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

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 virtually useless as analytical tools; they become too unwieldy to be of practical use. For these games we need a different solution technique. We present the analytical tools for handling such **continuous strategy** games in the first part of this chapter.

This chapter also takes up some broader matters relevant to behavior in simultaneous-move games and to the concept of Nash equilibrium. We review the empirical evidence on Nash equilibrium play that has been collected both from the laboratory and from real-life situations. We also present some theoretical criticisms of the Nash equilibrium concept and rebuttals of these criticisms. You will see that game-theoretic predictions are often a reasonable starting point for understanding actual behavior, with some caveats, such as the level of player.

1 PURE STRATEGIES THAT ARE CONTINUOUS VARIABLES

In Chapter 4 we developed the method of best-response analysis for finding all pure-strategy Nash equilibria of simultaneous-move games. Now we extend that method to games in which each player has available a continuous range of choices—for example, firms setting prices of their products. To calculate best responses in this type of game, we find, for each possible value of one firm's price, the value of the other firm's price that is best for it (maximizes its payoff). The continuity of the sets of strategies allows us to use algebraic formulas to show how strategies generate payoffs and to show the best responses as curves in a graph, with each player's price (or any other continuous strategy) on one of the axes. In such an illustration, the Nash equilibrium of the game occurs where the two curves meet. We develop this idea and technique by using two stories.

A. Price Competition

Our first story is set in a small town, Yuppie Haven, that has two restaurants, Xavier's Tapas Bar and Yvonne's Bistro. To keep the story simple, we suppose that each place has a set menu. Xavier and Yvonne have to set the prices of their respective menus. Prices are their strategic choices in the game of competing with each other; each bistro's goal is to set price to maximize profit, the payoff in this game. We suppose that they must get their menus printed separately without knowing the other's price, so the game has simultaneous moves.[1] Because prices can take any value within an (almost) infinite range, we start with general or algebraic symbols for them. We then find **best-response rules** that we use to solve

[1]In reality, the competition extends over time, so each can observe the other's past choices. This repetition of the game introduces new considerations, which we cover in Chapter 11.

the game and to determine equilibrium prices. Let us call Xavier's price P_x and Yvonne's price P_y.

In setting its price, each restaurant has to calculate the consequences for its profit. To keep things relatively simple, we put the two restaurants in a very symmetric relationship, but readers with a little more mathematical skill can do a similar analysis by using much more general numbers or even algebraic symbols. Suppose the cost of serving each customer is $8 for each restaurateur. Suppose further that experience or market surveys have shown that, when Xavier's price is P_x and Yvonne's price is P_y, the number of their respective customers, respectively Q_x and Q_y (measured in hundreds per month), are given by the equations[2]

$$Q_x = 44 - 2P_x + P_y,$$
$$Q_y = 44 - 2P_y + P_x.$$

The key idea in these equations is that, if one restaurant raises its price by $1 (say, Yvonne increases P_y by $1), its sales will go down by 200 per month (Q_y changes by -2) and those of the other restaurant will go up by 100 per month (Q_x changes by 1). Presumably, 100 of Yvonne's customers switch to Xavier's and another 100 stay at home.

Xavier's profit per week (in hundreds of dollars per week), call it Π_x—the Greek letter Π (pi) is the traditional economic symbol for profit—is given by the product of the net revenue per customer (price less cost or $P_x - 8$) and the number of customers served:

$$\Pi_x = (P_x - 8)\, Q_x = (P_x - 8)\,(44 - 2\,P_x + P_y).$$

By multiplying out and rearranging the terms on the right-hand side of the preceding expression, we can write profit as a function of increasing powers of P_x:

$$\Pi_x = -8(44 + P_y) + (16 + 44 + P_y)P_x - 2(P_x)^2$$
$$= -8(44 + P_y) + (60 + P_y)P_x - 2(P_x)^2.$$

Xavier sets his price P_x to maximize this payoff. Doing so for each possible level of Yvonne's price P_y gives us Xavier's best-response rule; we can then graph it.

Many simple illustrative examples where one real number (such as the price) is chosen to maximize another real number that depends on it (such as the profit or the payoff) have a similar form. (In mathematical jargon, we would describe the second number as a function of the first.) In the appendix to this

[2]Readers who know some economics will recognize that the equations linking quantities to prices are demand functions for the two products X and Y. The quantity demanded of each product is decreasing in its own price (demands are downward sloping) and increasing in the price of the other product (the two are substitutes).

chapter we develop a simple general technique for performing such maximization; you will find many occasions to use it. Here we just state the formula.

The function we want to maximize takes the general form

$$Y = A + BX - CX^2$$

where we have used the descriptor Y for the number we want to maximize and X for the number we want to choose to maximize that Y. In our specific example, profit, Π_x, would be represented by Y, and the price, P_x, by X. Similarly, although in any specific problem, the terms A, B, and C in the equation above would be known numbers, we have denoted them by general algebraic symbols so that our formula can be applied across a wide variety of similar problems. (The technical term for the terms A, B, and C is *parameters*, or *algebraic constants*.) Because most of our applications involve nonnegative X entities, such as prices, and the maximization of the Y entity, we require $B > 0$ and $C > 0$. Then the formula giving the choice of X to maximize Y in terms of the known parameters A, B, and C is simply $X = B/(2C)$. Observe that A does not appear in the formula, although it will of course affect the value of Y that results.

Comparing the general function in the equation above and the specific example of the profit function in the pricing game on the previous page, we have[3]

$$B = 60 + P_y \text{ and } C = 2.$$

Therefore Xavier's choice of price to maximize his profit will satisfy the formula $B/(2C)$ and will be

$$P_x = 15 + 0.25P_y.$$

This equation determines the value of P_x that maximizes Xavier's profit, given a particular value of Yvonne's price, P_y. In other words, it is exactly what we want, the rule for Xavier's best response.

Yvonne's best-response rule can be found similarly. Because the costs and sales of the two restaurants are entirely symmetric, the equation is obviously going to be

$$P_y = 15 + 0.25P_x.$$

Both rules are used in the same way to develop best-response graphs. If Xavier sets a price of 16, for example, then Yvonne plugs this value into her best-response rule to find $P_y = 15 + 0.25(16) = 19$; similarly, Xavier's best response to Yvonne's $P_y = 16$ is $P_x = 19$, and each restaurant's best response to the other's price of 4 is 16, that to 8 is 17, and so on.

[3]Although P_y, chosen by Yvonne, is a variable in the full game, here we are considering only a part of the game, namely Xavier's best response, where he regards Yvonne's choice as outside his control and therefore like a constant.

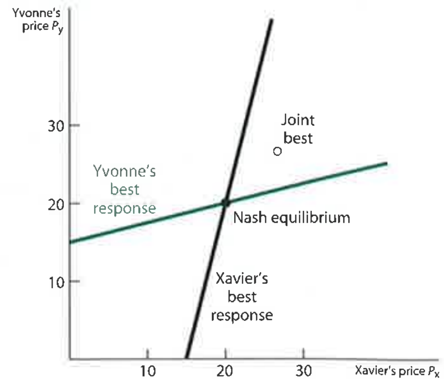

FIGURE 5.1 Best-Response Curves and Equilibrium in the Restaurant Pricing Game

Figure 5.1 shows the graphs of these two best-response relations. Owing to the special features of our example—namely, the linear relation between quantity sold and prices charged, and the constant cost of producing each meal—each of the two **best-response curves** is a straight line. For other specifications of demands and costs, the curves can be other than straight, but the method of obtaining them is the same: namely, first holding one restaurant's price (say, P_y) fixed and finding the value of the other's price (say, P_x) that maximizes the second restaurant's profit, and then the other way around.

The point of intersection of the two best-response curves is the Nash equilibrium of the pricing game between the two restaurants. That point represents the pair of prices, one for each firm, that are best responses to each other. The specific values for each restaurant's pricing strategy in equilibrium can be found algebraically by solving the two best-response rules jointly for P_x and P_y. We deliberately chose our example to make the equations linear, and the solution is easy. In this case, we simply substitute the expression for P_x into the expression for P_y to find

$$P_y = 15 + 0.25P_x = 15 + 0.25(15 + 0.25P_y) = 18.75 + 0.0625\,P_y.$$

This last equation simplifies to $P_y = 20$. Given the symmetry of the problem, it is simple to determine that $P_x = 20$ also.[4] Thus, in equilibrium, each restaurant charges \$20 for its menu and makes a profit of \$12 on each of the 2,400 customers

[4]Without this symmetry, the two best-response equations will be different, but given our other specifications, still linear. So it is not much harder to solve the nonsymmetric case. You will have a chance to do so in Exercise S2 at the end of this chapter.

[2,400 = (44−2 × 20 + 20) hundred] that it serves each month, for a total profit of $28,800 per month.

B. Some Economics of Oligopoly

Our main purpose in presenting the restaurant pricing example was to illustrate how the Nash equilibrium can be found in a game where the strategies are continuous variables, such as prices. But it is interesting to take a further look into this situation and to explain some of the economics behind pricing strategies and profits when a small number of firms (here just two) compete. In the jargon of economics, such competition is referred to as oligopoly, from the Greek words for "a small number of sellers."

Begin by observing that each of the two firm best-response curve slopes upward. Specifically, when one restaurant raises its price by $1, the other's best response is to raise its own price by 0.25, or 25 cents. When one restaurant raises its price, some of its customers switch to the other restaurant, and its rival can then best profit from them by raising its price part of the way. Thus a restaurant that raises its price is helping to increase the other's profit. In Nash equilibrium, where each restaurant chooses its price independently and out of concern for its own profit, it does not take into account this benefit that it conveys to the other. Could they get together and cooperatively agree to raise their prices, thereby raising both profits? Yes. Suppose the two restaurant charged $24 each. Then each would make a profit of $16 on each of the 2,000 customers [2,000 = (44 − 2 × 24 + 24) hundred] that it would serve each month, for a total profit of $32,000.

This pricing game is exactly like the prisoners' dilemma game presented in Chapter 4, but now the strategies are continuous variables. In the story in Chapter 4, the husband and wife were each tempted to cheat the other and confess to the police; but, when they both did so, both ended up with longer prison sentences (worse outcomes). In the same way, the more profitable price of $24 is not a Nash equilibrium. The separate calculations of the two restaurants will lead them to undercut such a price. Suppose that Yvonne somehow starts by charging $24. Using the best-response formula, we see that Xavier will then charge 15 + 0.25 × 24 = 21. Then Yvonne will come back with her best response to that: 15 + 0.25 × 21 = 20.25. Continuing this process, the prices of both will converge toward the Nash equilibrium price of $20.

But what price is jointly best for the two restaurants? Given the symmetry, suppose both charge the same price P. Then the profit of each will be

$$\Pi_x = \Pi_y = (P - 8)(44 - 2P + P) = (P - 8)(44 - P) = -352 + 52P - P^2.$$

The two can choose P to maximize this expression. Using the formula provided in Section 1.A, we see that the solution is $P = 52/2 = 26$. The resulting profit for each restaurant is $32,400 per month.

In the jargon of economics, such collusion to raise prices to the jointly optimal level is called a *cartel*. The high prices hurt consumers, and regulatory agencies of the U.S. government often try to prevent the formation of cartels and to make firms compete with one another. Explicit collusion over price is illegal, but it may be possible to maintain tacit collusion in a repeated prisoners' dilemma; we examine such repeated games in Chapter 11.[5]

Collusion need not always lead to higher prices. In the preceding example, if one restaurant lowers its price, its sales increase, in part because it draws some customers away from its rival because the products (meals) of the two restaurants are *substitutes* for each other. In other contexts, two firms may be selling products that are *complements* to each other—for example, hardware and software. In that case, if one firm lowers its price, the sales of both firms increase. In a Nash equilibrium, where the firms act independently, they do not take into account the benefit that a lower price of each would convey on the other. Therefore they keep prices higher than they would if they were able to coordinate their actions. Allowing them to cooperate would lead to lower prices and thus be beneficial to the consumers as well.

Competition need not always involve the use of prices as the strategic variables. For example, fishing fleets may compete to bring a larger catch to market; this is quantity competition as opposed to the price competition considered in this section. We consider quantity competition later in this chapter and in several of the end-of-chapter exercises.

C. Political Campaign Advertising

Our second example is one drawn from politics. It requires just a little more mathematics than we normally use, but we explain the intuition behind the calculations in words and with a graph.

Consider an election contested by two parties or candidates. Each is trying to win votes away from the other by advertising—either positive ads that highlight the good things about oneself or negative, attack ads that emphasize the bad things about the opponent. To keep matters simple, suppose the voters start entirely ignorant and unconcerned and are moved solely by the ads. (Many people would claim that this is a pretty accurate description of U.S. politics, but more advanced analyses in political science do recognize that there are informed and strategic voters. We address the behavior of such voters in detail in Chapter 16.) Even more simply, suppose the vote share of a party equals its share of the total campaign advertising that is done. Call the

[5]Firms do try to achieve explicit collusion when they think they can get away with it. An entertaining and instructive story of one such episode is in *The Informant*, by Kurt Eichenwald (New York: Broadway Books, 2000).

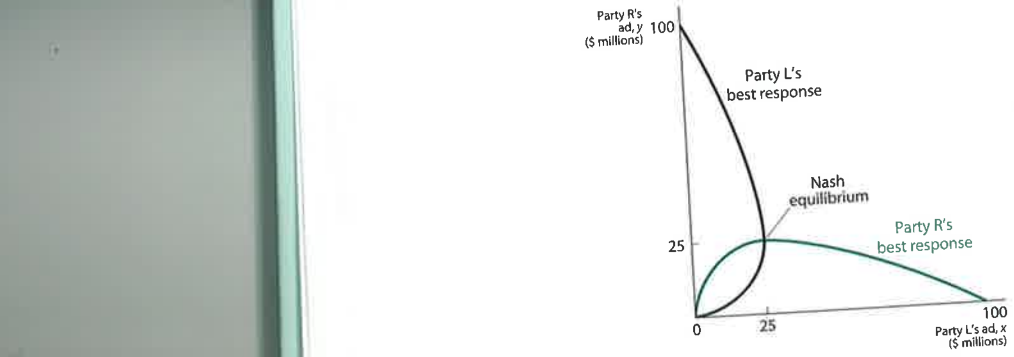

FIGURE 5.2 Best Responses and Nash Equilibrium in the Campaign Advertising Game

parties or candidates L and R; when L spends \$$x$ million on advertising and R spends \$$y$ million, L will get a share $x/(x + y)$ of the votes and R will get $y/(x + y)$. Once again, readers who get interested in this application can find more general treatments in specialized political science writings.

Raising money to pay for these ads includes a cost: money to send letters and make phone calls; time and effort of the candidates, party leaders, and activists; the future political payoff to large contributors; and possible future political costs if these payoffs are exposed and lead to scandals. For simplicity of analysis, let us suppose all these costs are proportional to the direct campaign expenditures x and y. Specifically, let us suppose that party L's payoff is measured by its vote percentage minus its advertising expenditure, $100x/(x + y) - x$. Similarly party R's payoff is $100y/(x + y) - y$.

Now we can find the best responses. Because we cannot do so without calculus, we derive the formula mathematically and then explain its general meaning intuitively, in words. For a given strategy x of party L, party R chooses y to maximize its payoff. The calculus first-order condition is found by holding x fixed and setting the derivative of $100y/(x + y) - y$ with respect to y equal to zero. It is $100x/(x + y)^2 - 1 = 0$, or $y = 10\sqrt{x} - x$. Figure 5.2 shows its graph and that of the analogous best-response function of party L—namely, $x = 10\sqrt{y} - y$.

Look at the best-response curve of party R. As the value of party L's x increases, party R's y increases for a while and then decreases. If the other party is advertising very little, then one's own ads have a high reward in the form of votes, and it pays to respond to a small increase in the other's expenditures by spending more oneself to compete harder. But if the other party already has a

massive expenditure, then one's own ads get only a small return in relation to their cost, so it is better to respond to the other's increase by scaling back.

As it happens, the two parties' best-response curves intersect at their peak points. Again, some algebraic manipulation of the equations for the two curves yields us exact values for the equilibrium values of x and y. You should verify that here x and y are each equal to 25, or \$25 million. (This is presumably a congressional election; Senate and presidential elections cost much more these days.)

As in the pricing game, we have a prisoners' dilemma. If both parties cut back on their ads in equal proportions, their vote shares would be entirely unaffected, but both would save on their expenditures and so both would have a larger payoff. Unlike a producers' cartel for substitute products that keeps prices high and hurts consumers, but like a producers' cartel for complements that leads to lower prices, a politicians' cartel to advertise less would probably benefit voters and society. We could all benefit from finding ways to resolve this particular prisoners' dilemma. In fact, Congress has been trying to do just that for several years and has imposed some partial curbs, but political competition seems too fierce to permit a full or lasting resolution.

What if the parties are not symmetrically situated? Two kinds of asymmetries can arise. One party (say, R) may be able to advertise at a lower cost, because it has favored access to the media. Or R's advertising dollars may be more effective than L's; for example, L's vote share may be $x/(x + 2y)$, while R's is $2y/(x + 2y)$.

In the first of these cases, R exploits its cheaper access to advertising by choosing a higher level of expenditures y for any given x for party L; that is, R's best-response curve in Figure 5.2 shifts upward. The Nash equilibrium shifts to the northwest along L's unchanged best-response curve. Thus R ends up advertising more, and L less, than before. It is as if the advantaged party uses its muscle and the disadvantaged party gives up to some extent in the face of this adversity.

In the second case, both parties' best-response curves shift in more complex ways. The outcome is that both spend equal amounts, but less than the 25 that they spent in the symmetric case. In our example where R's dollars are twice as effective as L's, it turns out that their common expenditure level is $200/9 = 22.2 < 25$. (Thus the symmetric case is the one of most intense competition.) When R's spending is more effective, it is also true that the best-response curves are asymmetric in such a way that the new Nash equilibrium, rather than being at the peak points of the two best-response curves, is on the downward part of L's best-response curve and on the upward part of R's best-response curve. That is to say, although both parties spend the same dollar amount, the favored party, R, spends more than the amount that would bring forth the maximum response from party L, and the underdog party, L, spends less than the amount that would bring forth the maximum response from party R. We include an optional exercise (Exercise U12) in this chapter that lets the mathematically advanced students derive these results.

