4

Simultaneous-Move Games: 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 choosing to do at the same time that the others are figuring out what she is choosing to do. This circularity makes the analysis of simultaneous-move games somewhat more intricate than the analysis of sequential-move games, but the analysis is not difficult. In this chapter, we will 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 Chapter 7. 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 for a product 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 then 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

[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.

		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

dimensions 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 blue.

Next we consider a second example with more of a story attached. Figure 4.2 represents a very simplified version of a single play in American football. Offense attempts 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). Defense can adopt one of three strategies to try to keep Offense at bay: a run defense, a pass defense, or a blitz of the quarterback.

[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, −2	5, −5	13, −13
	Short Pass	6, −6	5.6, −5.6	10.5, −10.5
	Medium Pass	6, −6	4.5, −4.5	1, −1
	Long Pass	10, −10	3, −3	−2, 2

FIGURE 4.2 A Single Play in American Football

Offense tries to gain yardage while 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 a gain of 4.5 yards, or +4.5.[3] Defense's "payoff" is a loss of 4.5 yards, or −4.5. The other cells similarly show our estimates of each team's gain or loss of yardage.

Note that the payoffs sum to 0 in every cell of this table: when the offense gains 5 yards, the defense loses 5 yards, and when the offense loses 2 yards, the defense gains 2 yards. This pattern is quite common in sports contexts, where the interests of the two sides are exactly the opposite of each other. As noted in Chapter 2, we call this a zero-sum (or sometimes constant-sum) game. You should remember that the definition of a zero-sum game is that the payoffs sum to the same *constant* across cells, whether that number is 0, 6, or 1,000. (In Section 4.7, we describe a game where the two players' payoffs sum to 100.) The key feature of any zero-sum game is that one player's loss is the other player's gain.

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 17. 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 8, 9, and 13. 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 S13 of this chapter and in Exercise S14 of Chapter 7.

		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 10. 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 Defense is choosing the Pass defense, then the best choice for Offense is Short Pass (payoff of 5.6 versus 5, 4.5, or 3). Conversely, if Offense is choosing the Short Pass, then Defense's best choice is the Pass defense—it holds 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 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 a systematic analysis is foolproof, but tedious and unmanageable except in simple games or with the use of a good computer program to check cells for equilibria. Luckily, there are 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 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 Chapter 7, 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 Chapter 7, 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 Chapter 7 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 Chapter 7, 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, while 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—this is the meaning of dominance—

but if each of them attributes to the other the same rationality as he or she practices, then both of them 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 *L.A. 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 10 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 choose 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. There is 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 10.

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 applies 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,

		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

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 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, we will 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, 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, Congress likes worst (payoff 1) the outcome with a balanced budget and high

[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.

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 a balanced budget 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 Congress is opting for a balanced budget (Fed's payoff 4 rather than 3), but it does better choosing high interest rates if it believes that 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 Congress does. If Congress believes that the Fed is choosing low interest rates, it does better for itself by choosing a budget deficit rather than a balanced budget (Congress's payoff 4 instead of 3). If Congress believes that the Fed is choosing high interest rates, again it does better for itself by choosing a budget deficit rather than a balanced budget (Congress's payoff 2 instead of 1). Choosing to run a budget deficit is then Congress's dominant strategy.

The choice for Congress is now clear. No matter what it believes the Fed is doing, 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 Congress will choose its dominant strategy (budget deficit) and therefore 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 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 9, 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

		COLIN	
		Left	Right
ROWENA	Up	0, 0	1, 1
	Down	1, 1	1, 1

FIGURE 4.6 Elimination of Weakly Dominated Strategies

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.

And now, 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, where we introduce Rowena as the row player and Colin as the column player.[8] For Rowena, Up is weakly dominated by Down; if Colin plays Left, then Rowena gets a better payoff by playing Down than by playing Up, and, if Colin plays Right, then Rowena gets the same payoff from her two strategies.

[8] We use these names in the hope that they will aid you in remembering which player chooses the row and which chooses the column. We acknowledge Robert Aumann, who shared the Nobel Prize with Thomas Schelling in 2005 (and whose ideas will be prominent in Chaper 9), for inventing this clever naming idea.

Similarly, for Colin, Right weakly dominates Left. Dominance solvability then tells us that (Down, Right) is a Nash equilibrium. That is true, but (Down, Left) and (Up, Right) also are Nash equilibria. Consider (Down, Left). When Rowena is playing Down, Colin cannot improve his payoff by switching to Right, and, when Colin is playing Left, Rowena'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 use other methods (such as the one described in the next section) 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 will 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 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.

Let's return to the game played by Row and Column and reproduce it as Figure 4.7. We'll first consider Row's responses. 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). Again, we show Row's best choices by circling the appropriate payoffs.

		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

Similarly, Column's best responses are shown by circling her payoffs: 3 (Middle as best response to Row's Top), 5 (Left to Row's High), 4 (Middle to Row's Low), and 7 (Right to Row's Bottom).[9] 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 in the Husband–Wife prisoners' dilemma shown in Figure 4.4 and the Congress–Federal Reserve game depicted in Figure 4.5 emerge from such an analysis.

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

[9] 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.

find a Nash equilibrium, then the game has no equilibrium in pure strategies. We address games of this type in Section 7 of this chapter. In Chapter 5, we will extend best-response analysis to games where the players' strategies are continuous variables—for example, prices or advertising expenditures. Moreover, we will construct best-response *curves* to help us find Nash equilibria, and we will see that such games are less likely—by virtue of the continuity of strategy choices—to have no equilibrium.

5 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 four 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 that the garden was no better when all three contributed 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 possible choices for her 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 of the other players 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 of them. 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 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, thereby providing 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 must 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 4.8. 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

TALIA chooses:

		Contribute: NINA Contribute	Contribute: NINA Don't	Don't Contribute: NINA Contribute	Don't Contribute: NINA 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 4.8 Street–Garden Game

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 Contribute. 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 11, 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 best-response analysis, as shown in Figure 4.9. Because each player has Don't Contribute as her dominant strategy, all of Emily's best responses are on her Don't Contribute row, all of Nina's best responses are on her Don't Contribute column, and all of

TALIA chooses:

Contribute

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

Don't Contribute

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

FIGURE 4.9 Best-Response Analysis in the Street–Garden Game

Talia's best responses are on her Don't Contribute page. The cell at the bottom right has all three best responses; therefore, it gives us the Nash equilibrium.

6 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.[10] They are attracted to each other and would like to continue the conversation, but they 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?

Figure 4.10 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. Best-response analysis quickly reveals that the

[10] The names come from the 1989 movie *When Harry Met Sally*, starring Meg Ryan and Billy Crystal, with its classic line "I'll have what she's having."

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

FIGURE 4.10 Pure Coordination

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.[11]

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 na-

[11] 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 stairs at TKTS in Times Square. The focal time remains twelve noon.