GAMES OF STRATEGY

THIRD EDITION

Avinash Dixit
Princeton University

Susan Skeath
Wellesley College

David Reiley
University of Arizona and Yahoo! Research

W. W. Norton & Company
New York • London

and $1 - x - y$ for Rudy; similarly, the probability of winning when Rich gives up is z for Kelly and $1 - z$ for Rudy. Further, suppose that Rich's chance of being picked by the jury is p if he has won immunity and has voted Rudy off the island; his chance of being picked is q if Kelly has won immunity and has voted Rudy off. Continue to assume that if Rudy wins immunity, he keeps Rich with probability 1, and that Rudy wins the game with probability 1 if he ends up in the final two. Note that in the example of Figure 3.11, we had $x = 0.45$, $y = 0.5$, $z = 0.9$, $p = 0.4$, and $q = 0.6$. (In general, the variables p and q need not sum to 1, though this happened to be true in Figure 3.11.)

(a) Find an algebraic formula, in terms of x, y, z, p, and q, for the probability that Rich wins the million dollars if he chooses Continue. (Note: Your formula might not contain all of these variables.)

(b) Find a similar algebraic formula for the probability that Rich wins the million dollars if he chooses Give Up. (Again, your formula might not contain all of the variables.)

(c) Use these results to find an algebraic inequality telling us under what circumstances Rich should choose Give Up.

(d) Suppose all the values are the same as in Figure 3.11 except for z. How high or low could z be so that Rich would still prefer to Give Up? Explain intuitively why there are some values of z for which Rich is better off choosing Continue.

(e) Suppose all the values are the same as in Figure 3.11 except for p and q. Assume that since the jury is more likely to choose a "nice" person who doesn't vote Rudy off, we should have $p > 0.5 > q$. For what values of the ratio (p/q) should Rich choose Give Up? Explain intuitively why there are some values of p and q for which Rich is better off choosing Continue.

4

Simultaneous-Move Games with Pure Strategies I: Discrete Strategies

Recall from Chapter 2 that games are said to have simultaneous moves if players must move without knowledge of what their rivals have chosen to do. It is obviously so if players choose their actions at exactly the same time. A game is also simultaneous when players choose their actions in isolation, with no information about what other players have done or will do, even if the choices are made at different hours of the clock. (For this reason, simultaneous-move games have *imperfect information* in the sense we defined in Chapter 2, Section 2.D.) This chapter focuses on games that have such purely simultaneous interactions among players. We consider a variety of types of simultaneous games, introduce a solution concept called Nash equilibrium for these games, and study games with one equilibrium, many equilibria, or no equilibrium at all.

Many familiar strategic situations can be described as simultaneous-move games. The various producers of television sets, stereos, or automobiles make decisions about product design and features without knowing what rival firms are doing about their own products. Voters in U.S. elections simultaneously cast their individual votes; no voter knows what the others have done when she makes her own decision. The interaction between a soccer goalie and an opposing striker during a penalty kick requires both players to make their decisions simultaneously—the goalie cannot afford to wait until the ball has actually been kicked to decide which way to go, because then it would be far too late.

When a player in a simultaneous-move game chooses her action, she obviously does so without any knowledge of the choices made by other players. She also cannot look ahead to how they will react to her choice, because they act simultaneously and do not know what she is choosing. Rather, each player must figure out what others are doing when what the others are doing is figuring out what this player is doing. This circularity makes the analysis of simultaneous-move games somewhat more intricate than that of sequential-move games, but the analysis is not difficult. In this chapter, we develop a simple concept of equilibrium for such games that has considerable explanatory and predictive power.

1 DEPICTING SIMULTANEOUS-MOVE GAMES WITH DISCRETE STRATEGIES

In Chapters 2 and 3, we emphasized that a strategy is a complete plan of action. But in a purely simultaneous-move game, each player can have at most one opportunity to act (although that action may have many component parts); if a player had multiple opportunities to act, that would be an element of sequentiality. Therefore there is no real distinction between strategy and action in simultaneous-move games, and the terms are often used as synonyms in this context. There is only one complication. A strategy can be a probabilistic choice from the basic actions initially specified. For example, in sports, a player or team may deliberately randomize its choice of action to keep the opponent guessing. Such probabilistic strategies are called **mixed strategies,** and we consider them in Chapters 7 and 8. In this chapter, we confine our attention to the basic initially specified actions, which are called **pure strategies.**

In many games, each player has available to her a finite number of discrete pure strategies—for example, Dribble, Pass, or Shoot in basketball. In other games, each player's pure strategy can be any number from a continuous range—for example, the price charged by a firm.[1] This distinction makes no difference to the general concept of equilibrium in simultaneous-move games, but the ideas are more easily conveyed with discrete strategies; solution of games with continuous strategies needs slightly more advanced tools. Therefore, in this chapter, we restrict the analysis to the simpler case of discrete pure strategies and take up continuously variable strategies in Chapter 5.

Simultaneous-move games with discrete strategies are most often depicted with the use of a **game table** (also called a **game matrix** or **payoff table**). The table is called the **normal form** or the **strategic form** of the game. Games with any number of players can be illustrated by using a game table, but its dimen-

[1]In fact, prices must be denominated in the minimum unit of coinage—for example, whole cents—and can therefore take on only a finite number of discrete values. But this unit is usually so small that it makes more sense to think of the price as a continuous variable.

sion must equal the number of players. For a two-player game, the table is two dimensional and appears similar to a spreadsheet. The row and column headings of the table are the strategies available to the first and second players, respectively. The size of the table, then, is determined by the numbers of strategies available to the players.[2] Each cell within the table lists the payoffs to all players that arise under the configuration of strategies that placed players into that cell. Games with three players require three-dimensional tables; we consider them later in this chapter.

We illustrate the concept of a payoff table for a simple game in Figure 4.1. The game here has no special interpretation; so we can develop the concepts without the distraction of a "story." The players are named Row and Column. Row has four choices (strategies or actions) labeled Top, High, Low, and Bottom; Column has three choices labeled Left, Middle, and Right. Each selection of Row and Column generates a potential outcome of the game. Payoffs associated with each outcome are shown in the cell corresponding to that row and that column. By convention, of the two payoff numbers, the first is Row's payoff and the second is Column's. For example, if Row chooses High and Column chooses Right, the payoffs are 6 to Row and 4 to Column. For additional convenience, we show everything pertaining to Row—player name, strategies, and payoffs—in black, and everything pertaining to Column in green.

Remember that, in some games, most notably in sports contexts, the interests of the two sides are exactly the opposite of each other. Then, for each combination of the players' choices, the payoffs of one can be obtained by reversing the sign of the payoffs to the other. As noted in Chapter 2, we call these **zero-sum** (or **constant-sum**) games.

		COLUMN		
		Left	Middle	Right
ROW	Top	3, 1	2, 3	10, 2
	High	4, 5	3, 0	6, 4
	Low	2, 2	5, 4	12, 3
	Bottom	5, 6	4, 5	9, 7

FIGURE 4.1 Representing a Simultaneous-Move Game in a Table

[2]If each firm can choose its price at any number of cents in a range that extends over a dollar, each has 100 distinct discrete strategies, and the table becomes 100 by 100. That is surely too unwieldy to analyze. Algebraic formulas with prices as continuous variables provide a simpler approach, not a more complicated one as some readers might fear. We develop this "Algebra is our friend" method in Chapter 5.

		DEFENSE		
		Run	Pass	Blitz
OFFENSE	Run	2	5	13
	Short Pass	6	5.6	10.5
	Medium Pass	6	4.5	1
	Long Pass	10	3	−2

FIGURE 4.2 Representing a Zero-Sum Simultaneous-Move Game in a Table

In zero-sum games, we can simplify the game table by showing the payoffs of just one player, generally the Row player. Those of the Column player are left implicit. Figure 4.2 shows an example of this shorthand notation for a very simplified version of a single play in (American) football. The team on the offense is attempting to move the ball forward to improve its chances of kicking a field goal. It has four possible strategies: a run and one of three different-length passes (short, medium, and long). The defense can adopt one of three strategies to try to keep the offense at bay: a run defense, a pass defense, or a blitz of the quarterback. The game is zero-sum; the offense tries to gain yardage while the defense tries to prevent it from doing so. Suppose we have enough information about the underlying strengths of the two teams to work out the probabilities of completing different plays and to determine the average gain in yardage that could be expected under each combination of strategies. For example, when Offense chooses the Medium Pass and Defense counters with its Pass defense, we estimate Offense's payoff to be 4.5 (yards).[3] Defense's "payoff" is the loss of 4.5 yards, or negative 4.5 yards (–4.5), but this number is not explicitly shown in the table. The other cells similarly show Offense's payoff, with Defense's payoff implicit and equal to the negative of whatever Offense receives.

2 NASH EQUILIBRIUM

To analyze simultaneous games, we need to consider how players choose their actions. Return to the game table in Figure 4.1. Focus on one specific outcome—

[3]Here is how the payoffs for this case were constructed. When Offense chooses the Medium Pass and Defense counters with its Pass defense, our estimate is that with probability 50% the pass will be completed for a gain of 15 yards, with probability 40% the pass will fall incomplete (0 yards), and with probability 10% the pass will be intercepted with a loss of 30 yards; this makes an average of $0.5 \times 15 + 0.4 \times 0 + 0.1 \times (-30) = 4.5$ yards. The numbers in the table were constructed by a small panel of expert neighbors and friends convened by Dixit on one fall Sunday afternoon. They received a liquid consultancy fee.

namely, the one where Row chooses Low and Column chooses Middle; payoffs there are 5 to Row and 4 to Column. Each player wants to pick an action that yields her the highest payoff, and in this outcome each indeed makes such a choice, given what her opponent chooses. Given that Row is choosing Low, can Column do any better by choosing something other than Middle? No, because Left would give her the payoff 2, and Right would give her 3, neither of which is better than the 4 she gets from Middle. Thus Middle is Column's **best response** to Row's choice of Low. Conversely, given that Column is choosing Middle, can Row do better by choosing something other than Low? Again no, because the payoffs from switching to Top (2), High (3), or Bottom (4) would all be no better than what Row gets with Low (5). Thus Low is Row's best response to Column's choice of Middle.

The two choices, Low for Row and Middle for Column, have the property that each is the chooser's best response to the other's action. If they were making these choices, neither would want to switch to anything different *on her own*. By the definition of a noncooperative game, the players are making their choices independently; therefore such unilateral changes are all that each player can contemplate. Because neither wants to make such a change, it is natural to call this state of affairs an equilibrium. This is exactly the concept of Nash equilibrium.

To state it a little more formally, a **Nash equilibrium**[4] in a game is a list of strategies, one for each player, such that no player can get a better payoff by switching to some other strategy that is available to her while all the other players adhere to the strategies specified for them in the list.

A. Some Further Explanation of the Concept of Nash Equilibrium

To understand the concept of Nash equilibrium better, we take another look at the game in Figure 4.1. Consider now a cell other than (Low, Middle)—say, the one where Row chooses High and Column chooses Left. Can this be a Nash equilibrium? No, because, if Column is choosing Left, Row does better to choose Bottom and get the payoff 5 rather than to choose High, which gives her only 4. Similarly, (Bottom, Left) is not a Nash equilibrium, because Column can do better by switching to Right, thereby improving her payoff from 6 to 7.

