

Stable Play: Nash Equilibria in Discrete Games with Two or Three Players

4

Nothing in the world is so powerful as an idea whose time has come.

—VICTOR HUGO

4.1 Defining Nash Equilibrium

DreamWorks initially threw down the gauntlet in the clash of the 'toon titans way back in June 2002, claiming a release date of November 5, 2004 for Sharkslayer. . . . The studio's staking out of the November date was seen as a slap at Disney, which has traditionally released its Pixar pictures that month. Disney . . . kicked up the brinkmanship factor, announcing that Sharkslayer or no, The Incredibles would also open on November 5. . . . DreamWorks [then] blinked [as it] moved the release date for its film . . . [to] October 1.¹

IN SPITE OF ITS REFERENCE to a nonlethal, passive fowl, Chicken is a dangerous game. In its classic form, it begins with two cars facing each other in duel-like fashion (and typically occupied by male teenagers in pursuit of testosterone-inspired adventures). As the cars come hurtling toward one another, each driver is frantically deciding whether to swerve to avoid a collision or to hang tough (hoping that the other will swerve). The goal is to avoid being the first to swerve, although if both hang tough, then the result is a mangled mess of metal and flesh. Chicken has been played in many contexts, including contests between movie executives (with release dates) and between the leaders of the United States and the former Soviet Union (with nuclear weapons). TABLE 4.1 lists a few other games of Chicken that have arisen in fact and fiction.

TABLE 4.1 Chicken in Action	
Mode	Description
Tractors	<i>Footloose</i> (1984, movie)
Bulldozers	Buster and Gob in <i>Arrested Development</i> (2004, TV)
Wheelchairs	Two old ladies with motorized wheelchairs in <i>Banzai</i> (2003, TV)
Snowmobiles	"[Two adult males] died in a head-on collision, earning a tie in the game of chicken they were playing with their snowmobiles." < www.seriouslyinternet.com/278.0.html >
Film release dates	Dreamworks and Disney-Pixar (2004)
Nuclear weapons	Cuban Missile Crisis (1963): "Since the nuclear stalemate became apparent, the Governments of East and West have adopted the policy which Mr. Dulles calls 'brinkmanship.' This is a policy adapted from a sport which, I am told, is practised by some youthful degenerates. This sport is called 'Chicken!' " (Bertrand Russell, <i>Common Sense and Nuclear Warfare</i> , 1959)

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, all strategies survive the iterative deletion of strictly dominated strategies (IDSDS, Section 3.4.2), in which case it won't help us solve this game. But don't abandon hope, because 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. To derive those beliefs, Chapter 3 used the assumption that rationality is common knowledge among the players. 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 the players to 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.

*You might be tempted to put large negative numbers for the strategy pair in which both participants choose *hang tough*, because this means certain injury and possible death. You can do so, but it'll make no difference to 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.

**While staying at the Three Village Inn in Stony Brook, New York for an annual game-theory conference, I was fortunate enough to have breakfast with John Nash. My impression of him was a gentle, thoughtful man who did not depart far from the usual misfitness of game theorists! *A Beautiful Mind* is a truly wonderful biography by Sylvia Nasar that portrays the brilliance of the man, the tragedy of mental illness, and the inspirational humanity of his wife.

➤ **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 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.

Each and every player acting in a manner to maximize his or her payoff, as described by Nash equilibrium, has the desirable property that it is stable in the sense that each player is content to do what she is doing, given what everyone else is 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 one player who will learn that his strategy is not the best one available, given what others are doing. He then will 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. Because *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	
		<i>Swerve</i>	<i>Hang tough</i>
Driver 1	<i>Swerve</i>	2,2	1,3
	<i>Hang tough</i>	3,1	0,0

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

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

- (*swerve, hang tough*). If driver 2 chooses *hang tough*, *swerve* is the best strategy for driver 1, because 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. Although Nash equilibrium tells us someone will swerve, it doesn't tell us which driver will swerve. This indeterminacy is not surprising, because there is nothing in the game to distinguish the two drivers. Any argument for driver 1's swerving (and driver 2's hanging tough) works equally well for driver 2's swerving (and driver 1's hanging tough).

Herman Kahn, who came to prominence as a military strategist at the RAND Corporation, described how one might act so as to credibly convey to the other driver that he'll have to be the one to swerve:

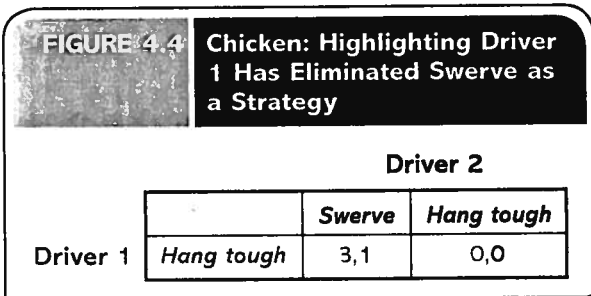
The "skilful" player may get into the car quite drunk, throwing whisky bottles out of the window to make it clear to everybody just how drunk he is. ... As soon as the car reaches high speed, he takes the steering wheel and throws it out of the window. If his opponent is watching, he has won. If his opponent is not watching, he has a problem. [On Escalation (1965), p.11]

On the subject of nuclear warfare, Kahn was one of the leading thinkers during the Cold War, which makes this quotation absolutely frightening! (It is also said that he was the inspiration for the title character of the 1964 dark comedy *Dr. Strangelove or: How I Learned to Stop Worrying and Love the Bomb*.)

What Kahn is saying is that perhaps the best way to play Chicken is to commit to not swerving by effectively eliminating *swerve* from your strategy set and, most important, 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 implemented 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."



riun
ciat
Our
emb
tion
Nas
garr
wha

4.

THE
cep
it is
can
devi
exp
the
con
acc
dec
mal
neo
we
in a
eral
this
of t

of v
equ
ide
eml
the
is r
Altl
libr

resj
rep
rep
pla
ers,
givi

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 we can use this concept 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 the 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 finding Nash equilibria. The process by which it is determined that a candidate strategy profile is or is not a Nash equilibrium can be summarized in the following mantra: *In a Nash equilibrium, unilateral deviation does not pay.* *Unilateral deviation* means that we go through the thought experiment in which one player chooses a strategy different from her strategy in the candidate strategy profile. When we perform this thought experiment, we consider each player in turn but always assume that the other players are acting according to the candidate strategy profile. The reason is that each player's decision-making process is presumed independent of other players' decision-making processes; thus, joint deviations, whereby two or more players simultaneously consider doing something different, are not considered. By *does not pay*, we mean that the alternative strategy in the thought experiment does not result in a higher payoff. For a strategy profile to be a Nash equilibrium, every unilateral deviation for a player must not make the player better off. Equipped with this mantra, you might find it useful to go back and think about our evaluation of the game of Chicken.

The plan is to warm up with a few simple games involving two players, each of whom has at most 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 them in Chapter 7.

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 that 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 not only the most widely examined game in game theory but has infiltrated our everyday lexicon. To take just a few examples from the popular press: "The Prisoner's Dilemma and the Financial Crisis" (*Washington Post*, October 29, 2009); "Lance Armstrong and the Prisoners' Dilemma of Doping in Professional Sports" (*Wired*, October 26, 2012); and "The Prisoner's Dilemma and Your Money" (*Christian Science Monitor*, June 5, 2013).

So what exactly is this Prisoners' Dilemma? Well, it all began when Albert Tucker described the following scenario before an audience of psychologists in 1950: 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, whereas 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 three 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 1 and 1 for cr (silence, testify) 2's payoff is 4. choose the sar profile—such other player.

INSIGHT

Second, *te* does, *testify* p it should be c rium must ha rium, the stra Because a do must be used

INSIGHT

Here is tl vidual intere although (*tes* jointly movi Therein lies 1

INSIGHT

SITUATION: A

All vehic the same —GENEV

We next look players have game that r which to dri (as in Engla just that we The gan verify that

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, whereas 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. Hence, if a symmetric strategy profile—such as (*testify*, *testify*)—is optimal for one player, it is also optimal for the other player.

◆ **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.

Second, *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.

◆ **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.

Here is the third and final point: The fact that all players act in their individual interests does not imply that they act in their collective interests. For although (*testify*, *testify*) is a Nash equilibrium, both players could do better by jointly moving to (*silence*, *silence*); each would raise his payoff from 2 to 3. Therein lies the dilemma that the prisoners face.

◆ **INSIGHT** Nash equilibrium ensures that each player is doing the best she can individually but does not ensure that the group of players are doing the best they can collectively.

► SITUATION: A COORDINATION GAME—DRIVING CONVENTIONS

All vehicular traffic proceeding in the same direction on any road shall keep to the same side of the road, which shall be uniform in each country for all roads.

