CHAPTER 12

The Theory of Games

If you had to be a pig, would you rather be a strong pig or a weak pig? Sometimes it pays to be weak.

The biologist John Maynard Smith reports an experiment where two pigs are kept in a box with a lever at one end and a food dispenser at the other.[1] When the lever is pushed, food appears at the dispenser.

If the weak pig pushes the lever, the strong pig waits by the dispenser and eats all the food. Even if the weak pig races to the dispenser and arrives before the food is gone, the strong pig pushes the weak pig away. The weak pig is smart enough to figure this out, so it never bothers pressing the lever in the first place.

On the other hand, if the strong pig pushes the lever, the weak pig waits by the dispenser and gets most of the food. But the strong pig can race to the dispenser and shove the weak pig aside before it has entirely finished eating and then help itself to the leftovers. This makes it worthwhile for the strong pig to push the lever.

The outcome is that the strong pig does all of the work, and the weak pig does most of the eating.

Strategic situations can yield surprising outcomes. The Prisoner's Dilemma of Chapter 11 provides one example; the pigs in a box provide another. In this chapter, we will study the **theory of games** (or **game theory** for short), which allows us to catalog many of those outcomes and to discuss both their positive and their normative aspects.

Theory of games or game theory

A system for studying strategic behavior.

12.1 Game Matrices

In this section, we will introduce game matrices and show how they can be used to systematically analyze strategic situations.

Game matrix

A diagram showing one player's strategy choices across the top, the other player's along the left side, and the corresponding outcomes in the appropriate boxes.

Pigs in a Box

Consider the pigs from the introduction to this chapter. We represent the pigs' dilemma by a **game matrix** as in Exhibit 12.1. Across the top we list the possible strategies of the

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

EXHIBIT 12.1 **Pigs in a Box**

		Strong Pig's Strategy	
		Push lever	Wait by dispenser
Weak Pig's Strategy	Push lever	Strong pig gets 90 calories Weak pig gets –10 calories	Strong pig gets 100 calories Weak pig gets –10 calories
	Wait by dispenser	Strong pig gets 15 calories Weak pig gets 75 calories	Strong pig gets 0 calories Weak pig gets 0 calories

The dispenser gives 100 calories worth of food, and it requires 10 calories to push the lever. If both pigs arrive at the dispenser simultaneously, only the strong pig eats. But if the weak pig waits at the dispenser while the strong pig pushes the lever, he can eat ¾ of the food before the strong pig arrives. The game matrix shows the pigs' rewards for each combination of strategies.

The lower left-hand box is the only Nash equilibrium. Starting from any other box, at least one of the pigs would want to change his strategy.

strong pig, who can either push the lever or wait by the food dispenser. Along the left side we list the possible strategies of the weak pig, who has the same options.

In each of the four boxes of the matrix we show the consequences of the pigs' behavior. We assume that the food dispenser yields 100 calories worth of food and that pushing the lever burns 10 calories. We assume also that pigs care only about calories (which is presumably why they are called pigs).

If both pigs decide to push the lever, then they both run to the dispenser, where the strong pig shoves the weak pig aside and eats all of the food. The net gain is 90 calories for the strong pig (100 calories worth of food minus 10 calories burned pushing the lever) and *minus* 10 calories for the weak pig, who pushes the lever and runs but gets no food. The upper left-hand box in the exhibit shows this outcome.

If the strong pig waits by the dispenser while the weak pig pushes the lever, the strong pig gets all 100 calories worth of food and the weak pig loses 10 calories, as shown in the upper right-hand box.

If the strong pig pushes the lever while the weak pig waits by the dispenser, the weak pig is able to consume 75 calories before the strong pig arrives and takes the remaining 25, leaving him with a net gain of 15 after subtracting the 10 that he burns by pushing the lever. This is the outcome in the lower left-hand box.

And finally, if both pigs wait by the dispenser, then nobody gets to eat anything at all, as indicated in the lower right-hand box.

Choosing Strategies

In the introduction to this chapter, we argued that the pigs will end up in the lower left-hand box, which is to say that the strong pig will push the lever while the weak pig waits by the dispenser and gets most of the food. Let us see how we can use the game matrix to reach this conclusion systematically.

When the strong pig selects a strategy, he decides which column of the matrix both pigs will occupy. When the weak pig selects a strategy, he decides on a row. There are four possible outcomes, represented by the four boxes of the game matrix. For each outcome, we can ask this question: If this *were* the outcome, would either pig want to change his mind? If one or both pigs *would* want to change their minds, then we can rule out that outcome as a possibility.

For example, suppose for the moment that we are in the upper left-hand box, where both pigs push the lever. If the strong pig changes his mind and waits by the dispenser, we move to the upper right-hand box, while if the weak pig changes his mind we move to the lower left-hand box. Would the strong pig want to change his mind? The answer is yes: By moving from the upper left to the upper right he gains 10 calories. *This is already enough to rule out the upper left-hand box.*

Would the weak pig want to change his mind? The answer is yes again: By moving from the upper left to the lower left he gains 10 calories (or more precisely, he avoids losing 10 calories). This by itself would *also* be enough to rule out the upper left. So the upper left is ruled out for each of two separate reasons: If that were the outcome, the strong pig would change his mind *and* the weak pig would change his mind.