[4]This concept is named for the mathematician and economist John Nash, who developed it in his doctoral dissertation at Princeton in 1949. Nash also proposed a solution to cooperative games, which we consider in Chapter 18. He shared the 1994 Nobel Prize in economics with two other game theorists, Reinhard Selten and John Harsanyi; we will treat some aspects of their work in Chapters 9, 10, and 14. Sylvia Nasar's biography of Nash, *A Beautiful Mind* (New York: Simon & Schuster, 1998), was the (loose) basis for a movie starring Russell Crowe. Unfortunately, the movie's attempt to explain the concept of Nash equilibrium fails. We explain this failure in Exercise S12 of this chapter and in Exercise S14 of Chapter 8.

		COLUMN		
		Left	Middle	Right
Row	Top	3, 1	2, 3	10, 2
	High	4, 5	3, 0	6, 4
	Low	2, 2	5, 4	12, 3
	Bottom	5, 6	5, 5	9, 7

FIGURE 4.3 Variation on Game of Figure 4.1 with a Tie in Payoffs

The definition of Nash equilibrium does not require equilibrium choices to be strictly better than other available choices. Figure 4.3 is the same as Figure 4.1 except that Row's payoff from (Bottom, Middle) is changed to 5, the same as that from (Low, Middle). It is still true that, given Column's choice of Middle, Row *could not do any better* than she does when choosing Low. So neither player has a reason to change her action when the outcome is (Low, Middle), and that qualifies it for a Nash equilibrium.[5]

More important, a Nash equilibrium does not have to be jointly best for the players. In Figure 4.1, the strategy pair (Bottom, Right) gives payoffs (9, 7), which are better for both players than the (5, 4) of the Nash equilibrium. However, playing independently, they cannot sustain (Bottom, Right). Given that Column plays Right, Row would want to deviate from Bottom to Low and get 12 instead of 9. Getting the jointly better payoffs of (9, 7) would require cooperative action that made such "cheating" impossible. We examine this type of behavior later in this chapter and in more detail in Chapter 11. For now, we merely point out the fact that a Nash equilibrium may not be in the joint interests of the players.

To reinforce the concept of Nash equilibrium, look at the football game of Figure 4.2. If the Defense is choosing the Pass defense, then the best choice for the Offense is Short Pass (payoff of 5.6 versus 5, 4.5, or 3). Conversely, if the Offense is choosing the Short Pass, then the Defense's best choice is the Pass defense—it holds the Offense down to 5.6 yards, whereas the Run defense and the Blitz would be expected to concede 6 and 10.5 yards, respectively. (Remember that the entries in each cell of a zero-sum game are the Row player's payoffs; therefore the best choice for the Column player is the one that yields the smallest number, not the largest.) In this game, the strategy combination (Short Pass, Pass defense) is a Nash equilibrium, and the resulting payoff to the Offense is 5.6 yards.

[5]But note that (Bottom, Middle) with the payoffs of (5, 5) is not itself a Nash equilibrium. If Row was choosing Bottom, Column's own best choice would not be Middle; she could do better by choosing Right. In fact, you can check all the other cells in the table to verify that none of them can be a Nash equilibrium.

How does one find Nash equilibria in games? One can always check every cell to see if the strategies that generate it satisfy the definition of a Nash equilibrium. Such **cell-by-cell-inspection,** or **enumeration,** is foolproof but tedious and unmanageable except in simple games or unless one is using a good computer program for finding equilibria. Luckily, there are many other methods, applicable to special types of games, that not only find Nash equilibria more quickly when they apply, but also give us a better understanding of the process of thinking by which beliefs and then choices are formed. We develop several such methods in later sections.

B. Nash Equilibrium As a System of Beliefs and Choices

Before we proceed with further study and use of the Nash equilibrium concept, we should try to clarify something that may have bothered some of you. We said that in a Nash equilibrium each player chooses her "best response" to the other's choice. But the two choices are made simultaneously. How can one *respond* to something that has not yet happened, at least when one does not *know* what has happened?

People play simultaneous-move games all the time and do make choices. To do so, they must find a substitute for actual knowledge or observation of the others' actions. Players could make blind guesses and hope that they turn out to be inspired ones, but luckily there are more systematic ways to try to figure out what the others are doing. One method is experience and observation—if the players play this game or similar games with similar players all the time, they may develop a pretty good idea of what the others do. Then choices that are not best will be unlikely to persist for long. Another method is the logical process of thinking through the others' thinking. You put yourself in the position of other players and think what they are thinking, which of course includes their putting themselves in your position and thinking what you are thinking. The logic seems circular, but there are several ways of breaking into the circle, and we demonstrate these ways by using specific examples in the sections that follow. Nash equilibrium can be thought of as a culmination of this process of thinking about thinking, where each player has correctly figured out the others' choice.

Whether by observation or logical deduction or some other method, you, the game player, acquire some notion of what the others are choosing in simultaneous-move games. It is not easy to find a word to describe the process or its outcome. It is not anticipation, nor is it forecasting, because the others' actions do not lie in the future but occur simultaneously with your own. The word most frequently used by game theorists is **belief.** This word is not perfect either, because it seems to connote more confidence or certainty than is intended; in fact, in Chapters 7 and 8, we allow for the possibility that beliefs are held with some uncertainty. But for lack of a better word, it will have to suffice.

This concept of belief also relates to our discussion of uncertainty in Chapter 2, Section 2.D. There we introduced the concept of strategic uncertainty. Even when all the rules of a game—the strategies available to all players and the payoffs for each as functions of the strategies of all—are known without any uncertainty external to the game, such as weather, each player may be uncertain about what actions the others are taking at the same time. Similarly, if past actions are not observable, each player may be uncertain about what actions the others took in the past. How can players choose in the face of this strategic uncertainty? They must form some subjective views or estimates about the others' actions. That is exactly what the notion of belief captures.

Now think of Nash equilibrium in this light. We defined it as a configuration of strategies such that each player's strategy is her best response to that of the others. If she does not know the actual choices of the others but has beliefs about them, in Nash equilibrium those beliefs would have to be correct—the others' actual actions should be just what you believe them to be. Thus we can define Nash equilibrium in an alternative and equivalent way: it is a set of strategies, one for each player, such that (1) each player has correct beliefs about the strategies of the others and (2) the strategy of each is the best for herself, given her beliefs about the strategies of the others.[6]

This way of thinking about Nash equilibrium has two advantages. First, the concept of "best response" is no longer logically flawed. Each player is choosing her best response, not to the as yet unobserved actions of the others, but only to her own already formed beliefs about their actions. Second, in Chapters 7 and 8, where we allow mixed strategies, the randomness in one player's strategy may be better interpreted as uncertainty in the other players' beliefs about this player's action. For now, we proceed by using both interpretations of Nash equilibrium in parallel.

You might think that formation of correct beliefs and calculation of best responses is too daunting a task for mere humans. We discuss some criticisms of this kind, as well as empirical and experimental evidence concerning Nash equilibrium, in Chapter 5 for pure strategies and Chapters 7 and 8 for mixed strategies. For now, we simply say that the proof of the pudding is in the eating. We develop and illustrate the Nash equilibrium concept by applying it. We hope that seeing it in use will prove a better way to understand its strengths and drawbacks than would an abstract discussion at this point.

[6]In this chapter we consider only Nash equilibria in pure strategies, namely the ones initially listed in the specification of the game, and not mixtures of two or more of them. Therefore in such an equilibrium, each player has certainty about the actions of the others; strategic uncertainty is removed. When we consider mixed strategy equilibria in Chapters 7 and 8, the strategic uncertainty for each player will consist of the probabilities with which the various strategies are played in the other players' equilibrium mixtures.

3 DOMINANCE

Some games have a special property that one strategy is uniformly better than or worse than another. When this is the case, it provides one way in which the search for Nash equilibrium and its interpretation can be simplified.

The well-known game of the **prisoners' dilemma** illustrates this concept well. Consider a story line of the type that appears regularly in the television program *Law and Order.* Suppose that a husband and wife have been arrested under the suspicion that they were conspirators in the murder of a young woman. Detectives Green and Lupo place the suspects in separate detention rooms and interrogate them one at a time. There is little concrete evidence linking the pair to the murder, although there is some evidence that they were involved in kidnapping the victim. The detectives explain to each suspect that they are both looking at jail time for the kidnapping charge, probably 3 years, even if there is no confession from either of them. In addition, the husband and wife are told individually that the detectives "know" what happened and "know" how one had been coerced by the other to participate in the crime; it is implied that jail time for a solitary confessor will be significantly reduced if the whole story is committed to paper. (In a scene common to many similar programs, a yellow legal pad and a pencil are produced and placed on the table at this point.) Finally, they are told that, if both confess, jail terms could be negotiated down but not as much as they would be if there were one confession and one denial.

Both husband and wife are then players in a two-person, simultaneous-move game in which each has to choose between confessing and not confessing to the crime of murder. They both know that no confession leaves them each with a 3-year jail sentence for involvement with the kidnapping. They also know that, if one of them confesses, he or she will get a short sentence of 1 year for cooperating with the police, while the other will go to jail for a minimum of 25 years. If both confess, they figure that they can negotiate for jail terms of 10 years each.

The choices and outcomes for this game are summarized by the game table in Figure 4.4. The strategies Confess and Deny can also be called Defect and Cooperate to capture their roles in the relationship between the *two players;* thus Defect

		WIFE	
		Confess (Defect)	Deny (Cooperate)
HUSBAND	Confess (Defect)	10 yr, 10 yr	1 yr, 25 yr
	Deny (Cooperate)	25 yr, 1 yr	3 yr, 3 yr

FIGURE 4.4 Prisoners' Dilemma

means to defect from any tacit arrangement with the spouse, and Cooperate means to take the action that helps the spouse (not cooperate with the cops).

Payoffs here are the lengths of the jail sentences associated with each outcome, so low numbers are better for each player. In that sense, this example differs from those of most of the games that we analyze, in which large payoffs are good rather than bad. We take this opportunity to alert you that "large is good" is not always true. When payoff numbers indicate players' rankings of outcomes, people often use 1 for the best alternative and successively higher numbers for successively worse ones. Also, in the table for a zero-sum game that shows only one player's bigger-is-better payoffs, smaller numbers are better for the other. In the prisoners' dilemma here, smaller numbers are better for both. Thus, if you ever write a payoff table where large numbers are bad, you should alert the reader by pointing it out clearly. And when reading someone else's example, be aware of the possibility.

Now consider the prisoners' dilemma game in Figure 4.4 from the husband's perspective. He has to think about what the wife will choose. Suppose he believes that she will confess. Then his best choice is to confess; he gets a sentence of only 10 years, when denial would have meant 25 years. What if he believes the wife will deny? Again, his own best choice is to confess; he gets only 1 year instead of the 3 that his own denial would bring in this case. Thus, in this special game, Confess is better than Deny for the husband *regardless of his belief about the wife's choice*. We say that, for the husband, the strategy Confess is a **dominant strategy** or that the strategy Deny is a **dominated strategy.** Equivalently, we could say that the strategy Confess *dominates* the strategy Deny or that the strategy Deny is *dominated* by the strategy Confess.