—GENEVA CONVENTION ON ROAD TRAFFIC (1949)

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 on 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

FIGURE 4.6 Driving Conventions

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

driving on the left, and the other has both driving on the right. If Louise drives on the left, then Thelma's best reply is to drive on the left and receive a payoff of 1 rather than receive -1 from driving on the right. The same argument verifies that Louise's driving on the left is best. In fact, because 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 best reply is to also drive on the right. 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 on; instead, it tells us only that they will settle on some standard. History shows that societies do settle on 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. ◀◀◀

INSIGHT In a coordination game, players have both an individual and common interest to make the same choice.

Society is full of coordination games including the most ubiquitous one of language. It doesn't really matter what sounds we attach to some object but only that we all agree on the sound associated with an object. That different countries settle on different equilibria can be particularly troubling in the case of dialects, because the same sound can be present in both dialects but mean very different things, and thus create the opportunity for awkward situations. For example, "blow off" means "not to show up for a meeting" in U.S. English but to "break wind" in UK English. *Graft* refers to "political corruption" in the United States but "hard work" in the UK. (Of course, corruption can be hard work.) A *bogey* is an unidentified aircraft in the United States but dried nasal mucus in the UK. You would bring forth very different responses in the two countries if you announced, "Incoming bogeys." And they can be far worse than that! If you don't believe me, go to Wikipedia and put "different meanings in American and British English" in the search field.

▷ SITUATION

In the
were
on th
playe
Co
they'r
Winn
they l
ing to
versa
try to
don't
when
calls
game
To
reply
there
ting
reply
necti
Winn
what
librit
a sec
is no
to wa
it is l

For

HC

*Ans

*This
ting. T
go to l

► 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.

Colleen is chatting on the phone with Winnie when 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. (Assume we are in the stone age of technology when there was no call waiting feature.) Obviously, if neither calls back, then they don't reconnect either. The strategic form game is shown in FIGURE 4.7.*

FIGURE 4.7 Telephone

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

To find the Nash equilibria, let us begin by deriving each caller's best reply. If Colleen expects Winnie to call back, then Colleen should wait, thereby receiving a payoff of 3 rather than 0 from also calling (and both getting a busy signal). If Colleen expects Winnie to wait, then Colleen's best reply is to call and get a payoff of 2 rather than 0 from waiting (and not connecting). Because this game is symmetric, the same analysis applies to Winnie. Thus, each person's best reply is to make a different choice from what she expects the other person to do. It is then clear that one Nash equilibrium is for Colleen to call back and Winnie to wait for Colleen's call, and a second Nash equilibrium is for Winnie to call back and Colleen to wait. It is not a Nash equilibrium for both to call back, because each would do better to wait if she expects the other person to call, nor for both to wait, because it is better for someone to take the initiative and call.

⊖ 4.1 CHECK YOUR UNDERSTANDING

For the game in FIGURE 4.8, find all Nash equilibria.*

FIGURE 4.8

		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

*Answers to Check Your Understanding are in the back of the book.

*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, because scissors break when trying to cut rock. If one person chooses *rock* and the other *paper*, then *paper* wins, because paper can be wrapped around rock. And if one person chooses *paper* and the other *scissors*, then *scissors* wins, because 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 4.9. Contrary to Bart's belief, *rock* is not a dominant strategy. Although *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*), because 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 next to (*paper, rock*), and although Bart's strategy of *paper* is a best reply, Lisa's is not, because *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 because 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," because, if strategy *y* is to have even a

FIGURE 4.9 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

chance of
For *y* to be
that is a b
Bart's Nas
ing optim
(which, re
best reply
we have sl
is, in fact,
we conclu
swoop w
Putting it
part of a

Lisa

To pu
part of a
he wants
scissors. I
would ch
which Li
chooses r
equilibria
in which
(Try it!)
game of

Rock
guessing
other pl
the othe
by playi
instead,
thinks t
Bart doe
rock. (U
Groenin

As it
tary con
sively in
Rock-Pa
help you

That
two lead
worth ir
sell his
on the l
play the
tive for
three s

chance of 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 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*. Because 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. The same trick can be used to show that there is no Nash equilibrium in which Lisa chooses *paper* and no equilibrium in which she chooses *scissors*. (Try it!) In this manner, we can prove 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 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 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*. For a notable instance in which scissors beat paper, go to www.smittenbybritain.com/scissors-beat-paper

Mickey: All right, rock beats paper! (Mickey smacks Kramer's hand for losing).

Kramer: I thought paper covered rock.

Mickey: Nah, rock flies right through paper.

Kramer: What beats rock?

Mickey: (looks at hand) Nothing beats rock.

—FROM THE EPISODE "THE STAND-IN" OF *SEINFELD* ◀◀◀

4.2 CHECK YOUR UNDERSTANDING

Two competitive siblings—Juan and Maria—are deciding when to show up at their mom's house for Mother's Day. They are simultaneously choosing between times of 8:00 A.M., 9:00 A.M., 10:00 A.M. and 11:00 A.M. The payoffs to a sibling are shown in TABLE 4.2 and depend on what time he or she shows up and whether he or she shows up first, second, or at the same time. (Note that this is not a payoff matrix.) For example, if Juan shows up at 9:00 A.M. and Maria shows up at 10:00 A.M. then Juan's payoff is 8 (because he is first) and Maria's payoff is 4 (because she is second). Payoffs have the property that each would like to show up first but would prefer to show up second rather than show up at the same time. (They really do not like one another.) Furthermore, conditional on showing up first or at the same time or second, each prefers to show up later in the morning; note that payoffs are increasing as we move down a column that is associated with arriving later in the morning.

Time/Order of Arrival	First	Same Time	Second
8:00 A.M.	7	-3	-
9:00 A.M.	8	-2	3
10:00 A.M.	9	-1	4
11:00 A.M.	-	0	5

Using the method deployed for Rock–Paper–Scissors, show that there is no Nash equilibrium. (Because you'll want to start getting used to solving games without a payoff matrix before you, I'd recommend trying to answer this question without constructing the payoff matrix. However, the answer in the back of the book does include the payoff matrix if you choose to work with it.)*

*Answers to Check Your Understanding are in the back of the book.

► SITUATION: 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 4.9: Players' payoffs always sum to the same number (which happens to be zero). For

exam
of th
gets
pair,
sum
num
sum
S
ber, i
in a l
and ;
sense
sum
C
oppo
strat
(left,
0. Th
their
C
tegi
want
(han
to he
the l
Colle
of th
prefe
com
diver
K
flict,
are N
the p
tions

4
For
Dile
FIG
*Ans

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, as in Figure 4.9, 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 makes Lisa worse off, and analogously, what makes Lisa better off makes 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 the driving conventions game. 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**, because 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, because 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 diverge—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. ◀◀◀


4.3 CHECK YOUR UNDERSTANDING

For the game in FIGURE 4.10, find all values for x and y such that it is a Prisoners' Dilemma. Find all values for x and y such that it is a coordination game.*

FIGURE 4.10

		Barbara	
		Alpha	Beta
Kenneth	Alpha	5,5	y,x
	Beta	x,y	3,3

*Answers to Check Your Understanding are in the back of the book.



BRIAN A. JACKSON/SHUTTERSTOCK

Corporate Leniency Programs

Scott Hammond (Deputy Assistant Attorney General, Antitrust Division, U.S. Department of Justice):

"The Antitrust Division's Corporate Leniency Program has been the Division's most effective investigative tool. Cooperation from leniency applicants has cracked more cartels than all other tools at our disposal combined." [Reference: "Cracking Cartels With Leniency Programs," OECD Competition Committee, Paris, October 18, 2005.]

A cartel is a collection of firms in a market that, instead of competing for customers' business, coordinate in raising prices and profits to the detriment of consumers. Although such behavior is illegal in many countries, a challenge to enforcing the law is that colluding firms know well enough to keep their activity hidden by, for example, meeting clandestinely and avoiding any written documentation. Once cartel members are caught and convicted, penalties in the United States are severe, with government fines and customer damages from private litigation imposed on corporations as well as fines and imprisonment for involved individuals. However, levying those penalties requires first finding cartels and then amassing the evidence to convict them.

The most important innovation in recent times for uncovering and prosecuting cartels is the corporate leniency program. In the United States, the first member of an unlawful cartel to come forward and cooperate with the Department of Justice (DOJ) in convicting the other cartel members will be absolved of all government

penalties. Relying on the adage "There is no honor among thieves," a leniency program is intended to undermine the mutual interest of raising profits with the dangling of the carrot of no government penalties. That such a carrot can go to only one firm inserts conflict into the relations among cartel members. After an important revision in the leniency program in 1993, applications increased twentyfold and resulted in many convictions. Other jurisdictions quickly took notice, and now more than 50 countries and unions have leniency programs. During the week before Spain's leniency program became active, cartel members were literally lining up outside the door of the competition authority on Calle del Barquillo to make sure they were the first from their cartel!