Next suppose that we are in the upper right-hand box. Would the strong pig want to change his mind and move to the upper left? No; he prefers the upper right, gaining 100 calories instead of 90. Would the weak pig want to change his mind and move to the lower right? Yes; he can then avoid losing 10 calories. So we rule out the upper right on the grounds that the weak pig would change his mind.

What about the lower right? The weak pig would not want to change rows, but the strong pig *would* want to change columns. Because the strong pig wants to change his mind, this outcome can also be ruled out.

Exercise 12.1 In the lower-right corner, how much would the weak pig lose by changing rows? How much would the strong pig gain by changing columns?

Finally, consider the lower left. Starting from here, the weak pig has the option to move up a box, reducing his calorie intake from 75 to -10; this option is not attractive. The strong pig has the option to move to the right, reducing his net calorie intake from 15 to 0; this is not attractive either. Neither pig changes his mind, and the pigs remain in the lower left-hand box.

Any outcome that survives this process of elimination is called a **Nash equilibrium** outcome. An outcome is a Nash equilibrium if neither player would want to deviate from it, taking his opponent's behavior as given. The phrase *taking his opponent's behavior as given* is an important one here. Starting in the lower left, the strong pig *would* want to deviate provided he thought that for some crazy reason the weak pig was going to deviate too and he could end up in the upper right. But as long as the strong pig assumes that the weak pig is going to stick to his strategy of waiting by the food dispenser, he has no desire to change his own strategy.

Nash equilibrium
An outcome from which neither player would want to deviate, taking the other player's behavior as given.

The Prisoner's Dilemma Revisited

The Prisoner's Dilemma of Chapter 11 is already represented by a game matrix, which we reproduce in Exhibit 12.2. We argued in Chapter 11 that the prisoners land in the upper left-hand box. Let us confirm this conclusion using the techniques we've just developed.

Suppose the prisoners were in the upper right-hand box, with B confessing and A not confessing. If B switches strategies, we move down a row, increasing B's prison

EXHIBIT 12.2 **The Prisoner's Dilemma**

		Prisoner A's Strategy	
		Confess	Not confess
Prisoner B's Strategy	Confess	A gets 5 years B gets 5 years	A gets 10 years B gets 1 year
	Not confess	A gets 1 year B gets 10 years	A gets 2 years B gets 2 years

The prisoners face the same dilemma as in Chapter 11. The only Nash equilibrium is in the upper left-hand corner; this is also the only square that is not Pareto-optimal.

term; therefore B does *not* want to switch. But if A switches strategies, then we move a column to the left, where A's prison term falls from 10 years to 5; therefore A *does* want to switch. Because at least one of the prisoners wants to switch, the upper right-hand box is *not* a Nash equilibrium.

Dangerous Curve

It's worth noting that the pigs in a box were out to *maximize* their calorie intake, while the prisoners are out to *minimize* their jail sentences. In all of the other examples of this chapter, the goal will be to maximize outcomes (as the pigs do) rather than to minimize them (as the prisoners do).

Exercise 12.2 Explain why the lower left-hand box and the lower right-hand box are not Nash equilibria. In each case, which prisoner wants to switch?

Dominant Strategies

In Chapter 11, we pointed out that Prisoner A would want to confess *regardless* of his beliefs about Prisoner B's behavior. If Prisoner B is known to be confessing (placing us in the top row), then Prisoner A has a choice between getting a sentence of 5 years by confessing or getting a sentence of 10 years by not confessing. If Prisoner B is known to be not confessing (placing us in the bottom row), then Prisoner A has a choice between getting a sentence of 1 year by confessing or getting a sentence of 2 years by not confessing. Either way, Prisoner A prefers to confess.

Dominant strategy
A strategy that a player would want to follow regardless of the other player's behavior.

Confessing in the Prisoner's Dilemma is called a **dominant strategy** for Prisoner A, because he would want to follow that strategy regardless of what Prisoner B was up to. Confessing is also a dominant strategy for Prisoner B. When both prisoners follow their dominant strategies, we reach the Nash equilibrium outcome where both confess.

Pigs in a Box Revisited

Sometimes a player has no dominant strategy. Let us return to the pigs of Exhibit 12.1. Should the strong pig push the lever or wait by the dispenser?

It depends on what he thinks the weak pig is doing. If the weak pig can be counted on to push the lever, then the strong pig should wait by the dispenser; but if the weak pig waits by the dispenser, then the strong pig should push the lever.

We can see this in the game matrix. If the weak pig pushes the lever we are in the first row. The strong pig can push (for a gain of 90) or wait (for a gain of 100); it is better to wait (that is, to choose the second column). But if the weak pig waits by the dispenser, we are in the second row. The strong pig can push (for a gain of 15) or wait (for a gain of 0); it is better to push (that is, to choose the first column).

Before the strong pig can choose his strategy, he'd like to know what the weak pig is going to do. This means that the strong pig has no dominant strategy. If he had a dominant strategy, he would not need to inquire about the weak pig's behavior before deciding on his own.

