 + 20 \times 1/2 = 55$ if row kicks right. In this case row will, of course, choose to kick left.

We can plot the two payoffs in Figure 29.3, which is analogous to the previous diagram. From column's viewpoint, it is the maximum of the two lines that is relevant, since this reflects row's optimal choice for each choice of q. Hence, the diagram depicts these lines in color. Just as before we can find the best q for column—the point where row's maximum payoff is minimized. This occurs where

$$50q + 80(1 - q) = 90q + 20(1 - q),$$

which implies $q = .6$.

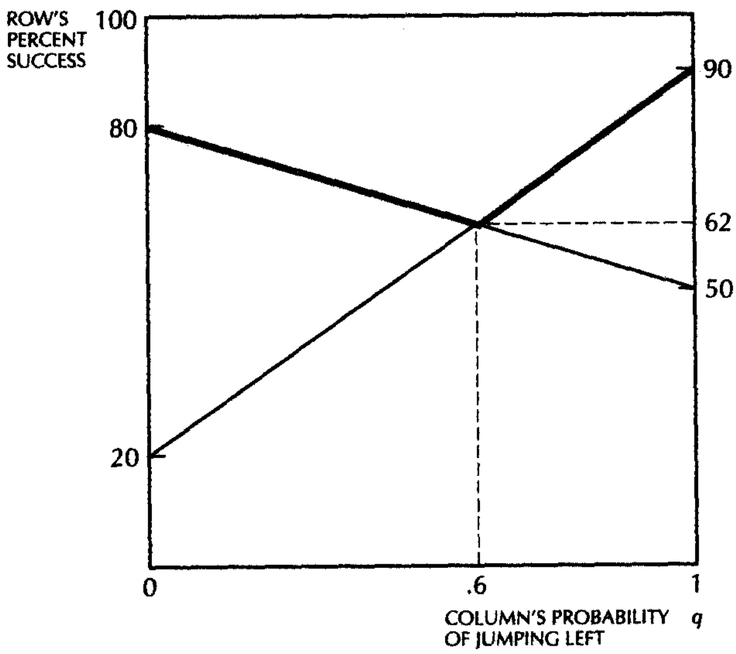

Column's strategy. The two lines show row's expected payoff as a function of q, the probability that column jumps to the left. Whatever q column chooses, row will try to maximize his own payoff.

We have now calculated the equilibrium strategies for each of the two players. Row should kick to the left with probability .7, and column should jump to the left with probability .6. These values were chosen so that row's payoffs and column's payoffs will be the same, whatever the other player does, since we found the values by equating the payoffs from the two strategies the opposing player could choose.

So when row chooses .7, column is indifferent between jumping left and jumping right, or, for that matter jumping left with any probability q. In particular, column is perfectly happy jumping left with probability .6.

Similarly, if column jumps left with probability .6, then row is indifferent between kicking left and kicking right, or any mixture of the two. In particular, he is happy to kick left with probability .7. Hence these choices are a Nash equilibrium: each player is optimizing, given the choices of the other.

In equilibrium row scores 62 percent of the time and fails to score 38 percent of the time. This is the best he can do, if the other player responds optimally.

What if column responds nonoptimally? Can row do better? To answer this question, we can use the best response curves introduced at the beginning of this chapter. We have already seen that when p is less than .7, column will want to jump left, and when p is greater than .7, column will want to jump right. Similarly when q is less than .6, row will want to kick left, and when q is greater than .6, row will want to kick right.

Figure 29.4 depicts these best response curves. Note that they intersect at the point where $p = .7$ and $q = .6$. The nice thing about the best response curves is that they tell each player what to do for every choice the other player makes, optimal or not. The only choice that is an optimal response to an optimal choice is where the two curves cross—the Nash equilibrium.

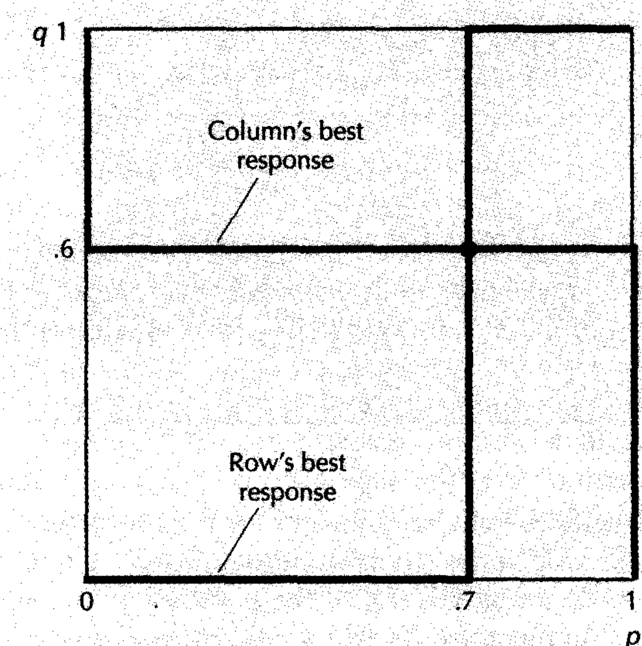

Best response curves. These are the best response curves for row and column, as a function of p, the probability that row kicks to the left, and q, the probability that column jumps to the left.

