

FIGURE 4.1 The Game of Chicken

		Driver 2	
		Swerve	Hang tough
Driver 1	Swerve	2,2	1,3
	Hang tough	3,1	0,0

FIGURE 4.1 provides a strategic form representation of Chicken.* Because neither player has a strictly dominated strategy, the iterative deletion of strictly dominated strategies (IDSDS, Section 3.4.2) won't help us solve this game. But don't forsake hope, as game theory has many more game-solving tricks to offer.

If you've either read the book or seen the movie *A Beautiful Mind*, then you know about the brilliant schizophrenic mathematician John Nash. In his

doctoral thesis at Princeton University, Dr. Nash made two striking game-theoretic advances—one of which became known as Nash equilibrium—that resulted in his winning the Nobel Prize in Economics more than 40 years later.

To understand what Nash equilibrium is and why it is an appropriate method for solving a game, let us return to the discussion of the previous chapter. In the context of a game, a player is rational when he chooses a strategy to maximize his payoff, given his beliefs about what other players will do. The tricky part is figuring out what is reasonable for a player to believe about the strategy another player will select. Chapter 3 used the assumption that rationality is common knowledge among the players to derive those beliefs. For example, if player 2 has a dominant strategy and player 1 believes that player 2 is rational, then player 1 believes that player 2 will use her dominant strategy. In this manner, we derived player 1's beliefs regarding player 2's strategy.

The approach of Nash equilibrium maintains the assumption that players are rational, but takes a different approach to nailing down beliefs. What Nash equilibrium does is require that each player's beliefs about other players' strategies be *correct*. For example, the strategy that player 1 conjectures that player 2 will use is exactly what player 2 actually does use. The definition of Nash equilibrium is then made up of two components:

- 1. Players are rational:** Each player's strategy maximizes his payoff, given his beliefs about the strategies used by the other players.
- 2. Beliefs are accurate:** Each player's beliefs about the strategies used by the other players are true.

Condition (1) is innocent enough; it's condition (2) that is tougher to swallow. It requires that players be effective prognosticators of the behavior of others. In some settings, that may be a reasonable assumption; in others, it may not. Combining the assumptions about behavior—that it is always rational—and beliefs—that they are always true—gives us the definition of Nash equilibrium.

+ DEFINITION 4.1 A strategy profile is a **Nash equilibrium** if each player's strategy maximizes his or her payoff, given the strategies used by the other players.

With n players, there are, then, n conditions that must be satisfied in order for a strategy profile to be a Nash equilibrium—one condition for each player

*You might be tempted to put large negative numbers for the strategy pair in which both participants choose *hang tough*, since this means certain injury and possible death. You can do so, but it'll make no difference as regards the solution. As long as the payoffs when both hang tough are less than all the other payoffs in the matrix, our conclusions regarding behavior will be the same. This condition reflects the property that what matters is the *ranking* of the payoffs, not their actual values.

which ensures that a player's strategy is optimal, given the other players' strategies. Thus, all players are simultaneously doing their best. A violation of one or more of those conditions means that a strategy profile is not a Nash equilibrium. Unlike the game of horseshoes, you don't come "close to a Nash equilibrium" by having all but one of the conditions satisfied; it's either all or nothing.

An appeal of Nash equilibrium as a solution concept is that it identifies strategy profiles that are stable in the sense that each player is content to do what she is doing, given what everyone else is doing. Consider, for instance, a strategy profile that is *not* a Nash equilibrium because, say, player 3's strategy is not best for her, given what the other players are up to. We would then expect player 3 to change her strategy once she discovers that it is not optimal. In contrast, a Nash equilibrium is not subject to such second-guessing, because players are happy with what they are doing.

To be more concrete on this point, imagine that players play the same game over and over. If they are not currently acting according to a Nash equilibrium, then, after one of the game's interactions, there will be a player who will learn that his strategy is not the best one available, given what others are doing. He will then have an incentive to change his strategy in order to improve his payoff. In contrast, if players are behaving according to a Nash equilibrium, they are satisfied with their actions after each round of interactions. Behavior generated by a Nash equilibrium is then expected to persist over time, and social scientists are generally interested in understanding persistent behavior (not necessarily because unstable behavior is uninteresting, but rather because it is just much harder to explain).

Hopefully having convinced you that Nash equilibrium is a worthy solution concept (and if not, bear with me), let's put it to use with the game of Chicken. We begin by considering the four strategy pairs and asking whether each is a Nash equilibrium.

- (*hang tough, hang tough*). If driver 2 chooses *hang tough*, then driver 1's payoff from *swerve* is 1 and from *hang tough* is 0. (See FIGURE 4.2.) Thus, driver 1 prefers to swerve (and live with a few clucking sounds from his friends) than to hang tough (and learn whether or not there is an afterlife). Thus, *hang tough* is not best for player 1, which means that player 1's Nash equilibrium condition is not satisfied. Hence, we can conclude that (*hang tough, hang tough*) is not a Nash equilibrium. (It is also true that driver 2's strategy of *hang tough* is not best for her either, but we've already shown this strategy pair is not a Nash equilibrium.)
- (*swerve, swerve*). If driver 1 chooses *swerve*, then driver 2's payoff from *swerve* is 2 and from *hang tough* is 3. (See FIGURE 4.3.) Driver 2 thus prefers *hang tough* if driver 1 is going to chicken out. Since *swerve* is not the best strategy for driver 2, (*swerve, swerve*) is not a Nash equilibrium either.

FIGURE 4.2 Chicken: Highlighting Driver 1's Payoffs when Driver 2 Chooses Hang Tough

		Driver 2	
		Swerve	Hang tough
Driver 1	Swerve	2,2	1,3
	Hang tough	3,1	0,0

FIGURE 4.3 Chicken: Highlighting Driver 2's Payoffs when Driver 1 Chooses Swerve

		Driver 2	
		Swerve	Hang tough
Driver 1	Swerve	2,2	1,3
	Hang tough	3,1	0,0

- (*swerve, hang tough*). If driver 2 chooses *hang tough*, *swerve* is the best strategy for driver 1, as it produces a payoff of 1 compared with 0 from hanging tough. Consequently, the requirement that driver 1's strategy is best for him is satisfied. Turning to driver 2, we see that *hang tough* is best for her, because it yields a payoff of 3, rather than 2 from *swerve*. The condition ensuring that driver 2's strategy is optimal for her is satisfied as well. Because each driver is choosing the best strategy, given what the other driver is expected to do, (*swerve, hang tough*) is a Nash equilibrium.
- (*hang tough, swerve*). By logic similar to that in the preceding case, this strategy pair is a Nash equilibrium, too.

Summing up, there are two Nash equilibria in this game: (*swerve, hang tough*) and (*hang tough, swerve*). Both predict that there will be no car crash and, furthermore, that one and only one driver will swerve. However, Nash equilibrium doesn't tell us which driver will swerve.

Perhaps the best way to play Chicken is to commit to not swerving by eliminating *swerve* from your strategy set and, most importantly, making this known to the other driver. FIGURE 4.4 illustrates what the game would look like if driver 1 were to eliminate *swerve* from his strategy set. The game now has only one Nash equilibrium: driver 1 hangs tough and driver 2 chickens out.

A tactic similar to that illustrated in Figure 4.4 was taken in a naval encounter about 20 years ago.

Let's listen in on the radio conversation between the two participants.²

1: "Please divert your course 15 degrees to the north to avoid a collision."

2: "Recommend that you change your course 15 degrees to the south to avoid a collision."

1: "This is the captain of a U.S. navy ship. I say again, divert your course."

2: "No, I say again, divert your course."

1: "This is the aircraft carrier Enterprise; we are a large warship of the U.S. navy. Divert your course now!"

2: "This is a lighthouse. Your call."

We have several tasks ahead of us in this chapter. Having defined Nash equilibrium, we want to learn how to solve games for Nash equilibria and begin to appreciate how this concept can be used to derive an understanding of human behavior. Our analysis commences in Section 4.2 with some simple two-player games that embody both the conflict and mutual interest that can arise in strategic situations. To handle more complicated games, the best-reply method for solving for Nash equilibria is introduced in Section 4.3 and is then applied to three-player games in Section 4.4. Finally, Section 4.5 goes a bit deeper into understanding what it means to suppose that players behave as described by a Nash equilibrium.