If an action is clearly best for a player, no matter what the others might be doing, then there is compelling reason to think that a rational player would choose it. And if an action is clearly bad for a player, no matter what the others might be doing, then there is equally compelling reason to think that a rational player would avoid it. Therefore dominance, when it exists, provides a compelling basis for the theory of solutions to simultaneous-move games.

A. Both Players Have Dominant Strategies

In the preceding prisoners' dilemma, dominance should lead the husband to choose Confess. Exactly the same logic applies to the wife's choice. Her own strategy Confess dominates her own strategy Deny; so she also should choose Confess. Therefore (Confess, Confess) is the outcome predicted for this game. Note that it is a Nash equilibrium. (In fact it is the only Nash equilibrium.) Each player is choosing his or her own best strategy.

In this special game, the best choice for each is independent of whether their beliefs about the other are correct—that is the meaning of dominance—but

each, attributing to the other the same rationality as he or she practices, should be able to form correct beliefs. And the actual action of each is the best response to the actual action of the other. Note that the fact that Confess dominates Deny for both players is completely independent of whether they are actually guilty, as in many episodes of *Law and Order*, or are being framed, as happened in the movie *LA Confidential*. It only depends on the pattern of payoffs dictated by the various sentence lengths.

Any game with the same general payoff pattern as that illustrated in Figure 4.4 is given the generic label "prisoners' dilemma." More specifically, a prisoners' dilemma has three essential features. First, each player has two strategies: to cooperate with one's rival (deny any involvement in the crime, in our example) or to defect from cooperation (confess to the crime, here). Second, each player also has a dominant strategy (to confess or to defect from cooperation). Finally, the dominance solution equilibrium is worse for both players than the nonequilibrium situation in which each plays the dominated strategy (to cooperate with rivals).

Games of this type are particularly important in the study of game theory for two reasons. The first is that the payoff structure associated with the prisoners' dilemma arises in many quite varied strategic situations in economic, social, political, and even biological competitions. This wide-ranging applicability makes it an important game to study and to understand from a strategic standpoint. The whole of Chapter 11 and sections in several other chapters deal with its study.

The second reason that prisoners' dilemma games are integral to any discussion of games of strategy is the somewhat curious nature of the equilibrium outcome achieved in such games. Both players follow conventional wisdom in choosing their dominant strategies, but the resulting equilibrium outcome yields them payoffs that are lower than they could have achieved if they had each chosen their dominated strategies. Thus the equilibrium outcome in the prisoners' dilemma is actually a bad outcome for the players. They could find another outcome that they both prefer to the equilibrium outcome; the problem is how to guarantee that someone will not cheat. This particular feature of the prisoners' dilemma has received considerable attention from game theorists who have asked an obvious question: What can players in a prisoners' dilemma do to achieve the better outcome? We leave this question to the reader momentarily, as we continue the discussion of simultaneous games, but return to it in detail in Chapter 11.

B. One Player Has a Dominant Strategy

When a rational player has a dominant strategy, she will use it, and the other player can safely believe this. In the prisoners' dilemma, it applied to both players. In some other games, it applies only to one of them. If you are playing in a game in which you do not have a dominant strategy but your opponent does, you can assume that she will use her dominant strategy and so you can choose your equilibrium action (your best response) accordingly.

We illustrate this case by using a game frequently played between the Congress, which is responsible for fiscal policy (taxes and government expenditures), and the Federal Reserve (Fed), which is in charge of monetary policy (primarily, interest rates).[7] In a version that simplifies the game to its essential features, the Congress's fiscal policy can have either a balanced budget or a deficit, and the Fed can set interest rates either high or low. In reality, the game is not clearly simultaneous; nor is who has the first move obvious if choices are sequential. We consider the simultaneous-move version here, and in Chapter 6 study how the outcomes differ for different rules of the game.

Almost everyone wants lower taxes. But there is no shortage of good claims on government funds: defense, education, health care, and so on. There are also various politically powerful special interest groups—including farmers and industries hurt by foreign competition—who want government subsidies. Therefore the Congress is under constant pressure both to lower taxes and to increase spending. But such behavior runs the budget into deficit, which can lead to higher inflation. The Fed's primary task is to prevent inflation. However, it also faces political pressure for lower interest rates from many important groups, especially homeowners who benefit from lower mortgage rates. Lower interest rates lead to higher demand for automobiles, housing, and capital investment by firms, and all this demand can cause higher inflation. The Fed is generally happy to lower interest rates but only so long as inflation is not a threat. And there is less threat of inflation when the government's budget is in balance. With all this in mind, we construct the payoff matrix for this game in Figure 4.5.

Congress likes best (payoff 4) the outcome with a budget deficit and low interest rates. This pleases all the immediate political constituents. It may entail trouble for the future, but political time horizons are short. For the same reason,

		FEDERAL RESERVE	
		Low interest rates	High interest rates
CONGRESS	Budget balance	3, 4	1, 3
	Budget deficit	4, 1	2, 2

FIGURE 4.5 Game of Fiscal and Monetary Policies

[7]Similar games are played in many other countries with central banks that have operational independence in the choice of monetary policy. Fiscal policies may be chosen by different political entities—the executive or the legislature—in different countries.

Congress likes worst (payoff 1) the outcome with a balanced budget and high interest rates. Of the other two outcomes, it prefers (payoff 3) the outcome with a balanced budget and low interest rates; this outcome pleases the important home-owning middle classes and, with low interest rates, less expenditure is needed to service the government debt, so the balanced budget still has room for many other items of expenditure or for tax cuts.

The Fed likes worst (payoff 1) the outcome with a budget deficit and low interest rates, because this combination is the most inflationary. It likes best (payoff 4) the outcome with a balanced budget and low interest rates, because this combination can sustain a high level of economic activity without much risk of inflation. Comparing the other two outcomes with high interest rates, the Fed prefers the one with budget balance because it reduces the risk of inflation.

We look now for dominant strategies in this game. The Fed does better by choosing low interest rates if it believes that the Congress is opting for budget balance (Fed's payoff 4 rather than 3), but it does better choosing high interest rates if it believes that the Congress is choosing to run a budget deficit (Fed's payoff 2 rather than 1). The Fed, then, does not have a dominant strategy. But the Congress does. If the Congress believes that the Fed is choosing low interest rates, it does better for itself by choosing a budget deficit rather than budget balance (Congress's payoff 4 instead of 3). If the Congress believes that the Fed is choosing high interest rates, again it does better for itself by choosing a budget deficit rather than budget balance (Congress's payoff 2 instead of 1). Choosing to run a budget deficit is then Congress's dominant strategy.

The choice for the Congress in the game is now clear. No matter what it believes the Fed is doing, the Congress will choose to run a budget deficit. The Fed can now take this choice into account when making its own decision. The Fed should believe that the Congress will choose its dominant strategy (budget deficit) and choose the best strategy for itself, given this belief. That means that the Fed should choose high interest rates.

In this outcome, each side gets payoff 2. But an inspection of Figure 4.5 shows that, just as in the prisoners' dilemma, there is another outcome—namely, a balanced budget and low interest rates—that can give both players higher payoffs—namely, 3 for the Congress and 4 for the Fed. Why is that outcome not achievable as an equilibrium? The problem is that Congress would be tempted to deviate from its stated strategy and sneakily run a budget deficit. The Fed, knowing this temptation and that it would then get its worst outcome (payoff 1), deviates also to its high interest rate strategy. In Chapters 6 and 10, we consider how the two sides can get around this difficulty to achieve their mutually preferred outcome. But we should note that, in most countries and at many times, the two policy authorities are indeed stuck in the bad outcome; the fiscal policy is too loose, and the monetary policy has to be tightened to keep inflation down.

C. Successive Elimination of Dominated Strategies

The games considered so far have had only two pure strategies available to each player. In such games, if one strategy is dominant, the other is dominated; so choosing the dominant strategy is equivalent to eliminating the dominated one. In larger games, some of a player's strategies may be dominated even though no single strategy dominates all of the others. If players find themselves in a game of this type, they may be able to reach an equilibrium by removing dominated strategies from consideration as possible choices. Removing dominated strategies reduces the size of the game, and then the "new" game may have another dominated strategy, for the same player or for her opponent, that can also be removed. Or the "new" game may even have a dominant strategy for one of the players. **Successive** or **iterated elimination of dominated strategies** uses this process of removal of dominated strategies and reduction in the size of a game until no further reductions can be made. If this process ends in a unique outcome, then the game is said to be **dominance solvable;** that outcome is the Nash equilibrium of the game, and the strategies that yield it are the equilibrium strategies for each player.

We can use the game of Figure 4.1 to provide an example of this process. Consider first Row's strategies. If any one of Row's strategies always provides worse payoffs for Row than another of her strategies, then that strategy is dominated and can be eliminated from consideration for Row's equilibrium choice. Here, the only dominated strategy for Row is High, which is dominated by Bottom; if Column plays Left, Row gets 5 from Bottom and only 4 from High; if Column plays Middle, Row gets 4 from Bottom and only 3 from High; and, if Column plays Right, Row gets 9 from Bottom and only 6 from High. So we can eliminate High. We now turn to Column's choices to see if any of them can be eliminated. We find that Column's Left is now dominated by Right (with similar reasoning, $1 < 2$, $2 < 3$, and $6 < 7$). Note that we could not say this before Row's High was eliminated; against Row's High, Column would get 5 from Left but only 4 from Right. Thus the first step of eliminating Row's High makes possible the second step of eliminating Column's Left. Then, within the remaining set of strategies (Top, Low, and Bottom for Row, and Middle and Right for Column), Row's Top and Bottom are both dominated by his Low. When Row is left with only Low, Column chooses his best response—namely, Middle.

The game is thus dominance solvable, and the outcome is (Low, Middle) with payoffs (5, 4). We identified this outcome as a Nash equilibrium when we first illustrated that concept by using this game. Now we see in better detail the thought process of the players that leads to the formation of correct beliefs. A rational Row will not choose High. A rational Column will recognize this, and thinking about how her various strategies perform for her against Row's remaining

strategies, will not choose Left. In turn, Row will recognize this, and therefore will not choose either Top or Bottom. Finally, Column will see through all this, and choose Middle.

Other games may not be dominance solvable, or successive elimination of dominated strategies may not yield a unique outcome. Even in such cases, some elimination may reduce the size of the game and make it easier to solve by using one or more of the techniques described in the following sections. Thus eliminating dominated strategies can be a useful step toward solving a large simultaneous-play game, even when their elimination does not completely solve the game.

Thus far in our consideration of iterated elimination of dominated strategies, all the payoff comparisons have been unambiguous. What if there are some ties? Consider the variation on the preceding game that is shown in Figure 4.3. In that version of the game, High (for Row) and Left (for Column) also are eliminated. And, at the next step, Low still dominates Top. But the dominance of Low over Bottom is now less clear-cut. The two strategies give Row equal payoffs when played against Column's Middle, although Low does give Row a higher payoff than Bottom when played against Column's Right. We say that, from Row's perspective at this point, Low *weakly* dominates Bottom. In contrast, Low *strictly* dominates Top, because it gives strictly higher payoffs than does Top when played against both of Column's strategies, Middle and Right, under consideration at this point.