D. General Method for Finding Nash Equilibria

Although the strategies (prices or campaign expenditures) and payoffs (profits or vote shares) in the two previous examples are specific to the context of competition between firms or political parties, the method for finding the Nash equilibrium of a game with continuous strategies is perfectly general. Here we state its steps so you can use it as a recipe for solving other games of this kind.

Suppose the players are numbered 1, 2, 3, . . . Label their strategies $x, y, z, \ldots$ in that order, and their payoffs by the corresponding upper-case letters $X, Y, Z, \ldots$ The payoff of each is in general a function of the choices of all; label the respective functions $F, G, H, \ldots$ Construct payoffs from the information about the game, and write them as

$$X = F(x, y, z, \ldots),\ Y = G(x, y, z, \ldots),\ Z = H(x, y, z, \ldots).$$

Using this general format to describe our example of price competition between two players (firms) makes the strategies x and y become the prices P_x and P_y. The payoffs X and Y are the profits Π_x and Π_y. The functions F and G are the quadratic formulas,

$$\Pi_x = -8(44 + P_y) + (16 + 44 + P_y)P_x - 2(P_x)$$

and similarly for Π_y.

In the general approach, player 1 regards the strategies of players 2, 3, . . . as outside his control, and chooses his own strategy to maximize his own payoff. Therefore for each given set of values of $y, z, \ldots$, player 1's choice of x maximizes $X = F(x, y, z, \ldots)$. If you use calculus, the condition for this maximization is that the derivative of X with respect to x holding $y, z, \ldots$ constant (the partial derivative) equals zero. For special functions, simple formulas are available, such as the one we stated and used above for the quadratic. And even if an algebra or calculus formulation is too difficult, computer programs can tabulate or graph best-response functions for you. Whatever method you use, you find an equation for player 1's optimal choice of x for given $y, z, \ldots$ that is player 1's best-response function. Similarly, you can find the best-response functions for each of the other players.

The best-response functions are equal in number to the number of the strategies in the game and can be solved simultaneously while regarding the strategy variables as the unknowns. The solution is the Nash equilibrium we seek. Some games may have multiple solutions, yielding multiple Nash equilibria. Other games may have no solution, requiring further analysis, such as inclusion of mixed strategies.

2 EMPIRICAL EVIDENCE CONCERNING NASH EQUILIBRIUM

In Chapter 3, when we considered empirical evidence on sequential-move games and rollback, we said that the evidence came from observations on games actually played in the world, as well as games deliberately constructed for testing the theory in the laboratory or the classroom. We pointed out the different merits and drawbacks of the two methods. Similar issues arise in securing and interpreting the evidence on simultaneous-move games.

Real-world games are played for substantial stakes by experienced players, who therefore have the knowledge and the incentives to employ good strategies. But these situations include many factors beyond those considered in the theory. Therefore, if the data do not bear out the predictions of the theory, we cannot tell whether the theory is wrong or whether some other factor is having an effect that overwhelms the strategic considerations.

Laboratory experiments can control for the other factors and therefore provide cleaner tests. But they bring in inexperienced players and provide them meager time and incentives to learn the game and play it well. Confronted with a new game, most of us would flounder and try things out at random. Thus the first several plays of the game in an experimental setting may represent this learning phase and not the equilibrium that experienced players would learn to play. Some experiments do control for inexperience and learning by discarding several initial plays from their data. But the learning phase may last longer than the one morning or one afternoon that is the typical limit of laboratory sessions.

A. Laboratory and Classroom Experiments

Researchers have conducted numerous laboratory experiments in the past three decades to test how people act when placed in certain interactive strategic situations. In particular, do they play Nash equilibrium strategies? Reviewing this work, Douglas Davis and Charles Holt conclude that, in relatively simple single-move games with a unique Nash equilibrium, that outcome "has considerable drawing power . . . after some repetitions with different partners."[6] But in more complex or repeated situations or when coordination is required because there are multiple Nash equilibria or when the calculations required for finding a Nash equilibrium are more complex, the theory's success is more mixed. We briefly consider the performance of Nash equilibrium in such circumstances.

[6] Douglas D. Davis and Charles A. Holt, *Experimental Economics* (Princeton: Princeton University Press, 1993), Ch. 2.

I. CHOOSING AMONG MULTIPLE EQUILIBRIA When there are multiple equilibria, players generally fail to coordinate unless they have some common cultural background (and this fact is common knowledge among them) that is needed for locating focal points. Thomas Schelling and David Kreps report on several experiments of coordination games.[7] Kreps played the following game between pairs of his students. One student was assigned Boston, and the other was assigned San Francisco. Each was given a list of nine other U.S. cities—Atlanta, Chicago, Dallas, Denver, Houston, Los Angeles, New York, Philadelphia, and Seattle—and asked to choose a subset of these cities. The two chose simultaneously and independently. If their choices divided up the nine cities completely and without any overlap between them, both got a prize. Otherwise, neither got anything. This game has numerous (512) Nash equilibria in pure strategies. But, when both players were Americans or long-time U.S. residents, more than 80% of the time they chose the division geographically; the student assigned Boston chose all the cities east of the Mississippi, and the student assigned San Francisco chose those west of the Mississippi. Such coordination was much less likely when one or both students were non-U.S. residents. In such pairs the choices were sometimes made alphabetically, but even then there was no clear dividing point.

II. REVELATION OF INNATE ALTRUISM OR PUBLIC-SPIRITEDNESS IN EXPERIMENTS One respect in which the behavior of players in some experimental situations does not often conform to the experimenter's predicted Nash equilibrium is that people seem to "err" on the side of niceness or fairness. Thus in prisoners' dilemma games, we observe "too much" cooperation and, in bargaining games, people concede "too much" to opponents. But the reason may not be any failure to calculate or to learn to play Nash equilibrium. It may instead be that the players' payoffs are different from those assumed by the experimenter.

As with observations of naturally occurring games, participants in experimental situations also may know some complexities of the situation better than the experimenter knows them. For example, the possibility of repetition or a separate ongoing relationship with the other player may affect their choices in this game. Or the players' value systems may have internalized some social norms of niceness and fairness that have proved useful in the larger social context and that therefore carry over to their behavior in the experimental game.[8]

[7]Thomas Schelling, *The Strategy of Conflict* (New York: Oxford University Press, 1960), pp. 54–58; David Kreps, *A Course in Microeconomic Theory* (Princeton: Princeton University Press, 1990), pp. 302–303, 414–415.

[8]The distinguished game theorist Jörgen Weibull argues this position in detail in "Testing Game Theory," in *Advances in Understanding Strategic Behaviour: Game Theory, Experiments and Bounded Rationality: Essays in Honour of Werner Güth*, ed. Steffen Huck (Basingstoke: Palgrave MacMillan, 2004).

These observations do not show any deficiency of the Nash equilibrium concept itself. However, they do warn us against using the concept under naive or mistaken assumptions about people's payoffs; it might be a mistake, for example, to assume that players are always driven by the selfish pursuit of money.

III. LEARNING FROM EXPERIENCE One game, often used in classrooms or laboratories, asks each participant to choose a number between 0 and 100. Typically, the players are handed cards on which to write their names and a choice, so this game is a simultaneous-move game. When the cards are collected, the average of the numbers is calculated. The person whose choice is closest to a specified fraction—say half—of the average is the winner. The rules of the game (this whole procedure) are announced in advance.

The Nash equilibrium of this game is for everyone to choose 0. In fact the game is dominance solvable. Even if everyone chooses 100, half of the average can never exceed 50; so, for each player, any choice above 50 is dominated by 50.[9] But all players should rationally figure this out, so the average can never exceed 50 and half of it can never exceed 25, and so any choice above 25 is dominated by 25. The iteration goes on until only 0 is left.

However, when a group actually plays this game for the first time, the winner is typically a player who has chosen a number just a little less than 25. This outcome seems to suggest that the winner assumes that everyone else will choose randomly (so their average is 50) and then chooses her own best response to that. The outcome is quite far from the Nash equilibrium.

What happens if the game is repeated? Our experience in classroom trials has been that the winning choice falls rapidly in successive plays. In Skeath's class, half the class played the game first while the other half watched, then the other half played, and finally the whole class played. In Dixit's class, the game was played by different groups of 10 students at a time. By the third round, the winner's choice was usually as low as 2 or 3.

How should one interpret this result? Critics would say that, unless the exact Nash equilibrium is reached, the theory is refuted. Indeed, they would argue, if you have good reason to believe that other players will not play their Nash equilibrium strategies, then your best choice is not your Nash equilibrium strategy either. If you can figure out how others will deviate from their Nash equilibrium strategies, then you should play your best response to what you believe they are choosing. Others would argue that theories in social science can never hope for the kind

[9]If you factor in your own choice, the calculation is strengthened. Suppose there are N players. In the "worst-case scenario" where all the other $(N - 1)$ players choose 100 and you choose x, the average is $[x + (N - 1)100]/N$. Then your best choice is half of this; so $x = [x + (N - 1)100]/(2N)$, or $x = 100(N - 1)/(2N - 1)$. If $N = 10$, then $x = 900/19 = 47$ (approximately). So any choice above 47 is dominated by 47. The same reasoning applies to the successive rounds.

of precise prediction that we expect in sciences such as physics and chemistry. If the observed outcomes are close to the Nash equilibrium, that is a vindication of the theory. In this case, the experiment not only produces such a vindication, but illustrates the process by which people gather experience and learn to play strategies close to Nash equilibrium. We sympathize with this latter viewpoint.

Interestingly, we have found that people learn somewhat faster by observing others play a game than while they play it themselves. This may be because, as observers, they are free to focus on the game as a whole and think about it analytically. Players' brains are occupied with the task of making their own choices and they are less able to take the broader perspective.

We should clarify the concept of gaining experience by playing the game. The quotation from Davis and Holt at the start of this section spoke of "repetitions with different partners." In other words, experience should be gained by playing the game frequently, but with different opponents each time. However, for any learning process to generate outcomes increasingly closer to the Nash equilibrium, the whole *population* of learners needs to be stable. If novices keep appearing on the scene and trying new experimental strategies, then the original group may unlearn what they had learned by playing against one another.

If a game is played repeatedly between two players or even among the same small group of known players, then any pair is likely to play each other repeatedly. In such a situation, the whole repeated game becomes a game in its own right. It can have very different Nash equilibria from those that simply repeat the Nash equilibrium of a single play. For example, tacit cooperation may emerge in repeated prisoners' dilemmas, owing to the expectation that any temporary gain from cheating will be more than offset by the subsequent loss of trust. If games are repeated in this way, then learning about them must come from playing whole sets of the repetitions frequently, against different partners each time.

B. Real-World Evidence

The predictions of game theory have been subjected to real-world empirical evidence in two distinct ways. One is to see if the theory can explain some observed phenomena in general terms. The other is to test statistically some implications of the theory against data. We briefly consider each in turn.

I. EXPLANATORY POWER The first approach uses game-theoretic reasoning to explain phenomena that are observed in reality, of which there are numerous successful examples. One of the earliest was in the area of international relations. Thomas Schelling pioneered the use of game theory to explain phenomena such as the escalation of arms races, even between countries that have no intention of attacking each other, and the credibility of deterrent threats. Subsequent applications in this area have included the questions of when and how a country

can credibly signal its resolve in diplomatic negotiation or in the face of a potential war. Game theory began to be used systematically in economics and business in the mid-1970s, and such applications continue to proliferate. We have space for only a couple of prominent examples.

The theory has helped us to understand when and how the established firms in an industry can make credible commitments to deter new competition—for example, to wage a destructive price war against any new entrant. The prisoners' dilemma game, in its one-time and repeated forms, has helped us to understand what kinds of industries will see fierce competition and exhibit low prices and what kinds will sustain tacit agreements to keep prices and profits high. More recently, game theory has become the tool of choice for the study of political systems and institutions within a country as well as for cross-country comparisons. For example, game theory has shown how voting and agenda setting in committees and elections can be strategically manipulated in pursuit of one's ultimate objectives. In this introductory book, we can develop only a few elementary examples of this kind. We already saw an example (price competition) in this chapter. More examples appear later, including a case study of the Cuban missile crisis and analyses of auctions, voting, and bargaining.[10]

Some critics remain unpersuaded by these successful applications of the theory. They claim that the same understanding of these phenomena can be obtained without using game theory, by basing one's analysis on previously known general principles of economics, political science, and so on. In one sense they are right. A few of these analyses existed before game theory came along. For example, the equilibrium of the interaction between two price-setting firms, which we developed in Section 1 of this chapter, was known in economics for more than a hundred years; one can think of Nash equilibrium as just a general formulation of that equilibrium concept for all games. Some theories of strategic voting date to the 18th century, and some notions of credibility can be found in history as far back as Thucydides' *Peloponnesian War*. However, what game theory does is to unify all these applications and thereby facilitate the development of new ones.

In the past 30 years, several new ideas and applications have been identified. For example, we now understand how different forms of auctions (English

[10]For those who would like to see more applications, here are some suggested sources. Thomas Schelling's *Strategy of Conflict* (New York: Oxford University Press, 1960) and *Arms and Influence* (New Haven: Yale University Press, 1966) are still required reading for all students of game theory. The classic textbook on game-theoretic treatment of industries is Jean Tirole, *Industrial Organization* (Cambridge: MIT Press, 1988). In political science, an early classic is William Riker, *Liberalism Against Populism* (San Francisco: W. H. Freeman, 1982). For surveys of more recent work, see several articles in *The Handbook of Game Theory*, ed. Robert J. Aumann and Sergiu Hart (Amsterdam: North-Holland, 1992, 1994, 2002), particularly Barry O'Neill, "Game Theory Models of Peace and War," in volume 2, and Kyle Bagwell and Asher Wolinsky, "Game Theory and Industrial Organization," and Jeffrey Banks, "Strategic Aspects of Political Systems," in volume 3.

and Dutch, sealed bid and open outcry) lead to differences in bidding strategies and in the seller's revenues. We understand how the existence of a second-strike capability reduces the fear of surprise attack. And we understand how governments can successfully manipulate fiscal and monetary policies to improve their chances of reelection even when the voters are sophisticated and aware of such attempts. If these examples were all amenable to previously known approaches, they would have been discovered long ago.

II. STATISTICAL TESTING The second approach to examining empirical evidence is quantitative and statistical. The general procedure in this work is to assume that Nash equilibrium prevails and to derive the implications of this assumption in the form of equations linking various magnitudes—the players' choices and outcomes—that may be observable in the situation being studied. These equations can then be estimated by using real data. In industrial economics, firms compete by choosing their quantities and prices as illustrated in the examples in this chapter; they also have other strategic choices at their disposal, including product quality, investment, R & D, and so on. While the choice of quantities or prices may be studied in a static context (at a given time), games of strategic competition in investment or R & D are dynamic. Numerous studies of both kinds of interactions have been carried out.[11] This work has produced encouraging results. Game-theoretic models, based on the Nash equilibrium concept and its dynamic generalizations, fit the data for many major industries, such as automobile manufacturers, reasonably well and give us a better understanding of the determinants of competition than the older analysis, which assumed perfect competition and estimated supply-and-demand curves.

In politics, the votes on various issues within legislatures are the outcome of the legislators' strategic interaction. The equilibrium of this game depends on the legislators' underlying preferences. Detailed voting records in the U.S. Congress are public information. On the basis of the relation between preferences and voting in a Nash equilibrium, these data can be used to infer the legislators' preferences. This method has been used with remarkable success by Keith Poole and Howard Rosenthal.[12] They find that U.S. politics can be

[11]A survey of static studies of prices and quantity competition is "Empirical Studies of Industries with Market Power," by Timothy F. Bresnahan, in *Handbook of Industrial Organization*, ed. Richard L. Schmalansee and Robert D. Willig, vol. 2 (Amsterdam: North-Holland, 1989). A general method for dynamic studies is developed in "A Framework for Applied Dynamic Analysis in Industrial Organization," by Ulrich Doraszelski and Ariel Pakes, in *Handbook of Industrial Economics*, ed. Mark Armstrong and Robert Porter, vol. 3 (Amsterdam: North-Holland, 2007).

[12]Keith Poole and Howard Rosenthal, "Patterns of Congressional Voting," *American Journal of Political Science*, vol. 35, no. 1 (February 1991), pp. 228–278, and *Congress: A Political-Economic History of Roll Call Voting* (New York: Oxford University Press, 1996).

adequately summarized by conflicts over issues in a two-dimensional space, one representing economic inequality and the other racial inequality.

Pankaj Ghemawat, a professor at the Harvard Business School, has developed a mixed mode of quantitative analysis, using case studies of individual firms or industries and statistical analysis of larger data samples.[13] His game-theoretic models are remarkably successful in improving our understanding of several initially puzzling business decisions on pricing, capacity, innovation, and so on. However, his work also brings out the need to construct models that have sufficiently rich detail to do justice to the circumstances of the firms or industries being analyzed. In a general and introductory textbook such as this one, we lack the space and eschew the more advanced techniques that are needed to construct such models. But we will set you on the way to further study that will bring these methods within your grasp. And in Chapter 15 we develop one such theory-based case study from the field of international politics to illustrate the method.

III. A REAL-WORLD EXAMPLE OF LEARNING We conclude by offering an interesting illustration of equilibrium and the learning process in a real-world game. The setting is outside the laboratory or classroom, where people play the game frequently and the stakes are high, creating strong motivation and good opportunities to learn. Stephen Jay Gould discovered this beautiful example.[14] Through most of the 20th century, the best batting averages recorded in a baseball season have been declining. In particular, the number of instances of a player averaging .400 or better used to be much more frequent than they are now. Devotees of baseball history often explain this decline by invoking nostalgia: "There were giants in those days." A moment's thought should make one wonder why there were no corresponding pitching giants who would keep batting averages low. But Gould demolishes such arguments in a more systematic way. He points out that we should look at all batting averages, not just the top ones. The worst batting averages are not as bad as they used to be; there are also many fewer .150 hitters in the major leagues than there used to be. He argues that this overall decrease in *variation* is a standardization or stabilization effect:

> When baseball was very young, styles of play had not become sufficiently regular to foil the antics of the very best. Wee Willie Keeler could "hit 'em where they ain't" (and compile an average of .432 in 1897) because fielders didn't yet know where they should be. Slowly, players moved toward *optimal* methods of positioning, fielding, pitching, and batting—and variation inevitably declined. The best [players] now met an opposition too finely honed to its own perfection to permit the extremes of achievement that characterized a more casual age. [emphasis added]

[13]Pankaj Ghemawat, *Games Businesses Play: Cases and Models* (Cambridge: MIT Press, 1997).