4.2 Classic Two-Player Games

THE MAIN OBJECTIVE OF this chapter is to get you comfortable both with the concept of Nash equilibrium and with deriving Nash equilibria. Let's warm up with a few simple games involving two players, each of whom has at most

FIGURE 4.4 Chicken when Driver 1 Has Eliminated Swerve as a Strategy

		Driver 2	
		<i>Swerve</i>	<i>Hang tough</i>
Driver 1	<i>Hang tough</i>	3,1	0,0

three strategies. As we'll see, a game can have one Nash equilibrium, several Nash equilibria, or no Nash equilibrium. The first case is ideal in that we provide a definitive statement about behavior. The second is an embarrassment of riches: we cannot be as precise as we'd like, but in some games there may be a way to select among those equilibria. The last case—when there is no Nash equilibrium—gives us little to talk about, at least at this point. Although in this chapter we won't solve games for which there is no Nash equilibrium, we'll talk extensively about how to handle that problem in Chapter 7.

You may be wondering whether there is an “easy-to-use” algorithm for solving Nash equilibria. Chapter 3, for example, presented an algorithm for finding strategies consistent with rationality being common knowledge: the iterative deletion of strictly dominated strategies (IDSDS). Unfortunately, there is no such method for solving Nash equilibrium. For finite games—that is, when there is a finite number of players and each player has a finite number of strategies—the only universal algorithm is **exhaustive search**, which means that one has to check each and every strategy profile and assess whether it is a Nash equilibrium. We will, however, present some shortcuts for engaging in exhaustive searches.

A useful concept in deriving Nash equilibria is a player's *best reply* (or best response). For each collection of strategies for the other players, a player's **best reply** is a strategy that maximizes her payoff. Thus, a player has not just one best reply, but rather a best reply for each configuration of strategies for the other players. Furthermore, for a given configuration of strategies for the other players, there can be more than one best reply if there is more than one strategy that gives the highest payoff.

+ DEFINITION 4.2 A **best reply** for player i to $(s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$ is a strategy that maximizes player i 's payoff, given that the other $n - 1$ players use strategies $(s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$.

A Nash equilibrium can be understood as a strategy profile which ensures that a player's strategy is a best reply to the other players' strategies, for each and every player. These are the same n conditions invoked by Definition 4.1, but we're just describing them a bit differently.

► SITUATION: PRISONERS' DILEMMA

During the time of Stalin, an orchestra conductor was on a train reading a musical score. Thinking that it was a secret code, two KGB officers arrested the conductor, who protested that it was just Tchaikovsky's Violin Concerto. The next day, the interrogator walks in and says, “You might as well confess, as we've caught your accomplice Tchaikovsky, and he's already talking.”

The Prisoners' Dilemma, which we previously considered under the guise of the opera *Tosca*, is the most widely examined game in game theory. Two members of a criminal gang have been arrested and placed in separate rooms for interrogation. Each is told that if one testifies against the other and the other does not testify, the former will go free and the latter will get three years of jail time. If each testifies against the other, they will both be sentenced to two years. If neither testifies against the other, each gets one year. Presuming that each player's payoff is higher when he receives a shorter jail sentence, the strategic form is presented in **FIGURE 4.5**.

FIGURE 4.5 The Prisoners' Dilemma

		Criminal 2	
		Testify	Silence
Criminal 1	Testify	2,2	4,1
	Silence	1,4	3,3

This game has a unique Nash equilibrium, which is that both players choose *testify*. Let us first convince ourselves that $(testify, testify)$ is a Nash equilibrium. If criminal 2 testifies, then criminal 1's payoff from also testifying is 2, while it is only 1 from remaining silent. Thus, the condition ensuring that criminal 1's strategy is optimal is satisfied. Turning to criminal 2, we see that, given that criminal 1 is to testify, she earns 2 from choosing *testify* and 1 from choosing *silence*. So, the condition ensuring that criminal 2's strategy is optimal is also satisfied. Hence, $(testify, testify)$ is a Nash equilibrium.

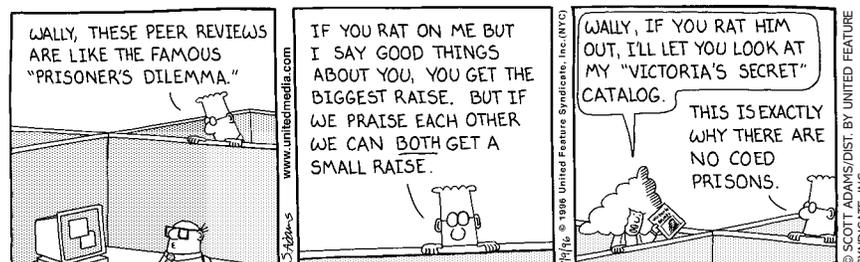
Let us make two further points. First, the Prisoners' Dilemma is an example of a *symmetric game*. A two-player game is **symmetric** if players have the same strategy sets, and if you switch players' strategies, then their payoffs switch. For example, if the strategy pair is $(testify, silence)$, then the payoffs are 4 for criminal 1 and 1 for criminal 2. If we switch their strategies so that the strategy pair is $(silence, testify)$, the payoffs switch: now criminal 1's payoff is 1 while criminal 2's payoff is 4. A trivial implication of the symmetric condition is that players who choose the same strategy will get the same payoff.

An important aspect of symmetric games is that if a symmetric strategy profile—such as $(testify, testify)$ —is optimal for one player, it is also optimal for the other player. If criminal 1's strategy is optimal (given what criminal 2 is doing), then it must also be the case that criminal 2's strategy is optimal (given what criminal 1 is doing). By the symmetry in the game and the consideration of a symmetric strategy profile, the equilibrium conditions for the players are identical. Thus, for a symmetric strategy profile in a symmetric game, either all of the Nash equilibrium conditions hold or none do. Now that we know this property, it is sufficient to show that if *testify* is optimal for criminal 1 (given that criminal 2 chooses *testify*), *testify* must also be optimal for criminal 2.

INSIGHT For a symmetric strategy profile in a symmetric game, if one player's strategy is a best reply, then all players' strategies are best replies.

Note also that *testify* is a dominant strategy. Regardless of what the other criminal does, *testify* produces a strictly higher payoff than *silence*. With a little thought, it should be clear that if a player has a dominant strategy, then a Nash equilibrium must have her using it. For a player's strategy to be part of a Nash equilibrium, the strategy must be optimal, given the strategies used by the other players. Because a dominant strategy is *always* the uniquely best strategy, then it surely must be used in a Nash equilibrium. It follows that if all players have dominant strategies—as in the Prisoners' Dilemma—the game has a unique Nash equilibrium in which those dominant strategies are used. Thus, $(testify, testify)$ is the unique Nash equilibrium in the Prisoners' Dilemma.

Dilbert/Scott Adams



INSIGHT If a player has a dominant strategy, a Nash equilibrium requires that the player use it. If all players have a dominant strategy, then there is a unique Nash equilibrium in which each player uses his or her dominant strategy.

► **SITUATION: A COORDINATION GAME—DRIVING CONVENTIONS**

We next look at an example of a **coordination game**, which has the property that players have a common interest in coordinating their actions. A coordination game that most adults engage in every day is the choice of the side of the road upon which to drive. It's not really that important whether everyone drives on the left (as in England and Japan) or on the right (as in the United States and Chile), but just that we agree to a standard.

The game between two drivers is represented in **FIGURE 4.6**. It is easy to verify that there are two Nash equilibria. One has both Thelma and Louise driving on the left, and the other has both driving on the right. If Louise drives on the left, then Thelma's payoff from doing the same is 1, while it is -1 from driving on the right. Thus, *left* is indeed best for Thelma, given that Louise is choosing *left*. The same argument verifies that Louise's driving on the left is best, given that Thelma drives on the left.

In fact, since this is a symmetric game, I can invoke the magic words—"by symmetry"—to conclude that Louise's strategy of *left* is optimal as well. This makes (*left, left*) a Nash equilibrium. An analogous argument allows us to conclude that (*right, right*) is a Nash equilibrium.

In contrast, (*left, right*) is not a Nash equilibrium. Given that Louise is driving on the right, Thelma's payoff from driving on the left is -1 , while she can do better by driving on the right and getting a payoff of 1. It is straightforward also to argue that (*right, left*) is not a Nash equilibrium.