29.5 Games of Coexistence

We have interpreted mixed strategies as randomization by the players. In the penalty kick game, if row's strategy is to play left with probability .7 and right with probability .3, then we think that row will "mix it up" and play left 70 percent of the time and right 30 percent of the time.

But there is another interpretation. Suppose that kickers and goalies are matched up at random and that 70 percent of the kickers always kick left and 30 percent always kick right. Then, from the goalie's point of view, it is just like facing a single player who randomizes with those probabilities.

This isn't all that compelling as a story for soccer games, but it is a reasonable story for animal behavior. The idea is that various kinds of behavior are genetically programmed and that evolution selects the mixtures of the population that are stable with respect to evolutionary forces. In recent years, biologists have come to regard game theory as an indispensible tool to study animal behavior.

The most famous game of animal interaction is the **hawk-dove game**. This doesn't refer to a game between hawks and doves (which would have a pretty predictable outcome) but rather to a game involving a single species that exhibits two kinds of behavior.

Think of a wild dog. When two wild dogs come across a piece of food, they have to decide whether to fight or to share. Fighting is the hawkish strategy: one will win and one will lose. Sharing is a dovish strategy: it works well when the other player is also dovish, but if the other player is hawkish, the offer to share is rejected and the dovish player will get nothing.

A possible set of payoffs is given in Table 29.8.

Hawk-dove game.

		Column: Hawk	Column: Dove
Row	Hawk	-2, -2	4, 0
Row	Dove	0, 4	2, 2

If both wild dogs play dove, they end up with (2, 2). If one plays hawk and the other plays dove, the hawkish player wins everything. But if both play hawk, each dog will be seriously injured.

It obviously can't be an equilibrium if everyone plays hawk, since if some dog played dove, it would end up with 0 rather than -2. And if all dogs played dove, it would pay someone to deviate and play hawk. So there will have to be some mixture of hawk types and dove types in equilibrium. What sort of mixture should we expect?

Suppose that the fraction playing hawk is p. Then a hawk will meet another hawk with probability p and meet a dove with probability $1-p$. The expected payoff to the hawk type will be

$$H = -2p + 4(1-p).$$

The expected payoff to the dove type will be

$$D = 2(1-p).$$

Suppose that the type that has the higher payoff reproduces more rapidly, passing its tendency to play hawk or dove on to its offspring. So if $H > D$, we would see the fraction of hawk types in the population increase, and if $H < D$, we would expect to see the number of dove types increase.

The only way the population can be in equilibrium is if the payoffs to each type are the same. This requires

$$H = -2p + 4(1-p) = 2(1-p) = D,$$

which solves for $p = 1/2$.

We have found that a 50-50 mixture of doves and hawks is an equilibrium. Is it stable, in some sense? We plot the payoffs to hawk and dove as a function of p, the fraction of the population playing hawk in Figure 29.5. Note that when $p > 1/2$, the payoff to playing hawk is less than that of playing dove, so we would expect to see the doves reproduce more rapidly, moving us back to the equilibrium 50-50 ratio. Similarly, when $p < 1/2$, the payoff to hawk is greater than the payoff to dove, leading the hawks to reproduce more rapidly.

This argument shows that not only is $p = 1/2$ an equilibrium but it is also stable under evolutionary forces. Considerations of this sort lead to a concept known as an **evolutionarily stable strategy** or an **ESS**.[1] Remarkably, an ESS turns out to be a Nash equilibrium, even though it was derived from quite different considerations.

The Nash equilibrium concept was designed to deal with calculating, rational individuals, each of whom is trying to devise a strategy appropriate for the best strategy the other player might choose. The ESS was designed to model animal behavior under evolutionary forces, where strategies that had greater fitness payoffs would reproduce more rapidly. But the ESS equilibria are also Nash equilibria, giving another argument for why this particular concept in game theory is so compelling.

[1] See John Maynard Smith, *Evolution and the Theory of Games*, (Cambridge University Press, 1982).