[14]"Losing the Edge," in *The Flamingo's Smile* (New York: Norton, 1985), pp. 215–229.

In other words, through a succession of adjustments of strategies to counter one another, the system settled down into its (Nash) equilibrium.

Gould marshals decades of hitting statistics to demonstrate that such a decrease in variation did indeed occur, except for occasional "blips." And indeed the blips confirm his thesis, because they occur soon after an equilibrium is disturbed by an externally imposed change. Whenever the rules of the game are altered (the strike zone is enlarged or reduced, the pitching mound is lowered, or new teams and many new players enter when an expansion takes place) or the technology changes (a livelier ball is used or perhaps, in the future, aluminum bats are allowed), the preceding system of mutual best responses is thrown out of equilibrium. Variation increases for a while as players experiment, and some succeed while others fail. Finally a new equilibrium is attained, and variation goes down again. That is exactly what we should expect in the framework of learning and adjustment to a Nash equilibrium.

We take up the evidence concerning mixed strategies in Chapter 8 and the evidence for some specific games or types of games—for example, the prisoners' dilemma, bargaining, and auctions—at appropriate points in later chapters. For now, the experimental and empirical evidence that we have presented should make you cautiously optimistic about using Nash equilibrium as a first approach or as the point of departure for your analysis. On the whole, we believe you should have considerable confidence in using the Nash equilibrium concept when the game in question is played frequently by players from a reasonably stable population and under relatively unchanging rules and conditions. When the game is new or is played just once and the players are inexperienced, you should use the equilibrium concept more cautiously and should not be surprised if the outcome that you observe is not the equilibrium that you calculate. But even then, your first step in the analysis should be to look for a Nash equilibrium; then you can judge whether it seems a plausible outcome and, if not, proceed to the further step of asking why not. Often the reason will be your misunderstanding of the players' objectives, not the players' failure to play the game correctly, given their true objectives.

3 CRITICAL DISCUSSION OF THE NASH EQUILIBRIUM CONCEPT

In addition to the critiques lodged against the Nash equilibrium concept by those who have examined the empirical evidence, there have also been theoretical criticisms of the concept. In this section, we briefly review some such criticisms and some rebuttals, in each case by using an example.[15] Some of the

[15]David M. Kreps, *Game Theory and Economic Modelling* (Oxford: Clarendon Press, 1990) gives an excellent in-depth discussion.

criticisms are mutually contradictory, and some can be countered by thinking of the games themselves in a better way. Others tell us that the Nash equilibrium concept by itself is not enough and suggest some augmentations or relaxations of it that have better properties. We develop one such alternative here and point to some others that appear in later chapters. We believe our presentation will leave you with renewed but cautious confidence in using the Nash equilibrium concept. But some serious doubts remain unresolved, indicating that game theory is not yet a settled science. Even this should give encouragement, not the opposite, to budding game theorists, because it shows that there is a lot of room for new thinking and new research in the subject. A totally settled science would be a dead science.

We begin by considering the basic appeal of the Nash equilibrium concept. Most of the games in this book are noncooperative, in the sense that every player takes her action independently. Therefore it seems natural to suppose that, if her action is not the best according to her own value system (payoff scale), given what everyone else does, then she will change it. In other words, it is appealing to suppose that every player's action will be the best response to the actions of all the others. Nash equilibrium has just this property of "simultaneous best responses"; indeed, that is its very definition. In any purported final outcome that is not a Nash equilibrium, at least one player could have done better by switching to a different action.

This consideration leads eminent game theorist Roger Myerson to rebut those criticisms of the Nash equilibrium that are based on the intuitive appeal of playing a different strategy. His rebuttal simply shifts the burden of proof onto the critic. "When asked why players in a game should behave as in some Nash equilibrium," he says, "my favorite response is to ask 'Why not?' and to let the challenger specify what he thinks the players should do. If this specification is not a Nash equilibrium, then . . . we can show that it would destroy its own validity if the players believed it to be an accurate description of each other's behavior."[16]

A. The Treatment of Risk in Nash Equilibrium

Some critics argue that the Nash equilibrium concept does not pay due attention to risk. In some games, people might find strategies different from their Nash equilibrium strategies to be safer and may therefore choose those strategies. We offer two examples of this kind. The first is due to John Morgan, an economics professor at the University of California, Berkeley; Figure 5.3 shows the game table.

Cell-by-cell inspection quickly reveals that this game has a unique Nash equilibrium—namely, (A, A), yielding the payoffs (2, 2). But you may think, as

[16]Roger Myerson, *Game Theory* (Cambridge: Harvard University Press, 1991), p. 106.

		COLUMN		
		A	B	C
ROW	A	2, 2	3, 1	0, 2
	B	1, 3	2, 2	3, 2
	C	2, 0	2, 3	2, 2

FIGURE 5.3 A Game with a Questionable Nash Equilibrium

did several participants in an experiment conducted by Morgan, that playing C has a lot of appeal, for the following reasons. It *guarantees* you the same payoff as you would get in the Nash equilibrium—namely, 2; whereas if you play your Nash equilibrium strategy A, you will get a 2 only so long as the other player also plays A. Why take that chance? What is more, if you think the other player might use this rationale for playing C, then you would be making a serious mistake by playing A; you would get only a 0 when you could have gotten your guaranteed 2 by playing C.

Myerson would respond, "Not so fast. If you really believe that the other player would think thus and would play C, then you should play B to get the payoff 3. And if you think the other person would think thus and so would play B, then your best response to B is A. And if you think the other would figure this out too, you should be playing your best response to A, namely A. Back to the Nash equilibrium!" As you can see, criticizing Nash equilibrium and rebutting the criticisms is itself something of an intellectual game, and quite a fascinating one.

The second example, due to David Kreps, is even more dramatic. The payoff matrix is in Figure 5.4. Before doing any theoretical analysis of this game, you should pretend that you are actually playing the game and that you are player A. Which of your two actions would you choose?

Keep in mind your answer to the preceding question and let us proceed to analyze the game. If we start by looking for dominant strategies, we see that player A has no dominant strategy but player B does. Playing Left guarantees B a payoff of 10, no matter what A does, versus the 9.9 that is gained by playing Right (also no matter what A does). Thus, player B should play Left. Given that player B is going to go Left, player A does better to go Down. The unique pure-strategy Nash equilibrium of this game is (Down, Left); each player achieves a payoff of 10 at this outcome.

The problem that arises here is that many people (but not all) would not, as player A, choose to play Down. (What did you choose?) This is true for those who have been students of game theory for years as well as for those who have

		B	
		Left	Right
A	Up	9, 10	8, 9.9
	Down	10, 10	−1000, 9.9

FIGURE 5.4 Disastrous Nash Equilibrium?

never heard of the subject. If A has *any* doubts about *either* B's payoffs *or* B's rationality, then it is a lot safer for A to play Up than to play her Nash equilibrium strategy of Down. What if A thought the payoff table was as illustrated in Figure 5.4 but in reality B's payoffs were the reverse—the 9.9s went with Left and the 10s went with Right? What if the 9.9s were only an approximation and the exact payoffs were actually 10.1? What if B was a player with a substantially different value system or was not a truly rational player who might choose the "wrong" action just for fun? Obviously, our assumptions of perfect information and rationality can really be crucial to the analysis that we use in the study of strategy. Doubts of players can alter equilibria from those that we would normally predict and can call the reasonableness of the Nash equilibrium concept into question.

However, the real problem with many such examples is not that the Nash equilibrium concept is inappropriate but that the examples choose to use it in an inappropriately simplistic way. In this example, if there are any doubts about B's payoffs, then this fact should be made an integral part of the analysis. If A does not know B's payoffs, the game is one of asymmetric information, and we do not develop the general techniques for studying such games until Chapter 9. But this particular example is a relatively simple game of that kind, and we can figure out its equilibrium very easily.

Suppose A thinks there is a probability p that B's payoffs from Left and Right are the reverse of those shown in Figure 5.4; so $(1 - p)$ is the probability that B's payoffs are as stated in that figure. Because A must take her action without knowing which is the case, she must choose her strategy to be "best on the average." In this game the calculation is simple, because in each case B has a dominant strategy; the only problem for A is that in the two different cases different strategies are dominant for B. With probability $(1 - p)$, B's dominant strategy is Left (the case shown in the figure) and, with probability p, it is Right (the opposite case). Therefore if A chooses Up, then with probability $(1 - p)$ he will meet B playing Left and so get a payoff of 9; with probability p, he will meet B playing Right and so get a payoff of 8. Thus A's statistical or probability-weighted average payoff from playing Up is $9(1 - p) + 8p$. Similarly, A's statistical average payoff from playing Down is $10(1 - p) - 1000p$. Therefore it is better for A to choose Up if

$$9(1 - p) + 8p > 10(1 - p) - 1000p, \quad \text{or} \quad p > 1/1009.$$

Thus, even if there is only a very slight chance that B's payoffs are the opposite of those in Figure 5.4, it is optimal for A to play Up. In this case, analysis based on rational behavior, when done correctly, contradicts neither the intuitive suspicion nor the experimental evidence after all.

In the preceding calculation, we supposed that, facing an uncertain prospect of payoffs, player A would calculate the statistical average payoffs from her different actions and would choose that action which yields her the highest statistical average payoff. This implicit assumption, though it serves the purpose in this example, is not without its own problems. For example, it implies that a person faced with two situations, one having a 50-50 chance of winning or losing \$10 and the other having a 50-50 chance of winning \$10,001 and losing \$10,000, should choose the second situation, because it yields a statistical average winning of 50 cents ($\frac{1}{2} \times 10{,}001 - \frac{1}{2} \times 10{,}000$), whereas the first yields 0 ($\frac{1}{2} \times 10 - \frac{1}{2} \times 10$). But most people would think that the second situation carries a much bigger risk and would therefore prefer the first situation. This difficulty is quite easy to resolve. In the Appendix to Chapter 7, we show how the construction of a scale of payoffs that is suitably nonlinear in money amounts enables the decision maker to allow for risk as well as return Then in Chapter 9, we show how the concept can be used for understanding how people respond to the presence of risk in their lives, for example, by arranging the sharing of risk with others, or through the provision of insurance.

B. Multiplicity of Nash Equilibria

Another criticism of the Nash equilibrium concept is based on the observation that many games have multiple Nash equilibria. Thus, the argument goes, the concept fails to pin down outcomes of games sufficiently precisely to give unique predictions. This argument does not automatically require us to abandon the Nash equilibrium concept. Rather, it suggests that if we want a unique prediction from our theory, we must add some criterion for deciding which one of the multiple Nash equilibria we want to select.

In Chapter 4, we studied many games of coordination with multiple equilibria. From among these equilibria, the players may be able to select one as a focal point if they have some common social, cultural, or historical knowledge or if the game has some deliberate or accidental features that enable their expectations to converge. Here is a very extreme example of multiplicity of Nash equilibria in a coordination game. Two players are asked to write down, simultaneously and independently, the share that each wants of a total prize of \$100. If the amounts that they write down add up to \$100 or less, each player is given what she wrote. If the two add up to more than \$100, neither gets anything. For any x, one player writing x and the other writing $(100 - x)$ is a Nash equilibrium. Thus the game has an (almost) infinite range of Nash equilibria. But in practice, 50 : 50 emerges as a focal point. This social norm of equality or fairness seems so deeply ingrained as to be almost an instinct; players who choose 50 say that it is the obvious answer. To be a true focal point, not only should it be obvious to each, but everyone should know that it is obvious to each, and everyone should know that . . . ; in other words, its obviousness should be common knowledge. That need not always be the case, as we see when we consider a situation in which one player is a woman from an enlightened and egalitarian society who believes that 50 : 50 is obvious and the other is a man from a patriarchal society who believes it is obvious that, in any matter of division, a man should get three times as much as a woman. Then each will do what is obvious to her or him, and they will end up with nothing, because neither's obvious solution is obvious as common knowledge to both.

The existence of focal points is often a matter of coincidence, and creating them where none exist is basically an art that requires a lot of attention to the historical and cultural context of a game and not merely its mathematical description. This bothers many game theorists, who would prefer the outcome to depend only on an abstract specification of a game—players and their strategies should be identified by numbers without any external associations. We disagree. We think that historical and cultural contexts are just as important to a game as its purely mathematical description, and, if such context helps in selecting a unique outcome from multiple Nash equilibria, that is all to the better.

In Chapter 6, we will see that sequential-move games can have multiple Nash equilibria. There, we introduce the requirement of *credibility* that enables us to select a particular equilibrium; it turns out that this one is in fact the rollback equilibrium of Chapter 3. In more complex games with information asymmetries or additional complications, other restrictions called **refinements** have been developed to identify and rule out Nash equilibria that are unreasonable in some way. In Chapter 9, we consider one such refinement process that selects an outcome called a *perfect Bayesian equilibrium*. The motivation for each refinement is often specific to a particular type of game. A refinement stipulates how players update their information when they observe what moves other players made or failed to make. Each such stipulation is often perfectly reasonable in its context, and in many games it is not difficult to eliminate most of the Nash equilibria and therefore to narrow down the ambiguity in prediction.

The opposite of the criticism that some games may have too many Nash equilibria is that some games may have none at all. We saw an example of this in Section 4.8 and said that, by extending the concept of strategy to random mixtures, Nash equilibrium could be restored. In Chapters 7 and 8 we explain and consider Nash equilibria in mixed strategies. In higher reaches of game theory, there are more esoteric examples of games that have no Nash equilibrium in mixed strategies either. However, this added complication is not relevant for the types of analysis and applications that we deal with in this book, so we do not attempt to address it here.

C. Requirements of Rationality for Nash Equilibrium

Remember that Nash equilibrium can be regarded as a system of the strategy choices of each player and the belief that each player holds about the other players' choices. In equilibrium, (1) the choice of each should give her the best payoff given her belief about the others' choices, and (2) the belief of each player should be correct—that is, her actual choices should be the same as what this player believes them to be. These seem to be natural expressions of the requirements of the mutual consistency of individual rationality. If all players have common knowledge that they are all rational, how can any one of them rationally believe something about others' choices that would be inconsistent with a rational response to her own actions?

To begin to address this question, we consider the three-by-three game in Figure 5.5. Cell-by-cell inspection quickly reveals that it has only one Nash equilibrium—namely, (R2, C2), leading to payoffs (3, 3). In this equilibrium, Row plays R2 because she believes that Column is playing C2. Why does she believe this? Because she knows Column to be rational, Row must simultaneously believe that Column believes that Row is choosing R2, because C2 would not be Column's best choice if she believed Row would be playing either R1 or R3. Thus, the claim goes, in any rational process of formation of beliefs and responses, beliefs would have to be correct.

The trouble with this argument is that it stops after one round of thinking about beliefs. If we allow it to go far enough, we can justify other choice combinations. We can, for example, rationally justify Row's choice of R1. To do so, we note that R1 is Row's best choice if she believes Column is choosing C3. Why does she believe this? Because she believes that Column believes that Row is playing R3. Row justifies this belief by thinking that Column believes that Row

		COLUMN		
		C1	C2	C3
ROW	R1	0, 7	2, 5	7, 0
	R2	5, 2	3, 3	5, 2
	R3	7, 0	2, 5	0, 7

FIGURE 5.5 Justifying Choices by Chains of Beliefs and Responses

believes that Column is playing C1, believing that Row is playing R1, believing in turn . . . This is a chain of beliefs, each link of which is perfectly rational.

Thus rationality alone does not justify Nash equilibrium. There are more sophisticated arguments of this kind that do justify a special form of Nash equilibrium in which players can condition their strategies on a publicly observable randomization device. But we leave that to more advanced treatments. In the next section, we develop a simpler concept that captures what is logically implied by the players' common knowledge of their rationality alone.

4 RATIONALIZABILITY

What strategy choices in games can be justified on the basis of rationality alone? In the matrix of Figure 5.5, we can justify any pair of strategies, one for each player, by using the same type of logic as that used in Section 3.C. In other words, we can justify any one of the nine logically conceivable combinations. Thus rationality alone does not give us any power to narrow down or predict outcomes at all. Is this a general feature of all games? No. For example, if a strategy is dominated, rationality alone can rule it out of consideration. And when players recognize that other players, being rational, will not play dominated strategies, iterated elimination of dominated strategies can be performed on the basis of common knowledge of rationality. Is this the best that can be done? No. Some more ruling out of strategies can be done, by using a property slightly stronger than being dominated in pure strategies. This property identifies strategies that are **never a best response.** The set of strategies that survive elimination on this ground are called **rationalizable,** and the concept itself is known as **rationalizability.**

Why introduce this additional concept, and what does it do for us? As for why, it is useful to know how far we can narrow down the possible outcomes of a game based on the players' rationality alone, without invoking correctness of expectations about the other player's actual choice. It is sometimes possible to figure out that the other player *will not* choose some available action or actions, even when it is not possible to pin down the single action that she *will* choose. As for what it achieves, that depends on the context. In some cases rationalizability may not narrow down the outcomes at all. This was so in the three-by-three example of Figure 5.5. In some cases it narrows down the possibilities to some extent, but not all the way down to the Nash equilibrium if the game has a unique one, or to the set of Nash equilibria if there are several. An example of such a situation is the four-by-four enlargement of the previous example, considered in Section 4.A. In some other cases, the narrowing down may go all the way to the Nash equilibrium; in these cases we have a more powerful justification for the Nash equilibrium that relies on rationality alone, without assuming correctness of expectations. The quantity competition example of Section 4.B is an example in which the rationalizability argument takes us all the way to the game's unique Nash equilibrium.

A. Applying the Concept of Rationalizability

Consider the game in Figure 5.6, which is the same as Figure 5.5 but with an additional strategy for each player.[17] We just indicated that nine of the strategy combinations that pick one of the first three strategies for each of the players can be justified by a chain of beliefs about each other's beliefs. That remains true in this enlarged matrix. But can R4 and C4 be justified in this way?

Could Row ever believe that Column would play C4? Such a belief would have to be justified by Column's beliefs about Row's choice. What might Column believe about Row's choice that would make C4 Column's best response? Nothing. If Column believes that Row would play R1, then Column's best choice is C1. If Column believes that Row will play R2, then Column's best choice is C2. If Column believes that Row will play R3, then C3 is Column's best choice. And, if Column believes that Row will play R4, then C1 and C3 are tied for her best choice. Thus C4 is never a best response for Column.[18] This means that Row, knowing Column to be rational, can never attribute to Column any belief about Row's choice that would justify Column's choice of C4. Therefore Row should never believe that Column would choose C4.