Nash equilibrium doesn't tell us which standard a population of drivers will settle upon; instead, it tells us only that they will settle upon some standard. History shows that societies do settle upon a driving convention, and which side of the road it is can vary across time and space. It is estimated that about 75% of all roads have the custom of driving on the right.³ Although today everyone conforms to a driving convention because it's the law, conventions developed long before they were legislated (and, indeed, long before automobiles came on the scene). Generally, the law just codified a custom that had developed on its own.

CONUNDRUM Suppose an American driver and an English driver are driving towards each other on neutral ground. Each is driving a car from her own country. All of this is common knowledge. On which side of the road will each drive? Will it turn into a game of Chicken? Is Nash equilibrium a good predictor?

► **SITUATION: A GAME OF COORDINATION AND CONFLICT—TELEPHONE**

In the driving conventions game, there were two Nash equilibria and the players were indifferent between them: driving on the right was just as good as driving on the left. Now let us consider a setting that also has two equilibria, but the players rank them differently.

FIGURE 4.6 Driving Conventions

		Louise	
		Drive left	Drive right
Thelma	Drive left	1, 1	-1, -1
	Drive right	-1, -1	1, 1

FIGURE 4.7 The Telephone Game

		Winnie	
		Call	Wait
Colleen	Call	0,0	2,3
	Wait	3,2	1,1

Colleen is chatting on the phone with Winnie and suddenly they're disconnected. Should Colleen call Winnie or should Winnie call Colleen? Colleen and Winnie are the players, and they have a strategy set composed of *call* and *wait*. Each is willing to call the other if that is what it takes to continue the conversation, but each would prefer the other to do so. If they both try to call back, then the other's phone is busy and thus they don't reconnect. Obviously, if neither calls back, then they don't reconnect either. The strategic form game is shown in **FIGURE 4.7**.*

The strategy pair (*call, call*) is not a Nash equilibrium, since, for example, Colleen earns 0 from *call*, but a higher payoff of 3 from *wait*. Thus, Colleen's strategy of *call* is not optimal should Winnie choose *call*. Nor is (*wait, wait*) a Nash equilibrium, since, for example, Winnie earns 1 from choosing *wait* and a higher payoff of 2 from choosing *call* should Colleen choose *wait*. Thus, Winnie's strategy is not optimal, given what Colleen is doing.

A strategy pair that is a Nash equilibrium has one player calling back while the other waits. Consider the strategy pair (*call, wait*), so that Colleen is to call. Given that Winnie is waiting for Colleen's call, Colleen prefers to call, as it delivers a payoff of 2 while waiting has a payoff of only 1. And if Winnie anticipates that Colleen will call her, then Winnie prefers to wait for Colleen's call. For Winnie, waiting delivers a payoff of 3, whereas should she call, the payoff is 0.

If the telephone game is symmetric (and it is), we can infer by symmetry from (*call, wait*) being a Nash equilibrium that (*wait, call*) is also a Nash equilibrium (where now it is Winnie who is to call). To see why, first convince yourself that this is a symmetric game. If the strategy pair is (*call, wait*), then Colleen gets a payoff of 2 and Winnie gets 3. If we switch the strategies so that the strategy pair is (*wait, call*), then Colleen now gets a payoff of 3 and Winnie gets 2. Furthermore, they have the same payoff when they choose the same strategy. This symmetry implies that the Nash equilibrium condition for Winnie at strategy pair (*wait, call*) is the same as the Nash equilibrium condition for Colleen at (*call, wait*), and the Nash equilibrium condition for Colleen at (*wait, call*) is the same as the Nash equilibrium condition for Winnie at (*call, wait*). Thus, if, at (*call, wait*), Colleen finds it optimal to choose *call*, then Winnie must find it optimal to choose *call* at strategy pair (*wait, call*); the conditions are exactly the same. Either both conditions hold or neither does. Similarly, if, at (*call, wait*), Winnie finds it optimal to choose *wait*, then Colleen must find it optimal to choose *wait* at (*wait, call*).

In sum, by virtue of the game being symmetric, either both (*wait, call*) and (*call, wait*) are Nash equilibria or neither are. Since we've shown that (*call, wait*) is a Nash equilibrium, then, by symmetry, (*wait, call*) is a Nash equilibrium.

CONUNDRUM If a game is symmetric, but the equilibrium is asymmetric, how do players coordinate? How would you coordinate in the telephone game?

*This is the same game as the well-known "Battle of the Sexes," though recast in a more gender-neutral setting. The original game was one in which the man wants to go to a boxing match and the woman wants to go to the opera. Both would prefer to do something together than to disagree.

► **SITUATION: AN OUTGUESSING GAME—ROCK–PAPER–SCISSORS**

Lisa: *Look, there's only one way to settle this: Rock–Paper–Scissors.*

Lisa's Brain: *Poor predictable Bart. Always picks rock.*

Bart's Brain: *Good ol' rock. Nothin' beats that!*

(Bart shows rock, Lisa shows paper)

Bart: *Doh!*

—FROM THE EPISODE “THE FRONT,” OF *THE SIMPSONS*.

How many times have you settled a disagreement by using Rock–Paper–Scissors? In case you come from a culture that doesn't use this device, here's what it's all about. There are two people, and each person moves his hands up and down four times. On the fourth time, each person comes down with either a closed fist (which signals her choice of *rock*), an open hand (signaling *paper*), or the middle finger and forefinger in the shape of *scissors* (no explanation required). The winner is determined as follows: If one person chooses *rock* and the other *scissors*, then *rock* wins, since scissors break when trying to cut rock. If one person chooses *rock* and the other *paper*, then *paper* wins, as paper can be wrapped around rock. And if one person chooses *paper* and the other *scissors*, then *scissors* wins, since scissors can cut paper. If the two players make identical choices, then it is considered a draw (or, more typically, they play again until there is a winner).

If we assign a payoff of 1 to winning, -1 to losing, and 0 to a draw, then the strategic form game is as described in **FIGURE 4.8**. Contrary to Bart's belief, *rock* is not a dominant strategy. While *rock* is the unique best reply against *scissors*, it is not the best reply against *paper*. In fact, there is no dominant strategy. Each strategy is a best reply against some strategy of the other player. *Paper* is the unique best reply against *rock*, *rock* is the unique best reply against *scissors*, and *scissors* is the unique best reply against *paper*.

Without any dominated strategies, the IDSDS won't get us out of the starting gate; all strategies survive the IDSDS. So, being good game theorists, we now pull Nash equilibrium out of our toolbox and go to work. After much hammering and banging, we chip away some of these strategy pairs. We immediately chip off $(rock, rock)$, as Bart ought to choose *paper*, not *rock*, if Lisa is choosing *rock*. Thus, $(rock, rock)$ now lies on the floor, having been rejected as a solution because it is not a Nash equilibrium. We turn to $(paper, rock)$, and while Bart's strategy of *paper* is a best reply, Lisa's is not, since *scissors* yields a higher payoff than *rock* when Bart is choosing *paper*. Hence, $(paper, rock)$ joins $(rock, rock)$ on the floor. We merrily continue with our work, and before we know it, the floor is a mess as everything lies on it! None of the nine strategy pairs is a Nash equilibrium.

You could check each of these nine strategy pairs and convince yourself that that claim is true, but let me offer a useful shortcut for two-player games. Suppose we ask whether Lisa's choice of some strategy, call it y , is part of a Nash equilibrium. (I say “part of,” since, to even have a chance at being a Nash equilibrium, there must also be a strategy for Bart.) For y to be part of a Nash equilibrium, Bart must choose a strategy (call it c) that is a best reply to Lisa's

FIGURE 4.8 Rock–Paper–Scissors

		Lisa		
		Rock	Paper	Scissors
Bart	Rock	0,0	-1,1	1,-1
	Paper	1,-1	0,0	-1,1
	Scissors	-1,1	1,-1	0,0

choosing y . Choosing such a strategy ensures that Bart's Nash equilibrium condition is satisfied. To ensure that Lisa is also acting optimally, we then need to derive her best reply to Bart's choosing c (which, recall, is his best reply to Lisa's choosing y). Now suppose that Lisa's best reply to c is actually y , which is the strategy we started with for Lisa. Then we have shown that y is indeed part of a Nash equilibrium and the equilibrium is, in fact, (c, y) . However, if Lisa's best reply to Bart's choosing c is not y , then we conclude that y is not part of *any* Nash equilibrium. In that case, in one fell swoop we've eliminated all strategy profiles involving Lisa's choosing y . Putting it pictorially, this is what we need to happen for Lisa's playing y to be part of a Nash equilibrium:

Lisa plays $y \rightarrow$ Bart's best reply to y is $c \rightarrow$ Lisa's best reply to c is y .