How attractive it is to apply for leniency—and thus how effective such a program is in inducing cartel members to come forward—depends on how much leniency reduces the penalties received. Although leniency negates all government penalties, a company in the United States still must pay customer damages. If convicted in a private litigation suit, the guilty firm must pay triple the amount of damages that it inflicted on its customers (a penalty referred to as "treble damages"). Generally, these damages vastly exceed government fines. For example, in the fine arts auction house cartel, Christie's avoided the \$45 million fine but still had to pay \$256 million in damages.

If the amount of penalty relief from leniency is low (as when damages are large relative to fines), then the cartel members face a coordination game for which there are two Nash equilibria: all apply and no one applies. If a cartel member believes other cartel members will apply, then failing to apply ensures conviction and full penalties, whereas applying gives a firm a shot at leniency. But if no one is expected to apply, then—because leniency reduces penalties only modestly—it is better not to apply and hope no one is caught. The situation is then as depicted in FIGURE 4.11 (where payoffs can be thought of as the negative of the amount of expected monetary penalties). This is a coordination game—in which each firm wants to do what it expects the other to do—but where one equilibrium is superior (in contrast to the driving conventions game). Firms want to coordinate on the "no one applies" equilibrium, whereas the DOJ wants them to be at the equilibrium for which they apply.

FIGURE

Christie's

So
lenien
faced
game
ment
domin
way to
that a
is exa

4.3 THE

AS THE CE
by adding
with you

Recal
maximiz
other pla
rium as a
egy is a l
are usin
best-repl
Nash equ
let's walk
in FIGURE

For e
best repl
replies—
which ex
ble strat
unique b
Jack's on
replies, c
his best r

Next,
ing her be
gies. If Ja

FIGURE 4.11 Leniency as a Coordination Game

		Sotheby's	
		Apply	Do not apply
Christie's	Apply	-7, -7	-6, -8
	Do not apply	-8, -6	-3, -3

FIGURE 4.12 Leniency as a Prisoners' Dilemma

		Sotheby's	
		Apply	Do not apply
Christie's	Apply	-5, -5	-2, -8
	Do not apply	-8, -2	-3, -3

So how can the government ensure a race for leniency? By changing the strategic situation faced by cartel members from a coordination game into a Prisoners' Dilemma! What the government wants to do is make applying for leniency a dominant strategy, as depicted in FIGURE 4.12. One way to do so is to increase the amount of penalties that are avoided from receiving leniency, and this is exactly what Congress did in 2004 when it

passed (and the President signed into law) the Antitrust Criminal Penalty Enforcement and Reform Act. Instead of being liable for treble customer damages, firms now are liable only for single damages. Furthermore, the avoided double damages are to be paid by the other cartel members. All this makes it more likely that applying for leniency is the optimal strategy irrespective of what other cartel members do.

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 we do, I'll share a useful shortcut with you for deriving Nash equilibria.

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 4.13.

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, which 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 4.14.

Next, perform the same exercise for 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

FIGURE 4.13 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

FIGURE 4.14 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

4.15 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

FIGURE 4.16 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

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 4.15.

Because a Nash equilibrium is a strategy pair in which each player's strategy is a best reply, we can identify Nash equilibria in Figure 4.15 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 to use the best-reply method 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 4.16. 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.

4.4 CHECK YOUR UNDERSTANDING

For the game in FIGURE 4.17, use the best-reply method to find the Nash equilibria.*

FIGURE 4.17

		Player 2			
		w	x	y	z
Player 1	a	4, 3	1, 1	0, 3	3, 3
	b	2, 4	3, 1	2, 2	3, 1
	c	0, 6	4, 4	3, 7	1, 2
	d	2, 3	3, 2	1, 0	0, 0

*Answers to Check Your Understanding are in the back of the book.

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 episode 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 wear T-shirts that spe a big "A," Kaitlyn wi who knows? They mi

Although they all tie new top just pur now an hour before home trying to deci should each wear?

In specifying the all wear their lettere sequence). This payc tops, which is 1. Fina girls do not yields a p doesn't look alluring

The strategic form selecting a row—eithe Kaitlyn chooses a col Alicia chooses a matri

Using the b method to solve th consider the situati by Lauren. If Alicia v T-shirt with E and wears hers with Lauren's best reply is part and wear the T: A. So we circle Laure of 2 in the cell associ If, instead, Kaitlyn c to wear her Bebe to Lauren. If Alicia w Bebe top and Kaitlyr then Lauren's best again to wear her Be we circle Lauren's pa Finally, if both of 1 two girls choose th tops, then Lauren does so as well, whi that we now circle payoff of 1 in that in

*This game is a 21st-century *Origins and Foundations of a stag* (rather than spell out is a bit too trendy for your which the stag is North Kor Each of those five countries they can work together and part of multiparty talks un com/2010/04/26/six-party-ta

opportunity to make a statement. While texting, 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 a 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!

Although they all like this idea, each is tempted to wear instead an attractive new top just purchased from their latest shopping expedition to 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 tops, 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 4.18.* 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.

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 4.19. 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 tops, then Lauren optimally does so as well, which means that we now circle Lauren’s payoff of 1 in that instance.

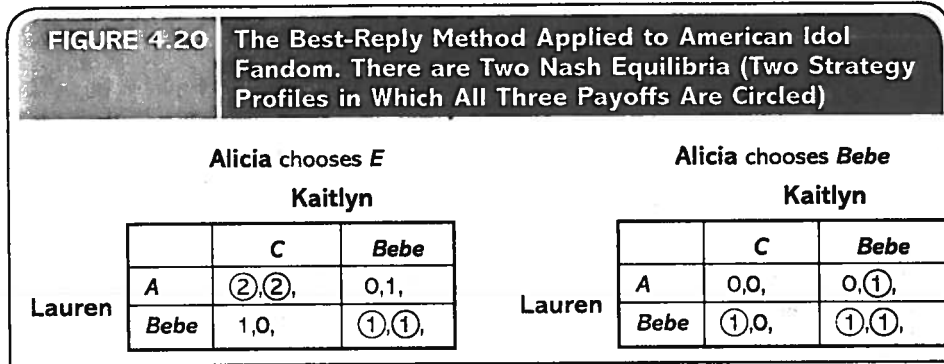
FIGURE 4.18 American Idol Fandom

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

FIGURE 4.19 The Best-Reply Method Applied to Lauren.

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

*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). If this example is a bit too trendy for your tastes—or if, like me, you don’t shop at Bebe—then consider instead a setting in which the stag is North Korea and the hunters are Japan, the United States, South Korea, Russia, and China. Each of those five countries can engage in bilateral talks with North Korea that will have limited success, or they can work together and hold six-party talks with a much greater chance of success. No one wants to be part of multiparty talks unless everyone else has joined in. For details, see asiansecurityblog.wordpress.com/2010/04/26/six-party-talks-as-a-game-theoretic-stag-hunt-1-n-korea-is-the-stag.



Performing this same exercise for Kaitlyn and Alicia, we end up with FIGURE 4.20 which has the best replies circled for all three of the players. Examining this figure, we find that there are two Nash equilibria—that is, two strategy profiles in which all three payoffs are circled—signifying that all three teens are choosing best replies. One equilibrium occurs when Lauren wears A, Kaitlyn wears C, and Alicia wears E, which, to the delight of Ace Young, spells out his name on *American Idol*. The other equilibrium occurs when each tosses her lettered T-shirt aside and instead wears that eye-catching top from Bebe.

Note that the players rank the Nash equilibria the same: all prefer wearing the T-shirts spelling out ACE to wearing Bebe shirts. The situation differs from that pertaining to the driving conventions game, in which players are indifferent among equilibria, and also from that in the Chicken and telephone games, in which players ranked equilibria differently.

INSIGHT When an individual can either work alone for personal benefit or work with others for a common benefit that yields a higher payoff for all—but only when everyone works together—then there are two Nash equilibria: one where everyone works together and one where everyone works alone.

4.5 CHECK YOUR UNDERSTANDING

For the game in FIGURE 4.21, find all Nash equilibria.*

FIGURE 4.21

		Player 3: I		
		Player 2		
Player 1	a	2,1,2	0,0,2	1,2,3
	b	0,3,1	2,2,4	3,1,0
	c	1,1,1	3,2,1	2,2,2

		Player 3: II		
		Player 2		
Player 1	a	2,1,3	1,0,3	1,0,4
	b	1,2,1	3,3,3	1,1,1
	c	1,2,1	1,0,0	2,1,2

*Answers to Check Your Understanding are in the back of the book.