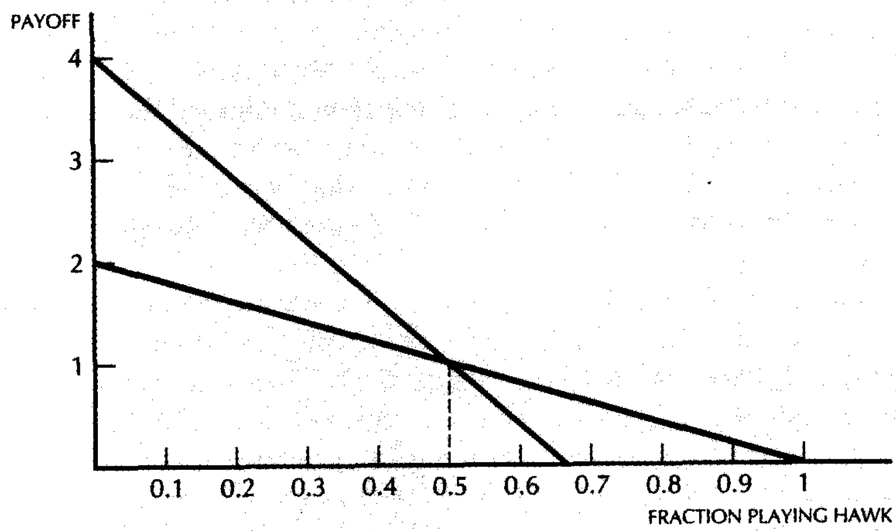

Payoffs in the hawk-dove game. The payoff to hawk is depicted in color; the payoff to dove is in black. When $p > 1/2$, the payoff to hawk is less than dove and vice versa, showing that the equilibrium is stable.

29.6 Games of Commitment

The previous examples involving games of cooperation and competition have been concerned with games with **simultaneous moves**. Each player had to make his or her choice without knowing what the other player was choosing (or had chosen). Indeed, games of coordination or competition can be quite trivial if one player knows the other's choices.

In this section we turn our attention to games with **sequential moves**. An important strategic issue that arises in such games is **commitment**. To see how this works, look back at the game of chicken described earlier in this chapter. We saw there that if one player could force himself to choose straight, the other player would optimally choose to swerve. In the assurance game, the outcome would be better for both players if one of them moved first.

Note that this committed choice must be both irreversible and observable by the other player. Irreversibility is part of what it means to be committed, while observability is crucial if the other player is going to be persuaded to change his or her behavior.

The Frog and the Scorpion

We begin with the fable of the frog and the scorpion. They were standing on the bank of the river, trying to figure out a way across. "I know," said the scorpion "I will climb on your back and you can swim across the river." The frog said, "But what if you sting me with your stinger?" The scorpion said, "Why would I do that? Then we would both die."

The frog found this convincing, so the scorpion climbed on his back and they started across the river. Halfway across, at the deepest point, the scorpion stung the frog. Writhing in pain, the frog cried out, "Why did you do that? Now we are both doomed!" "Alas," said the scorpion, as he sank into the river, "it is my nature."

Let's look at the frog and the scorpion from the viewpoint of game theory. Figure 29.6 depicts a sequential game with payoffs consistent with the story. Start at the bottom of the game tree. If the frog refuses the scorpion, both get nothing. Looking up one line, we see that if the frog carries the scorpion, he receives utility 5, for doing a good deed, and the scorpion receives a payoff of 3, for getting across the river. In line where the frog is stung, he receives a payoff of -10, and the scorpion gets a payoff of 5, representing the satisfaction from fulfilling his natural instincts.

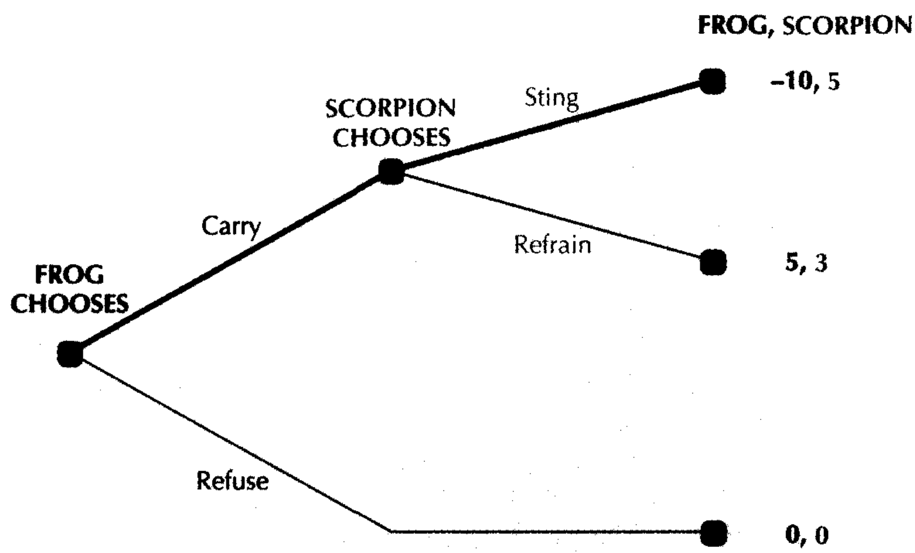

The frog and the scorpion. If the frog chooses to carry the scorpion, the scorpion will choose to sting him and both will die.

It is best to start with the final move of the game: the scorpion's choice of sting or refrain. Stinging has a higher payoff to the scorpion because "it is his nature" to sting. Hence the frog should rationally choose to refuse to carry the scorpion. Unfortunately, the frog didn't understand the scorpion's payoffs; apparently, he thought that the scorpion's payoffs looked something like those in Figure 29.7. Alas, this mistake was fatal for the frog.