To put this algorithm into action, let us ask whether Lisa's choosing *rock* is part of a Nash equilibrium. If Bart thinks that Lisa is going to choose *rock*, then he wants to choose *paper*. Now, if Bart chooses *paper*, then Lisa wants to choose *scissors*. Since this option is different from what we initially assumed that Lisa would choose, which was *rock*, we conclude that there is no Nash equilibrium in which Lisa chooses *rock*. Hence, none of the strategy profiles in which Lisa chooses *rock*—namely, $(rock, rock)$, $(paper, rock)$, and $(scissors, rock)$ —are Nash equilibria. Now let's do the same trick on the strategy *paper* for Lisa in order to determine whether her choosing *paper* is part of an equilibrium. If Lisa chooses *paper*, Bart's best reply is *scissors*, and Lisa's best reply to Bart's selection of *scissors* is *rock*, not *paper*. Hence, Lisa's using *paper* is not part of any Nash equilibrium, so we can eliminate $(rock, paper)$, $(paper, paper)$, and $(scissors, paper)$ as Nash equilibria. Finally, using the same method, we can show that Lisa's choosing *scissors* is not part of any Nash equilibrium. In this manner, we've proven that there is no Nash equilibrium for the game of Rock–Paper–Scissors.

Rock–Paper–Scissors is an example of an *outguessing game*. In an **outguessing game**, maximizing your payoff requires that you outguess the other player (or players). That is, you want to do what they don't expect. If the other player thinks that you're going to play strategy x , and she responds by playing b , then you don't want to play x in response to her playing b ; instead, you want to respond with something else. For example, if Lisa thinks that Bart is going to play *rock*, then she'll play *paper*, in which case Bart doesn't want to do as Lisa expects. Instead, he should play *scissors*, not *rock*. (Unfortunately, Bart isn't that smart, but you have to blame Matt Groening for that, not game theory.)

As it turns out, outguessing games arise in many situations. Sports and military conflicts are two prominent examples; we'll investigate them quite extensively in Chapter 7. However, be forewarned: if you intend to enter the USA Rock–Paper–Scissors League (yes, there is such a thing), game theory really can't help you design a winning strategy.

That Rock–Paper–Scissors is not just a kid's game was recently demonstrated by the two leading auction houses: Christie's and Sotheby's. The owner of an art collection worth in excess of \$20 million decided to determine which auction house would sell his collection—and, consequently, earn millions of dollars in commissions—on the basis of the outcome of a round of Rock–Paper–Scissors.⁴ Rather than play the game in the traditional way, however, with physical hand movements, an executive for Christie's and an executive for Sotheby's each wrote down one of the three strategies on a piece of paper. Christie's won, choosing *rock* to beat Sotheby's *scissors*.

CONFLICT AND MUTUAL INTEREST IN GAMES

Rock–Paper–Scissors is a game of *pure conflict*. What do I mean by that? Well, take note of an interesting property of the payoff matrix in Figure 4.8: Players' payoffs always sum to the same number (which happens to be zero). For example, if both Bart and Lisa choose *rock*, then each gets zero, so the sum of their payoffs is zero. If Bart chooses *paper* and Lisa chooses *rock*, then Bart gets 1 and Lisa gets -1 , which once again sums to zero. For every strategy pair, the sum of their payoffs is zero. This type of game is known as a **constant-sum game**, because the payoffs always sum to the same number. When that number happens to be zero, the game is called a **zero-sum game**.

So think about what this implies. Since payoffs must sum to the same number, if some strategy pair results in a higher payoff for Bart, then it must result in a lower payoff for Lisa. Thus, what makes Bart better off has to make Lisa worse off, and analogously, what makes Lisa better off has to make Bart worse off. It is in that sense that Rock–Paper–Scissors is a game of pure conflict. In fact, all constant-sum games have this property.

Contrast this game with driving conventions. Here we have the opposite of Rock–Paper–Scissors, in the sense that there is no conflict at all. A strategy pair that makes driver 1 better off—such as (*left, left*) compared with (*left, right*)—also makes driver 2 better off; they both get a payoff of 1 rather than 0. This is a game of **mutual interest**, as the rankings of strategy pairs by their payoffs coincides for the players.

Chicken and the telephone game lie between these two extremes. Those strategic settings do provide grounds for mutual interest. In Chicken, both players want to avoid (*hang tough, hang tough*); they both prefer (*swerve, hang tough*) and (*hang tough, swerve*). But there is also room for conflict, as they disagree as to how they rank (*swerve, hang tough*) and (*hang tough, swerve*); driver 1 prefers the latter and driver 2 prefers the former. Similarly, with the telephone game, both Colleen and Winnie agree that one of them calling is preferable to either both of them waiting or both calling, but they disagree as to who should call. Colleen prefers that it be Winnie, while Winnie prefers it to be Colleen. They share a common interest in coordinating on exactly one person calling, but their interests depart—they are in conflict—when it comes to who that person should be.

Knowing whether players' interests are entirely in conflict, partially in conflict, or entirely in common can provide some insight into which strategy profiles are Nash equilibria. So, when you come to a game, think about the interests of the players before launching into a robotic search for solutions. Your ruminations may offer some valuable shortcuts.

4.1 CHECK YOUR UNDERSTANDING

For the game in **FIGURE 4.9**, find all Nash equilibria.

FIGURE 4.9

		Player 2			
		w	x	y	z
Player 1	a	0,1	0,1	1,0	3,2
	b	1,2	2,2	4,0	0,2
	c	2,1	0,1	1,2	1,0
	d	3,0	1,0	1,1	3,1

4.3 The Best-Reply Method

AS THE CELEBRATED TV chef Emeril Lagasse would say, "Let's kick it up a notch!" by adding a third player to the mix. But before doing so, I'll share a useful shortcut with you for deriving Nash equilibria.

FIGURE 4.10 Jack and Diane's Strategies

		Diane		
		x	y	z
Jack	a	1,1	2,1	2,0
	b	2,3	0,2	2,1
	c	2,1	1,2	3,0

Recall that a player's best reply is a strategy that maximizes his payoff, given the strategies used by the other players. We can then think of a Nash equilibrium as a strategy profile in which each player's strategy is a best reply to the strategies the other players are using. Stemming from this perspective, the *best-reply method* offers a way of finding all of the Nash equilibria. Rather than describe it in the abstract, let's walk through the method for the two-player game in **FIGURE 4.10**.

For each strategy of Diane, we want to find Jack's best replies. If Diane uses x , then Jack has two best replies— b and c —each of which gives a payoff of 2 that exceeds the payoff of 1 from the other possible strategy, a . If Diane uses y , then Jack has a unique best reply of a . And if Diane uses z , then c is Jack's only best reply. To keep track of these best replies, circle those of Jack's payoffs associated with his best replies, as shown in **FIGURE 4.11**.

Next, perform the same exercise on Diane by finding her best replies in response to each of Jack's strategies. If Jack uses a , then both x and y are Diane's best replies. If Jack uses b , then Diane's best reply is x . Finally, if Jack uses c , then y is Diane's best reply. Circling the payoffs for Diane's best replies, we now have **FIGURE 4.12**.

FIGURE 4.11 Jack's Best Replies (Circled)

		Diane		
		x	y	z
Jack	a	1,1	②,1	2,0
	b	②,3	0,2	2,1
	c	②,1	1,2	③,0

FIGURE 4.12 Diane's and Jack's Best Replies (Circled). Two Strategy Pairs Are Nash Equilibria: (b,x) and (a,y)

		Diane		
		x	y	z
Jack	a	1,①	②,①	2,0
	b	②,③	0,2	2,1
	c	②,1	1,②	③,0

Since a Nash equilibrium is a strategy pair in which each player's strategy is a best reply, we can identify Nash equilibria in Figure 4.12 as those strategy pairs in which *both* payoffs in a cell are circled. Thus, (b,x) and (a,y) are Nash equilibria. We have just used the best-reply method to derive all Nash equilibria.