SITUATION: V

"Wasting
asking e
ARGUMENT
CANDIDATE

A company shareholder offers from also has a third choice B's offer, and ranking is lower rankings are

TABLE	
Shareholder	
1	
2	
3	

Assume implement choice. The votes, allocated shareholder required to all of her 2 The strategy applies, which

To derive plurality with B is the winner 100 votes (of the 27 d

The next in a cell in off is 1 (best first choice Substitution

When a rational possibility is used when that shareholder B, and shareholder be approved

► SITUATION: VOTING, SINCERE OR DEVIOS?

"Wasting your vote is voting for somebody that you don't believe in . . . I'm asking everybody watching this nationwide to waste your vote on me."—CLOSING ARGUMENTS BY THE LIBERTARIAN GARY JOHNSON IN A DEBATE AMONG THIRD-PARTY CANDIDATES DURING THE 2012 PRESIDENTIAL RACE.

A company has three shareholders. Shareholder 1 controls 25% of the shares, shareholder 2 controls 35%, and shareholder 3 controls 40%. The company has offers from two other companies, denoted *A* and *B*, to purchase it. The company also has a third option, which is to decline both offers. Shareholder 1 ranks the three choices, from the most to least preferred, as follows: accept *A*'s offer, accept *B*'s offer, and accept neither offer (which we'll denote option *C*). Shareholder 2's ranking is *B*, then *C*, then *A*; and shareholder 3's ranking is *C*, then *B*, then *A*. The rankings are summarized in TABLE 4.3.

Shareholders	1st Choice	2nd Choice	3rd Choice
1	<i>A</i>	<i>B</i>	<i>C</i>
2	<i>B</i>	<i>C</i>	<i>A</i>
3	<i>C</i>	<i>B</i>	<i>A</i>

Assume that a shareholder gets a payoff of 2 if his most preferred choice is implemented, a payoff of 1 for his second choice, and a payoff of 0 for his third 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 applies, 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 4.22 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 4.22. For example, if *B* is the winner, then shareholder 1's payoff is 1 (because *B* is his second choice), shareholder 2's payoff is 2 (since *B* is his first choice), and shareholder 3's payoff is 1 (because *B* is his second choice). Substitution, then, gives us FIGURE 4.23.

When we make statements about how these shareholders might vote, a natural possibility to consider is what political scientists call *sincere voting*. This term 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, because it received the most votes. But is sincere voting a Nash

FIGURE 4.22 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 4.23 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

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 for *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 is that the most preferred option may not be viable—that is, it won't win, regardless of how you vote. In that situation, a player should focus on those options that could prevail and vote for the one that is most preferred. In the case 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*. Because she prefers *B*, she ought to vote for *B* to make that option the winner. ◀◀◀

INSIGHT

Having let's see if equilibrium. Shareholder 1's sense, because votes, regardless. Thus, all the best replies her payoff shareholder associated shareholder votes for A vote for A the associated shareholder votes for replies are for shareholder

Now let all three profiles is a best reply best he or ers are doing equilibrium. If there are three play are Nash (C), (A, C, equilibria lead results in are voting in offer B' (A, C, C) le

We have of Nash eq what mak for sharel answer is alternative If he vote wins (tho (again wit

INSIGHT Use your vote where it can have the most beneficial impact. From among those options for which your vote can influence the outcome, vote for the most preferred option.

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, because 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 4.24 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 4.24.

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 other players are doing; in other words, it is a Nash equilibrium. Inspecting Figure 4.24, we see that there are five strategy profiles for which all three players are using best replies and thus are Nash equilibria: (*A*, *A*, *A*), (*B*, *B*, *B*), (*C*, *C*, *C*), (*A*, *C*, *C*), and (*B*, *B*, *C*). Note that the equilibria lead to different outcomes: (*A*, *A*, *A*) results in offer *A*'s being accepted, because all are voting for *A*. (*B*, *B*, *B*) and (*B*, *B*, *C*) result in offer *B*'s being accepted, and (*C*, *C*, *C*) and (*A*, *C*, *C*) lead to *C*'s being chosen.

We have rather robotically derived the set of Nash equilibria. Although this is useful, it is more important to understand what makes them equilibria. Consider equilibrium (*A*, *A*, *A*). Why is it optimal for shareholders 2 and 3 to vote for their least preferred alternative? The answer is that neither shareholder is pivotal, in that the outcome is the same—alternative *A* wins—regardless of how each votes. Now consider shareholder 2. If he votes for *A*, then *A* wins with 100 votes; if she votes for *B*, then *A* still wins (though now with only 65 votes); and if she votes for *C*, then *A* still wins (again with 65 votes). It is true that shareholders 2 and 3 could work together

FIGURE 4.24 Best-Reply Method Applied to the Voting Game. There Are Five Nash Equilibria.

3 votes for *A*

2

	<i>A</i>	<i>B</i>	<i>C</i>
<i>A</i>	2, 0, 0	2, 0, 0	2, 0, 0
<i>B</i>	2, 0, 0	1, 2, 1	2, 0, 0
<i>C</i>	2, 0, 0	2, 0, 0	0, 1, 2

3 votes for *B*

2

	<i>A</i>	<i>B</i>	<i>C</i>
<i>A</i>	2, 0, 0	1, 2, 1	1, 2, 1
<i>B</i>	1, 2, 1	1, 2, 1	1, 2, 1
<i>C</i>	1, 2, 1	1, 2, 1	0, 1, 2

3 votes for *C*

2

	<i>A</i>	<i>B</i>	<i>C</i>
<i>A</i>	2, 0, 0	0, 1, 2	0, 1, 2
<i>B</i>	0, 1, 2	1, 2, 1	0, 1, 2
<i>C</i>	0, 1, 2	0, 1, 2	0, 1, 2

1's Strategy	3's Strategy	2 Votes for B		2 Votes for A
A	A	0	=	0
A	B	2	>	0
A	C	1	=	0
B	A	2	>	0
B	B	2	=	2
B	C	2	>	1
C	A	0	=	0
C	B	2	=	2
C	C	1	=	1

to achieve higher payoffs. If they both vote for *B*, then *B* wins and shareholders 2 and 3 get payoffs of 2 and 1, respectively, which is better than 0 (which is what they get when *A* wins). But such coordination among players is not permitted. Nash equilibrium requires only that each player, acting *independently* of others, can do no better.*

Equilibrium (*A, A, A*) has another interesting property: shareholders 2 and 3 are using a weakly dominated strategy by voting for *A*. As shown in TABLE 4.4, voting for *A* is weakly dominated in voting for *B* for shareholder 2. For every strategy pair for shareholders 1 and 3, shareholder 2's payoffs in the column "2 votes for *B*" are at least as great as those in the column "2 votes for *A*," and in some of the rows the payoff is strictly greater. So, regardless of how shareholders 1 and 3 vote, voting for *B* gives shareholder 2 at least as high a payoff as does voting for *A*. Of course, when shareholders 1 and 3 vote for *A*—as they do at Nash equilibrium (*A, A, A*)—a vote for *A* and a vote for *B* result in the same payoff of 0 for shareholder 2, so he is acting optimally by voting for *A*. However, 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.

*Note that, in each of the five Nash equilibria, at least one shareholder is not pivotal. With (*B, B, B*) and (*C, C, C*), all three players are not pivotal, just as with (*A, A, A*). With (*A, C, C*), shareholder 1 is not pivotal, although shareholders 2 and 3 are pivotal, because each could ensure *A*'s winning if they voted for *A*. With (*B, B, C*), shareholder 3 is not pivotal, although shareholders 1 and 2 are.

As
an
th
A
ec
gc
piv
*An

► SIT
Sup
mar
your
But
degr
you
;
next
paig
they
don'
rent
weal
ple,
Com
Jian
I
trick
Who
of ef
her
one
E
exert
exert
1. If
by 4.
detr
sible
but t
E
playe
again
effor
S
prom

4.6 CHECK YOUR UNDERSTANDING

Assume that passage of option *A* or *B* requires not a plurality but 70% of the votes, and that otherwise the status quo, *C*, prevails. That is, if *A* receives 70% or more of the votes, then *A* is adopted; if *B* receives 70% or more, then *B* is adopted; and if *A* and *B* each receive less than 70% of the votes, then *C* is adopted. Find all Nash equilibria in which players do not use weakly dominated strategies. (Hint: Before going to the effort of deriving the payoff matrices, think about when each voter is pivotal, and who must vote for an option in order for that option to be adopted.)*

*Answers to Check Your Understanding are in the back of the book.