The weak pig, by contrast, *does* have a dominant strategy: He should wait by the dispenser regardless of how the strong pig behaves. If the strong pig pushes (choosing the first column), then the weak pig can push (for a gain of −10) or wait (for a gain of 75); it is better to wait (that is, to choose the second row). If the strong pig waits (choosing the second column), then the weak pig can push (for a gain of −10) or wait (for a gain of 0); it is still better to wait (that is, to choose the second row).

Dominant Strategies versus Nash Equilibria

When both players have dominant strategies, as in the Prisoner's Dilemma, there is one and only one Nash equilibrium. In the Nash equilibrium, both players play their dominant strategies.

But Nash equilibria can exist even when one or both players have no dominant strategy. In the "pigs in a box" example of Exhibit 12.1, the strong pig has no dominant strategy, but the lower-left corner is still a Nash equilibrium.

To keep track of the differences in these concepts, continue to focus on the pigs. We know that it is a dominant strategy for the weak pig to wait by the dispenser; in terms of the game matrix this means that the weak pig will always choose the second row.

Now suppose that we are in the lower-left box (where the strong pig is pushing the lever) and consider the following two questions:

1. Would the strong pig want to change strategies, given that he knows the weak pig will choose the second row?
2. Might the strong pig want to change strategies if he wasn't sure what the weak pig will do?

The answer to question 1 is no. Once the second row is chosen, the strong pig certainly prefers the first column to the second. Neither the strong pig nor the weak pig wants to change, so the lower left is a Nash equilibrium.

The answer to question 2 is yes. If the strong pig thought that the weak pig had (foolishly) chosen the first row, then he would want to switch to the second column. His choice of columns depends on what he thinks the weak pig is doing, so he has no dominant strategy.

The Battle of the Sexes

Exhibit 12.3 shows a game that is usually called the *Battle of the Sexes.*

Fred prefers to go to boxing matches and Ethel prefers to go to the opera, but they both like doing things together. If they go their separate ways, both are miserable. The

EXHIBIT 12.3 **The Battle of the Sexes**

		Fred's Strategy	
		Boxing match	Opera
Ethel's Strategy	Boxing match	Fred gets 5 Ethel gets 3	Fred gets 0 Ethel gets 0
	Opera	Fred gets 1 Ethel gets 1	Fred gets 3 Ethel gets 5

Fred likes boxing and Ethel likes opera, but they both like to be together. The upper-left and lower-right corners are Nash equilibria.

game matrix puts numerical values on Fred and Ethel's happiness (which economists sometimes call *utility*). If Fred goes to the opera while Ethel goes to the boxing match, they each earn zero units of utility; if Fred goes to the boxing match while Ethel goes to the opera, they each earn 1 unit.

But if Fred and Ethel attend the boxing match together, then Fred earns 5 units of utility while Ethel earns 3 just by being with Fred; if they attend the opera together, then Ethel earns 5 units of utility and Fred earns 3 just by being with Ethel.

Does Fred have a dominant strategy in this game? If he thinks that Ethel is going to the boxing match, he prefers to be at the boxing match, while if he thinks that Ethel is going to the opera, he prefers to be at the opera. This means that he has no dominant strategy. Neither does Ethel.

What about Nash equilibria? Suppose that Fred and Ethel both go to the boxing match (the upper left-hand corner). Would Fred want to switch to the opera, knowing that Ethel is going to the boxing match? The answer is no. And would Ethel want to switch to the opera knowing that Fred is going to the boxing match? No again. So this outcome is a Nash equilibrium.

The lower right-hand corner (both going to the opera) is also a Nash equilibrium. But the two outcomes where Fred and Ethel go their separate ways are *not* Nash equilibria, because in either of these situations both Fred and Ethel would want to switch.

Dangerous Curve

Suppose that Fred goes to the boxing match while Ethel goes to the opera (the lower left-hand box). Then, given Ethel's plans, Fred prefers to switch, and, given Fred's plans, Ethel prefers to switch. You might wonder whether Ethel would reason a little more deeply. "I know that as long as I am going to the opera, Fred will want to switch to the opera as well, so I think that I'll just head over to the opera and wait for him to follow along." It is true that Ethel might think this way, but such reasoning is not relevant to the question of whether this outcome is a Nash equilibrium. Given Fred's intention to attend the boxing match, Ethel does want to switch. This rules out the lower left-hand corner as a Nash equilibrium.

From the lower left-hand box (or from the upper right-hand box) both Fred and Ethel want to switch (each taking the other's behavior as given). This is more information than necessary to rule out these boxes as Nash equilibria; as long as at least one of Fred and Ethel wants to switch, the box is ruled out.

Dangerous Curve

The Copycat Game

Dot's brother Ditto is a copycat. If Dot watches television, Ditto wants to watch television, too. If Dot goes out to play in the yard, then so does Ditto.

Dot, on the other hand, always wants to be by herself. She's happy watching television as long as Ditto is out in the yard, and happy in the yard as long as Ditto is watching television.

The matrix in Exhibit 12.4 shows Dot and Ditto's game. As long as they are doing something together, Ditto gets 5 units of utility and Dot gets 0. As long as they are apart, Ditto gets 0 units of utility and Dot gets 5.