Before we explore how the best-reply method is used in three-player games, let's deploy it in Rock-Paper-Scissors. Marking each of Lisa's and Bart's best replies, we have **FIGURE 4.13**. For example, if Lisa chooses *rock*, then Bart's best reply is *paper*, so we circle Bart's payoff of 1 earned from the strategy pair (*paper*, *rock*). Note that no cell has two circled payoffs, indicating that there is no Nash equilibrium; this is the same result we derived earlier.

FIGURE 4.13 Best Replies (Circled) for Bart and Lisa's Game of Rock-Paper-Scissors. There Is No Nash Equilibrium

		Lisa		
		Rock	Paper	Scissors
Bart	Rock	0,0	-1,①	①,-1
	Paper	①,-1	0,0	-1,①
	Scissors	-1,①	①,-1	0,0

4.4 Three-Player Games

► SITUATION: AMERICAN IDOL FANDOM

Alicia, Kaitlyn, and Lauren are ecstatic. They've just landed tickets to attend this week's segment of *American Idol*. The three teens have the same favorite among the nine contestants that remain: Ace Young. They're determined to take this opportunity to make a statement. While IMing, they come up with a plan to wear T-shirts that spell out "ACE" in large letters. Lauren is to wear a T-shirt with an big "A," Kaitlyn with a "C," and Alicia with an "E." If they pull this stunt off, who knows—they might end up on national television! OMG!

While they all like this idea, each is tempted to wear instead an attractive new top just purchased from their latest shopping expedition at Bebe. It's now an hour before they have to leave to meet at the studio, and each is at home trying to decide between the Bebe top and the lettered T-shirt. What should each wear?

In specifying the strategic form of this game, we assign a payoff of 2 if they all wear their lettered T-shirts (and presumably remember to sit in the right sequence). This payoff is higher than the one they get from wearing the Bebe top, which is 1. Finally, wearing a lettered T-shirt when one or both of the other girls do not yields a payoff of 0, as the wearer realizes the worst of all worlds: she doesn't look alluring and they don't spell out ACE.

The strategic form is shown in **FIGURE 4.14**.^{*} Lauren's choice is represented as selecting a row—either wearing the T-shirt with the letter "A" or her Bebe top—while Kaitlyn chooses a column and Alicia chooses a matrix.

FIGURE 4.14 American Idol Fandom

		Alicia chooses <i>E</i>		Alicia chooses <i>Bebe</i>	
		Kaitlyn		Kaitlyn	
Lauren	<i>A</i>	2,2,2	0,1,0	0,0,1	0,1,1
	<i>Bebe</i>	1,0,0	1,1,0	1,0,1	1,1,1

Using the best-reply method to solve this game, consider the situation faced by Lauren. If Alicia wears her T-shirt with *E* and Kaitlyn wears hers with *C*, then Lauren's best reply is to do her part and wear the T-shirt with *A*. So, we circle Lauren's payoff of 2 in the cell associated with strategy profile (*A*, *C*, *E*), as shown in **FIGURE 4.15**. If, instead, Kaitlyn chooses *Bebe* and Alicia wears *E*, then Lauren's best reply is to wear her *Bebe* top and receive a payoff of 1, so we circle that payoff for Lauren. If Alicia wears her *Bebe* top and Kaitlyn wears *C*, then Lauren's best reply is again to wear her *Bebe* top, so we circle Lauren's payoff of 1. Finally, if both of the other two girls choose their *Bebe*

^{*}This game is a 21st-century teen girl version of the Stag Hunt game due to Jean-Jacques Rousseau in *On the Origins and Foundations of Inequality among Men* (1755). In that setting, hunters can work together to catch a stag (rather than spell out ACE) or hunt individually for hare (rather than wear a *Bebe* top).

choice. The three shareholders cast their votes simultaneously. There are 100 votes, allocated according to share ownership, so shareholder 1 has 25 votes, shareholder 2 has 35 votes, and shareholder 3 has 40 votes. Shareholders are required to allocate their votes as a bloc. For example, shareholder 1 has to cast all of her 25 votes for A , B , or C ; she cannot divvy them up among the projects. The strategy set for a player is then composed of A , B , and C . Plurality voting is used, which means that the alternative with the most votes is implemented.

To derive the payoff matrix, let us first determine how votes translate into a plurality winner. For example, if shareholders 1 and 2 vote for alternative B , then B is the winner, with either 60 votes (if shareholder 3 votes instead for A or C) or 100 votes (if 3 votes for B as well). **FIGURE 4.16** shows the plurality winner for each of the 27 different ways in which the three players can vote.

The next step is to substitute the associated payoff vector for each alternative in a cell in Figure 4.16. For example, if B is the winner, then shareholder 1's payoff is 1 (since B is his second choice), shareholder 2's payoff is 2 (since B is his first choice), and shareholder 3's payoff is 1 (since B is his second choice). Substitution, then, gives us **FIGURE 4.17**.

In making statements about how these shareholders might vote, a natural possibility to consider is what political scientists call *sincere voting*. The term

FIGURE 4.16 Plurality Winners

3 votes for A
2

	A	B	C
A	A	A	A
B	A	B	A
C	A	A	C

3 votes for B
2

	A	B	C
A	A	B	B
B	B	B	B
C	B	B	C

3 votes for C
2

	A	B	C
A	A	C	C
B	C	B	C
C	C	C	C

FIGURE 4.17 Strategic Form of the Voting Game

3 votes for A
2

	A	B	C
A	2,0,0	2,0,0	2,0,0
B	2,0,0	1,2,1	2,0,0
C	2,0,0	2,0,0	0,1,2

3 votes for B
2

	A	B	C
A	2,0,0	1,2,1	1,2,1
B	1,2,1	1,2,1	1,2,1
C	1,2,1	1,2,1	0,1,2

3 votes for C
2

	A	B	C
A	2,0,0	0,1,2	0,1,2
B	0,1,2	1,2,1	0,1,2
C	0,1,2	0,1,2	0,1,2

is used when a voter casts his vote for his first choice. In this case, it would mean that shareholder 1 casts her 25 votes for *A*, shareholder 2 casts his 35 votes for *B*, and shareholder 3 casts her 40 votes for *C*. As a result, choice *C* would be approved, since it received the most votes. But is sincere voting a Nash equilibrium? Is it optimal for shareholders to vote sincerely? Actually, no. Note that shareholder 1 prefers choice *B* over *C*. Given that shareholders 2 and 3 are voting sincerely, shareholder 1 can instead engage in (shall we call it) *devious voting* and vote for choice *B* rather than *A*. Doing so means that *B* ends up with 60 votes—being supported by both shareholders 1 and 2—and thus is approved. Shifting her votes from her most preferred alternative, *A*, to her next most preferred alternative, *B*, raises shareholder 1's payoff from 0 to 1. Hence, sincere voting is not a Nash equilibrium for this game.

Although it can be shown that it is always optimal to vote sincerely when there are only two alternatives on the ballot, it can be preferable to vote for something other than the most preferred option when there are three or more options, as we just observed. The intuition behind this assertion is that the most preferred option may not be viable—that is, it won't win, regardless of how you vote. In the event that your first choice can't win, it's reasonable to start think-

ing about which remaining choices could prevail, depending on various voting scenarios, and, among those choices, vote for the one that is most preferred. In the case we have just examined, with shareholders 2 and 3 voting for *B* and *C*, respectively, shareholder 1 can cause *B* to win (by casting her votes for *B*) or cause *C* to win (by casting her votes for either *A* or *C*). The issue, then, is whether she prefers *B* or *C*. Since she prefers *B*, she ought to use her votes strategically to make that option the winner.