A smart frog would figure out some way to make the scorpion commit to not stinging. He could, for example, tie his tail. Or he could hire a hit frog, who would retaliate against the scorpion's family. Whatever the strategy, the critical thing for the frog to do is to change the payoffs to the scorpion by making stinging more costly or refraining more rewarding.

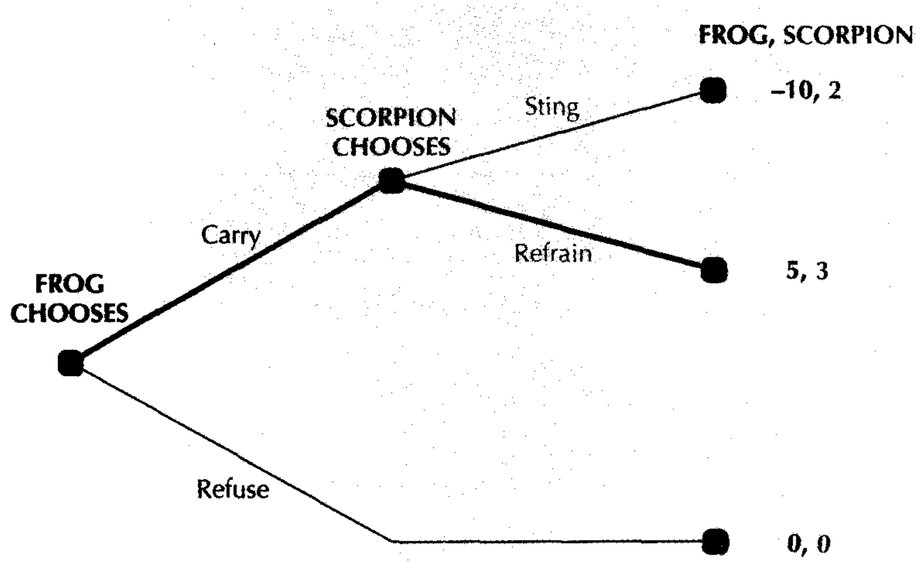

The frog and the scorpion. With these payoffs, if the frog chooses to carry the scorpion, the scorpion will not choose to sting him, and both will make it across the river safely.

The Kindly Kidnapper

Kidnapping for ransom is a big business in some parts of the world. In Columbia, it is estimated that there are over 2,000 kidnappings for ransom per year. In the former Soviet Union, kidnappings rose from 5 in 1992 to 105 in 1999. Many of the victims are Western businesspeople.

Some countries, such as Italy, have laws against paying ransom. The reasoning is that if the victim's family or employers can commit themselves not to pay ransom, then the kidnappers will have no motive to abduct the victim in the first place.

The problem is, of course, once a kidnapping has taken place, a victim's family will prefer to pay the kidnappers, even if it is illegal to do so. Hence penalties for paying ransom may not be effective as a commitment device.

Suppose some kidnappers abduct a hostage and then discover that they can't get paid. Should they release the hostage? The hostage, of course, promises not to reveal the identity of the kidnappers. But will he keep this promise? Once he is released, he has no incentive to do so—and every incentive to try to punish the kidnappers. Even if the kidnappers want to let the hostage go, they can't do so for fear of being identified.

Figure 29.8 depicts some possible payoffs. The kidnapper would feel bad about killing the hostage, receiving a payoff of -3. Of course, the hostage would feel even worse, receiving a payoff of -10. If the hostage is released, and refrains from identifying the kidnapper, the hostage gets a payoff of 3 and the kidnapper gets a payoff of 5. But if the hostage does identify the kidnapper, he gets a payoff of 5, leaving the kidnapper with a payoff of -5.

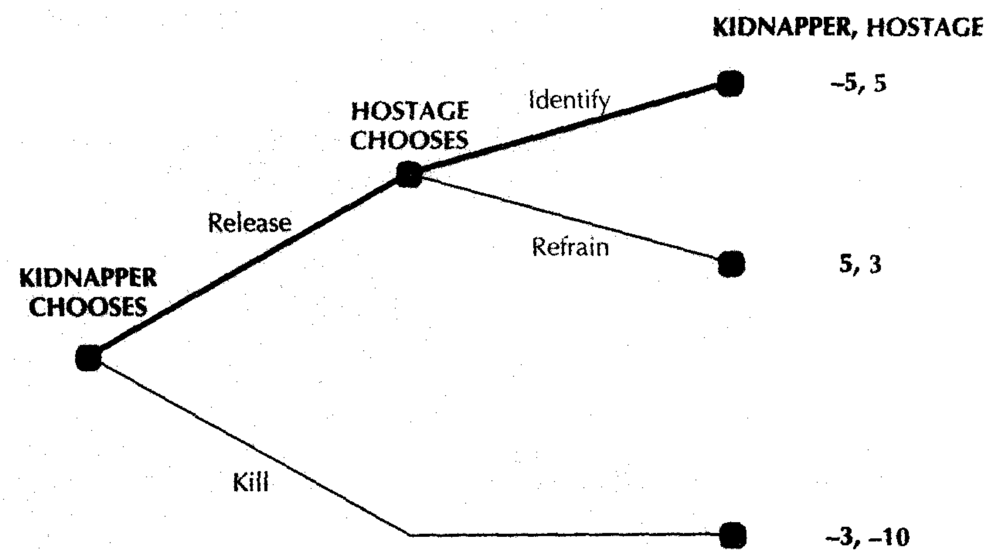

Figure 29.8 **Kidnap game.** The kidnapper would like to release the hostage, but if he does, the hostage will identify him.