We give a more precise definition of the distinction between strict and weak dominance in the Appendix to this chapter. Here, though, we provide a word of warning. Successive elimination of weakly dominated strategies can get rid of some Nash equilibria.

Consider the game illustrated in Figure 4.6. For Row, Up is weakly dominated by Down; if Column plays Left, then Row gets a better payoff by playing Down than by playing Up, and, if Column plays Right, then Row gets the same payoff from her two strategies. Similarly, for Column, Right weakly dominates Left. Dominance solvability then tells us that (Down, Right) is a Nash equilibrium.

		COLUMN	
		Left	Right
ROW	Up	0, 0	1, 1
	Down	1, 1	1, 1

FIGURE 4.6 Elimination of Weakly Dominated Strategies

That is true, but (Down, Left) and (Up, Right) also are Nash equilibria. Consider (Down, Left). When Row is playing Down, Column cannot improve her payoff by switching to Right, and, when Column is playing Left, Row's best response is clearly to play Down. A similar reasoning verifies that (Up, Right) also is a Nash equilibrium.

Therefore, if you use weak dominance to eliminate some strategies, it is a good idea to make a quick cell-by-cell check to see if you have missed any other equilibria. The iterated dominance solution seems to be a reasonable outcome to predict as the likely Nash equilibrium of this simultaneous-play game, but it is also important to consider the significance of multiple equilibria as well as of the other equilibria themselves. We address these issues in later chapters, taking up a discussion of multiple equilibria in Chapter 5 and the interconnections between sequential- and simultaneous-move games in Chapter 6.

4 BEST-RESPONSE ANALYSIS

Many simultaneous-move games have no dominant strategies and no dominated strategies. Others may have one or several dominated strategies, but iterated elimination of dominated strategies will not yield a unique outcome. In such cases, we need a next step in the process of finding a solution to the game. We are still looking for a Nash equilibrium in which every player does the best she can, given the actions of the other player(s), but we must now rely on subtler strategic thinking than the simple elimination of dominated strategies requires.

Here, we develop another systematic method for finding Nash equilibria that will prove very useful in later analysis. We begin without imposing a requirement of correctness of beliefs. We take each player's perspective in turn and ask the following question: For each of the choices that the other player(s) might be making, what is the best choice for this player? Thus we find the best responses

		COLUMN		
		Left	Middle	Right
ROW	Top	3, 1	2, ③	10, 2
	High	4, ⑤	3, 0	6, 4
	Low	2, 2	⑤, ④	⑫, 3
	Bottom	⑤, 6	4, 5	9, ⑦

FIGURE 4.7 Best Response Analysis

of each player to all available strategies of the others. In mathematical terms, we find each player's best-response strategy, depending on, or as a function of, the other players' available strategies.

Return to the game of Figure 4.1, reproduced as Figure 4.7, and consider Row first. If Column chooses Left, Row's best response is Bottom, yielding 5. We show this best response by circling that payoff in the game table. If Column chooses Middle, Row's best response is Low (also yielding 5). And, if Column chooses Right, Row's best choice is again Low (now yielding 12). As before, we show Row's best choices by circling the appropriate payoffs. Similarly, Column's best responses are shown by circling her payoffs 3 (Middle as best response to Row's Top), 5 (Left to High), 4 (Middle to Low), and 7 (Right to Bottom).[6] We see that one cell—namely, (Low, Middle)—has both its payoffs circled. Therefore the strategies Low for Row and Middle for Column are simultaneously best responses to each other. We have found the Nash equilibrium of this game. (Again.)

Best-response analysis is a comprehensive way of locating all possible Nash equilibria of a game. You should improve your understanding of it by trying it out on the other games that have been used in this chapter. The cases of dominance are of particular interest. If Row has a dominant strategy, that same strategy is her best response to all of Column's strategies; therefore her best responses are all lined up horizontally in the same row. Similarly, if Column has a dominant strategy, her best responses are all lined up vertically in the same column. You should see for yourself how the Nash equilibria of the preceding prisoners' dilemma and Congress–Fed games emerge from such a drawing.

There will be some games for which best-response analysis does not find a Nash equilibrium, just as dominance solvability sometimes fails. But in this case we can say something more specific than can be said when dominance fails. When best-response analysis of a discrete strategy game does not find a Nash equilibrium, then the game has no equilibrium in pure strategies. We address games of this type in Section 8 of this chapter. In Chapter 5, we extend best-response analysis to games where the players' strategies are continuous variables—for example, prices or advertising expenditures. There, we construct best-response *curves* to help us find Nash equilibria, and we see that such games are less likely—by virtue of the continuity of strategy choices—to have no equilibrium.

[6] Alternatively and equivalently, one could mark in some way the choices that are *not* made. For example, in Figure 4.3, Row will not choose Top, High, or Bottom as responses to Column's Right; one could show this by drawing slashes through Row's payoffs in these cases, respectively, 10, 6, and 9. When this is done for all strategies of both players, (Low, Middle) has both of its payoffs unslashed; it is then the Nash equilibrium of the game. The alternatives of circling choices that are made and slashing choices that are not made stand in a conceptually similar relation to each other, as do the alternatives of showing chosen branches by arrows and pruning unchosen branches for sequential-move games. We prefer the first alternative in each case, because the resulting picture looks cleaner and tells the story better.

5 THE MINIMAX METHOD FOR ZERO-SUM GAMES

For zero-sum games, an alternative to best-response analysis works by using the special logic of strict conflict that exists in such games. This approach, the **minimax method,** works only for zero-sum games and relies on a thought process that accounts for the fact that outcomes that are good for one player are, by definition, bad for the other. In this method, each player is assumed to choose her strategy by thinking: "Would this be the best choice for me, even if the other player found out that I was playing it?" She must then consider her opponent's best response to her chosen strategy. But in a zero-sum game, that best response is the worst one for her. In other words, each player believes that her opponent will choose an action that yields her the worst possible consequences of each of her own actions. Then acting on those beliefs she should choose the action that leads to the least-bad outcome.

This logic may seem extremely pessimistic, but it still relies on a type of best-response calculation and it is appropriate for finding the equilibrium of a zero-sum game. In equilibrium, each player is choosing her own best response, given her beliefs about what the other will do. In anticipating such best responses, each player will expect to receive the worst payoff associated with each action and will choose her own action accordingly. She is thus choosing her best payoff from among the set of worst payoffs.

Suppose the payoff table shows the row player's payoffs, and Row wants the outcome to be a cell with as high a number as possible. Then Column wants the outcome to be a cell with as low a number as possible. Using the pessimistic logic just described, Row figures that, for each of her rows, Column will choose the column with the lowest number in that row. Therefore Row should choose the row that gives her the highest among these lowest numbers, or the maximum among the minima—the **maximin** for short. Similarly, Column reckons that, for each of her columns, Row will choose the row with the largest number in that column. Then Column should choose the column with the smallest number among these largest ones, or the minimum among the maxima—the **minimax.** If Row's maximin value and Column's minimax value are in the same cell of the game table, then that outcome is a Nash equilibrium of the zero-sum game. This method of finding equilibria in zero-sum games should be called the maximin-minimax method, but it is called simply the *minimax method* for short. It will lead you to a Nash equilibrium in pure strategies if one exists.

To illustrate the minimax method, we use the football example of Figure 4.2. We already know the Nash equilibrium for that game, but now we obtain it by using the minimax method. We reproduce the game table in Figure 4.8, adding information that pertains to the minimax argument.

		DEFENSE			
		Run	Pass	Blitz	
OFFENSE	Run	2	5	13	min = 2
	Short Pass	6	5.6	10.5	min = 5.6
	Medium Pass	6	4.5	1	min = 1
	Long Pass	10	3	−2	min = −2
		max = 10	max = 5.6	max = 13	

FIGURE 4.8 The Minimax Method

Begin by finding the lowest number in each row (the offense's worst payoff from each strategy) and the highest number in each column (the defense's worst payoff from each strategy). The offense's worst payoff from Run is 2; its worst payoff from Short Pass, 5.6; its worst payoff from Medium Pass, 1; and its worst payoff from Long Pass, −2. We write the minimum for each row at the far right of that row. The defense's worst payoff from Run is 10; its worst payoff from Pass, 5.6; and its worst payoff from Blitz, 13. We write the maximum for each column at the bottom of that column.

The next step is to find the best of each player's worst possible outcomes, the largest row minimum and the smallest column maximum. The largest of the row minima is 5.6; so the offense can ensure itself a gain of 5.6 yards by playing the Short Pass; this is its maximin. The lowest of the column maxima is 5.6; so the defense can be sure of holding the offense down to a gain of 5.6 yards by deploying its Pass defense. This is the defense's minimax.

Looking at these two strategy choices, we see that the maximin and minimax values are found in the same cell of the game table. Thus the offense's maximin strategy is its best response to the defense's minimax and vice versa; we have found the Nash equilibrium of this game. That equilibrium entails the offense attempting a Short Pass while the defense defends against a Pass. A total of 5.6 yards will be gained by the offense (and given up by the defense).

The minimax method may fail to find an equilibrium in some zero-sum games. If so, then our conclusion is similar to that when best-response analysis fails: the game has no Nash equilibrium in pure strategies. We address this matter later in this chapter and examine mixed strategy equilibria in Chapters 7 and 8. And, to repeat, the minimax method cannot be applied to non-zero-sum games. In such games, your opponent's best is not necessarily your worst. Therefore the pessimistic assumption that leads you to choose the strategy that makes your minimum payoff as large as possible may not be your best strategy.

6 THREE PLAYERS

So far, we have analyzed only games between two players. All of the methods of analysis that have been discussed, however, can be used to find the pure-strategy Nash equilibria of any simultaneous-play game among any number of players. When a game is played by more than two players, each of whom has a relatively small number of pure strategies, the analysis can be done with a game table, as we did in the first five sections of this chapter.

In Chapter 3, we described a game among three players, each of whom had two pure strategies. The three players, Emily, Nina, and Talia, had to choose whether to contribute toward the creation of a flower garden for their small street. We assumed there that the garden when all three contributed was no better than when only two contributed and that a garden with just one contributor was so sparse that it was as bad as no garden at all. Now, let us suppose instead that the three players make their choices simultaneously and that there is a somewhat richer variety of possible outcomes and payoffs. In particular, the size and splendor of the garden will now differ according to the exact number of contributors; three contributors will produce the largest and best garden, two contributors will produce a medium garden, and one contributor will produce a small garden.

Suppose Emily is contemplating the possible outcomes of the street-garden game. There are six possibilities to consider. Emily can choose either to contribute or not to contribute when both Nina and Talia contribute or when neither of them contributes or when just one of them contributes. From her perspective, the best possible outcome, with a rating of 6, would be to take advantage of her good-hearted neighbors and to have both Nina and Talia contribute while she does not. Emily could then enjoy a medium-sized garden without putting up her own hard-earned cash. If both of the others contribute and Emily also contributes, she gets to enjoy a large, very splendid garden but at the cost of her own contribution; she rates this outcome second-best, or 5.