Having ascertained that sincere voting does not produce a Nash equilibrium, let's see if the best-reply method can derive a strategy profile that *is* a Nash equilibrium. Start with shareholder 1. If shareholders 2 and 3 vote for *A*, then shareholder 1's payoff is 2, whether she votes for *A*, *B*, or *C*. (This statement makes sense, since alternative *A* receives the most votes, regardless of how shareholder 1 votes.) Thus, all three strategies for shareholder 1 are best replies, and in **FIGURE 4.18** we've circled her payoff of 2 in the column associated with shareholder 2's choosing *A* and the matrix associated with shareholder 3's choosing *A*. If shareholder 2 votes for *B* and shareholder 3 votes for *A*, then shareholder 1's best reply is to vote for *A* or *C* (thereby ensuring that *A* wins); the associated payoff of 2 is then circled. If shareholder 2 votes for *C* and shareholder 3 votes for *A*, then, again, shareholder 1's best replies are *A* and *C*. Continuing in this manner for shareholder 1 and then doing the same for shareholders 2 and 3, we get **Figure 4.18**.

Now look for all strategy profiles in which all three payoffs are circled. Such a strategy profile is one in which each player's strategy is a best reply and thus each player is doing the best he or she can, given what the others players

FIGURE 4.18 Best-Reply Method Applied to the Voting Game. There Are Five Nash Equilibria

		3 votes for A		
		2		
		A	B	C
1	A	(2,0,0)	(2,0,0)	(2,0,0)
	B	(2,0,0)	1,(2,1)	(2,0,0)
	C	(2,0,0)	(2,0,0)	0,(1,2)

		3 votes for B		
		2		
		A	B	C
1	A	(2,0,0)	(1,2,1)	(1,2,1)
	B	1,(2,1)	(1,2,1)	(1,2,1)
	C	1,(2,1)	(1,2,1)	0,1,2)

		3 votes for C		
		2		
		A	B	C
1	A	(2,0,0)	0,(1,2)	(0,1,2)
	B	0,1,2)	(1,2,1)	(0,1,2)
	C	0,(1,2)	0,(1,2)	(0,1,2)

there are other votes by shareholders 1 and 3 (e.g., when one of them votes for A and the other for B) for which shareholder 2 does strictly better by voting for B rather than A .

We then find that a player using a weakly dominated strategy is not ruled out by Nash equilibrium. Though voting for B always generates at least as high a payoff for shareholder 2 as does voting for A (and, in some cases, a strictly higher payoff), as long as A gives the same payoff that voting for B does for the strategies that shareholders 1 and 3 are actually using, then A is a best reply and thereby consistent with Nash equilibrium.

INSIGHT A Nash equilibrium does not preclude a player's using a weakly dominated strategy.

► SITUATION: PROMOTION AND SABOTAGE

Suppose you are engaged in a contest in which the person with the highest performance wins a prize. Currently, you're in second place. What can you do to improve your chances of winning? One thought is to work hard to improve your performance. But what might prove more effective is engaging in a "dirty tricks" campaign to degrade the performance of the current front-runner. The goal is to end up on top, and that can be done either by clawing your way up or by dragging those ahead of you down.

Such destructive forms of competition arise regularly in the political arena. The next time the U.S. presidential primaries roll around, pay attention to the campaigning. Candidates who are behind will talk about not only what a good choice they are for President, but also what a bad choice the front-runner is. They generally don't waste their time denigrating the other candidates—just the one who is currently on top and thus is the "one to beat." It has been suggested that sabotage by weaker competitors has arisen as well in nondemocratic governments. For example, although Zhao Ziyang appeared destined to become the leader of the Chinese Communist Party after Deng Xiao-Ping died in 1997, two more minor figures—Jiang Zemin and Li Peng—took control instead. Sabotage may have been at work.

To explore when and how a front-runner can be dethroned through dirty tricks, consider a setting in which three players are competing for a promotion.⁵ Whoever has the highest performance is promoted. Each contestant has one unit of effort that she can allocate in three possible ways: she can use it to enhance her own performance (which we'll refer to as a "positive" effort) or to denigrate one of the two competing players (which we'll refer to as a "negative" effort).

Before the competition begins, player i 's performance equals v_i . If a player exerts a positive effort, then she adds 1 to her performance. If exactly one player exerts a negative effort against player i , then player i 's performance is reduced by 1. If both players go negative against her, then player i 's performance is reduced by 4. Hence, the marginal impact of a second person's being negative is more detrimental than the impact of one person's being negative. This idea seems plausible, since one person making negative remarks may be dismissed as a fabrication, but two people saying the same thing could be perceived as credible.

How effort affects performance is summarized in **TABLE 4.4**. For example, if player i exerts a positive effort and the other two players exert a negative effort

yields a strictly higher payoff, player 2 is satisfied with exerting a positive effort.

- **Player 3:** The situation of player 3 is identical to that of player 2. They face the same payoffs and are choosing the same strategy. Thus, if going positive is optimal for player 2, then it is optimal for player 3.
- In sum, all three players choosing a positive effort is a Nash equilibrium and results in the front-runner gaining the promotion.

In now considering a strategy profile in which some negative effort is exerted, let's think about the incentives of players and what might be a natural strategy profile. It probably doesn't make much sense for player 2 to think about denigrating player 3, because the "person to beat" is player 1, as she is in the lead at the start of the competition. An analogous argument suggests that player 3 should do the same. Player 1 ought to focus on improving her own performance, since she is in the lead and the key to winning is maintaining that lead.

Accordingly, let us consider the strategy profile in which player 1 promotes herself, while players 2 and 3 denigrate player 1. The resulting performance is -1 for player 1 (because her performance, which started at 2, is increased by 1 due to her positive effort and lowered by 4 due to the negative effort of the other two players) and 0 for players 2 and 3 (since no effort—positive or negative—is directed at them, so that their performance remains at its initial level). Because players 2 and 3 are tied for the highest performance, the payoffs are 0 for player 1 and $\frac{1}{2}$ each for players 2 and 3. Now let's see whether we have a Nash equilibrium:

- **Player 1:** Unfortunately for player 1, there's not much she can do about her situation. If she exerts a negative effort against player 2, then she lowers 2's performance to -1 and her own to -2 . Player 3's performance of 0 results in her own promotion, so player 1 still loses out. An analogous argument shows that player 1 loses if she engages instead in a negative effort targeted at player 3: now player 2 is the one who wins the promotion. Thus, there is no better strategy for player 1 than to exert a positive effort.
- **Player 2:** If, instead of denigrating player 1, player 2 goes negative against player 3, then player 1's performance is raised from -1 to 2, player 2's performance remains at 0, and player 3's performance is lowered from 0 to -1 . Since player 1 now wins, player 2's payoff is lowered from $\frac{1}{2}$ to 0, so player 2's being negative about player 1 is preferred to player 2's being negative about player 3. What about player 2's being positive? This does raise his performance to 1, so that he now outperforms player 3 (who still has a performance of 0), but it has also raised player 1's performance to 2, since only one person is being negative against her. Since player 1 has the highest performance, player 2's payoff is again 0. Thus, player 2's strategy of being negative against player 1 is strictly preferred to either player 2's being negative against player 3 or player 2's being positive.
- **Player 3:** By an argument analogous to that used for player 2, player 3's strategy of being negative against player 1 is optimal.
- In sum, player 1's going positive and players 2 and 3 denigrating player 1 is a Nash equilibrium. Doing so sufficiently lowers the performance of

player 1 (Zhao Ziyang?) such that the promotion goes to either player 2 (Jiang Zemin?) or player 3 (Li Peng?). The front-runner loses. As Wayne Campbell, of *Wayne's World*, would say, "Promotion . . . denied!"

The promotion game, then, has multiple Nash equilibria (in fact, there are many more than we've described), which can have very different implications. One equilibrium has all players working hard to enhance their performance, and the adage "Let the best person win" prevails. But there is a darker solution in which the weaker players gang up against the favorite and succeed in knocking her out of the competition. The promotion then goes to one of those weaker players. Perhaps the more appropriate adage in that case is the one attributed to baseball player and manager Leo Durocher: "Nice guys finish last".

4.2 CHECK YOUR UNDERSTANDING

For the game in **FIGURE 4.19**, find all Nash equilibria.

FIGURE 4.19

		Player 3: I					Player 3: II		
		Player 2					Player 2		
		x	y	z			x	y	z
Player 1	a	2,1,2	0,0,2	1,2,3	Player 1	a	2,1,3	1,0,3	1,0,4
	b	0,3,1	2,2,4	3,1,0		b	1,2,1	3,3,3	1,1,1
	c	1,1,1	3,2,1	2,2,2		c	1,2,1	1,0,0	2,1,2

4.5 Foundations of Nash Equilibrium

THUS FAR, WE'VE OFFERED two approaches to solving a game: iterative deletion of strictly dominated strategies and Nash equilibrium. It is natural to wonder how they are related, so we'll address this issue next. Then there is the matter of how a strategy is interpreted in the context of Nash equilibrium. As it turns out, a strategy plays double duty.