► 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 you can do that 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.⁶ 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, because one person's 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 4.5. For example, if player *i* exerts a positive effort and the other two players exert a negative effort against her, then her final performance is $v_i - 3$, going up 1 through her positive effort, but down 4 by the two units of negative effort directed against her.

Suppose players care intrinsically, not about performance, but rather about promotion. More specifically, a player's payoff is specified to be the probability

Amount of Positive Effort by i	Amount of Negative Effort Against i	Performance
0	0	v_i
0	1	$v_i - 1$
0	2	$v_i - 4$
1	0	$v_i + 1$
1	1	v_i
1	2	$v_i - 3$

that she is promoted. If a player ends up with a performance higher than those of the other two players, then she is promoted with probability 1, so her payoff is 1. If her performance is highest, but she is tied with one other player, then each has probability $\frac{1}{2}$ of being promoted, and thus each has a payoff of $\frac{1}{2}$. If all three players end up with the same performance, then each receives a payoff of $1/3$. Finally, if a player's performance is below that of another player, then her payoff is 0, since her probability of gaining the promotion is 0.

Assume that $v_1 = 2$ and $v_2 = 0 = v_3$ so that player 1 is the front-runner. To start the analysis, let's be a bit idealistic and consider the "no dirty tricks" strategy profile, in which each player exerts a positive effort, so that player i 's final performance is $v_i + 1$. This scenario translates into a final performance of 3 for player 1 (because she began with 2) and 1 for both players 2 and 3 (because each of them began with 0). Hence, player 1 is promoted. We see, then, that if all exert a positive effort in order to boost their own performances, then the player who was initially ahead will end up ahead and thus will be promoted. Let's now assess whether this is a Nash equilibrium:

- **Player 1:** First note that player 1's strategy is clearly optimal, because her payoff is 1 (recall that it is the probability of being promoted) and that is the highest feasible payoff. Thus, there can't be a strategy for player 1 that delivers a higher payoff.
- **Player 2:** Player 2's payoff from a positive effort is 0, because he is definitely not promoted: his performance of 1 falls short of player 1's performance of 3. Alternatively, he could exert a negative effort against player 3, but that isn't going to help, because the real competition for player 2 is player 1, and going negative against 3 doesn't affect 1's performance. The final alternative is for player 2 to exert a negative effort against player 1, in which case player 1's performance is 2 instead of 3, whereas player 2's performance is 0 instead of 1 (because he is no longer exerting a positive effort on his own behalf). In that case, player 2 is still not promoted. We then find that player 2 is indifferent among all three of his strategies, because all deliver a zero payoff. Thus, because there is no strategy that 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 and re:

In now c let's think al profile. It pr ing player 3 at the start should do th because she Accordir herself, and player 1 (bec positive effo and 0 for pl them, so tha 3 are tied fo players 2 an

- **Player** situati perfor in her shows at play no bet

- **Player** player forma Becau player negati raise h has a : becau: the hiq egy of 2's bei

- **Player** strateg

- In sur is a N player (Jiang Camp

The pro many more One equilib the adage " which the v

- In sum, all three players' choosing a positive effort is a Nash equilibrium and results in the front-runner's 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, because 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, because 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, and 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 (because 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 . Because 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, because only one person is being negative against her. Because 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 the baseball player and manager Leo Durocher: "Nice guys finish last." ◀◀◀

4.7 CHECK YOUR UNDERSTANDING

FIGURE 4.25 is the strategic form game for the Promotion and Sabotage Game. Use the best-reply method to find the Nash equilibria.*

FIGURE 4.25

Player 3: Positive

		Player 2		
		Positive	Negative - 1	Negative - 3
Player 1	Positive	1,0,0	1,0,0	1,0,0
	Negative - 2	1,0,0	$\frac{1}{2}, 0, \frac{1}{2}$	1,0,0
	Negative - 3	1,0,0	1,0,0	1,0,0

Player 3: Negative against player 1

		Player 2		
		Positive	Negative - 1	Negative - 3
Player 1	Positive	1,0,0	$0, \frac{1}{2}, \frac{1}{2}$	1,0,0
	Negative - 2	1,0,0	0,0,1	1,0,0
	Negative - 3	$\frac{1}{2}, \frac{1}{2}, 0$	0,1,0	1,0,0

Player 3: Negative against player 2

		Player 2		
		Positive	Negative - 1	Negative - 3
Player 1	Positive	1,0,0	1,0,0	1,0,0
	Negative - 2	1,0,0	1,0,0	1,0,0
	Negative - 3	1,0,0	1,0,0	1,0,0

*Answers to Check Your Understanding are in the back of the book.

4.5 Foundations of Nash Equilibrium

THUS FAR, WE'VE OFFERED two approaches to solving a game: iterative deletion of strictly dominated strategies (IDSDS) 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 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 a
against j
dominat
 x and pl
for playe
nated. T

Wha
player 3'
not, and
explains
why z st
rounds. I
works fo
gies surv
player ga

Altho
survive th
this chap
the strate
IDSDS sa
trast, onl
real fast!
which all
ria. Nash
strategy p



FIGURE
Nash equ
of the st
survive th
themselve
strategy
for any pa
sets could
for exam
equilibria
as those
which sur
in the Pr
Or the str
the IDSD
with the
profiles (a
Idol fando

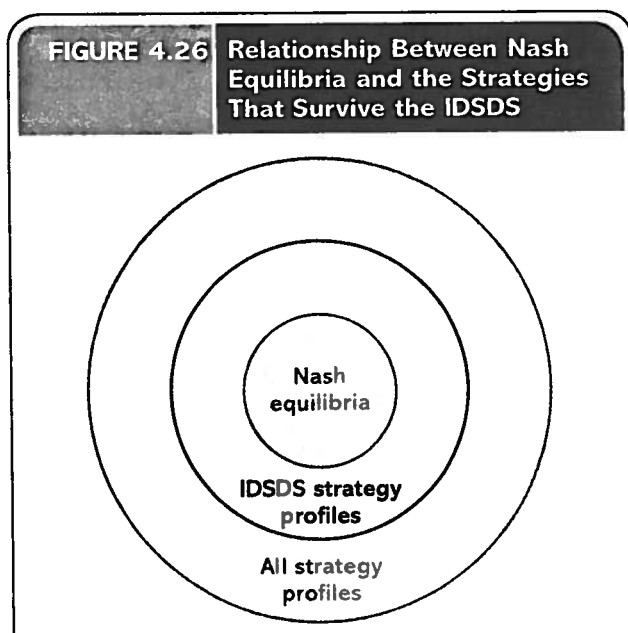
easy to argue that none are eliminated in the first round: Because x is a best reply against player 2's using y and player 3's using z , x is most definitely not strictly dominated. Analogously, because y is a best reply for player 2 when player 1 uses x and player 3 uses z , y is not strictly dominated. Finally, because 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 still is not strictly dominated for player 2 in the second round and why z still is 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.

Although 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, because 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, because 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 4.26 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 (as in the Prisoners' Dilemma). Or the strategies that survive the IDSDS might coincide with the set of all strategy profiles (as in the *American Idol* fandom game).



4.8 CHECK YOUR UNDERSTANDING

For the game in FIGURE 4.27, show that the set of Nash equilibria is smaller than the set of strategies that survive the IDSDS.*

FIGURE 4.27

Player 3: Up

		Player 2		
		Left	Center	Right
Player 1	Top	1,2,3	1,0,2	0,2,1
	Middle	4,4,1	1,6,3	3,7,1
	Bottom	3,3,3	2,3,2	1,2,4

Player 3: Down

		Player 2		
		Left	Center	Right
Player 1	Top	6,2,1	3,1,1	2,5,0
	Middle	3,3,0	5,2,2	4,6,0
	Bottom	2,4,1	1,3,1	4,2,1

*Answers to Check Your Understanding are in the back of the book.

4.5.2 The Definition of a Strategy, Revisited

To understand better 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 $s_j(i)$ —that 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 does 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
happen in

At a N
ture—in w
specify bel
egy both p_i
player's co

4.6 Fict

"The in
We has
—ALBE

THE APPEAL
ers accura
those conj
Although t
in the hec
will do? F
apply to se
Actual

For exampl
least for d
games fro
Driving C
game), wh
educated t
equilibrium
ent equilib

It is in
people use
Functiona
activity in
for domin
anterior c
neuroecon

As an e
simulated
who is rat
Chapter 1,
future beh
is that me
approach,
knowledge
abilities ar
how unso

*This section
read Section

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.