Now it is the hostage who has the commitment problem: how can he convince the kidnappers that he won't renege on his promise and reveal their identity?

The hostage needs to figure out a way to change the payoffs of the game. In particular, he needs to find a way to impose a cost on himself if he identifies the kidnappers.

Thomas Schelling, an economist at the University of Maryland who has worked extensively on strategic analysis in dynamic games, suggests that the hostage might have the kidnappers photograph him in some embarrassing act and leave them with the photos. This effectively changes the payoffs from his subsequently revealing the identity of the kidnappers, since they then have the option of revealing the embarrassing photograph.

This sort of strategy is known as an "exchange of hostages." In the Middle Ages, when two kings wanted to ensure a contract wouldn't be broken, they would exchange hostages such as family members. If either king broke the agreement, the hostages would be sacrificed. Neither wanted to sacrifice their family members, so each king would have an incentive to respect the terms of their contract.

In the case of the kidnapping, the embarrassing photo would impose costs on the hostage if it were released, thereby ensuring that he will stick to his agreement not to reveal the identity of the kidnappers.

When Strength Is Weakness

Our next example comes from the world of animal psychology. It turns out that pigs quickly establish dominance-subordinateness relations, in which the dominant pig tends to boss the subordinate pig around.

Some psychologists put two pigs, one dominant, one subordinate, in a long pen.[2] At one end of the pen was a lever that would release a portion of food to a trough located at the other end of the pen. The question of interest was this: which pig would push the lever and which would eat the food?

Somewhat surprisingly the outcome of the experiment was that the dominant pig pressed the lever, while the subordinate pig waited for the food. The subordinate pig then ate most of the food, while the dominant pig rushed as fast as it could to the trough end of the pen, ending up with only a few scraps. Table 29.9 depicts a game that illustrates the problem.

Pigs pressing levers.

		Dominant Pig	
		Don't press lever	Press lever
Subordinate Pig	Don't press lever	0, 0	4, 1
	Press lever	0, 5	2, 3

[2] The original reference is Baldwin and Meese, "Social Behavior in Pigs Studied by Means of Operant Conditioning," (*Animal Behavior*, (1979)). I draw on the description of John Maynard Smith, *Evolution and the Theory of Games* (Cambridge University Press, 1982).

The subordinate pig compares a payoff of (0, 4) to (0, 2) and concludes, sensibly enough, that pressing the lever is dominated by not pressing it. Given that the subordinate pig doesn't press the lever, the dominant pig has no choice but to do so.

If the dominant pig could refrain from eating all the food and reward the subordinate pig for pressing the lever, it could achieve a better outcome. The problem is that pigs have no contracts, and the dominant pig can't help being a hog!

As in the case of the kindly kidnapper, the dominant pig has a commitment problem. If he could only commit to not eating all the food, he would end up much better off.

Savings and Social Security

Commitment problems aren't limited to the animal world. They also show up in economic policy.

Saving for retirement is an interesting and timely example. Everyone gives lip service to the fact that saving is a good idea. Unfortunately, few people actually do it. Part of the reason for the reluctance to save is that individuals recognize that society won't let them starve, so there is a good chance they will be bailed out later on.

To formulate this in a game between the generations, let's consider two strategies for the older generation: save or squander. The younger generation likewise has two strategies: support their elders or save for their own retirement. A possible game matrix is shown in Table 29.10.

Intergenerational conflict over savings.

		Younger Generation	
		Support	Refrain
Older Generation	Save	3, −1	1, 0
	Squander	2, −1	−2, −2

If the older generation saves and the younger generation also supports them, the old folks end up with a utility level of 3 and the young folks end up with −1. If the older generation squanders and the younger generation supports them, the elders end up with a utility of 2 and the young folks end up with −1.

If the younger generation refrains from providing support to their elders and the older generation saves, the old folks get 1 and the young folks get 0. Finally, if the old folks squander and the young folks neglect them, each ends up with utility of -2, the old folks from starving and the young folks from having to watch.

It is not hard to see that there are two Nash equilibria in this game. If the old folks choose to save, then the young folks will choose optimally to neglect them. But if the old folks choose to squander, then it is optimal for the younger generation to support them. And of course, given that the younger generation will support their elders, it is optimal for their elders to squander!

However, this analysis ignores the time structure of the game: one of the (few) advantages of being old is that you get to move first. If we draw out the game tree, the payoffs become those in Figure 29.9.