Are there any Nash equilibria in this game? Consider the upper left-hand corner. If Dot and Ditto are both watching television, Ditto sees no reason to switch columns—but Dot wants to switch rows by going out to the yard. So the upper left-hand corner is not a Nash equilibrium. Neither is any other corner.

Exercise 12.3 Explain why the upper-right, lower-left, and lower-right corners are not Nash equilibria.

Nash Equilibrium as a Solution Concept

A **solution concept** is a rule for predicting how games will turn out when they are played. Nash equilibrium is one of the most popular solution concepts; that is, economists like to posit that when people play games, they end up in Nash equilibria. There are, however, some reasons to be uncomfortable with Nash equilibrium as a solution concept.

Solution concept
A rule for predicting how games will turn out.

EXHIBIT **12.4** **The Copycat Game**

		Dot's Strategy: Watch television	Dot's Strategy: Play in yard
Ditto's Strategy	Watch television	Dot gets 0 Ditto gets 5	Dot gets 5 Ditto gets 0
	Play in yard	Dot gets 5 Ditto gets 0	Dot gets 0 Ditto gets 5

Dot is happy as long as she is alone; Ditto is happy as long as he is with Dot. There is no Nash equilibrium in this game.

One problem is that some games, like the Battle of the Sexes, have more than one Nash equilibrium. There is no way to predict which Nash equilibrium is more likely.

Another problem is that some games, like the Copycat Game, have no Nash equilibrium at all. If Dot and Ditto start out watching television together, Dot will go out to the yard, whereupon Ditto will follow her out, whereupon Dot will come back in, whereupon Ditto will follow her in, whereupon. . . . There is nothing in the Nash equilibrium concept to tell us where this process will end.

Example: The Price of Car Insurance

A 19-year-old male who drives a five-year-old Chevrolet Caprice will pay about $1,800 a year for car insurance if he lives in Columbus, Ohio. That same 19-year-old male will pay about $2,500 if he lives in Detroit, $4,000 if he lives in Philadelphia, and $5,000 if he lives in Los Angeles! What can account for such enormous differences in price?

In a provocatively titled essay,[2] two economists have drawn attention to the "game" where each driver decides whether to buy insurance. They argue that observed price differences can be attributed to multiple Nash equilibria in this game.

Suppose, for example, that very few drivers buy insurance. Then insured drivers, when they have accidents, will usually have to collect from their own insurance companies—the other party will typically be uninsured. Therefore insurance becomes very expensive, so few drivers want to buy it. In other words, uninsured motorists cause high insurance prices, and high insurance prices cause uninsured motorists. This is an example of a Nash equilibrium: Everyone behaves rationally, taking everyone else's behavior as given.

On the other hand, suppose that most drivers buy insurance. Then insurance becomes cheaper and therefore, most drivers want to buy it. Again, we have a Nash equilibrium.

When a game has more than one Nash equilibrium, it's difficult to predict which of the equilibria will actually occur. But once an equilibrium is reached, it tends to remain stable. So if, for any reason, Columbus fell into the "bad" equilibrium while Philadelphia fell into the "good" equilibrium, it's not surprising that these equilibria would maintain themselves over time.

Example: Social Status

The average American earns almost $30,000 a year, according to official statistics, while the average citizen of Mali earns about $100. The latter figure is surely misleadingly low, but the fact remains that there are enormous differences in income across countries. No economist has succeeded in giving a complete account of those differences. Most partial explanations rely on differences in tastes (e.g., people with a strong preference for saving will be wealthier in the long run) and differences in available technology. But recently, a number of economists have pointed to the possibility of multiple equilibria.

---

[2] E. Smith and R. Wright, "Why is Automobile Insurance in Philadelphia So Damn Expensive?" *American Economic Review* 82 (1992), 756–772.

One intriguing story is that the relevant game is the mating game—the "game" in which people select marriage partners. To see how this can be relevant, let's imagine two stylized extremes.[3]

First, imagine a society where the richest people get the most desirable mates. In that society, people will be motivated to save, not just to acquire better mates for themselves, but also to acquire better mates for their children. And as long as all your neighbors play that strategy, you'll want to play it, too. In other words, we have a Nash equilibrium.

Now imagine a society where mates are allocated according to social status, which is inherited from your parents independent of wealth. In such a society, low-status people might try to attract high-status mates by acquiring a lot of wealth. But this strategy is discouraged if it dooms your children to even lower status. So *if* the "rules of the game" are that children of such "mixed marriages" have the lowest status of all, then there can be a Nash equilibrium in which people save very little.

Notice that even if the two societies are populated by identical people, their incomes will evolve very differently. A society that lands in either of the two equilibria will tend to remain there.

These highly stylized examples are far too simplistic to explain all the differences between the United States and Mali, but they do demonstrate that it's possible for multiple Nash equilibria to occur in this context and therefore that multiple equilibria might play an important role in understanding why some countries are so much wealthier than others.

Mixed Strategies

The Copycat Game has no Nash equilibrium. How might we expect Dot and Ditto to select their strategies in this game?