4.5.1 Relationship to Rationality Is Common Knowledge

To explore the relationship between those strategies which survive the iterative deletion of strictly dominated strategies (IDSDS) and Nash equilibria, let's start with an example. Consider a Nash equilibrium for a three-player game in which player 1 uses strategy x , player 2 uses strategy y , and player 3 uses strategy z . Do these strategies survive the IDSDS? It's pretty easy to argue that none are eliminated on the first round: Since x is a best reply against player 2's using y and player 3's using z , x is most definitely not strictly dominated. Analogously, since y is a best reply for player 2 when player 1 uses x and player 3 uses z , y is not strictly dominated. Finally, since z is a best reply for player 3 when players 1 and 2 use x and y , respectively, z is not strictly dominated. Thus, x , y , and z are not eliminated in the first round of the IDSDS.

What will happen in the second round? Although some of player 2's and player 3's strategies may have been eliminated in the first round, y and z were not, and that ensures that x is not strictly dominated. The same argument explains why y is still not strictly dominated for player 2 in the second round and why z is still not strictly dominated for player 3. Thus, x , y , and z survive two rounds. Like the Energizer Bunny, this argument keeps going and going . . . it works for every round! Thus, if (x, y, z) is a Nash equilibrium, then those strategies survive the IDSDS. Although we have demonstrated this property for a three-player game, the argument is general and applies to all games.

While every Nash equilibrium is consistent with IDSDS, can a strategy survive the IDSDS, but *not* be part of a Nash equilibrium? Absolutely, and in fact, this chapter is loaded with examples. In the *American Idol* fandom game, all of the strategies survive the IDSDS, since none are strictly dominated. Thus, the IDSDS says that any of the eight feasible strategy profiles could occur. In contrast, only two strategy profiles— (A, C, E) and $(Bebe, Bebe, Bebe)$ (try saying that real fast!)—are Nash equilibria. Another example is Rock-Paper-Scissors, in which all strategy profiles are consistent with IDSDS, but *none* are Nash equilibria. Nash equilibrium is a more stringent criterion than IDSDS, since fewer strategy profiles satisfy the conditions of Nash equilibrium.

INSIGHT All Nash equilibria satisfy the iterative deletion of strictly dominated strategies and thereby are consistent with rationality's being common knowledge. However, a strategy profile that survives the IDSDS need not be a Nash equilibrium.

FIGURE 4.20 Relationship Between Nash Equilibria and the Strategies That Survive the IDSDS

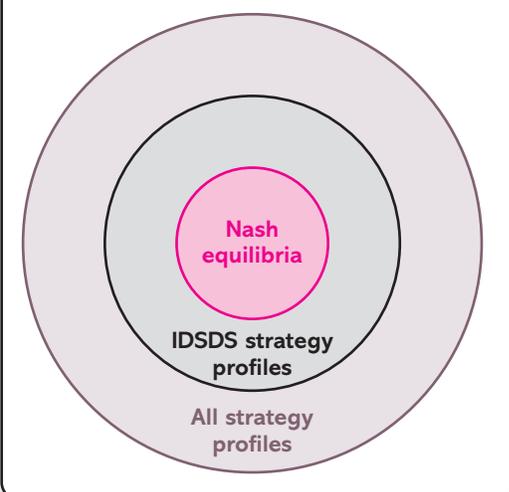


FIGURE 4.20 depicts how Nash equilibria are a subset of the strategy profiles that survive the IDSDS, which are themselves a subset of all strategy profiles. However, for any particular game, these sets could coincide, so that, for example, the set of Nash equilibria might be the same as those strategy profiles which survive the IDSDS, or the strategies that survive the IDSDS might coincide with the set of all strategy profiles.

4.5.2 The Definition of a Strategy, Revisited

To better understand the role of a strategy in the context of Nash equilibrium, think about specifying both a strategy for player i —denoted s_i and intended to be his decision rule—and a *conjecture* that player i holds regarding the strategy selected by player j —denoted $s_j(i)$ —which represents what i believes that j is going to play. A strategy profile (s'_1, \dots, s'_n) is then a Nash equilibrium if, for all i ,

1. s'_i maximizes player i 's payoff, given that he believes that player j will use $s_j(i)$, for all $j \neq i$.

2. $s_j(i) = s'_j$, for all $j \neq i$.

s'_i is then playing a dual role in a Nash equilibrium. As specified in condition 1, it is player i 's decision rule. In addition, as described in condition 2, s'_i is player j 's (accurate) conjecture as to what player i will do.

Recall from Section 2.3 that we required that a strategy specify what a player should do at every possible information set; that is, a strategy must specify behavior even at an information set that cannot be reached, given the prescribed behavior for some preceding information set. For example, in the kidnapping game, the kidnapper's strategy had to specify whether to release or kill the victim, even if at the initial node that strategy prescribed that he not perform the kidnapping. A strategy must meet this requirement because of the dual role of an equilibrium strategy. A player will have a conjecture as to how another player is going to behave, even if that player did not behave as predicted. For example, the victim's kin will have a conjecture as to whether the kidnapper will release or kill the victim, even if the kin originally predicted that the kidnapper would not perform the kidnapping. Just because a player did not behave as you expected doesn't mean that you don't have beliefs as to what will happen in the future.

At a Nash equilibrium, a strategy has two roles—decision rule and conjecture—in which case it's important that the strategy be fully specified; it must specify behavior at every information set for a player. A Nash equilibrium strategy both *prescribes*—being a player's decision rule—and *describes*—being another player's conjecture about that player's decision rule.

Summary

A rational player chooses a strategy that maximizes her payoff, given her beliefs about what other players are doing. Such an optimal strategy is referred to as a **best reply** to the conjectured strategies of the other players. If we furthermore suppose that these conjectures are accurate—that each player is correctly anticipating the strategy choices of the other players—then we have a **Nash equilibrium**. The appeal of Nash equilibrium is that it identifies a point of “mutual contentment” for all players. Each player is choosing a strategy that is best, given the strategies being chosen by the other players.

In many games, the iterative deletion of strictly dominated strategies (IDSDS) has no traction, because few, if any, strategies are strictly dominated. Nash equilibrium is a more selective criterion; thus, some games might have only a few Nash equilibria while having many more strategy profiles that survive the IDSDS. Generally, Nash equilibrium is a more useful solution concept, for that very reason. Nevertheless, as we found out by way of example, a game can have many Nash equilibria, a unique Nash equilibrium, or none at all.

In deriving the Nash equilibria for a game, one can approach the problem algorithmically, but also intuitively. The **best-reply method** was put forth as a procedure for deriving Nash equilibria, even though it can be cumbersome when players have many strategies to choose from. Intuition about the players' incentives can be useful in narrowing down the set of likely candidates for Nash equilibrium.

Games can range from **pure conflict** to ones where players have a **mutual interest**. **Constant-sum games** involve pure conflict, because something that makes one player better off must make other players worse off. An example is the children's game Rock–Paper–Scissors, which is also an example of an **out-guessing game** whereby each player is trying to do what the other players don't expect. At the other end of the spectrum are games in which the interests of the players coincide perfectly, so that what makes one player better off makes the others better off as well. This property describes driving conventions,

FIGURE PR 4.2 Modified Driving Conventions Game

	<i>Drive left</i>	<i>Drive right</i>	<i>Zigzag</i>
<i>Drive left</i>	1,1	-1,-1	0,0
<i>Drive right</i>	-1,-1	1,1	0,0
<i>Zigzag</i>	0,0	0,0	0,0

3. Return to the team project game in Chapter 3, and suppose that a frat boy is partnered with a sorority girl. The payoff matrix is shown in **FIGURE PR4.3**. Find all Nash equilibria.