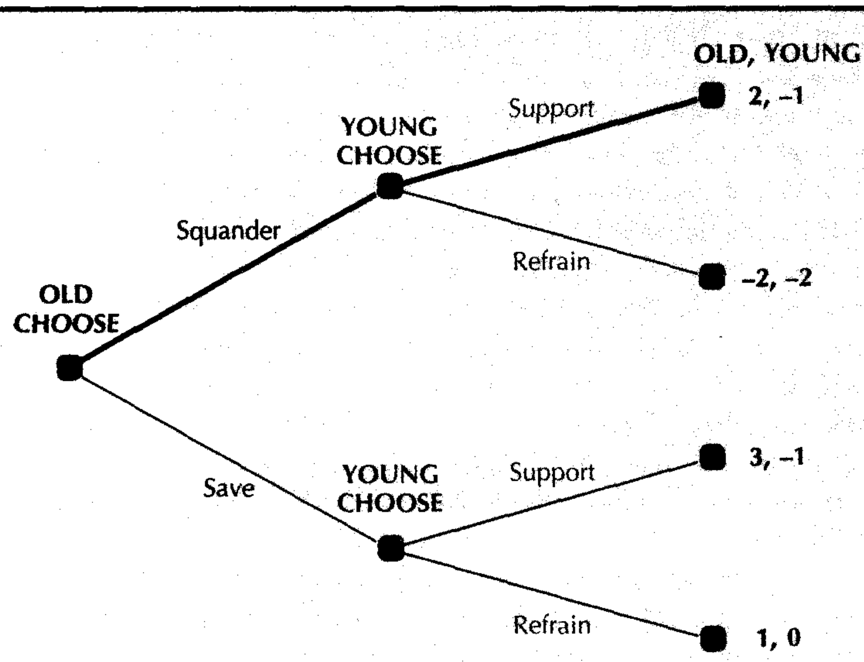

The savings game in extended form. Knowing that the younger generation will support them, the older generation chooses to squander. The subgame perfect equilibrium is (support, squander).

If the oldsters save, the youngsters will choose to neglect them, so the oldsters end up with a payoff of 1. If the oldsters squander, they know that the youngsters won't be able to bear watching them starve, so the oldsters end up with a payoff of 3. Hence the sensible thing for the oldsters to do is to squander, knowing they will be bailed out later on.

Of course, most developed countries now have a program like the U.S. Social Security program that forces each generation to save for retirement.

Hold Up

Consider the following strategic interaction. You hire a contractor to build a warehouse. After the plans are approved and the construction is almost done, you realize that the color is bad, so you ask the contractor to change the paint, which involves a trivial expense. The contractor comes back and says: "That change order will be \$1500, please."

You recognize that it will cost you at least that much to delay completion until you can find a painter, and you really do want the new color, so, muttering under your breath, you pay the cost. Congratulations, you have been held up!

Of course, contractors are not the only party at fault in this sort of game. The clients can "hold up" their payment as well, causing lots of grief for the contractor.

The game tree for the hold-up problem is depicted in Figure 29.10. We suppose that the value the owner places on having the new paint is \$1500 and that the actual cost of painting is \$200. Starting at the top of leaves of the tree, if the contractor charges \$1500, it will realize a profit of \$1300, and the client gets a net utility of zero.

If the client looks for another painter, it will cost him \$200 to pay the painter and, say, \$1400 in lost time. He gets the color he wants which is worth \$1500, buy has to pay \$1600 in direct costs and delay costs, leaving him with a net loss of \$100.

If the contractor charges the client the actual cost of \$200, he breaks even and the client gets a \$1500 value for \$200, leaving him with a net payoff of \$1300.

As can be seen, the optimal choice for the contractor is to extort the payment, and the optimal choice for the customer is to give in. But a sensible client will recognize that change orders will occur in any project. Because of this, the client will be reluctant to hire contractors with a reputation for extortion which is, of course, bad for the contractor.

How do firms solve the hold-up problem? The basic answer is contracts. Normally, contractors negotiate a contract specifying what kinds of change orders are appropriate and how their costs will be determined. Sometimes there are even arbitration or other dispute resolution procedures built into the contracts. A lot of time, energy, and money goes into writing contracts just to make certain that hold up won't occur.

But contracts aren't the only solution. Another way to solve the problem is through commitment. For example, the contractor might post a bond guaranteeing timely completion of the project. Again, there will generally be some objectively specified terms about what constitutes completion.