At the other end of the spectrum are the outcomes that arise when neither Nina nor Talia contributes to the garden. If that is the case, Emily would again prefer not to contribute, because she would foot the bill for a public garden that everyone could enjoy; she would rather have the flowers in her own yard. Thus, when neither other player is contributing, Emily ranks the outcome in which she contributes as a 1 and the outcome in which she does not as a 2.

In between these cases are the situations in which either Nina or Talia contributes to the flower garden but not both. When one of them contributes, Emily knows that she can enjoy a small garden without contributing; she also feels that the cost of her contribution outweighs the increase in benefit that she gets

TALIA chooses:		Contribute		Don't Contribute	
		NINA		NINA	
		Contribute	Don't	Contribute	Don't
EMILY	Contribute	5, 5, 5	3, 6, 3	3, 3, 6	1, 4, 4
EMILY	Don't	6, 3, 3	4, 4, 1	4, 1, 4	2, 2, 2

FIGURE 4.9 Street-Garden Game

from being able to increase the size of the garden. Thus she ranks the outcome in which she does not contribute, but still enjoys the small garden, as a 4 and the outcome in which she does contribute, to provide a medium garden, as a 3. Because Nina and Talia have the same views as Emily on the costs and benefits of contributions and garden size, each of them orders the different outcomes in the same way—the worst outcome being the one in which each contributes and the other two do not, and so on.

If all three women decide whether to contribute to the garden without knowing what their neighbors will do, we have a three-person simultaneous-move game. To find the Nash equilibrium of the game, we then need a game table. For a three-player game, the table must be three-dimensional and the third player's strategies correspond to the new dimension. The easiest way to add a third dimension to a two-dimensional game table is to add pages. The first page of the table shows payoffs for the third player's first strategy, the second page shows payoffs for the third player's second strategy, and so on.

We show the three-dimensional table for the street-garden game in Figure 4.9. It has two rows for Emily's two strategies, two columns for Nina's two strategies, and two pages for Talia's two strategies. We show the pages side by side so that you can see everything at the same time. In each cell, payoffs are listed for the row player first, the column player second, and the page player third; in this case, the order is Emily, Nina, Talia.

Our first test should be to determine whether there are dominant strategies for any of the players. In one-page game tables, we found this test to be simple; we just compared the outcomes associated with one of a player's strategies with the outcomes associated with another of her strategies. In practice this comparison required, for the row player, a simple check within columns of the single page of the table and vice versa for the column player. Here, we must check in both pages of the table to determine whether any player has a dominant strategy.

For Emily, we compare the two rows of both pages of the table and note that, when Talia contributes, Emily has a dominant strategy not to contribute, and, when Talia does not contribute, Emily also has a dominant strategy not to contribute. Thus the best thing for Emily to do, regardless of what either of the other players does, is not to contribute. Similarly, we see that Nina's dominant strategy—in both pages of the table—is not to contribute. When we check for a dominant strategy for Talia, we have to be a bit more careful. We must compare outcomes that keep Emily's and Nina's behavior constant, checking Talia's payoffs from choosing Contribute versus Don't. That is, we compare cells across pages of the table—the top-left cell in the first page (on the left) with the top-left cell in the second page (on the right), and so on. As for the first two players, this process indicates that Talia also has a dominant strategy not to contribute.

Each player in this game has a dominant strategy, which must therefore be her equilibrium pure strategy. The Nash equilibrium of the street-garden game entails all three players choosing not to contribute to the street garden and getting their second-worst payoffs; the garden is not planted, but no one has to contribute either.

Notice that this game is yet another example of a prisoners' dilemma. There is a unique Nash equilibrium in which all players receive a payoff of 2. Yet there is another outcome in the game—in which all three neighbors contribute to the garden—that for all three players yields higher payoffs of 5. Even though it would be beneficial to each of them for all to pitch in to build the garden, no one has the individual incentive to do so. As a result, gardens of this type are either not planted at all or paid for through tax dollars—because the town government can require its citizens to pay such taxes. In Chapter 12, we will encounter more such dilemmas of collective action and study some methods for resolving them.

The Nash equilibrium of the game can also be found using the cell-by-cell inspection method. For example, consider another cell in Figure 4.9—say, the one where Emily and Nina contribute but Talia does not, with the payoffs (3, 3, 6). When Emily considers changing her strategy, as the row player she can change only the row position of the game's outcome. Emily can move the outcome only from a given cell in a given row, column, and page to another cell in a different row but the same column and same page of the table. If she does that in this instance, she improves her payoff from 3 to 4. Similarly, Nina can change only the column position of the outcome, moving it to a cell in another column but in the same row and same page of the table. Doing so improves Nina's payoff from 3 to 4. Finally, Talia can change only the page position of the game's outcome. She can move the outcome to a different page, but the row and column positions must remain the same. Doing so would worsen Talia's payoff from 6 to 5. Because at least one player can do better by unilaterally changing her strategy, the cell that we examined cannot be the outcome of a Nash equilibrium.

TALIA chooses:

Contribute

		NINA: Contribute	NINA: Don't
EMILY	Contribute	5, 5, 5	3, 6, 3
	Don't	6, 3, 3	4, 4, 1

Don't Contribute

		NINA: Contribute	NINA: Don't
EMILY	Contribute	3, 3, 6	1, 4, 4
	Don't	4, 1, 4	2, 2, 2

FIGURE 4.10 Best-Response Analysis in the Street-Garden Game

We can also use the best-response method, as shown in Figure 4.10, by drawing circles around the best responses, as in Figure 4.7. Because each player has Don't as her dominant strategy, all of Emily's best responses are on her Don't rows, all of Nina's on her Don't columns, and all of Talia's on her Don't page. The cell at the bottom right has all three best responses; therefore it gives us the Nash equilibrium.

7 MULTIPLE EQUILIBRIA IN PURE STRATEGIES

Each of the games considered in preceding sections has had a unique pure-strategy Nash equilibrium. In general, however, games need not have unique Nash equilibria. We illustrate this result by using a class of games that have many applications. As a group, they may be labeled **coordination games.** The players in such games have some (but not always completely) common interests. But, because they act independently (by virtue of the nature of noncooperative games), the coordination of actions needed to achieve a jointly preferred outcome is problematic.

A. Will Harry Meet Sally? Pure Coordination

To illustrate this idea, picture two undergraduates, Harry and Sally, who meet in their college library. They are attracted to each other and would like to continue the conversation but have to go off to their separate classes. They arrange to meet for coffee after the classes are over at 4:30. Sitting separately in class, each realizes that in the excitement they forgot to fix the place to meet. There are two possible choices, Starbucks and Local Latte. Unfortunately, these locations are on opposite sides of the large campus; so it is not possible to try both. And Harry and Sally have not exchanged cell-phone numbers, so they can't send messages. What should each do?

		SALLY	
		Starbucks	Local Latte
HARRY	Starbucks	1, 1	0, 0
	Local Latte	0, 0	1, 1

FIGURE 4.11 Pure Coordination

Figure 4.11 illustrates this situation as a game and shows the payoff matrix. Each player has two choices—Starbucks and Local Latte. The payoffs for each are 1 if they meet and 0 if they do not. Cell-by-cell inspection shows at once that the game has two Nash equilibria, one where both choose Starbucks and the other where both choose Local Latte. It is important for both that they achieve one of the equilibria, but which one is immaterial because the two yield equal payoffs. All that matters is that they coordinate on the same action; it does not matter which action. That is why the game is said to be one of **pure coordination.**

But will they coordinate successfully? Or will they end up in different cafés, each thinking that the other has let him or her down? Alas, that risk exists. Harry might think that Sally will go to Starbucks because she said something about the class to which she was going and that class is on the Starbucks side of the campus. But Sally may have the opposite belief about what Harry will do. When there are multiple Nash equilibria, if the players are to select one successfully, they need some way to coordinate their beliefs or expectations about each other's actions.

The situation is similar to that of the heroes of the "Which tire?" game in Chapter 1, where we labeled the coordination device a **focal point.** In the present context, one of the two cafés may be generally known as the student hangout. But it is not enough that Harry knows this to be the case. He must know that Sally knows, and that she knows that he knows, and so on. In other words, their expectations must *converge* on the focal point. Otherwise Harry might be doubtful about where Sally will go because he does not know what she is thinking about where he will go; and similar doubts may arise at the third or fourth or higher level of thinking about thinking.[9]

[9]Thomas Schelling presented the classic treatment of coordination games and developed the concept of a focal point in his book *The Strategy of Conflict* (Cambridge: Harvard University Press, 1960); see pp. 54–58, 89–118. His explanation of focal points included the results garnered when he posed several questions to his students and colleagues. The best-remembered of these is "Suppose you have arranged to meet someone in New York City on a particular day, but have failed to arrange a specific place or time, and have no way of communicating with the other person. Where will you go and at what time?" Fifty years ago when the question was first posed, the clock at Grand Central Station was the usual focal place; now it might be the observation platform atop the Empire State Building or Times Square. The focal time remains twelve noon.

When one of us (Dixit) posed this question to students in his class, the freshmen generally chose Starbucks and the juniors and seniors generally chose the local café in the campus student center. These responses are understandable—freshmen, who have not been on campus long, focus their expectations on a nationwide chain that is known to everyone, whereas juniors and seniors have acquired the local habits, which they now regard as superior, and expect their peers to believe likewise.

If one café had an orange decor and the other a crimson decor, then in Princeton the former may serve as a focal point because orange is the Princeton color, whereas at Harvard crimson may be focal for the same reason. If one person is a Princeton student and the other a Harvard student, they may fail to meet at all, either because each thinks that his or her color "should" get priority or because each thinks that the other will be inflexible and so tries to accommodate him or her. More generally, whether players in coordination games can find a focal point depends on their having some commonly known point of contact, whether historical, cultural, or linguistic.

B. Will Harry Meet Sally? And Where? Assurance

Now change the game payoffs a little. The behavior of juniors and seniors suggests that our pair may not be quite indifferent about which café they both choose. The coffee may be better at one or the ambiance better at one. Or they may want to choose the one that is not the general student hangout, to avoid the risk of running into former boyfriends or girlfriends. Suppose they both prefer Local Latte; so the payoff of each is 2 when they meet there versus 1 when they meet at Starbucks. The new payoff matrix is shown in Figure 4.12.

Again, there are two Nash equilibria. But in this version of the game, each prefers the equilibrium where both choose Local Latte. Unfortunately, their mere liking of that outcome is not guaranteed to bring it about. First of all (and as always in our analysis), the payoffs have to be common knowledge—both have to know the entire payoff matrix, both have to know that both know, and so on. Such detailed knowledge about the game can arise if the two discussed and

		SALLY	
		Starbucks	Local Latte
HARRY	Starbucks	1, 1	0, 0
	Local Latte	0, 0	2, 2

FIGURE 4.12 Assurance

agreed on the relative merits of the two cafés but simply forgot to decide definitely to meet at Local Latte. Even then, Harry might think that Sally has some other reason for choosing Starbucks, or think that she thinks that he does, and so on. Without genuine **convergence of expectations** about actions, they may choose the worse equilibrium or, worse still, they may fail to coordinate actions and get 0 each.