FIGURE PR4.3 Team Project

		Sorority girl		
		<i>Low</i>	<i>Moderate</i>	<i>High</i>
Frat boy	<i>Low</i>	0,0	2,1	6,2
	<i>Moderate</i>	1,2	4,4	5,3
	<i>High</i>	2,6	3,5	3,4

4. Consider the two-player game illustrated in **FIGURE PR4.4**.
- For each player, derive those strategies which survive the iterative deletion of strictly dominated strategies.
 - Derive all strategy pairs that are Nash equilibria.

FIGURE PR4.4

		Player 2		
		<i>x</i>	<i>y</i>	<i>z</i>
Player 1	<i>a</i>	4,0	2,1	3,2
	<i>b</i>	2,2	3,4	0,1
	<i>c</i>	2,3	1,2	0,3

5. Consider the two-player game depicted in **FIGURE PR4.5**.
- Derive those strategies which survive the iterative deletion of strictly dominated strategies.
 - Derive all strategy pairs that are Nash equilibria.

FIGURE PR4.5

		Player 2		
		<i>x</i>	<i>y</i>	<i>z</i>
Player 1	<i>a</i>	1,2	1,2	0,3
	<i>b</i>	4,0	1,3	0,2
	<i>c</i>	3,1	2,1	1,2
	<i>d</i>	0,2	0,1	2,4

6. Return to the “white flight” game in Chapter 3. Now suppose that four of the eight homes are owned by one landlord, Donald Trump, and the other four are owned by a second landlord, John Jacob Astor. A strategy is the number of black families to whom to rent. Construct the payoff matrix and find the set of Nash equilibria. (Although you’re surely familiar with Donald Trump, John Jacob Astor has the noteworthy property of possibly being the first millionaire in U.S. history. Centuries before The Donald arrived on the real-estate scene in New York, Astor was wealthy beyond belief due to his New York City landholdings.)
7. Return to the kidnapping game, whose strategic form is shown in **FIGURE PR4.7**. Find all of the Nash equilibria.

FIGURE PR4.7 Kidnapping

		Vivica (kin of victim)	
		Pay ransom	Do not pay ransom
Guy (kidnapper)	Do not kidnap/Kill	3,5	3,5
	Do not kidnap/Release	3,5	3,5
	Kidnap/Kill	4,1	2,2
	Kidnap/Release	5,3	1,4

8. Queen Elizabeth has decided to auction off the crown jewels, and there are two bidders: Sultan Hassanal Bolkiah of Brunei and Sheikh Zayed Bin Sultan Al Nahyan of Abu Dhabi. The auction format is as follows: The Sultan and the Sheikh simultaneously submit a written bid. Exhibiting her well-known quiriness, the Queen specifies that the Sultan’s bid must be an odd number (in hundreds of millions of English pounds) between 1 and 9 (that is, it must be 1, 3, 5, 7, or 9) and the Sultan’s bid must be an even number between 2 and 10. The bidder who submits the highest bid wins the jewels and pays a price equal to his bid. (If you recall from Chapter 3, this is a first-price auction.) The winning bidder’s payoff equals his valuation of the item less the price he pays, while the losing bidder’s payoff is zero. Assume that the Sultan has a valuation of 8 (hundred million pounds) and the Sheikh has a valuation of 7.
 - a. In matrix form, write down the strategic form of this game.
 - b. Derive all Nash equilibria.
9. Find all of the Nash equilibria for the three-player game in **FIGURE PR4.9**.
10. Return to the game of promotion and sabotage in Section 4.4.
 - a. Determine whether the following strategy profile is a Nash equilibrium; (i) player 1 is negative against player 2, (ii) player 2 is negative against player 3, and (iii) player 3 is negative against player 1.
 - b. Find a Nash equilibrium in which player 2 wins the promotion with certainty (probability 1).
11. When there are multiple Nash equilibria, one approach to selecting among them is to eliminate all those equilibria which involve one or more players using a weakly dominated strategy. For the voting game in Figure 4.17, find all of the Nash equilibria that do not have players using a weakly dominated strategy.
12. Recall the example of Galileo Galilei and the Inquisition in Chapter 2. The strategic form of the game is reproduced in **FIGURE PR4.12**. Find all of the Nash equilibria.

FIGURE PR4.9

Player 3: A
Player 2

	<i>x</i>	<i>y</i>	<i>z</i>
Player 1 <i>a</i>	1,1,0	2,0,0	2,0,0
<i>b</i>	3,2,1	1,2,3	0,1,2
<i>c</i>	2,0,0	0,2,3	3,1,1

Player 3: B
Player 2

	<i>x</i>	<i>y</i>	<i>z</i>
Player 1 <i>a</i>	2,0,0	0,0,1	2,1,2
<i>b</i>	1,2,0	1,2,1	1,2,1
<i>c</i>	0,1,2	2,2,1	2,1,0

Player 3: C
Player 2

	<i>x</i>	<i>y</i>	<i>z</i>
Player 1 <i>a</i>	2,0,0	0,1,2	0,1,2
<i>b</i>	0,1,1	1,2,1	0,1,2
<i>c</i>	3,1,2	0,1,2	1,1,2

FIGURE PR4.12

Pope Urban VIII: *Refer*
Inquisitor

	<i>Torture</i>	<i>Do not torture</i>
Galileo <i>C/C</i>	3,4,5	3,4,5
<i>C/DNC</i>	3,4,5	3,4,5
<i>DNC/C</i>	1,5,4	4,2,2
<i>DNC/DNC</i>	2,1,1	4,2,2

Pope Urban VIII: *Do Not Refer*
Inquisitor

	<i>Torture</i>	<i>Do not torture</i>
Galileo <i>C/C</i>	5,3,3	5,3,3
<i>C/DNC</i>	5,3,3	5,3,3
<i>DNC/C</i>	5,3,3	5,3,3
<i>DNC/DNC</i>	5,3,3	5,3,3

13. Find all of the Nash equilibria for the three-player game shown in FIGURE PR4.13.

FIGURE PR4.13

		Player 3: A		
		Player 2		
		x	y	z
Player 1	a	2,0,4	1,1,1	1,2,3
	b	3,2,3	0,1,0	2,1,0
	c	1,0,2	0,0,3	3,1,1

		Player 3: B		
		Player 2		
		x	y	z
Player 1	a	2,0,3	4,1,2	1,1,2
	b	1,3,2	2,2,2	0,4,3
	c	0,0,0	3,0,3	2,1,0

4.6 Appendix: Nash Equilibrium

CONSIDER A GAME WITH n players: $1, 2, \dots, n$. Let S_i denote player i 's strategy set, and read $s'_i \in S_i$ as "strategy s'_i is a member of S_i ." Let S_{-i} be composed of all $(n - 1)$ -tuples of strategies for the $n - 1$ players other than player i , and let $V_i(s'_i, s'_{-i})$ be the payoff of player i when his strategy is s'_i and the other players use $s'_{-i} = (s'_1, \dots, s'_{i-1}, s'_{i+1}, \dots, s'_n)$. Then for all $i = 1, 2, \dots, n$, a strategy profile (s_1^*, \dots, s_n^*) is a Nash equilibrium if s_i^* maximizes player i 's payoff, given that the other players use strategies $(s_1^*, \dots, s_{i-1}^*, s_{i+1}^*, \dots, s_n^*)$. More formally, (s_1^*, \dots, s_n^*) is a Nash equilibrium if and only if for all $i = 1, 2, \dots, n$,

$$V_i(s_1^*, \dots, s_n^*) \geq V_i(s_1^*, \dots, s_{i-1}^*, s_i, s_{i+1}^*, \dots, s_n^*) \text{ for all } s_i \in S_i.$$

REFERENCES

1. "Disney Out-'Sharks' DreamWorks," by Marcus Errico (Aug. 18, 2003) <www.eonline.com>. *Sharkslayer* later had its title changed to *Shark Tale*.
2. A transcript of this conversation was reportedly released by the U.S. Chief of Naval Operations <www.unwind.com/jokes-funnies/militaryjokes/gamechicken.shtml>.
3. <www.brianlucas.ca/roadside/>.
4. Carol Vogel, "Rock, Paper, Payoff: Child's Play Wins Auction House an Art Sale," *New York Times*, Apr. 29, 2005.
5. This game is based on Kong-Pin Chen, "Sabotage in Promotion Tournaments," *Journal of Law, Economics, and Organization*, 19 (2003), 119–40. That paper also provides the anecdote regarding the Chinese Communist Party leadership.