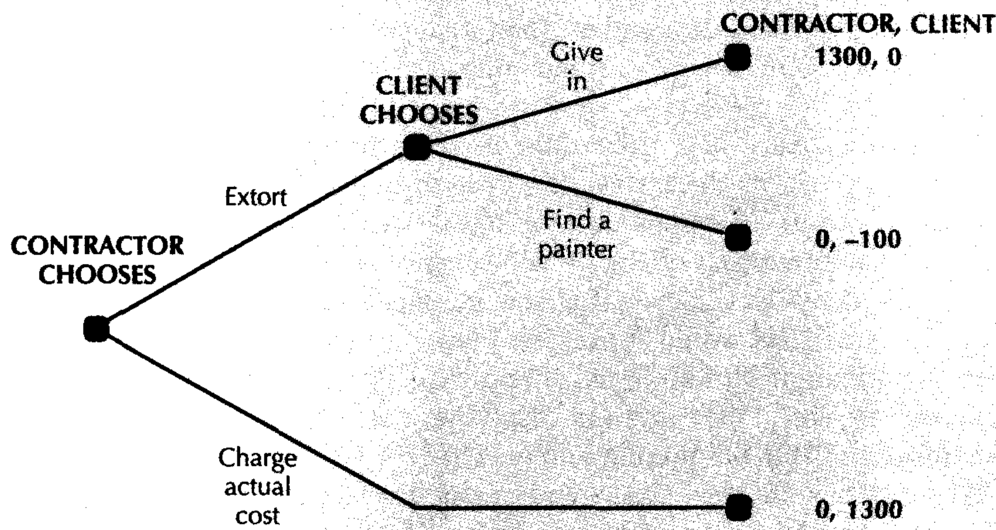

The hold-up problem. The contractor charges a high price for the change since the client has no alternative.

Another important factor is reputation. Obviously, a contractor who persistently tries to extort his customers will get a bad reputation. He won't be hired again by this customer, and he certainly won't get good recommendations. This reputation effect can be examined in a repeated game context in which hold up today will cost the contractor in the future.

29.7 Bargaining

The classical bargaining problem is divide the dollar. Two players have a dollar that they want to divide between them. How do they do it?

The problem, as stated, has no answer since there is too little information to construct a reasonable model. The challenge in modeling bargaining is to find some other dimensions on which the players can negotiate.

One solution, the **Nash bargaining model**, takes an axiomatic approach by specifying certain properties that a reasonable bargaining solution should have and then proving that there is only one outcome that satisfies these axioms.

The outcome ends up depending on how risk averse the players are and what will happen if no bargain is made. Unfortunately, a full treatment of this model is beyond the scope of this book.

An alternative approach, the **Rubinstein bargaining model**, looks at a sequence of choices and then solves for the subgame perfect equilibrium. Luckily the basic insight of this model is easy to illustrate in simple cases.

Two players, Alice and Bob, have \$1 to divide between them. They agree to spend at most three days negotiating over the division. The first day, Alice will make an offer, Bob either accepts or comes back with a counteroffer the next day, and on the third day Alice gets to make one final offer. If they cannot reach an agreement in three days, both players get zero.

We assume Alice and Bob differ in their degree of impatience: Alice discounts payoffs in the future at a rate of α per day, and Bob discounts payoffs at a rate of β per day. Finally, we assume that if a player is indifferent between two offers, he will accept the one that is most preferred by his opponent. This idea is that the opponent could offer some arbitrarily small amount that would make the player strictly prefer one choice and that this assumption allows us to approximate such an "arbitrarily small amount" by zero. It turns out that there is a unique subgame perfect equilibrium of this bargaining game.

We start our analysis at the end of the game, right before the last day. At this point Alice can make a take-it-or-leave-it offer to Bob. Clearly, the optimal thing for Alice to do at this point is to offer Bob the smallest possible amount that he would accept, which, by assumption, is zero. So if the game actually lasts three days, Alice would get \$1 and Bob would get zero (i.e., an arbitrarily small amount).

Now go back to the previous move, when Bob gets to propose a division. At this point Bob should realize that Alice can guarantee herself \$1 on the next move by simply rejecting his offer. A dollar next period is worth α to Alice this period, so any offer less than α would be sure to be rejected. Bob certainly prefers $1-\alpha$ now to zero next period, so he should rationally offer α to Alice, which Alice will then accept. So if the game ends on the second move, Alice gets α and Bob gets $1-\alpha$.

Now move to the first day. At this point Alice gets to make the offer and he realizes that Bob can get $1-\alpha$ if he simply waits until the second day. Hence Alice must offer a payoff that has at least this present value to Bob in order to avoid delay. Thus she offers $\beta(1-\alpha)$ to Bob. Bob finds this (just) acceptable and the game ends. The final outcome is that the game ends on the first move with Alice receiving $1-\beta(1-\alpha)$ and Bob receiving $\beta(1-\alpha)$.

The first panel in Figure 29.11 illustrates this process for the case where $\alpha=\beta<1$. The outermost diagonal line shows the possible payoff patterns on the first day, namely, all payoffs of the form $x_A+x_B=1$. The next diagonal line moving toward the origin shows the present value of the payoffs if the game ends in the second period: $x_A+x_B=\alpha$. The diagonal line closest to the origin shows the present value of the payoffs if the game ends in the third period; the equation for this line is $x_A+x_B=\alpha^2$. The right-angled path depicts the minimum acceptable divisions each period, leading up to the final subgame perfect equilibrium. The second panel in Figure 29.11 shows how the same process might look with more stages in the negotiation.