Note that, although C4 is never a best response, it is not dominated by any of C1, C2, and C3. For Column, C4 does better than C1 against Row's R3, better than

		COLUMN			
		C1	C2	C3	C4
ROW	R1	0, 7	2, 5	7, 0	0, 1
	R2	5, 2	3, 3	5, 2	0, 1
	R3	7, 0	2, 5	0, 7	0, 1
	R4	0, 0	0, −2	0, 0	10, −1

FIGURE 5.6 Rationalizable Strategies

[17]This example is taken from the original article that developed the concept of rationalizability. See Douglas Bernheim, "Rationalizable Strategic Behavior," *Econometrica*, vol. 52, no. 4 (July 1984), pp. 1007–1028. See also Andreu Mas-Colell, Michael Whinston, and Jerry Green, *Microeconomic Theory* (New York: Oxford University Press, 1995), pp. 242–245.

[18]Note that in each case the best choice is strictly better than C4 for Column. Thus C4 is never even tied for a best response. We can distinguish between weak and strong senses of never being a best response just as we distinguished between weak and strong dominance. Here, we have the strong sense.

C2 against Row's R4, and better than C3 against Row's R1. If a strategy *is* dominated, it also can never be a best response. Thus "never a best response" is a more general concept than "dominated." Eliminating strategies that are never a best response may be possible even when eliminating dominated strategies is not. So eliminating strategies that are never a best response can narrow down the set of possible outcomes more than can elimination of dominated strategies.[19]

The elimination of "never best response" strategies can also be carried out iteratively. Because a rational Row can never believe that a rational Column will play C4, a rational Column should foresee this. Because R4 is Row's best response only against C4, Column should never believe that Row will play R4. Thus R4 and C4 can never figure in the set of rationalizable stratgies. The concept of rationalizability does allow us to narrow down the set of possible outcomes of this game to this extent.

If a game has a Nash equilibrium, it is rationalizable and in fact can be sustained by a simple one-round system of beliefs, as we saw in Section 3.C. But more generally, even if a game does not have a Nash equilibrium, it may have rationalizable outcomes. Consider the two-by-two game obtained from Figure 5.5 or Figure 5.6 by retaining just the strategies R1 and R3 for Row and C1 and C3 for Column. It is easy to see that it has no Nash equilibrium in pure strategies. But all four outcomes are rationalizable with the use of exactly the chain of beliefs, constructed earlier, that went around and around these strategies.

Thus the concept of rationalizability provides a possible way of solving games that do not have a Nash equilibrium. And more important, the concept tells us how far we can narrow down the possibilities in a game on the basis of rationality alone.

B. Rationalizability Can Take Us All the Way to Nash Equilibrium

In some games, iterated elimination of never-best-response strategies can narrow things down all the way to Nash equilibrium. Note we said *can*, not *must*. But if it does, that is useful because in these games we can strengthen the case for Nash equilibrium by arguing that it follows purely from the players' rational thinking about each other's thinking. Interestingly, one class of games that can be solved in this way is very important in economics. This class consists of competition between firms that choose the quantities that they produce, knowing that the total quantity that is put on the market will determine the price.

We illustrate a game of this type in the context of a small coastal town. It has two fishing boats that go out every evening and return the following morning

[19]In Chapter 8, we will see that in two-player games, a strategy that is never a best response can be dominated by a *mixture* of the other strategies. Therefore, in two-player games that allow mixed strategies, the two kinds of elimination become equivalent.

to put their night's catch on the market. The game is played out in an era before modern refrigeration, so all the fish has to be sold and eaten the same day. Fish are quite plentiful in the ocean near the town, so the owner of each boat can decide how much to catch each night. But each knows that, if the total that is brought to the market is too large, the glut of fish will mean a low price and low profits.

Specifically, we suppose that, if one boat brings R barrels and the other brings S barrels of fish to the market, the price P (measured in ducats per barrel) will be $P = 60 - (R + S)$. We also suppose that the two boats and their crews are somewhat different in their fishing efficiency. Fishing costs the first boat 30 ducats per barrel and the second boat 36 ducats per barrel.

Now we can write down the profits of the two boat owners, U and V, in terms of their strategies R and S:

$$U = [(60 - R - S) - 30]R = (30 - S)R - R^2,$$
$$V = [(60 - R - S) - 36]S = (24 - R)S - S^2.$$

With these payoff expressions, we construct best-response curves and find the Nash equilibrium. As in our price competition example from Section 1, each player's payoff is a quadratic function of his own strategy, holding the strategy of the other player constant. Therefore the same mathematical methods we develop there and in the Appendix can be applied.

The first boat's best response R should maximize U for each given value of the other boat's S. With the use of calculus, this means that we should differentiate U with respect to R, holding S fixed, and set the derivative equal to zero, which gives

$$(30 - R) - 2R = 0; \quad \text{so} \quad R = 15 - S/2.$$

The noncalculus approach uses the result that the U-maximizing value of $R = B/(2C)$ where in this case $B = 30 - S$ and $C = 1$. This gives $R = (30 - S)/2$, or $R = 15 - S/2$.

Similarly, the best-response equation of the second boat is found by choosing S to maximize V for each fixed R, yielding

$$S = (24 - R)/2; \quad \text{so} \quad S = 12 - R/2.$$

The Nash equilibrium is found by solving the two best-response equations jointly for R and S, which is easy to do. So we just state the results:[20] quantities are $R = 12$ and $S = 6$; price is $P = 42$; and profits are $U = 144$ and $V = 36$.

[20] Although they are incidental to our purpose, some interesting properties of the solution are worth pointing out. The quantities differ because the costs differ; the more efficient (lower-cost) boat gets to sell more. The cost and quantity differences together imply even bigger differences in the resulting profits. The cost advantage of the first boat over the second is only 20%, but it makes four times as much profit as the second boat.

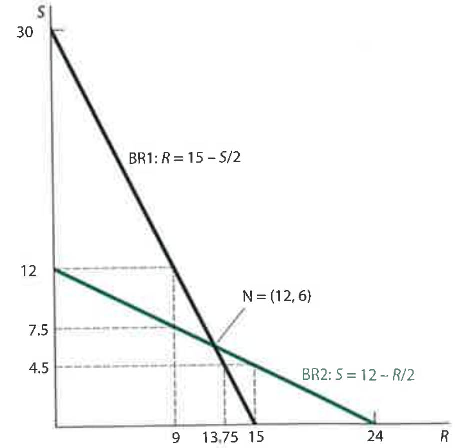

FIGURE 5.7 Nash Equilibrium Through Rationalizability

Figure 5.7 shows the two fishermen's best-response curves (labeled BR1 and BR2 with the equations displayed) and the Nash equilibrium (labeled N with its coordinates displayed) at the intersection of the two curves. Figure 5.7 also shows how the players' beliefs about each other's choices can be narrowed down by iteratively eliminating strategies that are never best responses.

What values of S can the first owner rationally believe the second owner will choose? That depends on what the second owner thinks the first owner will produce. But no matter what this might be, the whole range of the second owner's best responses is between 0 and 12. So the first owner cannot rationally believe that the second owner will choose anything else; all negative choices of S (obviously) and all choices of S greater than 12 (less obviously) are eliminated. Similarly, the second owner cannot rationally think that the first owner will produce anything less than 0 or greater than 15.

Now take this to the second round. When the first owner has restricted the second owner's choices of S to the range between 0 and 12, her own choices of R are restricted to the range of best responses to S's range. The best response to $S = 0$ is $R = 15$, and the best response to $S = 12$ is $R = 15 - 12/2 = 9$. Because BR1 has a negative slope throughout, the whole range of R allowed at this round of thinking is between 9 and 15. Similarly, the second owner's choice of S is restricted to the range of best responses to R between 0 and 15—namely, values between $S = 12$ and $S = 12 - 15/2 = 4.5$. Figure 5.7 shows these restricted ranges on the axes.

The third round of thinking narrows the ranges further. Because R must be at least 9 and BR2 has a negative slope, S can be at most the best response to 9—namely, $S = 12 - 9/2 = 7.5$. In the second round, S was already shown to be at least 4.5. Thus S is now restricted to be between 4.5 and 7.5. Similarly, because S must be at least 4.5, R can be at most $15 - 4.5/2 = 12.75$. In the second round, R was shown to be at least 9, so now it is restricted to the range from 9 to 12.75.

This succession of rounds can be carried on as far as you like, but it is already evident that the successive narrowing of the two ranges is converging on the Nash equilibrium, $R = 12$ and $S = 6$. Thus the Nash equilibrium is the only outcome that survives the iterated elimination of strategies that are never best responses.[21] We know that in general the rationalizability argument need not narrow down the outcomes of a game to its Nash equilibria, so this is a special feature of this example. Actually, the process works for an entire class of games; it will work for any game that has a unique Nash equilibrium at the intersection of downward-sloping best-response curves.[22]

This argument should be carefully distinguished from an older one based on a succession of best responses. The old reasoning proceeded as follows. Start at any strategy for one of the players—say, $R = 18$. Then the best response of the other is $S = 12 - 18/2 = 3$. The best response of R to $S = 3$ is $R = 15 - 3/2 = 13.5$. In turn, the best response of S to $R = 13.5$ is $12 - 13.5/2 = 5.25$. Then, in its turn, the best R against this S is $R = 15 - 5.25/2 = 12.375$. And so on.

The chain of best responses in the old argument also converges to the Nash equilibrium. But the argument is flawed. The game is played once with simultaneous moves. It is not possible for one player to respond to what the other player has chosen, then have the first player respond back again, and so on. If such dynamics of actual play were allowed, would the players not foresee that the other is going to respond and so do something different in the first place?

The rationalizability argument is different. It clearly incorporates the fact that the game is played only once and with simultaneous moves. All the thinking regarding the chain of best responses is done in advance, and all the successive rounds of thinking and responding are purely conceptual. Players are not responding to actual choices but are merely calculating those choices that will never be made. The dynamics are purely in the minds of the players.

[21] This example can also be solved by iteratively eliminating dominated strategies, but proving dominance is harder and needs more calculus, whereas the never-best-response property is obvious from Figure 5.7, so we use the simpler argument.

[22] A similar argument works with upward-sloping best-response curves, such as those in the pricing game of Figure 5.1, for narrowing the range of best responses starting at low prices. Narrowing from the higher end is possible only if there is some obvious starting point. This starting point might be a very high price that can never be exceeded for some externally enforced reason—if, for example, people simply do not have the money to pay prices beyond a certain level.

SUMMARY

When players in a simultaneous-move game have a continuous range of actions to choose, best-response analysis yields mathematical *best-response rules* that can be solved simultaneously to obtain Nash equilibrium strategy choices. The best-response rules can be shown on a diagram in which the intersection of the two curves represents the Nash equilibrium. Firms choosing price or quantity from a large range of possible values and political parties choosing campaign advertising expenditure levels are examples of games with *continuous strategies*.

The results of laboratory tests of the Nash equilibrium concept show that a common cultural background is essential for coordinating in games with multiple equilibria. Repeated play of some games shows that players can learn from experience and begin to choose strategies that approach Nash equilibrium choices. Further, predicted equilibria are accurate only when the experimenters' assumptions match the true preferences of players. Real-world applications of game theory have helped economists and political scientists, in particular, to understand important consumer, firm, voter, legislature, and government behaviors.

Theoretical criticisms of the Nash equilibrium concept have argued that the concept does not adequately account for risk, that it is of limited use because many games have multiple equilibria, and that it cannot be justified on the basis of rationality alone. In many cases, a better description of the game and its payoff structure or a refinement of the Nash equilibrium concept can lead to better predictions or fewer potential equilibria. The concept of *rationalizability* relies on the elimination of strategies that are *never a best response* to obtain a set of *rationalizable* outcomes. When a game has a Nash equilibrium, that outcome will be rationalizable, but rationalizability also allows one to predict equilibrium outcomes in games that have no Nash equilibria.

KEY TERMS

best-response curves (137)
best-response rule (134)
continuous strategy (134)
never a best response (157)
rationalizability (157)
rationalizable (157)
refinement (155)

SOLVED EXERCISES

S1. In the political campaign advertising game in Section 1.B, party L chooses an advertising budget, x (millions of dollars), and party R similarly chooses a budget, y (millions of dollars). We showed there that the best-response rules in that game are $y = 10\sqrt{x} - x$, for party R, and $x = 10\sqrt{y} - x$ for party L.

(a) What is party R's best response if party L spends $16 million?
(b) Use the specified best-response rules to verify that the Nash equilibrium advertising budgets are $x = y = 25$, or $25 million.

S2. The bistro game illustrated in Figure 5.1 defines demand functions for meals at Xavier's (Q_x) and Yvonne's (Q_y) as $Q_x = 44 - 2P_x + P_y$, and $Q_y = 44 - 2P_y + P_x$. Profits for each firm depend in addition on their costs of serving each customer. Suppose here that Yvonne's is able to reduce its costs to a mere $2 per customer by completely eliminating the wait staff (customers pick up their orders at the counter, and a few remaining employees bus the tables). Xavier's continues to incur a cost of $8 per customer.

(a) Recalculate the best-response rules and the Nash equilibrium prices for the two firms, given the change in the cost conditions.
(b) Graph the two best-response curves and describe the differences between your graph and Figure 5.1. In particular, which curve has moved and by how much? Explain why these changes occurred in the diagram.

S3. Yuppietown has two food stores, La Boulangerie, which sells bread, and La Fromagerie, which sells cheese. It costs $1 to make a loaf of bread and $2 to make a pound of cheese. If La Boulangerie's price is P_1 dollars per loaf of bread and La Fromagerie's price is P_2 dollars per pound of cheese, their respective weekly sales, Q_1 thousand loaves of bread and Q_2 thousand pounds of cheese, are given by the following equations:

$$Q_1 = 14 - P_1 - 0.5P_2, \quad Q_2 = 19 - 0.5P_1 - P_2.$$

(a) For each store, write its profit as a function of P_1 and P_2 (in the exercises that follow, we will call this "the profit function" for brevity). Then find their respective best-response rules. Graph the best-response curves, and find the Nash equilibrium prices in this game.
(b) Suppose that the two stores collude and set prices jointly to maximize the sum of their profits. Find the joint profit-maximizing prices for the stores.
(c) Provide a short intuitive explanation for the differences between the Nash equilibrium prices and those that maximize joint profit. Why is joint profit maximization not a Nash equilibrium?
(d) In this problem, bread and cheese are mutual *complements*. They are often consumed together; that is why a drop in the price of one increases the sales of the other. The products in our bistro example in Section 1.A are *substitutes* for each other. How does this difference explain the differences among your findings for the best-response rules, the Nash equilibrium prices, and the joint profit-maximizing prices in this question, and the corresponding entities in the bistro example in the text?

S4. The game illustrated in Figure 5.3 has a unique Nash equilibrium in pure strategies. However, all nine outcomes in that game are rationalizable. Confirm this assertion, explaining your reasoning for each outcome.

S5. For the game presented in Exercise S5 in Chapter 4, what are the rationalizable strategies for each player? Explain your reasoning.

S6. Section 4.B of this chapter describes a fishing game played in a small coastal town. When the response rules for the two boats have been derived, rationalizability can be used to justify the Nash equilibrium in the game. In the description in the text, we take the process of narrowing down strategies that can never be best responses through three rounds. By the third round, we know that R (the number of barrels of fish brought home by boat 1) must be at least 9, and that S (the number of barrels of fish brought home by boat 2) must be at least 4.5. The narrowing process in that round restricted R to the range between 9 and 12.75 while restricting S to the range between 4.5 and 7.5. Take this process of narrowing through one additional (fourth) round and show the reduced ranges of R and S that are obtained at the end of the round.

S7. Two carts selling coconut milk (from the coconut) are located at 0 and 1, 1 mile apart on the beach in Rio de Janeiro. (They are the only two coconut-milk carts on the beach.) The carts—Cart 0 and Cart 1—charge prices p_0 and p_1, respectively, for each coconut. One thousand beachgoers buy coconut milk, and these customers are uniformly distributed along the beach between carts 0 and 1. Each beachgoer will purchase one coconut milk in the course of her day at the beach and, in addition to the price, each will incur a transport cost of 0.5 times d^2, where d is the distance (in miles) from her beach blanket to the coconut cart. In this system, Cart 0 sells to all of the beachgoers located between 0 and x, and Cart 1 sells to all of the beachgoers located between x and 1, where x is the location of the beachgoer who pays the same total price if she goes to 0 or 1. Location x is then defined by the expression:

$$p_0 + 0.5x^2 = p_1 + 0.5(1 - x)^2.$$

The two carts will set their prices to maximize their bottom-line profit figures, B; profits are determined by revenue (the cart's price times its number of customers) and cost (the carts each incur a cost of $0.25 per coconut times the number of coconuts sold).

(a) For each cart, determine the expression for the number of customers served as a function of p_0 and p_1. (Recall that Cart 0 gets the customers between 0 and x, or just x, while Cart 1 gets the customers between x and 1, or $1 - x$. That is, cart 0 sells to x customers, where x is measured in thousands, and cart 1 sells to $(1 - x)$ thousand.)

(b) Write the profit functions for the two carts. Find the two best-response rules for each cart as a function of their rival's price.

(c) Graph the best-response rules, and then calculate (and show on your graph) the Nash equilibrium price level for coconut milk on the beach.

S8. Crude oil is transported across the globe in enormous tanker ships called Very Large Crude Carriers (VLCCs). By 2001, more than 92% of all new VLCCs were built in South Korea and Japan. Assume that the price of new VLCCs (in millions of dollars) is determined by the function $P = 180 - Q$, where $Q = q_{Korea} + q_{Japan}$. (That is, assume that only Japan and Korea produce VLCCs, so they are a duopoly.) Assume that the cost of building each ship is \$30 million in both Korea and Japan. That is, $c_{Korea} = c_{Japan} = 30$, where the per-ship cost is measured in millions of dollars.

(a) Write the profit functions for each country in terms of q_{Korea} and q_{Japan} and either c_{Korea} or c_{Japan}. Find each country's best-response function.

(b) Using the best-response functions found in part (a), solve for the Nash equilibrium quantity of VLCCs produced by each country per year. What is the price of a VLCC? How much profit is made in each country?