4.6 Fictitious Play and Convergence to Nash Equilibrium*

"The intuitive mind is a sacred gift and the rational mind is a faithful servant. We have created a society that honors the servant and has forgotten the gift."

—ALBERT EINSTEIN

THE APPEAL OF NASH EQUILIBRIUM is that once players are there—that is, once players accurately conjecture what others will do and choose the optimal strategy for those conjectures—there is no incentive for any player to do something different. Although that is all well and good, it leaves unaddressed a key presumption: How in the heck does each player come to have an accurate forecast of what others will do? From whence comes this gift of foresight? Is Nash equilibrium to only apply to soothsayers and psychics?

Actually, in some games, prescient prognostication may not be that difficult. For example, the IDSDS is an algorithm that produces an accurate conjecture, at least for dominance solvable games (such as the Existence of God and Doping games from Chapter 3). But what about a coordination game (such as the Driving Conventions game) or a Battle of the Sexes (such as the Telephone game), which have multiple Nash equilibria? Even if players have been similarly educated to expect a Nash equilibrium, it may still be unclear as to *which* Nash equilibrium will prevail. Players may then miscoordinate if each expects a different equilibrium to occur. As a result, play isn't consistent with either equilibrium.

It is interesting that some recent evidence from neuroeconomics indicates that people use different parts of the brain depending on the type of game being played. Functional magnetic resonance imaging (fMRI) showed relatively more brain activity in the middle frontal gyrus, the inferior parietal lobule, and the precuneus for dominance solvable games, and relatively more brain activity in the insula and anterior cingulate cortex for coordination games.⁷ Although this is fascinating, neuroeconomics is a long way from addressing the question before us.

As an avenue for figuring out what other players will do, the IDSDS is based on *simulated introspection*, whereby a player puts himself in the shoes of another player who is rational, believes other players are rational, and so forth. If you recall from Chapter 1, experiential learning is another method for forming beliefs regarding the future behavior of other players and involves using past play to predict behavior. It is that method which we'll explore here. Contrary to the simulated introspection approach, it will require neither much sophistication on the part of a player nor even knowledge of other players' payoffs and strategy sets. However, it does involve probabilities and expectations and, therefore, it will take a sophisticated reader to explain how unsophisticated players will end up at Nash equilibrium.

*This section uses probabilities and expectations. The reader unfamiliar with these concepts should first read Section 7.2.

FIGURE 4.28 Driving Conventions

		Louise	
		Drive left	Drive right
Thelma	Drive left	1, 1	-1, -1
	Drive right	-1, -4	2, 2

To give players an opportunity to learn about what other players will do, assume they play the game not once but over and over. There are many ways to learn from past play, and here we'll describe perhaps the simplest method for doing so. Referred to as *fictitious play*, it has a player believe that other players will act today as they've acted in the past. If a player has always chosen strategy z , then we believe she'll choose strategy z today. What about if the player has chosen various strategies in the past? We then assume that the more frequently a strategy was chosen in the past, the more likely it will be chosen today. More specifically (and here is where probabilities enter), the probability that player 1 assigns to player 2 of choosing z in the current period is specified to equal the percentage of past periods that player 2 chose z . Given those beliefs as to what player 2 will do today, player 1 chooses a strategy to maximize her expected payoff.

To make all of this more concrete, consider the modified version of the Driving Conventions game in FIGURE 4.28.

There are still two Nash equilibria, but the equilibrium in which everyone drives on the right is preferred to the equilibrium in which everyone drives on the left.*

Suppose Thelma and Louise have interacted for 17 periods and that, over those periods, Louise has chosen Left 12 times, and Thelma has chosen Left 10 times. What does fictitious play imply about behavior in their eighteenth encounter? The first step in answering that question is to derive a driver's beliefs regarding the other driver's strategy. Given that Louise has chosen Left in 71% ($\approx 12/17$) of past encounters, Thelma believes there is a 71% chance that Louise will choose Left and a 29% chance she'll choose Right. Given these beliefs, Thelma's expected payoff from driving on the left is $0.71 \times 1 + 0.29 \times (-1) = 0.42$, which follows from the fact that there is a 71% chance that both will be driving on the left and a 29% chance they'll fail to coordinate because Louise is driving on the right and Thelma is driving on the left. By analogous argument, Thelma's expected payoff from driving on the right is $0.71 \times (-1) + 0.29 \times 2 = -0.13$. Hence, Thelma will choose to drive on the left in order to receive the higher expected payoff of 0.42 compared with -0.13 . Let's turn to Louise, who thinks there is a 59% ($= 10/17$) chance that Thelma will choose Left, which means that Louise's expected payoff from driving on the left is $0.59 \times 1 + 0.41 \times (-1) = 0.18$ and on the right is $0.59 \times (-1) + 0.41 \times 2 = 0.23$. Hence, Louise will optimally choose Right.

What we have just described is how fictitious play operates in a single period—but remember that it is an ongoing process. Come encounter number 19, Louise's beliefs will change in response to the observed play of Thelma in encounter number 18. As argued above, Thelma will have chosen Left, in which case Louise now thinks there is a 61% ($= 11/18$) chance that Thelma will drive on the left. Hence, the expected payoff to Louise from Left is $0.61 \times 1 + 0.39 \times (-1) = .22$

*There is evidence that accidents are less likely when a driving convention coincides with a driver's dominant hand. That is, right-handers have lower accident rates when people drive on the right than when they drive on the left. Given that about 85 to 90% of people are right handed, it is then natural for the (Right, Right) equilibrium to yield a higher payoff than the (Left, Left) equilibrium.

and from Thelma. We period with Louise encourages beliefs in the past and this probability. And if the answer



If we play has players' given account two states to argue condition time." If, past period closer to times, so encounter ($= 93/100$ forth. Thus being account case, is Left strategy. Louise then fictitious and will choose having account must be a

Okay, when ing the same equilibrium question, it p denote the left. For example $p = 0.71$ to for driving $p \times 1 + (1$ Hence, Left

$p \times$

and from Right is $0.61 \times (-1) + 0.39 \times 2 = 0.17$, so she'll drive on the left. Thelma goes through an analogous process in determining her play.

We are not really interested in what fictitious play produces in any given period—indeed, it may produce different choices over time (as we just observed with Louise who drove on the right in encounter number 18 and on the left in encounter number 19)—but rather where it ends up. If each player forms her beliefs as to the choices of other players on the basis of how they've played in the past and, given those beliefs, acts to maximize her expected payoff, where will this process end? Will it settle down so that players are making the same choices? And if it does settle down, does their play form a Nash equilibrium? We can answer the second question definitively:

◆ INSIGHT If fictitious play settles down in the sense that each player is choosing the same strategy over time, then players' strategies form a Nash equilibrium.

If we give it a bit of thought, this Insight should not be surprising. Fictitious play has each player using an optimal strategy concerning beliefs about other players' strategies. Nash equilibrium has each player using an optimal strategy, given accurate beliefs as to other players' strategies. The only difference in those two statements is that Nash equilibrium involves *accurate* beliefs. Thus, we need to argue only that fictitious play eventually results in accurate beliefs. The key condition in this Insight is that "each player is choosing the same strategy over time." If, eventually, Louise is always driving on the left, then the percentage of past periods for which she chose left will steadily increase and get closer and closer to 1. Picking up where we left off, Thelma had chosen Left 12 out of 19 times, so Louise thinks there is a 63% chance that Thelma will choose Left in encounter 20. If Thelma keeps choosing left, then Louise's beliefs will rise to 93% ($= 93/100$) by encounter 101, to 99.3% by 1,001, to 99.93% by 10,001, and so forth. Thus, if Thelma always chooses Left, then Louise's beliefs will converge on being accurate. We then see that Louise is choosing a strategy (which, in this case, is Left) that maximizes her payoff, given her accurate beliefs about Thelma's strategy. Let's turn to Thelma. If Louise is choosing the same strategy over time, then fictitious play implies that Thelma eventually will have accurate beliefs and will choose a strategy that is optimal given those beliefs. With both players having accurate beliefs and choosing an optimal strategy given those beliefs, it must be a Nash equilibrium, for that is its very definition.