It is natural to let the horizon go to infinity and ask what happens in the

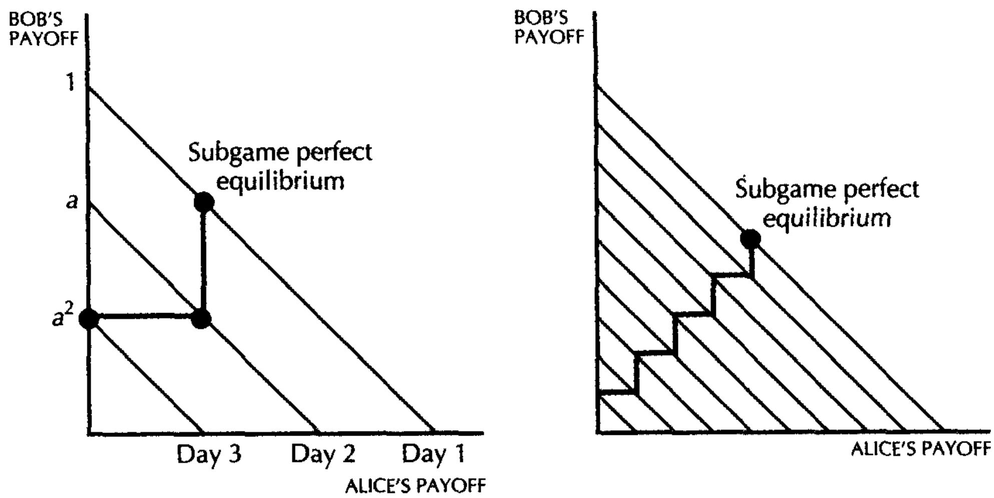

A bargaining game. The heavy line connects together the equilibrium outcomes in the subgames. The point on the line that is furthest out is the subgame perfect equilibrium.

infinite game. It turns out that the subgame perfect equilibrium division is

$$\text{Payoff to Alice} = \frac{1-\beta}{1-\alpha\beta}$$
$$\text{Payoff to Bob} = \frac{\beta(1-\alpha)}{1-\alpha\beta}.$$

Note that if $\alpha = 1$ and $\beta < 1$, then Alice receives the entire payoff.

The Ultimatum Game

The Rubinstein bargaining model is so elegant that economists rushed to test it in the laboratory. They found, alas, that elegance does not imply accuracy. Naive subjects (i.e., noneconomics majors) aren't very good at looking ahead more than one or two steps, if that.

In addition, there are other factors that cause problems. To see this, let us examine a one-step version of the bargaining model described above. Alice and Bob still have \$1 to divide between them. Alice proposes a division, and, if Bob agrees, the game ends. The question is, what should Alice say?

According to the theory, she should propose something like 99 cents for Alice, 1 cent for Bob. Bob, figuring that 1 cent is better than nothing, accepts, and Alice goes home happy that she studied economics.

Unfortunately, it doesn't work out like that. A more likely outcome is that Bob, disgusted by the paltry 1 cent, says "No way," and Alice ends up with nothing. Alice, recognizing this possibility, will tend to sweeten the offer. In actual experiments, the average offer for U.S. undergraduates is about 45 cents, and this offer tends to be accepted most of the time.

The offering players are behaving rationally, in the sense that the 45 cent offer is pretty close to maximizing the expected payoff, given the observed frequency of rejection. It is the receiving players who behave differently than the theory predicts, since they reject small offers, even though this makes them worse off.

There are many proposed explanations for this. One view is that too small an offer violates **social norms** of behavior. Indeed, economists have found quite significant cross-cultural differences in behavior in ultimatum games. Another, not inconsistent view, is that receivers get some utility payoff from hurting the offerers, in retaliation for the small offer. After all, if all you are losing is a penny, the satisfaction of striking back at the other player is pretty attractive by comparison. We will the ultimatum game in more detail in the next chapter.

Summary

1. A player's best response function gives the optimal choice for him as a function of the choices the other player(s) might make.

2. A Nash equilibrium in a two-person game is a pair of strategies, one for each player, each of which is a best response to the other.

3. A mixed strategy Nash equilibrium involves randomizing among several strategies.

4. Common games of coordination are the battle of the sexes, where both players want to do the same thing rather than different things; the prisoner's dilemma, where the dominant strategy ends up hurting both players; the assurance game, where both players want to cooperate as long as they think the other will cooperate; and chicken, where players want to avoid doing the same thing.

5. A two-person zero-sum game is one where the payoffs to one player are the negative of the payoffs to the other.

6. Evolutionary games are concerned with outcomes that are stable under population reproduction.

7. In sequential games, players move in turn. Each player therefore has to reason about what the other will do in response to his or her choices.

8. In many sequential games, commitment is an important issue. Finding ways to force commitment to play particular strategies can be important.

REVIEW QUESTIONS

1. In a two-person Nash equilibrium, each player is making a best response to what? In a dominant strategy equilibrium, each player is making a best response to what?

2. Look at the best responses for row and column in the section on mixed strategies. Do these give rise to best response functions?

3. If both players make the same choice in a coordination game, all will be well.

4. The text claims that row scores 62 percent of the time in equilibrium. Where does this number come from?

5. A contractor says that he intends to "low-ball the bid and make up for it on change orders." What does he mean?