(c) Labor costs in Korean shipyards are actually much lower than in their Japanese counterparts. Assume now that the cost per ship in Japan is \$40 million and that in Korea it is only \$20 million. Given $c_{Korea} = 20$ and $c_{Japan} = 40$, what is the market share of each country (ie, the percentage of ships that each country sells relative to the total number sold)? What are the profits for each country?

S9. Extending the previous problem, suppose China decides to enter the VLCC construction market. The duopoly now becomes a triopoly, so that although price is still $P = 180 - Q$, quantity is now given by $Q = q_{Korea} + q_{Japan} + q_{China}$. Assume that all three countries have a per-ship cost of \$30 million: $c_{Korea} = c_{Japan} = c_{China} = 30$.

(a) Write the profit functions for each of the three countries in terms of q_{Korea}, q_{Japan}, and q_{China}, and c_{Korea}, c_{Japan}, or c_{China}. Find each country's best-response rule.

(b) Using your answer to part (a), find the quantity produced, the market share captured (see Exercise S8, part (c)), and the profits earned by each country. This will require the solution of three equations in three unknowns.

(c) What happens to the price of a VLCC in the new triopoly relative to the duopoly situation in Exercise S8, part (b)? Why?

S10. Monica and Nancy have formed a business partnership to provide consulting services in the golf industry. They each have to decide how much effort to put into the business. Let m be the amount of effort put into the business by Monica, and n be the amount of effort put in by Nancy.

The joint profits of the partnership are given by $4m + 4n + mn$, in tens of thousands of dollars, and the two partners split these profits equally. However, they must each separately incur the costs of their own effort; the cost to Monica of her effort is m^2, while the cost to Nancy of her effort is n^2 (both measured in tens of thousands of dollars). Each partner must make her effort decision without knowing what effort decision the other player has made.

(a) If Monica and Nancy each put in effort of $m = n = 1$, then what are their payoffs?

(b) If Monica puts in effort of $m = 1$, then what is Nancy's best response?

(c) What is the Nash equilibrium to this game?

S11. Nash equilibrium through rationalizability can be achieved in games with upward-sloping best-response curves if the rounds of eliminating never-best-response strategies begin with the smallest possible values. Consider the pricing game between Xavier's Tapas Bar and Yvonne's Bistro that is illustrated in Figure 5.1. Use Figure 5.1 and the best-response rules from which it is derived to begin rationalizing the Nash equilibrium in that game. Start with the lowest possible prices for the two firms and describe (at least) two rounds of narrowing the set of rationalizable prices toward the Nash equilibrium.

S12. A professor presents the following game to Elsa and her forty-nine classmates. Each of them simultaneously and privately writes down a number between zero and 100 on a piece of paper, and they all hand in their numbers. The professor then computes the mean of these numbers and defines X to be the mean of the students' numbers. The student who submits the number closest to two-thirds of X wins \$50. If multiple students tie, they split the prize equally.

(a) Show that choosing the number 80 is a dominated strategy.

(b) What would the set of best responses be for Elsa if she knew that all of her classmates would submit the number 40? That is, what is the range of numbers for which each number in the range is closer to the winning number than 40?

(c) What would the set of best responses be for Elsa is she knew that all of her classmates would submit the number 10?

(d) Find a symmetric Nash equilibrium to this game. That is, what number is a best response to everyone else submitting that same number?

(e) Which strategies are rationalizable in this game?

UNSOLVED EXERCISES

U1. Diamond Trading Company (DTC), a subsidiary of De Beers, is the dominant supplier of high quality diamonds for the wholesale market. For simplicity, assume that DTC has a monopoly on wholesale diamonds. The quantity that DTC chooses to sell thus has a direct impact on the wholesale price of diamonds. Let the wholesale price of diamonds (in hundreds of dollars) be given by the following inverse demand function: $P = 120 - Q_{DTC}$. Assume that DTC has a cost of 12 (hundred dollars) per high-quality diamond.

(a) Write DTC's profit function in terms of Q_{DTC} and solve for DTC's profit-maximizing quantity. What will be the wholesale price of diamonds at that quantity? What will DTC's profit be?

Frustrated with DTC's monopoly, several diamond mining interests and large retailers collectively set up a joint venture called Adamantia to act as a competitor to DTC in the wholesale market for diamonds. The wholesale price is now given by $P = 120 - Q_{DTC} - Q_{ADA}$. Assume that Adamantia has a cost of 12 (hundred dollars) per high-quality diamond.

(b) Write the best-response functions for both DTC and Adamantia. What quantity does each wholesaler supply to the market in equilibrium? What wholesale price do these quantities imply? What will the profit of each supplier be in this duopoly situation?

(c) Describe the differences in the market for wholesale diamonds under the duopoly of DTC and Adamantia relative to the monopoly of DTC. What happens to the quantity supplied in the market and the market price when Adamantia enters? What happens to the collective profit of DTC and Adamantia?

U2. There are two movie theaters in the town of Harkinsville: Modern Multiplex, a first-run theater, and Sticky Shoe, which shows movies that have been out for a while at a cheaper price. The demand for Modern Multiplex is given by: $Q_{MM} = 14 - P_{MM} + P_{SS}$, while the demand for Sticky Shoe is: $Q_{SS} = 8 - 2P_{SS} + P_{MM}$, where prices are in dollars and quantities are measured in hundreds of moviegoers. Modern Multiplex has a per-customer cost of \$4, while Sticky Shoe has a per-customer cost of only \$2.

(a) From the demand equations alone, what indicates whether Modern Multiplex and Sticky Shoe offer services that are substitutes or complements?

(b) Write the profit function for each theater in terms of P_{SS} and P_{MM}. Find each theater's best-response rule.

(c) Find the Nash equilibrium price, quantity, and profit for each theater.

(d) What would each theater's price, quantity, and profit be if the two decided to collude to maximize joint profits in this market? Why isn't the collusive outcome a Nash equilibrium?

U3. Fast forward a decade beyond the situation in Exercise S3. Yuppietown's demand for bread and cheese has decreased, and the town's two food stores, La Boulangerie and La Fromagerie, have been bought out by a third company: L'Épicerie. It still costs \$1 to make a loaf of bread and \$2 to make a pound of cheese, but the quantities of bread and cheese sold (Q_1 and Q_2 respectively, measured in thousands) are now given by the equations:

$$Q_1 = 8 - P_1 - 0.5P_2, \quad Q_2 = 16 - 0.5P_1 - P_2.$$

Again, P_1 is the price in dollars of a loaf of bread, and P_2 is the price in dollars of a pound of cheese.

(a) Initially, L'Épicerie runs La Boulangerie and La Fromagerie as if they were separate firms, with independent managers who each try to maximize their own profit. What are the Nash equilibrium quantities, prices, and profits for the two divisions of L'Épicerie, given the new quantity equations?

(b) The owners of L'Épicerie think that they can make more total profit by coordinating the pricing strategies of the two Yuppietown divisions of their company. What are the joint-profit-maximizing prices for bread and cheese under collusion? What quantities do La Boulangerie and La Fromagerie sell of each good, and what is the profit that each division earns separately?

(c) In general, why might companies sell some of their goods at prices below cost? That is, explain a rationale of loss leaders, using your answer from part (b) as an illustration.

U4. The coconut-milk carts from Exercise S7 set up again the next day. Nearly everything is exactly the same as Exercise S7: the carts are in the same locations, the number and distribution of beachgoers is identical, and the demand of the beachgoers for exactly one coconut milk each is unchanged. The only difference is that it is a particularly hot day, so that now each beachgoer incurs a higher transport cost of $0.6d^2$. Again, Cart 0 sells to all of the beachgoers located between 0 and x, and Cart 1 sells to all of the beachgoers located between x and 1, where x is the location of the beachgoer who pays the same total price if she goes to 0 or 1. However, now location x is defined by the expression:

$$p_0 + 0.6x^2 = p_1 + 0.6(1 - x)^2.$$

Again, each cart has a cost of \$0.25 per coconut sold.

(a) For each cart, determine the expression for the number of customers served as a function of p_0 and p_1. (Recall that Cart 0 gets the customers between 0 and x, or just x, while Cart 1 gets the customers between x

and 1, or $1 - x$. That is, Cart 0 sells to x customers, where x is measured in thousands, and Cart 1 sells to $(1 - x)$ thousand.)

(b) Write out profit functions for the two carts and find the two best-response rules.

(c) Calculate the Nash equilibrium price level for coconuts on the beach. How does this price compare with the price found in Exercise S7? Why?

U5. The game illustrated in Figure 5.4 has a unique Nash equilibrium in pure strategies. Find that Nash equilibrium, and then show that it is also the unique rationalizable outcome in that game.

U6. What are the rationalizable strategies of the game "Evens or Odds" from Exercise S11 in Chapter 4?

U7. In the fishing-boat game of Section 4, we showed how it is possible for there to be a uniquely rationalizable outcome in continuous strategies that is also a Nash equilibrium. However, this is not always the case; there may be many rationalizable strategies, and not all of them will necessarily be part of a Nash equilibrium.

Returning to the political advertising game of Exercise S1, find the set of rationalizable strategies for party L. (Due to their symmetric payoffs, the set of rationalizable strategies will be the same for party R.) Explain your reasoning.

U8. Intel and AMD, the primary producers of computer central processing units (CPUs), compete with one another in the mid-range chip category (among other categories). Assume that global demand for mid-range chips depends on the quantity that the two firms make, so that the price (in dollars) for mid-range chips is given by $P = 210 - Q$, where $Q = q_{Intel} + q_{AMD}$ and where the quantities are measured in millions. Each mid-range chip costs Intel $60 to produce. AMD's production process is more streamlined; each chip costs them only $48 to produce.

(a) Write the profit function for each firm in terms of q_{Intel} and q_{AMD}. Find each firm's best-response rule.

(b) Find the Nash equilibrium price, quantity, and profit for each firm.

(c) **(Optional)** Suppose Intel acquires AMD, so that it now has two separate divisions with two different production costs. The merged firm wishes to maximize total profits from the two divisions. How many chips should each division produce? (Hint: You may need to think carefully about this problem, rather than blindly applying mathematical techniques.) What is the market price and the total profit to the firm?

U9. Return to the VLCC triopoly game of Exercise S9. In reality, the three countries do not have identical production costs. China has been gradually entering the VLCC construction market for several years, and its production costs started out rather high due to lack of experience.

(a) Solve for the triopoly quantities, market shares, price, and profits for the case where the per-ship costs are $20 million for Korea, $40 million for Japan, and $60 million for China ($c_{Korea} = 20$, $c_{Japan} = 40$, and $c_{China} = 60$).

After it gains experience and adds production capacity, China's per-ship cost will decrease dramatically. Since labor is even cheaper in China than Korea, eventually the per-ship cost will be even lower in China than it is in Korea.

(b) Repeat part (a) with the adjustment that China's per-ship cost is $16 million ($c_{Korea} = 20$, $c_{Japan} = 40$, and $c_{China} = 16$).

U10. Return to the story of Monica and Nancy from Exercise S10. After some additional professional training, Monica is more productive on the job, so that the joint profits of their company are now given by $5m + 4n + mn$, in tens of thousands of dollars. Again, m is the amount of effort put into the business by Monica, n is the amount of effort put in by Nancy, and the costs are m^2 and n^2 to Monica and Nancy respectively (in tens of thousands of dollars).

The terms of their partnership still require that the joint profits be split equally, despite the fact that Monica is more productive. Assume that their effort decisions are made simultaneously.

(a) What is Monica's best response if she expects Nancy to put in an effort of $n = \frac{4}{3}$?

(b) What is the Nash equilibrium to this game?

(c) Compared to the old Nash equilibrium found in Exercise S10, part (c), does Monica now put in more, less, or the same amount of effort? What about Nancy?

(d) What are the final payoffs to Monica and Nancy in the new Nash equilibrium (after splitting the joint profits and accounting for their costs of effort)? How do they compare to the payoffs to each of them under the old Nash equilibrium? In the end, who receives more benefit from Monica's additional training?

U11. A professor presents a new game to Elsa and her forty-nine classmates (similar to the situation in Exercise S12). As before, each of the students simultaneously and privately writes down a number between zero and 100 on a piece of paper, and the professor computes the mean of these numbers and calls it X. This time the student who submits the number closest to $(\frac{2}{3}) \times (X + 9)$ wins $50. Again, if multiple students tie, they split the prize equally.

(a) Find a symmetric Nash equilibrium to this game. That is, what number is a best response to everyone else submitting the same number?

(b) Show that choosing the number 5 is a dominated strategy. (Hint: what would class average X have to be for the target number to be 5?)

(c) Show that choosing the number 90 is a dominated strategy.
(d) What are all of the dominated strategies?
(e) Suppose Elsa believes that none of her classmates will play the dominated strategies found in part (d). Given these beliefs what strategies are never a best response for Elsa?
(f) Which strategies do you think are rationalizable in this game? Explain your reasoning.

U12. **(Optional—requires calculus)** Recall the political campaign advertising example from Section 1.B concerning parties L and R. In that example, when L spends $\$x$ million on advertising and R spends $\$y$ million, L gets a share $x/(x + y)$ of the votes and R gets $y/(x + y)$. We also mentioned that two types of asymmetries can arise between the parties in that model. One party—say, R—may be able to advertise at a lower cost or R's advertising dollars may be more effective in generating votes than L's. To allow for both possibilities, we can write the payoff functions of the two parties as

$$V_L = \frac{x}{x + ky} - x \quad \text{and} \quad V_R = \frac{ky}{x + ky} - cy, \text{ where } k > 0 \text{ and } c > 0.$$

These payoff functions show that R has an advantage in the relative effectiveness of its ads when k is high and that R has an advantage in the cost of its ads when c is low.

(a) Use the payoff functions to derive the best-response functions for R (which chooses y) and L (which chooses x).
(b) Use your calculator or your computer to graph these best-response functions when $k = 1$ and $c = 1$. Compare the graph with the one for the case in which $k = 1$ and $c = 0.8$. What is the effect of having an advantage in the cost of advertising?
(c) Compare the graph from part (b), when $k = 1$ and $c = 1$ with the one for the case in which $k = 2$ and $c = 1$. What is the effect of having an advantage in the effectiveness of advertising dollars?
(d) Solve the best-response functions that you found in part (a), jointly for x and y, to show that the campaign advertising expenditures in Nash equilibrium are:

$$x = \frac{ck}{(c + k)^2} - x \quad \text{and} \quad y = \frac{k}{(c + k)^2}.$$

(e) Let $k = 1$ in the equilibrium spending-level equations and show how the two equilibrium spending levels vary with changes in c (i.e., interpret the signs of dx/dc and dy/dc). Then let $c = 1$ and show how the two equilibrium spending levels vary with changes in k (i.e., interpret the signs of dx/dk and dy/dk). Do your answers support the effects that you observed in parts (b) and (c) of this exercise?

Appendix: Finding a Value to Maximize a Function

Here we develop in a simple way the method for choosing a variable X to obtain the maximum value of a variable that is a function of it, say $Y = F(X)$. Our applications will mostly be to cases where the function is quadratic, such as $Y = A + BX - CX^2$. For such functions we derive the formula $X = B/(2C)$ that was stated and used in the chapter. We develop the general idea using calculus, and then offer an alternative approach that does not use calculus but applies only to the quadratic function.*

The calculus method tests a value of X for optimality by seeing what happens to the value of the function for other values on either side of X. If X does indeed maximize $Y = F(X)$, then the effect of increasing or decreasing X should be a drop in the value of Y. Calculus gives us a quick way to perform such a test.

Figure 5A.1 illustrates the basic idea. It shows the graph of a function $Y = F(X)$, where we have used a function of the type that fits our application, even though the idea is perfectly general. Start at any point P with coordinates (X, Y) on the graph. Consider a slightly different value of X, say $(X + h)$. Let k be the resulting change in $Y = F(X)$, so the point Q with coordinates $(X + h, Y + k)$ is also on the graph. The slope of the chord joining P to Q is the ratio k / h. If this ratio is positive, then h and k have the same sign; as X increases, so does Y. If the ratio is negative, then h and k have opposite signs; as X increases, Y decreases.

If we now consider smaller and smaller changes h in X, and the corresponding smaller and smaller changes k in Y, the chord PQ will approach the tangent to the graph at P. The slope of this tangent is the limiting value of the ratio k / h. It is called the derivative of the function $Y = F(X)$ at the point X. Symbolically, it is written as $F'(X)$ or dY/dX. Its sign tells us whether the function is increasing or decreasing at precisely the point X.

For the quadratic function in our application, $Y = A + BX - CX^2$ and

$$Y + k = A + B(X + h) - C(X + h)^2.$$

Therefore, we can find an expression for k as follows:

$$k = [A + B(X + h) - C(X + h)^2] - [A + BX - CX^2]$$

$$= Bh - C[(X + h)^2 - X^2] = Bh - C[X^2 + 2Xh + h^2 - X^2]$$

$$= (B - 2CX)\,h - Ch^2.$$

*Needless to say, we give only the briefest, quickest treatment, leaving out all issues of functions that don't have derivatives, functions that are maximized at an extreme point of the interval over which they are defined, and so on. Some readers will know all we say here; some will know much more. Others who want to find out more should refer to any introductory calculus textbook.

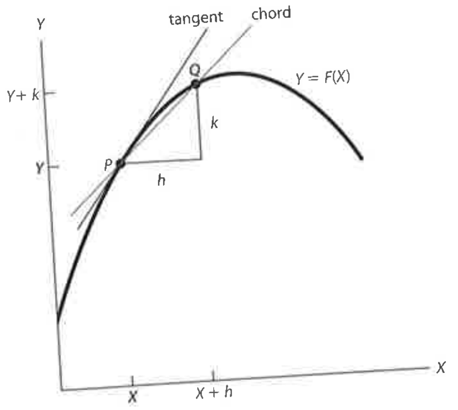

FIGURE 5A.1 Derivative of a Function Illustrated

Then $k/h = (B - 2CX) - Ch$. In the limit as h goes to zero, $k/h = (B - 2CX)$. This last expression is then the derivative of our function.

Now we use the derivative to find a test for optimality. Figure 5A.2 illustrates the idea. The point M yields the highest value of $Y = F(X)$. The function increases as we approach the point M from the left and decreases after we have passed to the right of M. Therefore the derivative $F'(X)$ should be positive for values of X smaller than M and negative for values of X larger than M. By continuity, the derivative precisely at M should be zero. In ordinary language, the graph of the function should be flat where it peaks.