If Ditto can predict Dot's behavior, he will simply mimic it; therefore, it is important for Dot to keep Ditto off guard. One way for her to do this is to flip a coin. On heads, she watches television and on tails she plays in the yard. Because her behavior is now totally unpredictable, Ditto can do no better than to flip his own coin and hope that it lands the same way Dot's does.

Notice that it is important to both Dot and Ditto that their coins be *fair coins,* with heads and tails equally probable. If Dot's coin is weighted so that she is more likely to watch television than to play outside, then Ditto will throw his coin away and watch television, giving him a better than even chance to win the game. And likewise, if Ditto's coin is weighted, then Dot has an opportunity to discard her own coin and follow a strategy that puts the odds on her side.

The Copycat Game is quite symmetric, in the sense that there is always a "winner" with 5 utilities and a "loser" with 0. In a game with less symmetry, Dot and Ditto might prefer to flip weighted coins, sacrificing some unpredictability in exchange for improving the chances of their preferred outcomes. We can view each possible weighting as an alternative strategy. (That is, "flip a fair coin" is one strategy; "flip a coin that comes up heads two-thirds of the time" is another; "flip a coin that comes up heads three-fourths

[3] The example to follow is based on H. Cole, G. Maulath, and A. Postlewaite, "Social Norms, Savings Behavior and Growth," *Journal of Political Economy* 100 (1992), 1092–1125.

Mixed strategy
A strategy that involves a random choice among pure strategies.

Pure strategy
A single choice of row (or column) in the game matrix.

of the time" is still another.) We call these options **mixed strategies,** as opposed to the **pure strategies** illustrated in Exhibit 12.3. If mixed strategies are allowed, then it is possible to prove under quite general circumstances that a Nash equilibrium must exist.

Mixed Strategies in Sports

In the international tournaments organized by the World Rock Paper Scissors Society (yes, that's a real organization), nobody ever consistently plays "Rock." Instead, the best players are the least predictable players. In Nash equilibrium, everyone plays a mixed strategy—1/3 "Rock," 1/3 "Paper," and 1/3 "Scissors."

Exercise 12.4 Explain why a strategy consisting of ½ "Rock," ¼ "Paper," and ¼ "Scissors" cannot be part of a Nash equilibrium.

Mixed strategies are common in more traditional sports as well. In baseball, pitchers want to be unpredictable—the pitcher who always throws a fastball will always face a batter who is prepared for a fastball. A football team that always passes will always face a defense that's prepared for a pass. In soccer, a kicker who always aims his penalty kicks in the same direction will always face a goalie who dives in that direction.

Recently, two economists[4] examined the strategies of championship tennis players. To keep it simple, they assumed that the server has just two options: Serve to the receiver's left or to the receiver's right. And the receiver has just two options: Prepare to receive the serve on the left or on the right. The payoffs depend on the particular strengths of particular players, so the associated game matrix depends on who's playing. The economists estimated the game matrices for various players, computed the Nash equilibrium mixed strategies, and examined the players' actual play. Their conclusion: The evidence is very strong that players do play just as the theory predicts.

Pareto Optima

Nash equilibrium is a *positive* (as opposed to normative) concept; it is designed to predict what *will* happen as opposed to enabling us to discuss what *ought* to happen. In this section, we will discuss the normative side of game theory.

Look again at Fred and Ethel, who played the Battle of the Sexes Game in Exhibit 12.3; this game is reproduced in Exhibit 12.5. In Exhibit 12.5, each of the four outcomes has been labeled with a letter (from A through D) for easy reference.

Fred and Ethel disagree about the desirability of the various outcomes; for example, Fred thinks outcome A is better than outcome D, while Ethel thinks just the opposite. But there are certain things they both agree on. For example, both agree that outcome C (where Fred and Ethel each get 1) is better than outcome B (where they both get 0).

Pareto improvement or Pareto-preferred
A change to which nobody objects.

Because Fred and Ethel are unanimous in this judgment, we say that moving from B to C is a **Pareto improvement,** or that C is **Pareto-preferred** to B. In general, a change is a Pareto improvement if nobody objects to it.[5]

[4] Mark Walker and John Wooders, "Minimax Play at Wimbledon," *American Economic Review* 91 (2002): 1521–1538.

[5] In some books, the phrase *Pareto improvement* is reserved for a change to which nobody objects *and* at least one person prefers.

EXHIBIT **12.5** **The Battle of the Sexes Revisited**

		Fred's Strategy	
		Boxing match	Opera
Ethel's Strategy	Boxing match	(A) Fred gets 5 Ethel gets 3	(B) Fred gets 0 Ethel gets 0
	Opera	(C) Fred gets 1 Ethel gets 1	(D) Fred gets 3 Ethel gets 5

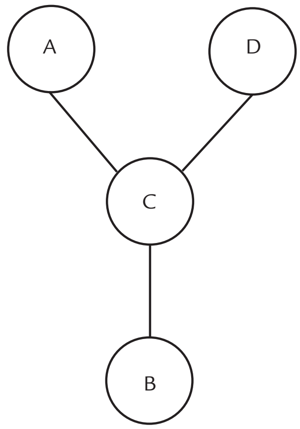

The tree shows that outcomes A and D are Pareto-preferred to C and B, and C is Pareto-preferred to B. A and D are Pareto optima, because nothing sits above them in the tree.