To repeat, players in the game illustrated in Figure 4.12 can get the preferred equilibrium outcome only if each has enough certainty or assurance that the other is choosing the appropriate action. For this reason, such games are called **assurance games.**[10]

In many real-life situations of this kind, such assurance is easily obtained, given even a small amount of communication between the players. Their interests are perfectly aligned; if one of them says to the other, "I am going to Local Latte," the other has no reason to doubt the truth of this statement and will follow to get the mutually preferred outcome. That is why we had to construct the story with the two students isolated in different classes with no means of communication. If the players' interests conflict, truthful communication becomes more problematic. We examine this problem further when we consider strategic manipulation of information in games in Chapter 9.

In larger groups, communication can be achieved by scheduling meetings or by making announcements. These devices work only if everyone knows that everyone else is paying attention to them, because successful coordination requires the desired equilibrium to be a focal point. The players' expectations must converge on it; everyone should know that everyone knows that . . . everyone is choosing it. Many social institutions and arrangements play this role. Meetings where the participants sit in a circle facing inward ensure that everyone sees everyone else paying attention. Advertisements during the Super Bowl, especially when they are proclaimed in advance as major attractions, ensure each viewer that many others are viewing them also. That makes such ads especially attractive to companies making products that are more desirable for any one buyer when many others are buying them, too; such products include those produced by the computer, telecommunication, and Internet industries.[11]

[10]The classic example of an assurance game usually offered is the stag hunt described by the 18th-century French philosopher Jean-Jacques Rousseau. Several people can successfully hunt a stag, thereby getting a large quantity of meat if they collaborate. If any one of them is sure that all of the others will collaborate, he also stands to benefit by joining the group. But if he is unsure whether the group will be large enough, he will do better to hunt for a smaller animal, a hare, on his own. However, it can be argued that Rousseau believed that each hunter would prefer to go after a hare regardless of what the others were doing, which would make the stag hunt a multiperson prisoners' dilemma, not an assurance game. We discuss this example in the context of collective action in Chapter 12.

[11]Michael Chwe develops this theme in *Rational Ritual: Culture, Coordination, and Common Knowledge* (Princeton: Princeton University Press, 2001).

		SALLY	
		Starbucks	Local Latte
HARRY	Starbucks	2, 1	0, 0
	Local Latte	0, 0	1, 2

FIGURE 4.13 Battle of the Sexes

C. Will Harry Meet Sally? And Where? Battle of the Sexes

Now introduce another complication to the café-choice game. Both players want to meet but prefer different cafés. So Harry might get a payoff of 2 and Sally a payoff of 1 from meeting at Starbucks, and the other way around from meeting at Local Latte. This payoff matrix is shown in Figure 4.13.

This game is called the **battle of the sexes.** The name derives from the story concocted for this payoff structure by game theorists in the sexist 1950s. A husband and wife were supposed to choose between going to a boxing match and a ballet, and (presumably for evolutionary genetic reasons) the husband was supposed to prefer the boxing match and the wife the ballet. The name has stuck and we will keep it, but our example—where either could easily have some non-gender-based reason to prefer either of the two cafés—should make it clear that it has no necessarily sexist connotations.

What will happen in this game? There are still two Nash equilibria. If Harry believes that Sally will choose Starbucks, it is best for him to do likewise, and the other way around. For similar reasons, Local Latte also is a Nash equilibrium. To achieve either of these equilibria and avoid the outcomes where the two go to different cafés, the players need a focal point, or convergence of expectations, exactly as in the pure-coordination and assurance games. But the risk of coordination failure is greater in the battle of the sexes. The players are initially in quite symmetric situations, but each of the two Nash equilibria gives them asymmetric payoffs; and their preferences between the two outcomes are in conflict. Harry prefers the outcome where they meet in Starbucks, and Sally prefers to meet in Local Latte. They must find some way of breaking the symmetry.

In an attempt to achieve his or her preferred equilibrium, each player may try to act tough and follow the strategy leading to the better equilibrium. In Chapter 10, we consider in detail such advance devices, called strategic moves, that players in such games can adopt to try to achieve their preferred outcomes. Or each may try to be nice, leading to the unfortunate situation where Harry goes to Local Latte because he wants to please Sally, only to find that she has chosen to please him and gone to Starbucks, like the couple choosing Christmas presents for each other in O. Henry's short story titled "The Gift of the Magi." Alternatively, if the game is repeated, successful coordination may be negotiated and maintained as an equilibrium. For example, the two can arrange to alternate between the cafés. In Chapter 11, we examine such tacit cooperation in repeated games in the context of a prisoners' dilemma.

D. Will James Meet Dean? Chicken

Our final example in this section is a slightly different kind of coordination game. In this game, the players want to avoid, not choose, actions with the same labels. Further, the consequences of one kind of coordination failure are far more drastic than those of the other kind.

The story comes from a game that was supposedly played by American teenagers in the 1950s. Two teenagers take their cars to opposite ends of Main Street, Middle-of-Nowhere, USA, at midnight and start to drive toward each other. The one who swerves to prevent a collision is the "chicken," and the one who keeps going straight is the winner. If both maintain a straight course, there is a collision in which both cars are damaged and both players injured.[12]

The payoffs for **chicken** depend on how negatively one rates the "bad" outcome—being hurt and damaging your car in this case—against being labeled chicken. As long as words hurt less than crunching metal, a reasonable payoff table for the 1950s version of chicken is found in Figure 4.14. Each player most prefers to win, having the other be chicken, and each least prefers the crash of the two cars. In between these two extremes, it is better to have your rival be chicken with you (to save face) than to be chicken by yourself.

This story has four essential features that define any game of chicken. First, each player has one strategy that is the "tough" strategy and one that is the "weak" strategy. Second, there are two pure-strategy Nash equilibria. These are the outcomes in which exactly one of the players is chicken, or weak. Third, each player strictly prefers that equilibrium in which the other player chooses

[12]A slight variant was made famous by the 1955 James Dean movie *Rebel Without a Cause*. There, two players drove their cars in parallel, very fast, toward a cliff. The first to jump out of his car before it went over the cliff was the chicken. The other, if he left too late, risked going over the cliff in his car to his death. The characters in the film referred to this as a "chicky game." In the mid-1960s, the British philosopher Bertrand Russell and other peace activists used this game as an analogy for the nuclear arms race between the United States and the USSR, and the game theorist Anatole Rapoport gave a formal game-theoretic statement. Other game theorists have chosen to interpret the arms race as a prisoners' dilemma or as an assurance game. For a review and interesting discussion, see Barry O'Neill, "Game Theory Models of Peace and War," in *The Handbook of Game Theory*, vol. 2, ed. Robert J. Aumann and Sergiu Hart (Amsterdam: North Holland, 1994), pp. 995–1053.

		DEAN	
		Swerve (Chicken)	Straight (Tough)
JAMES	Swerve (Chicken)	0, 0	−1, 1
	Straight (Tough)	1, −1	−2, −2

FIGURE 4.14 Chicken

chicken, or weak. Fourth, the payoffs when both players are tough are very bad for both players. In games such as this one, the real game becomes a test of how to achieve one's preferred equilibrium.

We are now back in a situation similar to that discussed for the battle-of-the-sexes game. One expects most real-life chicken games to be even worse as battles than most battles of the sexes—the benefit of winning is larger, as is the cost of the crash, and so all the problems of conflict of interest and asymmetry between the players are aggravated. Each player will want to try to influence the outcome. It may be the case that one player will try to create an aura of toughness that everyone recognizes so as to intimidate all rivals.[13] Another possibility is to come up with some other way to convince your rival that you will not be chicken, by making a visible and irreversible commitment to going straight. (In Chapter 10, we consider just how to make such commitment moves.) In addition, both players also want to try to prevent the bad (crash) outcome if at all possible.

As with the battle of the sexes, if the game is repeated, tacit coordination is a better route to a solution. That is, if the teenagers played the game every Saturday night at midnight, they would have the benefit of knowing that the game had both a history and a future when deciding their equilibrium strategies. In such a situation, they might logically choose to alternate between the two equilibria, taking turns being the winner every other week. (But if the others found out about this deal, both players would lose face.)

There is one final point, arising from these coordination games, that must be addressed. The concept of Nash equilibrium requires each player to have the correct belief about the other's choice of strategy. When we look for Nash equilibria in pure strategies, the concept requires each to be confident about the other's choice. But our analysis of coordination games shows that thinking about the other's choice in such games is fraught with strategic uncertainty. How can

[13]Why would a potential rival play chicken against someone with a reputation for never giving in? The problem is that participation in chicken, as in lawsuits, is not really voluntary. Put another way, choosing whether to play chicken is itself a game of chicken. As Thomas Schelling says, "If you are publicly invited to play chicken and say you would rather not, then you have just played [and lost]" (*Arms and Influence*, New Haven: Yale University Press, 1965, p. 118).

we incorporate such uncertainty in our analysis? In Chapter 7, we introduce the concept of a mixed strategy, where actual choices are made randomly among the available actions. This approach generalizes the concept of Nash equilibrium to situations where the players may be unsure about each other's actions.

8 NO EQUILIBRIUM IN PURE STRATEGIES

Each of the games considered so far has had at least one Nash equilibrium in pure strategies. Some of these games, such as those in Section 7, had more than one equilibrium, whereas games in earlier sections had exactly one. Unfortunately, not all games that we come across in the study of strategy and game theory will have such easily definable outcomes in which players always choose one particular action as an equilibrium strategy. In this section, we look at games in which there is not even one pure-strategy Nash equilibrium—games in which none of the players would consistently choose one strategy as that player's equilibrium action.

A simple example of a game with no equilibrium in pure strategies is that of a single point in a tennis match. Imagine a match between the two all-time best women players—Martina Navratilova and Chris Evert.[14] Navratilova at the net has just volleyed a ball to Evert on the baseline, and Evert is about to attempt a passing shot. She can try to send the ball either down the line (DL; a hard, straight shot) or crosscourt (CC; a softer, diagonal shot). Navratilova must likewise prepare to cover one side or the other. Each player is aware that she must not give any indication of her planned action to her opponent, knowing that such information will be used against her. Navratilova would move to cover the side to which Evert is planning to hit or Evert would hit to the side that Navratilova is not planning to cover. Both must act in a fraction of a second, and both are equally good at concealing their intentions until the last possible moment; therefore their actions are effectively simultaneous, and we can analyze the point as a two-player simultaneous-move game.

Payoffs in this tennis-point game are given by the fraction of times a player wins the point in any particular combination of passing shot and covering play.

[14]For those among you who remember only the latest phenom who shines for a couple of years and then burns out, here are some amazing facts about these two, who were at the top levels of the game for almost two decades and ran a memorable rivalry all that time. Navratilova was a left-handed serve-and-volley player. In grand-slam tournaments, she won 18 singles titles, 31 doubles, and 7 mixed doubles. In all tournaments, she won 167, a record. Evert, a right-handed baseliner, had a record win-loss percentage (90% wins) in her career and 150 titles, of which 18 were for singles in grand slam tournaments. She probably invented (and certainly popularized) the two-handed backhand that is now so common. From 1973 to 1988, the two played each other 80 times, and Navratilova ended up with a slight edge, 43–37.