In our quadratic example, the derivative is: $F'(X) = B - 2CX$. Our optimality test implies that the function is optimized when this is zero, or at $X = B/(2C)$. This is exactly the formula given in the chapter.

One additional check needs to be performed. If we turn the whole figure upside down, M is the minimum value of the upside-down function, and at this trough the graph will also be flat. So for a general function $F(X)$, setting $F'(X) = 0$ might yield an X that gives its minimum rather than the maximum. How do we distinguish the two possibilities?

At a maximum, the function will be increasing to its left and decreasing to its right. Therefore the derivative will be positive for values of X smaller than the purported maximum, and negative for larger values. In other words, the derivative, itself regarded as a function of X, will be decreasing at this point. A decreasing function has a negative derivative. Therefore the derivative of the derivative, what is called the second derivative of the original function, written as $F''(X)$ or d^2Y/dX^2,

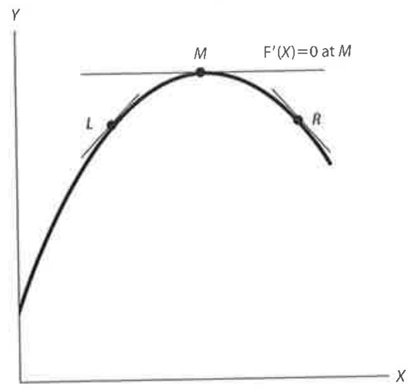

FIGURE 5A.2 Optimum of a Function

should be negative at a maximum. Similar logic shows that the second derivative should be positive at a minimum; that is what distinguishes the two cases.

For the derivative $F'(X) = B - 2CX$ of our quadratic example, applying the same h, k procedure to $F'(X)$ as we did to $F(X)$ shows $F''(X) = -2C$. This is negative so long as C is positive, which we assumed when stating the problem in the chapter. The test $F'(X) = 0$ is called the first-order condition for maximization of $F(X)$, and $F''(X) < 0$ is the second-order condition.

To fix the idea further, let us apply it to the specific example of Xavier's best response that we considered in the chapter. We had the expression

$$\Pi_x = -8(44 + P_y) + (16 + 44 + P_y)P_x - 2(P_x)^2.$$

This is a quadratic function of P_x (holding the other restaurant's price, P_y, fixed). Our method gives its derivative:

$$\frac{d\Pi_x}{dP_x} = (60 + p_y) - 4p_x.$$

The first-order condition for P_x to maximize Π_x is that this derivative should be zero. Setting it equal to zero and solving for P_x gives the same equation as derived in Section 5.1.A. (The second-order condition is $d^2\Pi_x/dP^2_x < 0$, which is satisfied because the second-order derivative is just –4.)

We hope you will regard the calculus method as simple enough and that you will have occasion to use it again in a few places later, for example, in Chapter 12 on collective action. But if you find it too difficult, here is a noncalculus alternative method that works for quadratic functions. Rearrange terms to write the function as

$$
\begin{aligned}
Y &= A + BX - CX^2 \\
&= A + B^2/(4C) - B^2/(4C) + BX - CX^2 \\
&= A + \frac{B^2}{4C} - C\left(\frac{B^2}{4C^2} - 2\frac{B}{C} + X^2\right) \\
&= A + \frac{B^2}{4C} - C\left(\frac{B}{2C} - X\right)^2.
\end{aligned}
$$

In the final form of the expression, X appears only in the last term, where a square involving it is being subtracted (remember $C > 0$). The whole expression is maximized when this subtracted term is made as small as possible, which happens when $X = B/(2C)$. Voila!

This method of "completing the square" works for quadratic functions and therefore will suffice for most of our uses. It also avoids calculus. But we must admit it smacks of magic. Calculus is more general and more methodical. It repays a little study many times over.

6

Combining Sequential and Simultaneous Moves

In Chapter 3 we considered games of purely sequential moves; Chapters 4 and 5 dealt with games of purely simultaneous moves. We developed concepts and techniques of analysis appropriate to the pure game types—trees and rollback equilibrium for sequential moves, payoff tables and Nash equilibrium for simultaneous moves. In reality, however, many strategic situations contain elements of both types of interaction. Also, although we used game trees (extensive forms) as the sole method of illustrating sequential-move games and game tables (strategic forms) as the sole method of illustrating simultaneous-move games, we can use either form for any type of game.

In this chapter, we examine many of these possibilities. We begin by showing how games that combine sequential and simultaneous moves can be solved by combining trees and payoff tables and by combining rollback and Nash equilibrium analysis in appropriate ways. Then we consider the effects of changing the nature of the interaction in a particular game. Specifically, we look at the effects of changing the rules of a game to convert sequential play into simultaneous play and vice versa and of changing the order of moves in sequential play. This topic gives us an opportunity to compare the equilibria found by using the concept of rollback, in a sequential-move game, with those found by using the Nash equilibrium concept, in the simultaneous version of the same game. From this comparison, we extend the concept of Nash equilibria to sequential-play games. It turns out that the rollback equilibrium is a special case, usually called a refinement, of these Nash equilibria.

1 GAMES WITH BOTH SIMULTANEOUS AND SEQUENTIAL MOVES

As mentioned several times thus far, most real games that you will encounter will be made up of numerous smaller components. Each of these components may entail simultaneous play or sequential play, so the full game requires you to be familiar with both. The most obvious examples of strategic interactions containing both sequential and simultaneous parts are those between two (or more) players over an extended period of time. You may play a number of different simultaneous-play games against your roommate, for example, in the course of a week; your play, as well as hers, in previous situations is important in determining how each of you decide to act in the next "round." Also, many sporting events, interactions between competing firms in an industry, and political relationships are sequentially linked series of simultaneous-move games. Such games are analyzed by combining the tools presented in Chapter 3 (trees and rollback) and in Chapters 4 and 5 (payoff tables and Nash equilibria).[1] The only difference is that the actual analysis becomes more complicated as the number of moves and interactions increases.

A. Two-Stage Games and Subgames

Our main illustrative example for such situations includes two would-be telecom giants, CrossTalk and GlobalDialog. Each can choose whether to invest $10 billion in the purchase of a fiber-optic network. They make their investment decisions simultaneously. If neither chooses to make the investment, that is the end of the game. If one invests and the other does not, then the investor has to make a pricing decision for its telecom services. It can choose either a high price, which will attract 60 million customers, from each of whom it will make an operating profit of $400, or a low price, which will attract 80 million customers, from each of whom it will make an operating profit of $200. If both firms acquire fiber-optic networks and enter the market, then their pricing choices become a second simultaneous-move game. Each can choose either the high or the low price. If both choose the high price, they will split the total market equally; so each will get 30 million customers and an operating profit of $400 from each. If both choose the low price, again they will split the total market equally; so each will get 40 million customers and an operating profit of $200 from each. If one chooses the high price and the other the low price, then the low-price

[1]Sometimes the simultaneous part of the game will have equilibria in mixed strategies, when the tools we develop in Chapters 7 and 8 will be required. We mention this possibility in this chapter where relevant and give you an opportunity to use such methods in exercises for the later chapters.

firm will get all the 80 million customers at that price, and the high-price firm will get nothing.

The interaction between CrossTalk and GlobalDialog forms a two-stage game. Of the four combinations of the simultaneous-move choices at the first (investment) stage, one ends the game, two lead to a second-stage (pricing) decision by just one player, and the fourth leads to a simultaneous-move (pricing) game at the second stage. We show this game pictorially in Figure 6.1.

Regarded as a whole, Figure 6.1 illustrates a game tree, but one that is more complex than the trees in Chapter 3. You can think of it as an elaborate "tree house" with multiple levels. The levels are shown in different parts of the same two-dimensional figure, as if you are looking down at the tree from a helicopter positioned directly above it.

The first-stage game is represented by the payoff table in the top-left quadrant of Figure 6.1. You can think of it as the first floor of the tree house. It has four "rooms." The room in the northwest corner corresponds to the "Don't invest" first-stage moves of both firms. If the firms' decisions take the game to this room, there are no further choices to be made, so we can think of it being like a terminal node of a tree in Chapter 3 and show the payoffs in the cell of the table;

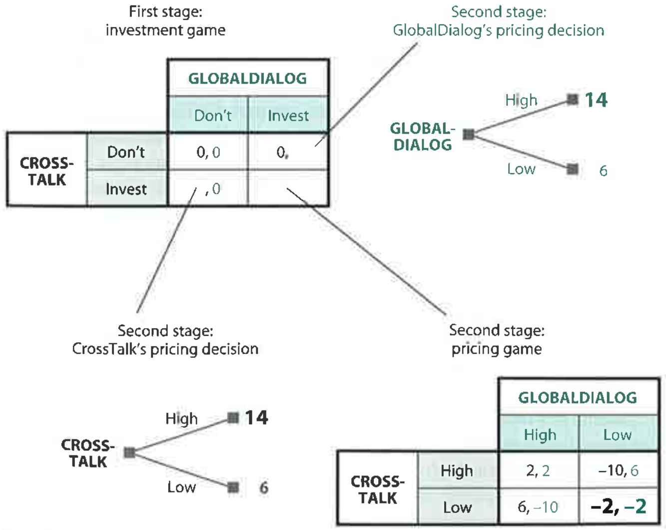

FIGURE 6.1 Two-Stage Game Combining Sequential and Simultaneous Moves

both firms get 0. However, all of the other combinations of actions for the two firms lead to rooms that lead to further choices; so we cannot yet show the payoffs in those cells. Instead, we show branches leading to the second floor. The northeast and southwest rooms show only the payoff to the firm that has not invested; the branches leading from each of these rooms take us to single-firm pricing decisions in the second stage. The southeast room leads to a multiroom second-floor structure within the tree house, which represents the second-stage pricing game that is played if both firms have invested in the first stage. This second-floor structure has four rooms corresponding to the four combinations of the two firms' pricing moves.

All of the second-floor branches and rooms are like terminal nodes of a game tree, so we can show the payoffs in each case. Payoffs here consist of each firm's operating profits minus the previous investment costs; payoff values are written in billions of dollars.

Consider the branch leading to the southwest corner of Figure 6.1. The game arrives in that corner if CrossTalk is the only firm that has invested. Then, if it chooses the high price, its operating profit is $\$400 \times 60$ million = \$24 billion; after subtracting the \$10 billion investment cost, its payoff is \$14 billion, which we write as 14. In the same corner, if CrossTalk chooses the low price, then its operating profit is $\$200 \times 80$ million = \$16 billion, yielding the payoff 6 after accounting for its original investment. In this situation, GlobalDialog's payoff is 0, as shown in the southwest room of the first floor of our tree. Similar calculations for the case in which GlobalDialog is the only firm to invest give us the payoffs shown in the northeast corner of Figure 6.1; again, the payoff of 0 for CrossTalk is shown in the northeast room of the first-stage game table.

If both firms invest, both play the second-stage pricing game illustrated in the southeast corner of the figure. When both choose the high price in the second stage, each gets operating profit of $\$400 \times 30$ million (half of the market), or \$12 billion; after subtracting the \$10 billion investment cost, each is left with a net profit of \$2 billion, or a payoff of 2. If both firms choose the low price in the second stage, each gets operating profit of $\$200 \times 40$ million = \$8 billion, and, after subtracting the \$10 billion investment cost, each is left with a net loss of \$2 billion, or a payoff of −2. Finally, if one firm charges the high price and the other firm the low price, then the low-price firm has operating profit of $\$200 \times 80$ million = \$16 billion, leading to the payoff 6, while the high-price firm gets no operating profit and simply loses its \$10 billion investment, for a payoff of −10.

As with any multistage game in Chapter 3, we must solve this game backward, starting with the second-stage game. In the two single-firm decision problems, we see at once that the high-price policy yields the higher payoff. We highlight this by showing that payoff in a larger-size type.

The second-stage pricing game has to be solved by using methods developed in Chapter 4. It is immediately evident, however, that this game is a prisoners' dilemma. Low is the dominant strategy for each firm; so the outcome is

the room in the southeast corner of the second-stage game table; each firm gets payoff −2.[2] Again, we show these payoffs in a larger type size to highlight the fact that they are the payoffs obtained in the second-stage equilibrium.

Rollback now tells us that each first-stage configuration of moves should be evaluated by looking ahead to the equilibrium of the second-stage game (or the optimum second-stage decision) and the resulting payoffs. We can therefore substitute the payoffs that we have just calculated into the previously empty or partly empty rooms on the first floor of our tree house. This substitution gives us a first floor with known payoffs, shown in Figure 6.2.

Now we can use the methods of Chapter 4 to solve this simultaneous-move game. You should immediately recognize the game in Figure 6.2 as a chicken game. It has two Nash equilibria, each of which entails one firm choosing Invest and the other choosing Don't. The firm that invests makes a huge profit; so each firm prefers the equilibrium in which it is the investor while the other firm stays out. In Chapter 4, we briefly discussed the ways in which one of the two equilibria might get selected. We also pointed out the possibility that each firm might try to get its preferred outcome, with the result that both of them invest and both lose money. Indeed, this is what seems to have happened in the real-life play of this game. In Chapter 7, we investigate this type of game further, showing that it has a third Nash equilibrium, in mixed strategies.

Analysis of Figure 6.2 shows that the first-stage game in our example does not have a unique Nash equilibrium. This problem is not too serious, because we can leave the solution ambiguous to the extent that was done in the preceding paragraph. Matters would be worse if the second-stage game did not have a unique equilibrium. Then it would be essential to specify the precise process by which an outcome gets selected so that we could figure out the second-stage payoffs and use them to roll back to the first stage.

The second-stage pricing game shown in the table in the bottom-right quadrant of Figure 6.1 is one part of the complete two-stage game. However, it is also

		GLOBALDIALOG	
		Don't	Invest
CROSSTALK	Don't	0, 0	0, 14
	Invest	14, 0	−2, −2

FIGURE 6.2 First-Stage Investment Game (After Substituting Rolled-Back Payoffs from the Equilibrium of the Second Stage)

[2]As is usual in a prisoners' dilemma, if the firms could successfully collude and charge high prices, both could get the higher payoff of 2. But this outcome is not an equilibrium, because each firm is tempted to cheat to try to get the much higher payoff of 6.

both firms get 0. However, all of the other combinations of actions for the two firms lead to rooms that lead to further choices; so we cannot yet show the payoffs in those cells. Instead, we show branches leading to the second floor. The northeast and southwest rooms show only the payoff to the firm that has not invested; the branches leading from each of these rooms take us to single-firm pricing decisions in the second stage. The southeast room leads to a multiroom second-floor structure within the tree house, which represents the second-stage pricing game that is played if both firms have invested in the first stage. This second-floor structure has four rooms corresponding to the four combinations of the two firms' pricing moves.

All of the second-floor branches and rooms are like terminal nodes of a game tree, so we can show the payoffs in each case. Payoffs here consist of each firm's operating profits minus the previous investment costs; payoff values are written in billions of dollars.

Consider the branch leading to the southwest corner of Figure 6.1. The game arrives in that corner if CrossTalk is the only firm that has invested. Then, if it chooses the high price, its operating profit is $400 × 60 million = $24 billion; after subtracting the $10 billion investment cost, its payoff is $14 billion, which we write as 14. In the same corner, if CrossTalk chooses the low price, then its operating profit is $200 × 80 million = $16 billion, yielding the payoff 6 after accounting for its original investment. In this situation, GlobalDialog's payoff is 0, as shown in the southwest room of the first floor of our tree. Similar calculations for the case in which GlobalDialog is the only firm to invest give us the payoffs shown in the northeast corner of Figure 6.1; again, the payoff of 0 for CrossTalk is shown in the northeast room of the first-stage game table.

If both firms invest, both play the second-stage pricing game illustrated in the southeast corner of the figure. When both choose the high price in the second stage, each gets operating profit of $400 × 30 million (half of the market), or $12 billion; after subtracting the $10 billion investment cost, each is left with a net profit of $2 billion, or a payoff of 2. If both firms choose the low price in the second stage, each gets operating profit of $200 × 40 million = $8 billion, and, after subtracting the $10 billion investment cost, each is left with a net loss of $2 billion, or a payoff of −2. Finally, if one firm charges the high price and the other firm the low price, then the low-price firm has operating profit of $200 × 80 million = $16 billion, leading to the payoff 6, while the high-price firm gets no operating profit and simply loses its $10 billion investment, for a payoff of −10.

As with any multistage game in Chapter 3, we must solve this game backward, starting with the second-stage game. In the two single-firm decision problems, we see at once that the high-price policy yields the higher payoff. We highlight this by showing that payoff in a larger-size type.

The second-stage pricing game has to be solved by using methods developed in Chapter 4. It is immediately evident, however, that this game is a prisoners' dilemma. Low is the dominant strategy for each firm; so the outcome is

the room in the southeast corner of the second-stage game table; each firm gets payoff −2.[2] Again, we show these payoffs in a larger type size to highlight the fact that they are the payoffs obtained in the second-stage equilibrium.

Rollback now tells us that each first-stage configuration of moves should be evaluated by looking ahead to the equilibrium of the second-stage game (or the optimum second-stage decision) and the resulting payoffs. We can therefore substitute the payoffs that we have just calculated into the previously empty or partly empty rooms on the first floor of our tree house. This substitution gives us a first floor with known payoffs, shown in Figure 6.2.

Now we can use the methods of Chapter 4 to solve this simultaneous-move game. You should immediately recognize the game in Figure 6.2 as a chicken game. It has two Nash equilibria, each of which entails one firm choosing Invest and the other choosing Don't. The firm that invests makes a huge profit; so each firm prefers the equilibrium in which it is the investor while the other firm stays out. In Chapter 4, we briefly discussed the ways in which one of the two equilibria might get selected. We also pointed out the possibility that each firm might try to get its preferred outcome, with the result that both of them invest and both lose money. Indeed, this is what seems to have happened in the real-life play of this game. In Chapter 7, we investigate this type of game further, showing that it has a third Nash equilibrium, in mixed strategies.

Analysis of Figure 6.2 shows that the first-stage game in our example does not have a unique Nash equilibrium. This problem is not too serious, because we can leave the solution ambiguous to the extent that was done in the preceding paragraph. Matters would be worse if the second-stage game did not have a unique equilibrium. Then it would be essential to specify the precise process by which an outcome gets selected so that we could figure out the second-stage payoffs and use them to roll back to the first stage.