Similarly, outcomes A and D are both Pareto improvements over B; nobody would object to a move from B to A or from B to D. A move from A to D is *not* a Pareto improvement, because Fred would object, and a move from D to A is not a Pareto improvement, because Ethel would object.

To the right of the game matrix in Exhibit 12.5, we have arranged the four outcomes in a "tree," where upward movements represent Pareto improvements. A, C, and D are all Pareto improvements over B, so A, C, and D all sit higher than B in the tree. Likewise, A and D both sit above C. But A sits neither above nor below D, because A is not a Pareto improvement over D and D is not a Pareto improvement over A.

We say that an outcome is **Pareto-optimal** if nothing sits above it in the tree. In this example, outcomes A and D are Pareto-optimal. From a normative point of view, we can think of outcomes that are *not* Pareto-optimal as "bad" outcomes. Outcome C, for example, is "bad" in the sense that both Fred and Ethel would prefer to climb higher in the tree, though they might disagree about whether it would be better to climb to A or to D.

Pareto-optimal
An outcome that allows no possibility of a Pareto improvement.

Exhibit 12.6 revisits the pigs in a box from Exhibit 12.1. Here outcome B is Pareto-preferred to outcome A and outcome C is Pareto-preferred to outcome D, but there are no other instances of Pareto improvements. Thus, the "tree" breaks into two pieces, one of which shows B above A and one of which shows C above D. The Pareto-optimal outcomes are at the tops of the trees: B and C.

Exercise 12.5 Explain why B is not Pareto-preferred to C or D. Explain why C is not Pareto-preferred to A or B.

Exercise 12.6 Build a tree for the Prisoner's Dilemma of Exhibit 12.2, keeping in mind that in this game, a shorter prison sentence is better than a long one. What are the Pareto optima in this game?

EXHIBIT 12.6 Pigs in a Box Revisited

		Strong Pig's Strategy	
		Push lever	Wait by dispenser
Weak Pig's Strategy	Push lever	(A) Strong pig gets 90 calories Weak pig gets −10 calories	(B) Strong pig gets 100 calories Weak pig gets −10 calories
	Wait by dispenser	(C) Strong pig gets 15 calories Weak pig gets 75 calories	(D) Strong pig gets 0 calories Weak pig gets 0 calories

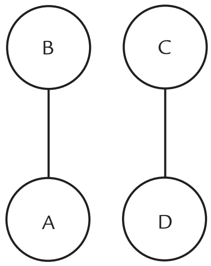

B is Pareto-preferred to A, and C is Pareto-preferred to D. B and C are the Pareto optima.

Pareto Optima versus Nash Equilibria

The pigs in Exhibit 12.1 have two Pareto optima (the lower left and upper right) but only one Nash equilibrium (the lower left). The Nash equilibrium happens to be one of the Pareto optima. But this is not always the case.

Consider the Prisoner's Dilemma of Exhibit 12.2. Here we have already seen that the only Nash equilibrium occurs in the upper left. This outcome is not Pareto-optimal, because a shift to the lower right would benefit both prisoners. In fact, the Nash equilibrium is the only outcome that is not Pareto-optimal.

Exercise 12.7 Explain why the upper-right box in the Prisoner's Dilemma is Pareto-optimal.

Exercise 12.8 Explain why the lower-left box in the Prisoner's Dilemma is Pareto-optimal.

In the Battle of the Sexes (Exhibit 12.3), both of the Nash equilibria (in the upper left and lower right) are Pareto-optimal. Starting in the upper left, any other square would be worse for Fred, and starting in the lower right, any other square would be worse for Ethel. Neither of the other two squares is Pareto-optimal.

Exercise 12.9 Explain why neither of the other two squares is Pareto-optimal.

12.2 Sequential Games

You have probably played the game of "scissors, paper, rock." Each player chooses one of three strategies (scissors, paper, or rock) and then the winner is determined by these rules: Scissors "cut" paper, paper "covers" rock, and rock "smashes" scissors.

Usually both players are required to choose their strategies simultaneously. There is a good reason for this. If players took turns, the second player would always win. Once you know what your opponent is doing, it is easy to choose a strategy that will defeat him.

On the other hand, there are games where it pays to go first instead of second. Consider the Battle of the Sexes (Exhibit 12.3), where Fred and Ethel disagree about where to spend the evening but want above all to be together. If Fred moves first, by going to the boxing match and waiting for Ethel to follow along, then she is sure to do so, giving Fred his most preferred outcome. If Ethel moves first by going to the opera, Fred follows her and Ethel wins.

In the games of Section 12.1, we have always assumed that both players must choose their strategies simultaneously. In this section, we will assume instead that there is a first player, who chooses a column in the game matrix, and then a second player, who chooses a row. This will require a new way of thinking about the outcome. We will illustrate the new method with some examples.