		NAVRATILOVA	
		DL	CC
EVERT	DL	50	80
	CC	90	20

FIGURE 4.15 No Equilibrium in Pure Strategies

Given that a down-the-line passing shot is stronger than a crosscourt shot and that Evert is more likely to win the point when Navratilova moves to cover the wrong side of the court, we can work out a reasonable set of payoffs. Suppose Evert is successful with a down-the-line passing shot 80% of the time if Navratilova covers crosscourt; she is successful with the down-the-line shot only 50% of the time if Navratilova covers down the line. Similarly, Evert is successful with her crosscourt passing shot 90% of the time if Navratilova covers down the line. This success rate is higher than when Navratilova covers crosscourt, in which case Evert wins only 20% of the time.

Clearly, the fraction of times that Navratilova wins this tennis point is just the difference between 100% and the fraction of time that Evert wins. Thus the game is zero-sum (more precisely, constant-sum, because the two payoffs sum to 100), and we can represent all the necessary information in the payoff table with just the payoff to Evert in each cell. Figure 4.15 shows the payoff table and the fraction of time that Evert wins the point against Navratilova in each of the four possible combinations of their strategy choices.

The rules for solving simultaneous-move games tell us to look first for dominant or dominated strategies and then to try minimax (in that this is a zero-sum game) or use cell-by-cell inspection to find a Nash equilibrium. It is a useful exercise to verify that no dominant strategies exist here. Going on to cell-by-cell inspection, we start with the choice of DL for both players. From that outcome, Evert can improve her success from 50% to 90% by choosing CC instead. But then Navratilova can hold Evert down to 20% by choosing CC. After this, Evert can raise her success again to 80% by making her shot DL, and Navratilova in turn can do better with DL. In every cell, one player always wants to change her play, and we cycle through the table endlessly without finding an equilibrium.

An important message is contained in the absence of a Nash equilibrium in this game and similar ones. What is important in games of this type is not what players should do, but what players should *not* do. In particular, each player should neither always nor systematically pick the same shot when faced with this situation. If either player engages in any determinate behavior of that type, the other can take advantage of it. (So if Evert consistently went crosscourt with her passing shot, Navratilova would learn to cover crosscourt every time and would thereby reduce Evert's chances of success with her crosscourt shot.) The most reasonable thing for players to do here is to act somewhat unsystematically, hoping for the element of surprise in defeating their opponents. An unsystematic approach entails choosing each strategy part of the time. (Evert should be using her weaker shot with enough frequency to guarantee that Navratilova cannot predict which shot will come her way. She should not, however, use the two shots in any set pattern, because that, too, would cause her to lose the element of surprise.) This approach, in which players randomize their actions, is known as mixing strategies and is the focus of Chapters 7 and 8. The game illustrated in Figure 4.15 may not have an equilibrium in pure strategies, but it can still be solved by looking for an equilibrium in mixed strategies, as we do in Chapter 7, Section 1.

SUMMARY

In simultaneous-move games, players make their strategy choices without knowledge of the choices being made by other players. Such games are illustrated by *game tables*, where cells show payoffs to each player and the dimensionality of the table equals the number of players. Two-person *zero-sum games* may be illustrated in shorthand with only one player's payoff in each cell of the game table.

Nash equilibrium is the solution concept used to solve simultaneous-move games; such an equilibrium consists of a set of strategies, one for each player, such that each player has chosen her best response to the other's choice. Nash equilibrium can also be defined as a set of strategies such that each player has correct *beliefs* about the others' strategies and strategies are best for each player given beliefs about the other's strategies. Nash equilibria can be found by using *cell-by-cell inspection*, through a search for *dominant strategies*, by *successive elimination of dominated strategies*, or with *best-response analysis*. Zero-sum games can also be solved by using the *minimax method*.

There are many classes of simultaneous games. *Prisoners' dilemma* games appear in many contexts. Coordination games, such as *assurance*, *chicken*, and *battle of the sexes*, have multiple equilibria, and the solution of such games requires players to achieve coordination by some means. If a game has no equilibrium in *pure strategies*, we must look for an equilibrium in *mixed strategies*, the analysis of which is presented in Chapters 7 and 8.

KEY TERMS

assurance game (114)
battle of the sexes (115)
belief (95)
best response (93)
best-response analysis (105)
cell-by-cell inspection (95)
chicken (116)
constant-sum game (91)
convergence of expectations (114)
coordination game (111)
dominance solvable (102)
dominant strategy (98)
dominated strategy (98)
enumeration (95)
focal point (112)
game matrix (90)
game table (90)
iterated elimination of dominated strategies (102)
maximin (106)
minimax (106)
minimax method (106)
mixed strategy (90)
Nash equilibrium (93)
normal form (90)
payoff table (90)
prisoners' dilemma (97)
pure coordination game (112)
pure strategy (90)
strategic form (90)
successive elimination of dominated strategies (102)
zero-sum game (91)

SOLVED EXERCISES

S1. "If a player has a dominant strategy in a simultaneous-move game, then she is sure to get her best possible outcome." True or false? Explain and give an example of a game that illustrates your answer.

S2. Find all Nash equilibria in pure strategies for the following zero-sum games. First check for dominant strategies. If neither player has a dominant strategy, use iterated elimination of dominated strategies to find the Nash equilibrium.

(a)

		COLUMN	
		Left	Right
ROW	Up	4	3
	Down	2	1

(b)

		COLUMN	
		Left	Right
ROW	Up	3	2
	Down	4	1

(c)

		COLUMN		
		Left	Middle	Right
ROW	Up	1	2	5
	Straight	2	4	3
	Down	1	3	3

(d)

		COLUMN			
		North	South	East	West
ROW	Up	6	7	5	6
	High	7	3	4	5
	Low	8	6	3	2
	Down	3	5	4	5

S3. Use the minimax method to find the Nash equilibria for the games in Exercise S2.

S4. Find all Nash equilibria in pure strategies in the following non-zero-sum games. Describe the steps that you used in finding the equilibria.

(a)

		COLUMN	
		Left	Right
ROW	Up	2, 4	1, 0
	Down	6, 5	4, 2

(b)

		COLUMN	
		Left	Right
ROW	Up	1, 1	0, 1
	Down	1, 0	1, 1

(c)

		COLUMN		
		Left	Middle	Right
ROW	Up	0, 1	9, 0	2, 3
	Straight	5, 9	7, 3	1, 7
	Down	7, 5	10, 10	3, 5

(d)

		COLUMN		
		West	Center	East
ROW	North	2, 3	8, 2	7, 4
	Up	3, 0	4, 5	6, 4
	Down	10, 4	6, 1	3, 9
	South	4, 5	2, 3	5, 2

S5. Consider the following table:

		COLUMN			
		North	South	East	West
ROW	Earth	1, 3	3, 1	0, 2	1, 1
	Water	1, 2	1, 2	2, 3	1, 1
	Wind	3, 2	2, 1	1, 3	0, 3
	Fire	2, 0	3, 0	1, 1	2, 2

(a) Does either Row or Column have a dominant strategy? Explain why or why not.
(b) Use iterated elimination of dominated strategies to reduce the game as much as possible. Give the order in which the eliminations occur and give the reduced form of the game.
(c) Is this game dominance solvable? Explain why or why not.
(d) State the Nash equilibrium (or equilibria) of this game.

S6. An old lady is looking for help crossing the street. Only one person is needed to help her; more are okay but no better than one. You and I are the two people in the vicinity who can help; we have to choose simultaneously whether to do so. Each of us will get pleasure worth a 3 from her success (no matter who helps her). But each one who goes to help will bear a cost of 1, this being the value of our time taken up in helping. If neither player helps, the payoff for each player is zero. Set this up as a game. Write the payoff table, and find all pure-strategy Nash equilibria.

S7. A university is contemplating whether to build a new lab or a new theater on campus. The science faculty would rather see a new lab built, and the humanities faculty would prefer a new theater. However, the funding for the project (whichever it may turn out to be) is contingent on unanimous support from the faculty. If there is disagreement, neither project will go forward, leaving each group with no new building and their worst payoff. The meetings of the two separate faculty groups on which proposal to support occur simultaneously, with payoffs given in the following table:

		HUMANITIES FACULTY	
		Lab	Theater
SCIENCE FACULTY	Lab	4, 2	0, 0
	Theater	0, 0	1, 5

(a) What are the pure-strategy Nash equilibria of this game?
(b) Which game described in this chapter is most similar to this game? Explain your reasoning.

S8. Suppose two game-show contestants, Alex and Bob, each separately select one of three doors numbered 1, 2, and 3. Both players get dollar prizes if their choices match, as indicated in the following table.

		B		
		1	2	3
A	1	10, 10	0, 0	0, 0
	2	0, 0	15, 15	0, 0
	3	0, 0	0, 0	15, 15

(a) What are the Nash equilibria of this game? Which, if any, is likely to emerge as the (focal) outcome? Explain.
(b) Consider a slightly changed game in which the choices are again just numbers, but the two cells with (15, 15) in the table become (25, 25). What is the expected (average) payoff to each player if each flips a coin to decide whether to play 2 or 3? Is this better than focusing on both choosing 1 as a focal equilibrium? How should you account for the risk that Alex might do one thing while Bob does the other?

S9. Marta has three sons: Arturo, Bernardo, and Carlos. She discovers a broken lamp in her living room and knows that one of her sons must have broken it at play. In reality, Carlos was the culprit, but Marta doesn't know this. She cares more about finding out the truth than she does about punishing the child who broke the lamp, so Marta announces that her sons are to play the following game.

Each child will write down his name on a piece of paper and write down either "Yes, I broke the lamp," or "No, I didn't break the lamp." If at least one child claims to have broken the lamp, she will give the normal allowance of $2 to each child who claims to have broken the lamp, and $5 to each child who claims not to have broken the lamp. If all three children claim not to have broken the lamp, none of them receives any allowance (each receives $0).

(a) Write down the game table. Make Arturo the row player, Bernardo the column player, and Carlos the page player.
(b) Find all the Nash equilibria of this game.
(c) There are multiple Nash equilibria of this game. Which one would you consider to be a focal point?

S10. Consider a game in which there is a prize worth $30. There are three contestants, Larry, Curly, and Moe. Each can buy a ticket worth $15 or $30 or not buy a ticket at all. They make these choices simultaneously and independently. Then, knowing the ticket-purchase decisions, the game organizer awards the prize. If no one has bought a ticket, the prize is not awarded. Otherwise, the prize is awarded to the buyer of the highest-cost ticket if there is only one such player or is split equally between two or three if there are ties among the highest-cost ticket buyers. Show this game in strategic form, using Larry as the Row player, Curly as the Column player, and Moe as the Page player. Find all pure-strategy Nash equilibria.

S11. Anne and Bruce would like to rent a movie, but they can't decide what kind of movie to get: Anne wants to rent a comedy, and Bruce wants to watch a drama. They decide to choose randomly by playing "Evens or Odds." On the count of three, each of them shows one or two fingers. If the sum is even, Anne wins and they rent the comedy; if the sum is odd, Bruce wins and they rent the drama. Each of them earns a payoff of 1 for winning and 0 for losing "Evens or Odds."