The second-stage pricing game shown in the table in the bottom-right quadrant of Figure 6.1 is one part of the complete two-stage game. However, it is also

		GLOBALDIALOG	
		Don't	Invest
CROSSTALK	Don't	0, 0	0, 14
	Invest	14, 0	−2, −2

FIGURE 6.2 First-Stage Investment Game (After Substituting Rolled-Back Payoffs from the Equilibrium of the Second Stage)

[2]As is usual in a prisoners' dilemma, if the firms could successfully collude and charge high prices, both could get the higher payoff of 2. But this outcome is not an equilibrium, because each firm is tempted to cheat to try to get the much higher payoff of 6.

a full-fledged game in its own right, with a fully specified structure of players, strategies, and payoffs. To bring out this dual nature more explicitly, it is called a **subgame** of the full game.

More generally, a subgame is the part of a multimove game that begins at a particular node of the original game. The tree for a subgame is then just that part of the tree for the full game that takes this node as its root, or initial, node. A multimove game has as many subgames as it has decision nodes.

B. Configurations of Multistage Games

In the multilevel game illustrated in Figure 6.1, each stage consists of a simultaneous-move game. However, that may not always be the case. Simultaneous and sequential components may be mixed and matched in any way. We give two more examples to clarify this point and to reinforce the ideas introduced in the preceding section.

The first example is a slight variation of the CrossTalk–GlobalDialog game. Suppose one of the firms—say, GlobalDialog—has already made the $10 billion investment in the fiber-optic network. CrossTalk knows of this investment and now has to decide whether to make its own investment. If CrossTalk does not invest, then GlobalDialog will have a simple pricing decision to make. If CrossTalk invests, then the two firms will play the second-stage pricing game already described. The tree for this multistage game has conventional branches at the initial node and has a simultaneous-move subgame starting at one of the nodes to which these initial branches lead. The complete tree is shown in Figure 6.3.

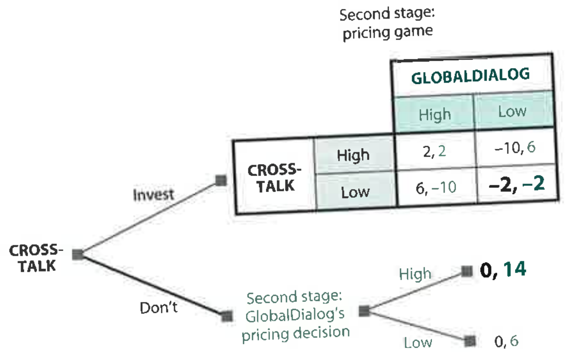

FIGURE 6.3 Two-Stage Game When One Firm Has Already Invested

When the tree has been set up, it is easy to analyze the game. We show the rollback analysis in Figure 6.3 by using large type for the equilibrium payoffs that result from the second-stage game or decision and a thicker branch for CrossTalk's first-stage choice. In words, CrossTalk figures out that, if it invests, the ensuing prisoners' dilemma of pricing will leave it with payoff −2, whereas staying out will get it 0. Thus it prefers the latter. GlobalDialog gets 14 instead of the −2 that it would have gotten if CrossTalk had invested, but CrossTalk's concern is to maximize its own payoff and not to ruin GlobalDialog deliberately.

This analysis does raise the possibility, though, that GlobalDialog may try to get its investment done quickly before CrossTalk makes its decision so as to ensure its most preferred outcome from the full game. And CrossTalk may try to beat GlobalDialog to the punch in the same way. In Chapter 10, we study some methods, called strategic moves, that may enable players to secure such advantages.

Our second example comes from football. Before each play, the coach for the offense chooses the play that his team will run; simultaneously, the coach for the defense sends his team out with instructions on how they should align themselves to counter the offense. Thus these moves are simultaneous. Suppose the offense has just two alternatives, a safe play and a risky play, and the defense may align itself to counter either of them. If the offense has planned to run the risky play and the quarterback sees the defensive alignment that will counter it, he can change the play at the line of scrimmage. And the defense, hearing the change, can respond by changing its own alignment. Thus we have a simultaneous-move game at the first stage, and one of the combination of choices of moves at this stage leads to a sequential-move subgame. Figure 6.4 shows the complete tree.

This is a zero-sum game, and we show only the offense's payoffs, measured in the number of yards that they expect to gain. The safe play gets 2 yards, even if the defense is ready for it; if the defense is not ready for it, the safe play does not do much better, gaining 6 yards. The risky play, if it catches the defense unready to cover it, gains 30 yards. But if the defense is ready for the risky play, the offense loses 10 yards. We show this payoff of −10 for the offense at the terminal node where the offense does not change the play. If the offense changes the play (back to safe), the payoffs are 2 if the defense responds and 6 if it does not; these payoffs are the same as those that arise when the offense plans the safe play from the start.

We show the chosen branches in the sequential subgame as thick lines in Figure 6.4. It is easy to see that, if the offense changes its play, the defense will respond to keep the offense's gain to 2 rather than 6 and that the offense should change the play to get 2 rather than −10. Rolling back, we should put the resulting payoff, 2, in the bottom-right cell of the simultaneous-move game of the first stage. Then we see that this game has no Nash equilibrium in pure strategies.

First stage:
coaches choose alignment

		DEFENSE TO COVER	
		Safe	Risky
OFFENSE TO PLAY	Safe	2	6
	Risky	30	

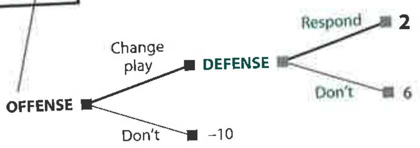

FIGURE 6.4 Simultaneous-Move First Stage Followed by Sequential Moves

The reason is the same as that in the tennis game of Chapter 4, Section 8; one player (defense) wants to match the moves (align to counter the play that the offense is choosing) while the other (offense) wants to unmatch moves (catch the defense in the wrong alignment). In Chapter 7, we show how to calculate the mixed-strategy equilibrium of such a game. It turns out that the offense should choose the risky play with probability 1/8, or 12.5%.

2 CHANGING THE ORDER OF MOVES IN A GAME

The games considered in preceding chapters were presented as either sequential or simultaneous in nature. We used the appropriate tools of analysis to predict equilibria in each type of game. In Section 1 of this chapter, we discussed games with elements of both sequential and simultaneous play. These games required both sets of tools to find solutions. But what about games that could be played either sequentially or simultaneously? How would changing the play of a particular game and thus changing the appropriate tools of analysis alter the expected outcomes?

The task of turning a sequential-play game into a simultaneous one requires changing only the timing or observability with which players make their choices of moves. Sequential-move games become simultaneous if the players cannot observe moves made by their rivals before making their own choices. In that case, we would analyze the game by searching for a Nash equilibrium rather than for a rollback equilibrium. Conversely, a simultaneous-move game could become sequential if one player were able to observe the other's move before choosing her own.

Any changes to the rules of the game can also change its outcomes. Here, we illustrate a variety of possibilities that arise owing to changes in different types of games.

A. Changing Simultaneous-Move Games into Sequential-Move Games

I. NO CHANGE IN OUTCOME Certain games have the same outcomes in the equilibria of both simultaneous and sequential versions and regardless of the order of play in the sequential-play game. This result generally arises only when both or all players have dominant strategies. We show that it holds for the prisoners' dilemma.

Consider the prisoners' dilemma game of Chapter 4, in which a husband and wife are being questioned regarding their roles in a crime. The simultaneous version of that game, reproduced in Figure 6.5a, can be redrawn as either of the sequential-play games shown in Figure 6.5b and c. As in Figure 4.4, the payoff numbers indicate years in jail; so low numbers are better than high ones. In Figure 6.5b, Husband chooses his strategy before Wife does; so she knows what he has chosen before making her own choice; in Figure 6.5c the roles are reversed.

The Nash equilibrium of the prisoners' dilemma game in Figure 6.5a is for each player to confess (or to defect from cooperating with the other). Using rollback to solve the sequential versions of the game, illustrated in Figure 6.5b and c, we see that the second player does best to confess if the first has confessed (10 rather than 25 years in jail) and the second player also does best to confess if the first has denied (1 year rather than 3 years of jail). Given these choices by the second player, the first player does best to confess (10 rather than 25 years in jail). The equilibrium entails 10 years of jail for both players regardless of which player moves first. Thus, the equilibrium is the same in all three versions of this game.

II. FIRST-MOVER ADVANTAGE A first-mover advantage may emerge when the rules of a game are changed from simultaneous to sequential play. At a minimum, if the simultaneous-move version has multiple equilibria, the sequential-move version enables the first mover to choose his preferred outcome. We illustrate such a situation with the use of chicken, the game in which two teenagers drive toward each other in their cars, both determined not to swerve. We reproduce the strategic form of Figure 4.14 in Figure 6.6a and two extensive forms, one for each possible ordering of play, in Figure 6.6b and c.

Under simultaneous play, the two outcomes in which one player swerves (is "chicken") and the others goes straight (is "tough") are both pure-strategy Nash

(a) Simultaneous play

		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

(b) Sequential play: Husband moves first

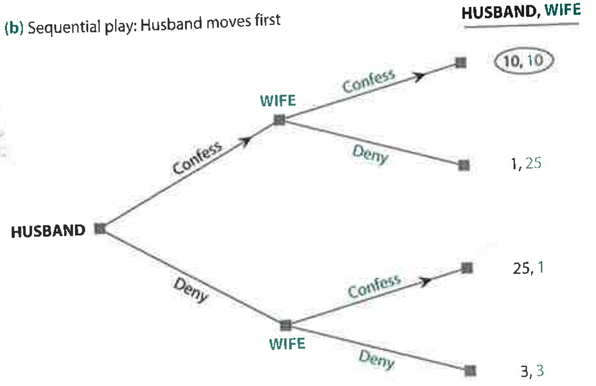

(c) Sequential moves: Wife moves first

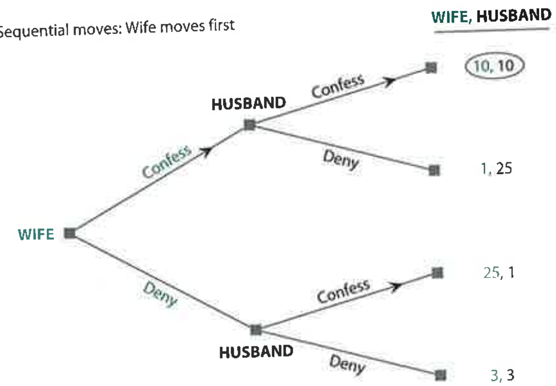

FIGURE 6.5 Three Versions of the Prisoners' Dilemma Game

(a) Simultaneous play

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

(b) Sequential play: James moves first

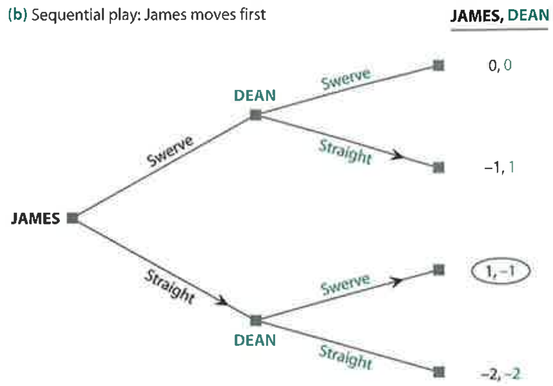

(c) Sequential play: Dean moves first

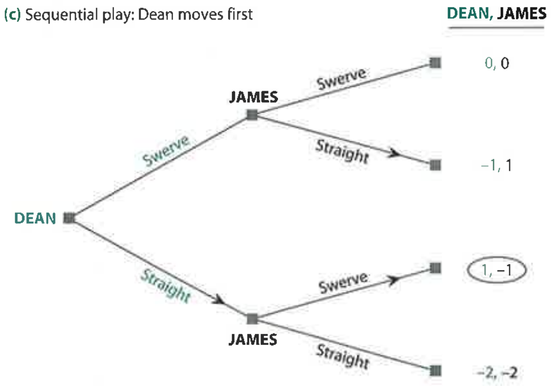

FIGURE 6.6 Chicken in Simultaneous- and Sequential-Play Versions

equilibria. Without specification of some historical, cultural, or other convention, neither has a claim to be a focal point. Our analysis in Chapter 4 suggested that coordinated play could help the players in this game, perhaps through an agreement to alternate between the two equilibria.

When we alter the rules of the game to allow one of the players the opportunity to move first, there are no longer two equilibria. Rather, we see that the second mover's equilibrium strategy is to choose the action opposite that chosen by the first mover. Rollback then shows that the first mover's equilibrium strategy is Straight. We see in Figure 6.6b and c that allowing one person to move first and to be observed making the move results in a single rollback equilibrium in which the first mover gets a payoff of 1, while the second mover gets a payoff of −1. The actual play of the game becomes almost irrelevant under such rules, which may make the sequential version uninteresting to many observers. Although teenagers might not want to play such a game with the rule change, the strategic consequences of the change are significant.

III. SECOND-MOVER ADVANTAGE In other games, a second-mover advantage may emerge when simultaneous play is changed into sequential play. This can be illustrated using the tennis game of Chapter 4. Recall that, in that game, Evert is planning the location of her return while Navratilova considers where to cover. The version considered earlier assumed that both players were skilled at disguising their intended moves until the very last moment so that they moved at essentially the same time. If Evert's movement as she goes to hit the ball belies her shot intentions, however, then Navratilova can react and move second in the game. In the same way, if Navratilova leans toward the side that she intends to cover before Evert actually hits her return, then Evert is the second mover. Figure 6.7 shows all three possibilities. The simultaneous-move version is Figure 4.15 reproduced as Figure 6.7a; the two orderings of the sequential-play game are Figure 6.7b and c.

The simultaneous-play version of this game has no equilibrium in pure strategies. In each ordering of the sequential version, however, there is a unique rollback equilibrium outcome; the equilibrium differs, depending on who moves first. If Evert moves first, then Navratilova chooses to cover whichever direction Evert chooses and Evert opts for a down-the-line shot. Each player is expected to win the point half the time in this equilibrium. If the order is reversed, Evert chooses to send her shot in the opposite direction from that which Navratilova covers; so Navratilova should move to cover crosscourt. In this case, Evert is expected to win the point 80% of the time. The second mover does better by being able to respond optimally to the opponent's move.

We return to the simultaneous version of this game in Chapter 7. There we show that it does have a Nash equilibrium in mixed strategies. In that equilibrium, Evert succeeds on average 62% of the time. Her success rate in the

(a) Simultaneous play

		NAVRATILOVA	
		DL	CC
EVERT	DL	50	80
	CC	90	20

(b) Sequential play: Evert moves first

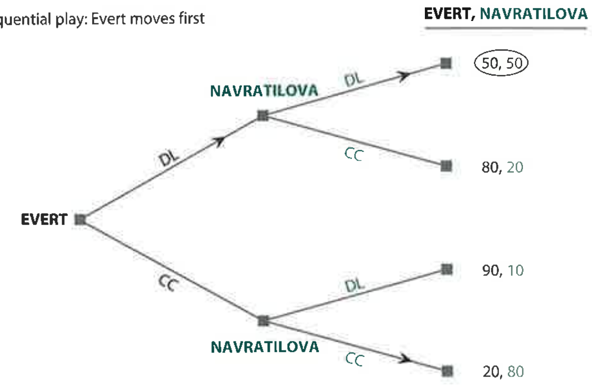

(c) Sequential play: Navratilova moves first

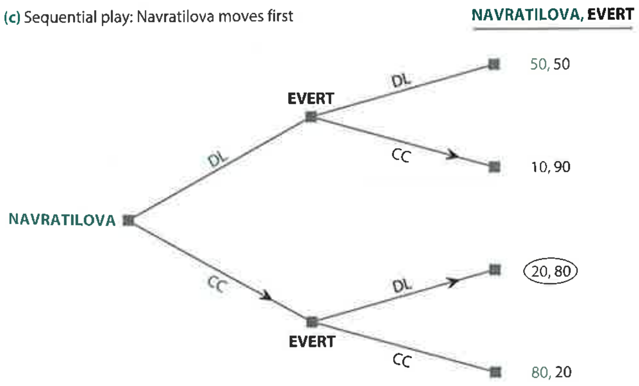

FIGURE 6.7 Tennis Game in Simultaneous- and Sequential-Play Versions

mixed-strategy equilibrium of the simultaneous game is thus better than the 50% that she gets by moving first but is worse than the 80% that she gets by moving second, in the two sequential-move versions.

IV. BOTH PLAYERS MAY DO BETTER That a game may have a first-mover or a second-mover advantage, which is suppressed when moves have to be simultaneous but emerges when an order of moves is imposed, is quite intuitive. Somewhat more surprising is the possibility that both players may do better under one set of rules of play than under another. We illustrate this possibility by using the game of monetary and fiscal policies played by the Federal Reserve and the Congress. In Chapter 4, we studied this game with simultaneous moves; we reproduce the payoff table (Figure 4.5) as Figure 6.8a and show the two sequential-move versions as Figure 6.8b and c. For brevity, we write the strategies as Balance and Deficit instead of Budget Balance and Budget Deficit for the Congress and as High and Low instead of High Interest Rates and Low Interest Rates for the Fed.

In the simultaneous-move version, the Congress has a dominant strategy (Deficit), and the Fed, knowing this, chooses High, yielding payoffs of 2 to both players. Almost the same thing happens in the sequential version where the Fed moves first. The Fed foresees that, for each choice it might make, the Congress will respond with Deficit. Then High is the better choice for Fed, yielding 2 instead of 1.

But the sequential-move version where the Congress moves first is different. Now the Congress foresees that, if it chooses Deficit, the Fed will respond with High, whereas, if it chooses Balance, the Fed will respond with Low. Of these two developments, the Congress prefers the latter, where it gets payoff 3 instead of 2. Therefore the rollback equilibrium with this order of moves is for the Congress to choose a balanced budget and the Fed to respond with low interest rates. The resulting payoffs, 3 for the Congress and 4 for the Fed, are better for both players than those of the other two versions.