An Oligopoly Problem

Kodak and Fuji produce photographic film. Suppose that there are no other significant firms in this industry, so that Kodak and Fuji constitute an oligopoly. Industrywide profits depend on industrywide output according to the following table:

Quantity (rolls of film per day)	Profits (dollars per day)
100	32
125	35
150	30
175	21
200	10

Moreover, the profits are divided in proportion to the firms' output. Thus, if one firm produces 100 rolls of film while the other produces 75 rolls (a ratio of 4 to 3), then the $21 profit is divided in the same ratio ($12 for one firm and $9 for the other).

Exhibit 12.7 shows the game matrix, where each company can produce either 50, 75, or 100 rolls of film.

The outcome of this game depends very much on how the game is played. Suppose first that the companies are able to collude, maximizing their joint profits and splitting them afterward. Then they will produce 125 rolls of film, for the maximum possible profit of $35.

Suppose instead that each company takes its rival's output as given and chooses its own output accordingly. In the language of game theory, this means that the companies achieve a Nash equilibrium in Exhibit 12.7. In the language of Chapter 11, we called the same thing a *Cournot equilibrium*. A Cournot equilibrium is nothing but a Nash equilibrium in a game where each company chooses its quantity.

EXHIBIT 12.7 **An Oligopoly Problem**

Fuji's Strategy \ Kodak's Strategy	50	75	100
50	Kodak gets 16 Fuji gets 16	Kodak gets 21 Fuji gets 14	Kodak gets 20 Fuji gets 10
75	Kodak gets 14 Fuji gets 21	Kodak gets 15 Fuji gets 15	Kodak gets 12 Fuji gets 9
100	Kodak gets 10 Fuji gets 20	Kodak gets 9 Fuji gets 12	Kodak gets 5 Fuji gets 5

The only Nash equilibrium is in the center square, where Kodak and Fuji each earn profits of $15. But if the game is played sequentially and Kodak moves first, then Kodak announces a policy of producing 100 rolls of film. Fuji's best response is to produce 50, leading to the upper right-hand square.

In Exhibit 12.7, the only Nash equilibrium is the center square. If each firm makes 75 rolls of film, neither wants to deviate. Kodak recognizes that dropping its output to 50 rolls would lower its profits from $15 to $14 and raising its output to 100 rolls would lower its profits from $15 to $12. Fuji recognizes the same thing.

Exercise 12.10 Explain why no other square in Exhibit 12.7 is a Nash equilibrium.

But now let's change the rules of the game. Suppose that Kodak is able to announce its output before Fuji gets to make a move. Now what will Kodak do?

Kodak needs to think through the consequences of each possible strategy. Suppose that Kodak produces 50 rolls of film (committing itself to the first column). Fuji will then pick its favorite square in the first column, producing 75 rolls for a profit of $21 (beating $16 and $20 in the other squares). Kodak ends up with $14 profit.

Suppose instead that Kodak produces 75 rolls of film (committing itself to the second column). Fuji will then pick its favorite square in the second column, producing 75 rolls for a profit of $15 (beating $14 and $12 in the other squares). Kodak ends up with $15 profit.

Suppose instead that Kodak produces 100 rolls of film (committing itself to the third column). Fuji will then pick its favorite square in the third column, producing 50 rolls for a profit of $10 (beating $9 and $5). Kodak ends up with $20 profit.

Stackelberg equilibrium

An equilibrium concept that arises when one player announces his strategy before the other.

Among these choices, Kodak likes the last one best. So Kodak announces that it will produce 100 rolls. Fuji responds by producing 50, and the game ends in the upper right-hand square, where Kodak earns twice what Fuji earns.

The outcome we have just described is called a **Stackelberg equilibrium.** A Stackelberg equilibrium occurs when one player commits to a strategy at the outset, accounting for the fact that the second player will choose an optimal response.

The Importance of Commitment

Suppose that Kodak announces it will produce 100 rolls of film and Fuji responds by producing 50 rolls as in the Stackelberg equilibrium of Exhibit 12.7. Once Fuji has agreed to produce only 50 rolls, Kodak wants to deviate. It is better for Kodak to produce 75 rolls for a profit of $21 than 100 rolls for a profit of $20.

So if Kodak moves first and Fuji moves second, then Kodak wants to change its move. If Kodak does change its move, and if Fuji foresees this, then Fuji goes ahead with plans to produce not 50 rolls of film but 75. (After all, Kodak will eventually place it in the middle column, where Fuji's optimal strategy is not 50 but 75.) The firms end up at the Nash equilibrium in the center instead of the Stackelberg equilibrium in the upper right. Kodak's profits fall from $20 to $15.

This means that Kodak is better off if it can commit itself to producing 100 rolls and assure Fuji that it is never going to back down from that commitment. This might surprise you. You might think that a firm is better off leaving itself some flexibility to deal with unforeseen contingencies. But that is not always so.

Consider the game of chicken, where two cretins drive their cars directly at each other until one of them loses by swerving. If you can absolutely guarantee that you will never swerve, you are a sure winner at this game. If you leave yourself the leeway to swerve in case your opponent is crazier than you are, then your opponent will have an incentive to *become* crazier than you are and you are liable to lose. The way to win the game of chicken is to disable your steering column and make sure your opponent is aware of it.

Summary

Strategic situations can be represented by game matrices, showing the outcome that results from each combination of strategies that the players can choose.