(a) Draw the game table for "Evens or Odds."

(b) Demonstrate that this game has no Nash equilibrium in pure strategies.

S12. In the film *A Beautiful Mind*, John Nash and three of his graduate-school colleagues find themselves faced with a dilemma while at a bar. There are four brunettes and a single blonde available for them to approach. Each young man wants to approach and win the attention of one of the young women. The payoff to each of winning the blonde is 10; the payoff of winning a brunette is 5; the payoff from ending up with no girl is zero. The catch is that if two or more young men go for the blonde, she rejects all of them, and then the brunettes also reject the men because they don't want to be second choice. Thus, each player gets a payoff of 10 only if he is the sole suitor for the blonde.

(a) First consider a simpler situation in which there are only two young men instead of four. (There are two brunettes and one blonde, but these women merely respond in the manner just described and are not active players in the game.) Show the playoff table for the game, and find all of the pure-strategy Nash equilibria of the game.

(b) Now show the (three-dimensional) table for the case in which there are three young men (and three brunettes and one blonde who are not active players). Again, find all of the Nash equilibria of the game.

(c) Without the use of a table, give all of the Nash equilibria for the case in which there are four young men (as well as four brunettes and a blonde).

(d) **(Optional)** Use your results to parts (a), (b), and (c) to generalize your analysis to the case in which there are n young men. Do not attempt to write down an n-dimensional payoff table; merely find the payoff to one player when k of the others choose Blonde and $(n - k - 1)$ choose Brunette, for $k = 0, 1, \ldots (n - 1)$. Can the outcome specified in the movie as the Nash equilibrium of the game—that all of the young men choose to go for brunettes—ever really be a Nash equilibrium of the game?

UNSOLVED EXERCISES

U1. Find all Nash equilibria in pure strategies for the zero-sum games in the following tables by checking for dominant strategies and using iterated dominance.

(a)

		COLUMN	
		Left	Right
ROW	Up	1	4
	Down	2	3

(b)

		COLUMN	
		Left	Right
ROW	Up	1	2
	Down	4	3

(c)

		COLUMN		
		Left	Middle	Right
ROW	Up	5	3	2
	Straight	6	4	3
	Down	1	6	2

(d)

		COLUMN		
		Left	Middle	Right
ROW	Up	5	1	3
	Straight	6	1	2
	Down	1	0	0

U2. Use the minimax method to find the Nash equilibria for the games in Exercise U1.

U3. Find all Nash equilibria in pure strategies in the following non-zero-sum games. Describe the steps that you used in finding the equilibria.

(a)

		COLUMN	
		Left	Right
ROW	Up	3, 1	4, 2
	Down	5, 2	2, 3

(b)

		COLUMN	
		Left	Right
ROW	Up	0, 0	0, 0
	Down	0, 0	1, 1

(c)

		COLUMN		
		Left	Middle	Right
ROW	Up	2, 9	5, 5	6, 2
	Straight	6, 4	9, 2	5, 3
	Down	4, 3	2, 7	7, 1

(d)

		COLUMN		
		Left	Middle	Right
ROW	Up	5, 3	7, 2	2, 1
	Straight	1, 2	6, 3	1, 4
	Down	4, 2	6, 4	3, 5

U4. Use successive elimination of dominated strategies to solve the following game. Explain the steps you followed. Show that your solution is a Nash equilibrium.

		COLUMN		
		Left	Middle	Right
ROW	Up	4, 3	2, 7	0, 4
	Down	5, 0	5, −1	−4, −2

U5. Find all of the pure-strategy Nash equilibria for the following game. Describe the process that you used to find the equilibria. Use this game to explain why it is important to describe an equilibrium by using the strategies employed by the players, not merely by the payoffs received in equilibrium.

		COLUMN		
		Left	Center	Right
ROW	Up	1, 2	2, 1	1, 0
	Level	0, 5	1, 2	7, 4
	Down	−1, 1	3, 0	5, 2

U6. Consider the following game table:

		COLUMN		
		Left	Center	Right
ROW	Top	4, ___	___, 2	3, 1
	Middle	3, 5	2, ___	2, 3
	Bottom	___, 3	3, 4	4, 2

(a) Complete the payoffs of the game table above so that Column has a dominant strategy. State which strategy is dominant and explain why. (Note: there are many equally correct answers.)

(b) Complete the payoffs of the game table above so that neither player has a dominant strategy, but also so that each player does have a dominated strategy. State which strategies are dominated and explain why. (Again, there are many equally correct answers.)

U7. The game known as the *Battle of the Bismarck Sea* (named for that part of the southwestern Pacific Ocean separating the Bismarck Archipelago from Papua–New Guinea) summarizes a well-known game actually played in a naval engagement between the United States and Japan during World War II. In 1943, a Japanese admiral was ordered to move a convoy of ships to New Guinea; he had to choose between a rainy northern route and a sunnier southern route, both of which required three days' sailing time. The Americans knew that the convoy would sail and wanted to send bombers after it, but they did not know which route it would take. The Americans had to send reconnaissance planes to scout for the convoy, but they had only enough reconnaissance planes to explore one route at a time. Both the Japanese and the Americans had to make their decisions with no knowledge of the plans being made by the other side.

If the convoy was on the route that the Americans explored first, they could send bombers right away; if not, they lost a day of bombing. Poor weather on the northern route would also hamper bombing. If the Americans explored the northern route and found the Japanese right away, they could expect only two (of three) good bombing days; if they explored the northern route and found that the Japanese had gone south, they could also expect two days of bombing. If the Americans chose to explore the southern route first, they could expect three full days of bombing if they found the Japanese right away but only one day of bombing if they found that the Japanese had gone north.

(a) Illustrate this game in a game table.

(b) Identify any dominant strategies in the game and solve for the Nash equilibrium.

U8. Two players, Jack and Jill, are put in separate rooms. Then each is told the rules of the game. Each is to pick one of six letters: G, K, L, Q, R, or W. If the two happen to choose the same letter, both get prizes as follows:

Letter	G	K	L	Q	R	W
Jack's Prize	3	2	6	3	4	5
Jill's Prize	6	5	4	3	2	1

If they choose different letters, each gets zero. This whole schedule is revealed to both players, and both are told that both know the schedules, and so on.

(a) Draw the table for this game. What are the Nash equilibria in pure strategies?

(b) Can one of the equilibria be a focal point? Which one? Why?

U9. Three friends (Julie, Kristin, and Larissa) independently go shopping for dresses for their high-school prom. On reaching the store, each girl sees only three dresses worth considering: one black, one lavender, and one yellow. Each girl furthermore can tell that her two friends would consider the same set of three dresses, because all three have somewhat similar tastes.

Each girl would prefer to have a unique dress, so a girl's utility is zero if she ends up purchasing the same dress as at least one of her friends. All three know that Julie strongly prefers black to both lavender and yellow, so she would get a utility of 3 if she were the only one wearing the black dress, and a utility of 1 if she were either the only one wearing the lavender dress or the only one wearing or the yellow dress. Similarly, all know that Kristin prefers lavender and secondarily prefers yellow, so her utility would be 3 for uniquely wearing lavender, 2 for uniquely wearing yellow, and 1 for uniquely wearing black. Finally, all know that Larissa prefers yellow and secondarily prefers black, so she would get 3 for uniquely wearing yellow, 2 for uniquely wearing black, and 1 for uniquely wearing lavender.

(a) Provide the game table for this three-player game. Make Julie the Row player, Kristin the Column player, and Larissa the Page player.

(b) Identify any dominated strategies in this game, or explain why there are none.

(c) What are the pure-strategy Nash equilibria in this game?

U10. Bruce, Colleen, and David are all getting together at Bruce's house on Friday evening to play their favorite game, Monopoly. They all love to eat sushi while they play. They all know from previous experience that two orders of sushi are just the right amount to satisfy their hunger. If they wind up with less than two orders, they all end up going hungry and don't enjoy the evening. More than two orders would be a waste, because they can't manage to eat a third order and the extra sushi just goes bad. Their favorite restaurant, Fishes in the Raw, packages its sushi in such large containers that each individual person can feasibly purchase at most one order of sushi. Fishes in the Raw offers takeout, but unfortunately doesn't deliver.

Suppose that each player enjoys \$20 worth of utility from having enough sushi to eat on Friday evening, and \$0 from not having enough to eat. The cost to each player of picking up an order of sushi is \$10.

Unfortunately, the players have forgotten to communicate about who should be buying sushi this Friday, and none of the players has a cell phone, so they must each make independent decisions of whether to buy (B) or not buy (N) an order of sushi.

(a) Write down this game in strategic form.

(b) Find all the Nash equilibria in pure strategies.

(c) Which equilibrium would you consider to be a focal point? Explain your reasoning.

U11. Roxanne, Sara, and Ted all love to eat cookies, but there's only one left in the package. No one wants to split the cookie, so Sara proposes the following extension of "Evens or Odds" (see Exercise S11) to determine who gets to eat it. On the count of three, each of them will show one or two fingers, they'll add them up, and then divide the sum by 3. If the remainder is zero Roxanne gets the cookie, if the remainder is 1 Sara gets it, and if it is 2 Ted gets it. Each of them receives a payoff of 1 for winning (and eating the cookie) and zero otherwise.

(a) Represent this three-player game in normal form, with Roxanne as the Row player, Sara as the Column player, and Ted as the Page player.

(b) Find all the pure-strategy Nash equilibria of this game. Is this game a fair mechanism for allocating cookies? Explain why or why not.

U12. **(Optional)** Construct the payoff matrix for your own two-player game that satisfies the following requirements. First, each player should have three strategies. Second, the game should not have any dominant strategies. Third, the game should not be solvable using minimax. Fourth, the game should have exactly two pure-strategy Nash equilibria. Provide your game matrix, and then demonstrate that all of the above conditions are true.

5

Simultaneous-Move Games with Pure Strategies II: Continuous Strategies and III: Discussion and Evidence

THE DISCUSSION OF SIMULTANEOUS-MOVE GAMES in Chapter 4 focused on games in which each player had a discrete set of actions from which to choose. Discrete strategy games of this type include sporting contests in which a small number of well-defined plays can be used in a given situation—soccer penalty kicks, in which the kicker can choose to go high or low, to a corner or the center, for example. Other examples include coordination and prisoners' dilemma games in which players have only two or three available strategies. Such games are amenable to analysis with the use of a game table, at least for situations with a reasonable number of players and available actions.

Many simultaneous-move games differ from those considered so far; they entail players choosing strategies from a wide range of possibilities. Games in which manufacturers choose prices for their products, philanthropists choose charitable contribution amounts, or contractors choose project bid levels are examples in which players have a virtually infinite set of choices. Technically, prices and other dollar amounts do have a minimum unit, such as a cent, and so there is actually only a finite and discrete set of price strategies. But in practice the unit is very small, and allowing the discreteness would require us to give each player too many distinct strategies and make the game table too large; therefore it is simpler and better to regard such choices as continuously variable real numbers. When players have such a large range of actions available, game tables become