Okay, we've shown that if players' behavior settles down so that each is choosing the same strategy over time, then they will be acting according to a Nash equilibrium. What about the other question: Will play settle down? To tackle this question, it'll be useful to know exactly how beliefs influence optimal choice. Let p denote the probability that a driver assigns to the other driver of driving on the left. For example, after 17 encounters, the preceding analysis had Thelma assign $p = 0.71$ to Louise, and Louise assign $p = 0.59$ to Thelma. What must p be in order for driving on the left to be optimal? The expected payoff from choosing Left is $p \times 1 + (1 - p) \times (-1)$ and from choosing Right is $p \times (-1) + (1 - p) \times 2$. Hence, Left is preferred only when:

$$p \times 1 + (1 - p) \times (-1) > p \times (-1) + (1 - p) \times 2, \text{ or } p > 0.6$$

Thus, if the chances of Thelma's driving on the left exceed 60%, then it is best for Louise to drive on the left. Likewise, if the chances of Thelma's driving on the left are less than 60%, then it is best for Louise to drive on the right. (Louise is indifferent between these two options when $p = 0.6$). Armed with that information, let us show that the play of Thelma and Louise will indeed settle down.

As we already noted for encounter 18, Thelma has $p = 0.71$ for Louise, in which case Thelma optimally chooses Left, whereas Louise has $p = 0.59$, who optimally chooses Right. In encounter 19, Thelma then lowers p from 0.71 to 0.67, because Louise chose Right in the previous encounter (and thus the frequency with which Louise chose Left went from 12/17 to 12/18). But because p still exceeds 0.6, Thelma still chooses Left. Let's turn to Louise, who has raised p from 0.59 to 0.61 between encounters 18 and 19 and, therefore, switches from Right to Left because p now exceeds 0.6. Because both chose Left in encounter 19, each will assign even greater chances to the other driver's choosing Left in encounter 20; hence, it will be optimal to choose Left again. From here on, both drivers choose Left, which leads to a higher probability that the other will choose Left, which reinforces even more the optimality of choosing Left. In this case, it is clear that the process does converge. Although Thelma and Louise were not at a Nash equilibrium in encounter 18, they were at encounter 19 and every encounter thereafter.

Although fictitious play results in Thelma and Louise's converging on a Nash equilibrium in this particular example, is that generally true? For the game in Figure 4.28, fictitious play always does converge, though where it converges to depends on initial play. For example, suppose that after 17 encounters, both Thelma and Louise had instead chosen Left half of the time. Because then $p = 0.5$ (which is less than 0.6), both will choose Right in encounter 18. This will lower p for encounter 19, which means that both will choose Right again. They will keep driving on the right and keep lowering the probability they assign to the other driver's driving on the left. As a result, play converges to the "drive on the right" Nash equilibrium.

More generally, fictitious play converges for some but not all games. Although it may be disappointing to end on such an ambiguous note, it is still impressive that such a simple learning rule could ultimately result in Nash equilibrium.

INSIGHT For some games, fictitious play eventually results in players' using strategies associated with a Nash equilibrium.

We've shown how unsophisticated players can learn their way to a Nash equilibrium. But what about moving from one equilibrium to another equilibrium? Though Nash equilibria are stable in the sense that no individual wants to change, players collectively may desire change. For example, if Thelma and Louise are at the Nash equilibrium in which the convention is driving on the left, they instead would prefer to be at the Nash equilibrium with a right-side convention, because that has a payoff of 2 rather than 1. Of course, as long as Thelma thinks Louise will drive on the left, then so will she. To manage a move from one

equilibrium which governs driving on

4.9
Suppose Louise had they end
*Answers to C

Summary

A rational about what best reply suppose that paying the rium. The contentme the strateg

In man has no trac librium is Nash equi IDSDS. Fo tion conce a unique N

In deri algorithmi procedure when play incentives Nash equil

Games interest. C makes one the childre guessing expect. At players coi others bett nation ga there are Telephone standing t another pl likely to be

equilibrium to another requires coordinated action among players, something at which governments can be effective—as they were, for example, in Austria in 1938, in Sweden in 1967, and in Myanmar in 1970, when they switched from driving on the left to driving on the right.

4.9 CHECK YOUR UNDERSTANDING

Suppose that, after 17 encounters, Thelma has driven on the left 14 times and Louise has driven on the left 8 times. If Thelma and Louise use fictitious play, will they end up with a convention of driving on the right or on the left?*

*Answers to Check Your Understanding are in the back of the book.

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. For that very reason, Nash equilibrium generally is a more useful solution concept. Nevertheless, as we found out by way of example, a game can have a unique Nash equilibrium, many Nash equilibria, or none at all.

In deriving the Nash equilibria for a game, we 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. One 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, a **coordination game** in which players simply want to choose the same action. Then there are games that combine conflict and mutual interest, such as the Telephone game, Chicken, and *American Idol* fandom. In these games, understanding the incentives of a player—how best a player should react to what another player is going to do—can provide insight into what strategy profiles are likely to be Nash equilibria.

EXERCISES

1. One of the critical moments early on in the *The Lord of the Rings* trilogy is the meeting in Rivendell to decide who should take the One Ring to Mordor. Gimli the Dwarf won't hear of an Elf doing it, whereas Legolas (who is an Elf) feels similarly about Gimli. Boromir (who is a Man) is opposed to either of them taking charge of the Ring. And then there is Frodo the Hobbit, who has the weakest desire to take the Ring but knows that someone must throw it into the fires of Mordor. In modeling this scenario as a game, assume there are four players: Boromir, Frodo, Gimli, and Legolas. (There were more, of course, including Aragorn and Elrond, but let's keep it simple.) Each of them has a preference ordering, shown in the following table, as to who should take on the task of carrying the One Ring.

Preference Rankings for The Lord of the Rings

Person	First	Second	Third	Fourth	Fifth
Boromir	Boromir	Frodo	No one	Legolas	Gimli
Gimli	Gimli	Frodo	No one	Boromir	Legolas
Legolas	Legolas	Frodo	No one	Gimli	Boromir
Frodo	Legolas	Gimli	Boromir	Frodo	No one

Of the three non-Hobbits, each prefers to take on the task himself. Each would prefer that other than themselves and Frodo, no one should take the Ring. As for Frodo, he doesn't really want to do it and prefers to do so only if no one else will. The game is one in which all players simultaneously make a choice among the four people. Only if they all agree—a unanimity voting rule is put in place—is someone selected; otherwise, no one takes on this epic task. Find all symmetric Nash equilibria.

2. Consider a modification of driving conventions, shown in the figure below, in which each player has a third strategy: to zigzag on the road. Suppose that if a player chooses *zigzag*, the chances of an accident are the same whether the other player drives on the left, drives on the right, or zigzags as well. Let that payoff be 0, so that it lies between -1, the payoff when a collision occurs for sure, and 1, the payoff when a collision does not occur. Find all Nash equilibria.

Modified Driving Conventions Game

	Drive left	Drive right	Zigzag
Drive left	1,1	-1,-1	0,0
Drive right	-1,-1	1,1	0,0
Zigzag	0,0	0,0	0,0

3. Return to is partne Nash equ

4. Consider

- a. For e:
deleti
- b. Deriv

5. Consider

- a. Deriv
domi
- b. Deriv

6. Return to eight hor owned b black far Nash equ has the r tory. Cen Astor wa

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 below. Find all Nash equilibria.

Team Project

		Sorority girl		
		Low	Moderate	High
Frat boy	Low	0,0	2,1	6,2
	Moderate	1,2	4,4	5,3
	High	2,6	3,5	3,4

4. Consider the two-player game illustrated here.

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

- a. For each player, derive those strategies which survive the iterative deletion of strictly dominated strategies.
 b. Derive all strategy pairs that are Nash equilibria.
5. Consider the two-player game depicted here.

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

- a. Derive those strategies which survive the iterative deletion of strictly dominated strategies.
 b. Derive all strategy pairs that are Nash equilibria.
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. (Surely you're 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 below. Find all of the Nash equilibria.

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 Hassan al Bolkuh 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 quirkiness, 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 that 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, whereas the losing bidder's payoff is 0. Assume that the Sultan has a valuation of 8 (hundred million pounds) and that the Sheikh has a valuation of 7.
- In matrix form, write down the strategic form of this game.
 - Derive all Nash equilibria.
9. Find all of the Nash equilibria for the three-player game here.

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

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

10. 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 4.23, find all of the Nash equilibria that do not have players using a weakly dominated strategy.
11. Recall the example of Galileo Galilei and the Inquisition in Chapter 2. The strategic form of the game is reproduced here. Find all of the Nash equilibria.