A Nash equilibrium is an outcome from which neither player would deviate, taking the other's behavior as given. A game can have one Nash equilibrium, no Nash equilibrium, or many Nash equilibria.

A dominant strategy is a strategy that a player would want to adopt regardless of his beliefs about the other player's strategy choice. The Prisoner's Dilemma is an example of a game where both players have dominant strategies.

One outcome is a Pareto improvement over another if it makes at least one player better off without making any player worse off. An outcome is Pareto-optimal if it allows no Pareto improvements.

There can be Nash equilibria that are not Pareto-optimal, and there can be Pareto optima that are not Nash equilibria.

When games are played sequentially instead of simultaneously, the Nash equilibrium is no longer a natural solution concept. Instead, we use the Stackelberg equilibrium, where the first player calculates the second player's responses to each of his possible strategies and then chooses the strategy that will yield him the best outcome. In a sequential game, it can be advantageous to go first or advantageous to go second, depending on the particular game.

In some games it is important to be able to commit to following a strategy even if better options become available. By committing, you can sometimes convince your opponent to behave in ways that are advantageous to you.

Author Commentary

www.cengage.com/economics/landsburg

AC1. For more information on multiple equilibria in the market for car insurance, see this article.

AC2. Read this article to learn about some slightly outdated applications of game theory to American politics.

AC3. This article is about the social status game.

Problem Set

The problems in this problem set refer to the following game matrices. In each case, Jack chooses "left or right" and Jill chooses "up or down." The outcomes show how many buckets of water are rewarded.

I.

		Jack's Strategy	
		Left	Right
Jill's Strategy	Up	Jack gets 1 Jill gets 1	Jack gets 4 Jill gets 2
	Down	Jack gets 2 Jill gets 4	Jack gets 3 Jill gets 3

II.

		Jack's Strategy	
		Left	Right
Jill's Strategy	Up	Jack gets 1 Jill gets 1	Jack gets 2 Jill gets 4
	Down	Jack gets 4 Jill gets 2	Jack gets 3 Jill gets 3

III.

		Jack's Strategy	
		Left	Right
Jill's Strategy	Up	Jack gets 1 Jill gets 1	Jack gets 4 Jill gets 4
	Down	Jack gets 2 Jill gets 2	Jack gets 3 Jill gets 3

IV.

Jill's Strategy	Jack's Strategy: Left	Jack's Strategy: Right
Up	Jack gets 2 Jill gets 2	Jack gets 4 Jill gets 1
Down	Jack gets 1 Jill gets 4	Jack gets 3 Jill gets 3

V.

Jill's Strategy	Jack's Strategy: Left	Jack's Strategy: Right
Up	Jack gets 1 Jill gets 3	Jack gets 3 Jill gets 1
Down	Jack gets 4 Jill gets 2	Jack gets 2 Jill gets 4

VI .

Jill's Strategy	Jack's Strategy: Left	Jack's Strategy: Right
Up	Jack gets 2 Jill gets 2	Jack gets 1 Jill gets 1
Down	Jack gets 1 Jill gets 1	Jack gets 3 Jill gets 3

VII.

Jill's Strategy	Jack's Strategy: Left	Jack's Strategy: Right
Up	Jack gets 2 Jill gets 3	Jack gets 1 Jill gets 1
Down	Jack gets 1 Jill gets 1	Jack gets 3 Jill gets 2

VIII.

Jill's Strategy	Jack's Strategy: Left	Jack's Strategy: Right
Up	Jack gets 12 Jill gets 8	Jack gets 9 Jill gets 8
Down	Jack gets 15 Jill gets 7	Jack gets 14 Jill gets 10

1. In each game above, identify all of the Nash equilibria.
2. In each game above, identify all of the Pareto optima.
3. In each game above, does Jack have a dominant strategy? Does Jill?
4. In each game above, what happens if Jack goes first?
5. In each game above, what happens if Jill goes first?
6. For each game above, create a reasonable story (like those that go with the exhibits in the text) that might lead to these numbers appearing in the matrix.
7. Create a "tree" showing which outcomes are Pareto-preferred to which in the Kodak–Fuji game of Exhibit 12.7.
8. Can you find examples of games (either among those that have appeared in the chapter or by creating them yourself) with the following characteristics?
   - **a.** There are no Nash equilibria.
   - **b.** There is exactly one Nash equilibrium, but it is not Pareto-optimal.
   - **c.** There is more than one Nash equilibrium, but none of them is Pareto-optimal.
   - **d.** There is more than one Nash equilibrium, and all of them are Pareto-optimal.
   - **e.** There is more than one Nash equilibrium, and some are Pareto-optimal while others are not.
9. Can there be a game with no Pareto optimum?
10. Suppose that the games of Exhibits 12.1, 12.2, 12.3, and 12.4 were played as sequential games. In each case, suppose that the player who chooses a column goes first. What are the outcomes of these games? Now suppose that the player who chooses a row goes first. In which cases do the outcomes change?
11. **True or False:** In a sequential game where the second player has a dominant strategy, he will always adopt that strategy.
12. **True or False:** In a sequential game where the first player has a dominant strategy, he will always adopt that strategy.