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

12. Find all of the Nash equilibria for the three-player game shown below.

Player 3: A

Player 2

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

Player 3: B

Player 2

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

13. On Friday night, Elton and his partner Rodney are deciding where to go for dinner. The choices are Indian, Korean, and Mexican. Elton most likes Indian food and most dislikes Mexican food, whereas Mexican is Rodney's favorite and Indian is his least favorite. Each cares about the food but also about dining together. As long as the food is Indian or Korean, Elton prefers to go to the restaurant he thinks Rodney will choose. However, he abhors Mexican food and would choose to dine alone at either the Indian or the Korean restaurant rather than joining Rodney at the Mexican place. As long as the food is Mexican or Korean, Rodney will decide to go where he thinks Elton will choose. However, Rodney is allergic to some of the Indian spices and prefers dining alone to eating Indian food. Both of them are at their separate workplaces and must simultaneously decide on a restaurant. Find all Nash equilibria.
14. Let us return to Juan and María from CYU 4.2 but modify their preferences. It is still the case that they are competitive and are deciding whether to show up at their mom's house at 8:00 A.M., 9:00 A.M., 10:00 A.M., or 11:00 A.M. But now they don't mind waking up early. Assume that the payoff is 1 if he or she shows up before the other sibling, it is 0 if he or she shows up after the other sibling, and it is -1 if they show up at the same time. The time of the morning does not matter. Find all Nash equilibria.
15. Two companies are deciding at what point to enter a market. The market lasts for four periods and companies simultaneously decide whether to enter in period 1, 2, 3, or 4, or not enter at all. Thus, the strategy set of a company is $\{1, 2, 3, 4, \text{do not enter}\}$. The market is growing over time, which is reflected in growing profit from being in the market. Assume that the profit received by a monopolist in period t (where a monopoly means that only one company has entered) is $10 \times t - 15$, whereas each duopolist (so both have entered) would earn $4 \times t - 15$. A company earns zero profit for any period that it is not in the market. For example, if company 1 entered in period 2 and company 2 entered in period 3, then company 1 earns zero profit in period 1; $5 (= 10 \times 2 - 15)$ in period 2; $-3 (= 4 \times 3 - 15)$ in period 3; and $1 (= 4 \times 4 - 15)$ in period 4, for a total payoff of 3. Company 2 earns zero profit in periods 1 and 2, -3 in period 3, and 1 in period 4, for a total payoff of -2 .
- Derive the payoff matrix.
 - Derive a company's best reply for each strategy of the other company.
 - Find the strategies that survive the IDSDS.
 - Find the Nash equilibria.
16. Consider an odd type of student who prefers to study alone except when the group is large. We have four of these folks: Melissa, Josh, Samina, and Wei. Melissa and Josh are deciding between studying in the common room in their dorm (which we will denote D) and the library (denoted L). Samina and Wei are choosing between the library and the local cafe (denoted C). If someone is the only person at a location, then his or her payoff is 6. If he or she is one of two people at a location, then the payoff is 2. If he or she is one of three people, then the payoff is 1. If all four end up together, then the payoff is 8.
- Is it a Nash equilibrium for Melissa and Josh to study in the common room and for Samina and Wei to study in the cafe?
 - Is it a Nash equilibrium for Josh to study in the common room, Samina to study in the cafe, and Melissa and Wei to study in the library?
 - Find the Nash equilibria.
17. Four political candidates are deciding whether to enter into a primary election. Each candidate has a certain probability of being elected, and each candidate has a certain probability of being elected if the other candidates do not enter. The candidates are: Candidate 1, Candidate 2, Candidate 3, and Candidate 4. The candidates are deciding whether to enter the primary election. Find all Nash equilibria.
18. Player 1 and Player 2 are deciding whether to enter a game. Player 1's payoff is 1 if he enters and Player 2's payoff is 2 if she enters. If both enter, the payoff is 0 for both. If neither enters, the payoff is 0 for both. Find all Nash equilibria.
- What is the payoff matrix?
 - For each strategy of the other player, what is the best response?
 - Is there a Nash equilibrium?
 - Is there a dominant strategy?
 - Is there a dominant strategy equilibrium?
 - Is there a mixed strategy Nash equilibrium?
19. It is the Scottish Highland Games. The village of Pitlochry is holding a village competition. The competition is a battle between the English forces and the Scottish forces. If the English forces win, then the English forces will receive a payoff of 10. If the Scottish forces win, then the Scottish forces will receive a payoff of 10. If the battle is a draw, then both sides will receive a payoff of 5. The English forces are deciding whether to enter the competition. The Scottish forces are deciding whether to enter the competition. Find all Nash equilibria.

17. Four political candidates are deciding whether or not to enter a race for an elected office where the decision depends on who else is throwing his or her hat into the ring. Suppose candidate A prefers not to enter if candidate B is expected to enter; otherwise, A prefers to enter. Candidate B prefers not to enter if she expects either candidate A and/or candidate D to enter; otherwise, she prefers to enter. Candidate C prefers not to enter if he expects candidate A to enter and prefers to enter in all other cases. Candidate D prefers not to enter if either candidate B and/or C are expected to enter; otherwise, she prefers to enter. If we assume that their choices are consistent with Nash equilibrium, who will enter the race?
18. Player 1 chooses a value for x from the set $\{0,1,2,3\}$. Once x has been selected, players 2 and 3 observe the value of x and then play the following game, where the first payoff in a cell is for player 1, the second payoff is for player 2, and the third payoff is for player 3:

		Player 3		
		a_3	b_3	c_3
Player 2	a_2	$1-x, 5, 5$	$1-x, 3, 2$	$5-x, 3, 2x$
	b_2	$1-x, 2, 3$	$3-x, x, x$	$5-x, x, 2x$
	c_2	$5-x, 2x, 3$	$5-x, 2x, x$	$5-x, 1.1x, 1x$

- What is the strategy set for player 1? for player 2? for player 3?
 - For each value of x , find the Nash equilibria for the game played between players 2 and 3.
 - Is there a Nash equilibrium in which $x = 0$?
 - Is there a Nash equilibrium in which $x = 1$?
 - Is there a Nash equilibrium in which $x = 2$?
 - Is there a Nash equilibrium in which $x = 3$?
19. It is thirteenth-century Scotland, and the English hordes are raiding the Scottish countryside. The villages of Aviemore, Braetongue, Glenfinnan, Pitlochry, and Shildaig are deciding whether to adopt a defensive or an offensive strategy. A defensive strategy means constructing barriers around a village, which will protect the people and homes from the English but will mean losing livestock and crops. If a village adopts a defensive strategy, its payoff is 30. Alternatively, a village can take an offensive strategy and fight the English. In addition, if two or more villages choose to fight, they will join forces. If some villages choose to fight and succeed in defeating the English, then the payoff (for those villages which chose to fight) is 100, but if they lose the battle then the payoff is -50 (again, for those villages which chose to fight). The probability of winning depends on how many villages unite to form an army. Assume the English are defeated with probability $(n/(n+3))$, where n is the number of villages that choose to fight, and they lose the battle with probability $(3/(n+3))$. Thus, the expected payoff to a village from this offensive strategy is $((n/(n+3))) \times 100 - ((3/(n+3))) \times 50$.
- Find the Nash equilibria.
- Now suppose that the payoff for losing the battle is 0, not 50, so the expected payoff to a village that chooses to fight is $((n/(n+3))) \times 100$, where again n is the total number of villages that unite to fight the English.

For another change in assumption, suppose a village that chooses the defensive strategy receives a payoff of 6 (not 30) but, in addition, realizes the expected benefit of defeating the English. Thus, if m other villages choose to fight, then the payoff to a village that chose the defensive strategy is $6 + ((m/(m + 3))) \times 100$.

b. Find the Nash equilibria.

20. Consider the telephone game in Figure 4.7. After 10 periods, Colleen has chosen Call eight times, and Winnie has chosen it seven times. What Nash equilibrium is predicted by fictitious play to occur eventually? (Note: If a player is indifferent between two strategies, then assume she chooses Call.)

4.7 Appendix: Formal Definition of 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 for 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 $(s_1^*, \dots, s_{i-1}^*, s_{i+1}^*, \dots, s_n^*)$. In more formal terms, (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 (August 18, 2003) www.eonline.com. *Sharkslayer* later had its title changed to *Shark Tale*.
2. Though this story may be apocryphal, 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. "Cracking Cartels With Leniency Programs," OECD Competition Committee, Paris, October 18, 2005.
6. 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.
7. Wen-Jui Kuo, Tomas Sjöström, Yu-Ping Chen, Yen-Hsiang Wang, and Chen-Ying Huang, "Intuition and Deliberation: Two Systems for Strategizing in the Brain," *Science*, 324 (2009): 519–22.

Stable
Nash
n-Pla

Man's
dimens

5.1 Int

IN THIS CH
array of s
players, b
I mean th
Whatever
librium lo
the numb

Before
it is helpf
players ha
players als
ric. Then, i
distinct p
equilibrium
symmetric
symmetric
asymmetri

Two in
make iden
Tipping re
when mon
shouldn't b
"fit in." In
your frienc
you. In a g
balance" so
the term "t
world caus
retailers, s
book, migh
erty, extren
strategy—a
tipping at v