The difference between the two outcomes is even more surprising because the better outcome obtained in Figure 6.8c results from the Congress choosing Balance, which is its dominated strategy in Figure 6.8a. To resolve the apparent paradox, one must understand more precisely the meaning of dominance. For Deficit to be a dominant strategy, it must be better than Balance from the Congress's perspective for *each given* choice of the Fed. This type of comparison between Deficit and Balance is relevant in the simultaneous-move game because there the Congress must make a decision without knowing the Fed's choice. The Congress must think through, or formulate a belief about, the Fed's action, and choose its best response to that. In our example, this best response is always Deficit for the Congress. The concept of dominance is also relevant with sequential moves if the Congress moves second, because then it knows what the

(a) Simultaneous moves

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

(b) Sequential moves: Fed moves first

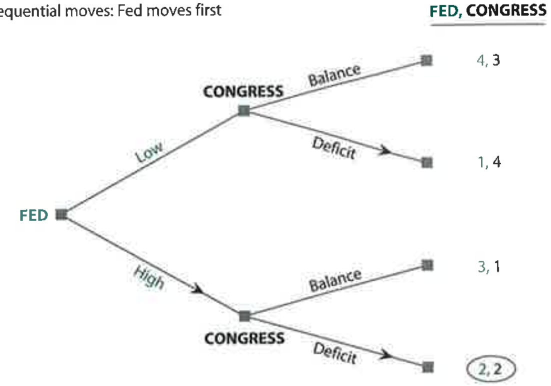

(c) Sequential moves: Congress moves first

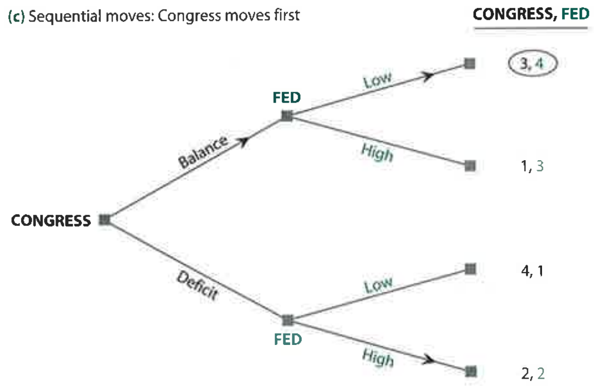

FIGURE 6.8 Three Versions of the Monetary–Fiscal Policy Game

Fed has already done and merely picks its best response, which is always Deficit. However, if the Congress moves first, it cannot take the Fed's choice as *given*. Instead, it must recognize how the Fed's second move will be affected by its own first move. Here it knows that the Fed will respond to Deficit with High and to Balance with Low. The Congress is then left to choose between these two alternatives; its most preferred outcome of Deficit and Low becomes irrelevant because it is precluded by the Fed's response.

The idea that dominance may cease to be a relevant concept for the first mover reemerges in Chapter 10. There we consider the possibility that one player or the other may deliberately change the rules of a game to become the first mover. Players can alter the outcome of the game in their favor in this way.

Suppose that the two players in our current example could choose the order of moves in the game. In this case, they would agree that the Congress should move first. Indeed, when budget deficits and inflation threaten, the chairs of the Federal Reserve in testimony before various congressional committees often offer such deals; they promise to respond to fiscal discipline by lowering interest rates. But it is often not enough to make a verbal deal with the other player. The technical requirements of a first move—namely, that it be observable to the second mover and not reversible thereafter—must be satisfied. In the context of macroeconomic policies, it is fortunate that the legislative process of fiscal policy in the United States is both very visible and very slow, whereas monetary policy can be changed quite quickly in a meeting of the Federal Reserve Board. Therefore the sequential play where the Congress moves first and the Fed moves second is quite realistic.

B. Other Changes in the Order of Moves

The preceding section presented various examples in which the rules of the game were changed from simultaneous play to sequential play. We saw how and why such rule changes can change the outcome of a game. The same examples also serve to show what happens if the rules are changed in the opposite direction, from sequential to simultaneous moves. Thus, if a first- or a second-mover advantage exists with sequential play, it can be lost under simultaneous play. And if a specific order benefits both players, then losing the order can hurt both.

The same examples also show us what happens if the rules are changed to reverse the order of play while keeping the sequential nature of a game unchanged. If there is a first-mover or a second-mover advantage, then the player who shifts from moving first to moving second may benefit or lose accordingly, with the opposite change for the other player. And if one order is in the common interests of both, then an externally imposed change of order can benefit or hurt them both.

3 CHANGE IN THE METHOD OF ANALYSIS

Game trees are the natural way to display sequential-move games, and payoff tables the natural representation of simultaneous-move games. However, each technique can be adapted to the other type of game. Here we show how to translate the information contained in one illustration to an illustration of the other type. In the process, we develop some new ideas that will prove useful in subsequent analysis of games.

A. Illustrating Simultaneous-Move Games by Using Trees

Consider the game of the passing shot in tennis as originally described in Chapter 4, where the action is so quick that moves are truly simultaneous, as shown in Figure 6.7a. But suppose we want to show the game in extensive form—that is, by using a tree. We show how this can be done in Figure 6.9.

To draw the tree in the figure, we must choose one player—say, Evert—to make her choice at the initial node of the tree. The branches for her two choices, DL and CC, then end in two nodes, at each of which Navratilova makes her choices. However, because the moves are actually simultaneous, Navratilova must choose without knowing what Evert has picked. That is, she must choose without knowing whether she is at the node following Evert's DL branch or the one following Evert's CC branch. Our tree diagram must in some way show this lack of information on Navratilova's part.

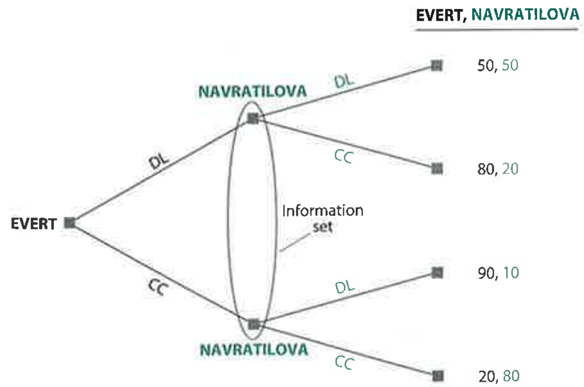

FIGURE 6.9 Simultaneous-Move Tennis Game Shown in Extensive Form

We illustrate Navratilova's strategic uncertainty about the node from which her decision is being made by drawing an oval to surround the two relevant nodes. (An alternative is to connect them by a dotted line; a dotted line is used to distinguish it from the solid lines that represent the branches of the tree.) The nodes within this oval or balloon are called an **information set** for the player who moves there. Such a set indicates the presence of imperfect information for the player; she cannot distinguish between the nodes in the set, given her available information (because she cannot observe the row player's move before making her own). As such, her strategy choice from within a single information set must specify the same move at all the nodes contained in it. That is, Navratilova must choose either DL at both the nodes in this information set or CC at both of them. She cannot choose DL at one and CC at the other, as she could in Figure 6.7b, where the game had sequential moves and she moved second.

Accordingly, we must adapt our definition of strategy. In Chapter 3, we defined a strategy as a complete plan of action, specifying the move that a player would make at each *node* where the rules of the game specified that it was her turn to move. We should now more accurately redefine a strategy as a complete plan of action, specifying the move that a player would make at each *information set* at whose nodes the rules of the game specify that it is her turn to move.

The concept of an information set is also relevant when a player faces external uncertainty about some conditions that affect his decision, rather than about another player's moves. For example, a farmer planting a crop is uncertain about the weather during the growing season, although he knows the probabilities of various alternative possibilities from past experience or meteorological forecasts. We can regard the weather as a random choice of an outside player, Nature, who has no payoffs but merely chooses according to known probabilities.[3] We can then enclose the various nodes corresponding to Nature's moves into an information set for the farmer, constraining the farmer's choice to be the same at all of these nodes. Figure 6.10 illustrates this situation.

Using the concept of an information set, we can formalize the concepts of perfect and imperfect information in a game, which we introduced in Chapter 2 (Section 2.D). A game has perfect information if it has neither strategic nor external uncertainty, which will happen if it has no information sets enclosing two or more nodes. Thus a game has perfect information if all of its information sets consist of singleton nodes.

[3]Some people believe that Nature is actually a malevolent player who plays a zero-sum game with us, so its payoffs are higher when ours are lower. For example, it is more likely to rain if we have forgotten to bring an umbrella. We understand such thinking, but it does not have real statistical support.

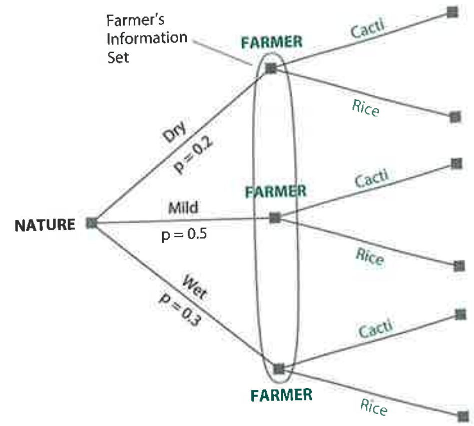

FIGURE 6.10 Nature and Information Sets

Although this representation is conceptually simple, it does not provide any simpler way of solving the game. Therefore we use it only occasionally, where it conveys some point more simply. Some examples of game illustrations using information sets can be found later in Chapters 9 and 15.

B. Showing and Analyzing Sequential-Move Games in Strategic Form

Consider now the sequential-move game of monetary and fiscal policy from Figure 6.8c, in which the Congress has the first move. Suppose we want to show it in normal or strategic form—that is, by using a payoff table. The rows and the columns of the table are the strategies of the two players. We must therefore begin by specifying the strategies.

For the Congress, the first mover, listing its strategies is easy. There are just two moves—Balance and Deficit—and they are also the two strategies. For the second mover, matters are more complex. Remember that a strategy is a complete plan of action, specifying the moves to be made at each node where it is a player's turn to move. Because the Fed gets to move at two nodes (and because we are supposing that this game actually has sequential moves and so the two nodes are not confounded into one information set) and can choose either Low or High at each node, there are four combinations of its choice patterns. These combinations are (1) Low if Balance, High if Deficit (we write this as "L if B, H if

		FED			
		L if B, H if D	H if B, L if D	Low always	High always
CONGRESS	Balance	3, 4	1, 3	3, 4	1, 3
	Deficit	2, 2	4, 1	4, 1	2, 2

FIGURE 6.11 Sequential-Move Game of Monetary and Fiscal Policy in Strategic Form

D" for short); (2) High if Balance, Low if Deficit ("H if B, L if D" for short); (3) Low always; and (4) High always.

We show the resulting two-by-four payoff matrix in Figure 6.11. The last two columns are no different from those for the two-by-two payoff matrix for the game under simultaneous-move rules (Figure 6.8a). This is because, if the Fed is choosing a strategy in which it makes the same move always, it is just as if the Fed were moving without taking into account what the Congress had done; it is as if their moves were simultaneous. But calculation of the payoffs for the first two columns, where the Fed's second move does depend on the Congress's first move, needs some care.

To illustrate, consider the cell in the first row and the second column. Here the Congress is choosing Balance, and the Fed is choosing "H if B, L if D." Given Congress's choice, the Fed's actual choice under this strategy is High. Then the payoffs are those for the Balance and High combination—namely, 1 for Congress and 3 for the Fed.

Cell-by-cell inspection quickly shows that the game has two pure-strategy Nash equilibria, which we show by shading the cells gray. One is in the top-left cell, where the Congress's strategy is Balance and the Fed's is "L if B, H if D," and so the Fed's actual choice is L. This outcome is just the rollback equilibrium of the sequential-move game. But there is another Nash equilibrium in the bottom-right cell, where the Congress chooses Deficit and the Fed chooses "High always." As always in a Nash equilibrium, neither player has a clear reason to deviate from the strategies that lead to this outcome. The Congress would do worse by switching to Balance, and the Fed could do no better by switching to any of its other three strategies, although it could do just as well with "L if B, H if D."

The sequential-move game, when analyzed in its extensive form, produced just one rollback equilibrium. But when analyzed in its normal or strategic form, it has two Nash equilibria. What is going on?

The answer lies in the different nature of the logic of Nash and rollback analyses. Nash equilibrium requires that neither player have a reason to deviate, given the strategy of the other player. However, rollback does not take the strategies of later movers as given. Instead, it asks what would be optimal to do if the opportunity to move actually arises.

In our example, the Fed's strategy of "High always" does not satisfy the criterion of being optimal if the opportunity to move actually arises. If the Congress chose Deficit, then High is indeed the Fed's optimal response. However, if the Congress chose Balance and the Fed had to respond, it would want to choose Low, not High. So "High always" does not describe the Fed's optimal response in all possible configurations of play and cannot be a rollback equilibrium. But the logic of Nash equilibrium does not impose such a test, instead regarding the Fed's "High always" as a strategy that the Congress could legitimately take as given. If it does so, then Deficit is the Congress's best response. And, conversely, "High always" is one best response of the Fed to the Congress's Deficit (although it is tied with "L if B, H if D"). Thus the pair of strategies "Deficit" and "High always" are mutual best responses and constitute a Nash equilibrium, although they do not constitute a rollback equilibrium.

We can therefore think of rollback as a further test, supplementing the requirements of a Nash equilibrium and helping to select from among multiple Nash equilibria of the strategic form. In other words, it is a refinement of the Nash equilibrium concept.

To state this idea somewhat more precisely, recall the concept of a subgame. At any one node of the full game tree, we can think of the part of the game that begins there as a subgame. In fact, as successive players make their choices, the play of the game moves along a succession of nodes, and each move can be thought of as starting a subgame. The equilibrium derived by using rollback corresponds to one particular succession of choices in each subgame and gives rise to one particular path of play. Certainly, other paths of play are consistent with the rules of the game. We call these other paths **off-equilibrium paths,** and we call any subgames that arise along these paths **off-equilibrium subgames,** for short.

With this terminology, we can now say that the equilibrium path of play is itself determined by the players' expectations of what would happen if they chose a different action—if they moved the game to an off-equilibrium path and started an off-equilibrium subgame. Rollback requires that all players make their best choices in *every* subgame of the larger game, whether or not the subgame lies along the path to the ultimate equilibrium outcome.

Strategies are complete plans of action. Thus a player's strategy must specify what she will do in each eventuality, or each and every node of the game, whether on or off the equilibrium path, where it is her turn to act. When one such node arrives, only the plan of action starting there—namely, the part of the full strategy that pertains to the subgame starting at that node—is pertinent. This part is called the **continuation** of the strategy for that subgame. Rollback requires that the equilibrium strategy be such that its continuation in every subgame is optimal for the player whose turn it is to act at that node, whether or not the node and the subgame lie on the equilibrium path of play.

Return to the monetary policy game with the Congress moving first, and consider the second Nash equilibrium that arises in its strategic form. Here the path of play is for the Congress to choose Deficit and the Fed to choose High. On the equilibrium path, High is indeed the Fed's best response to Deficit. The Congress's choice of Balance would be the start of an off-equilibrium path. It leads to a node where a rather trivial subgame starts—namely, a decision by the Fed. The Fed's purported equilibrium strategy "High always" asks it to choose High in this subgame. But that is not optimal; this second equilibrium is specifying a nonoptimal choice for an off-equilibrium subgame.

In contrast, the equilibrium path of play for the Nash equilibrium in the upper-left corner of Figure 6.11 is for the Congress to choose Balance and the Fed to follow with Low. The Fed is responding optimally on the equilibrium path. The off-equilibrium path would have the Congress choosing Deficit, and the Fed, given its strategy of "L if B, H if D," would follow with High. It is optimal for the Fed to respond to Deficit with High, so the strategy remains optimal off the equilibrium path, too.

The requirement that continuation of a strategy remain optimal under all circumstances is important because the equilibrium path itself is the result of players' thinking strategically about what would happen if they did something different. A later player may try to achieve an outcome that she would prefer by threatening the first mover that certain actions would be met with dire responses or by promising that certain other actions would be met with nice responses. But the first mover will be skeptical of the **credibility** of such threats and promises. The only way to remove that doubt is to check if the stated responses would actually be optimal if the need arose. If the responses are not optimal, then the threats or promises are not credible, and the responses would not be observed along the equilibrium path of play.

The equilibrium found by using rollback is called a **subgame-perfect equilibrium (SPE).** It is a set of strategies (complete plans of action), one for each player, such that, at every node of the game tree, whether or not the node lies along the equilibrium path of play, the continuation of the same strategy in the subgame starting at that node is optimal for the player who takes the action there. More simply, an SPE requires players to use strategies that constitute a Nash equilibrium in every subgame of the larger game.

In fact, as a rule, in games with finite trees and perfect information, where players can observe every previous action taken by all other players so that there are no multiple nodes enclosed in one information set, rollback finds the unique

(except for trivial and exceptional cases of ties) subgame-perfect equilibrium of the game. Consider: if you look at any subgame that begins at the last decision node for the last player who moves, the best choice for that player is the one that gives her the highest payoff. But that is precisely the action chosen with the use of rollback. As players move backward through the game tree, rollback eliminates all unreasonable strategies, including incredible threats or promises, so that the collection of actions ultimately selected is the SPE. Therefore, for the purposes of this book, subgame perfectness is just a fancy name for rollback. At more advanced levels of game theory, where games include complex information structures and information sets, subgame perfectness becomes a richer notion.

4 THREE-PLAYER GAMES

We have restricted the discussion so far in this chapter to games with two players and two moves each. But the same methods also work for some larger and more general examples. We now illustrate this by using the street-garden game of Chapter 3. Specifically, we (1) change the rules of the game from sequential to simultaneous moves and then (2) keep the moves sequential but show

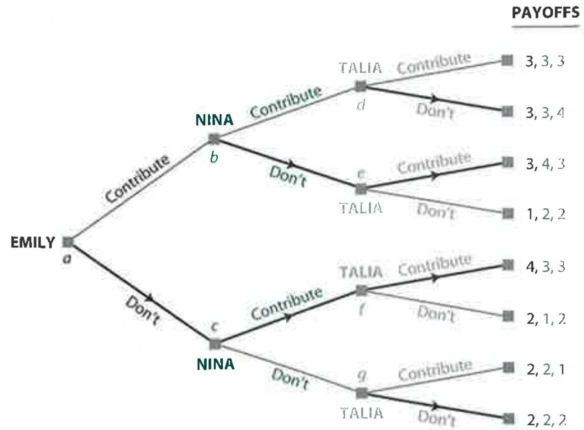

FIGURE 6.12 The Street-Garden Game with Sequential Moves

and analyze the game in its strategic form. First we reproduce the tree of that sequential-move game (Figure 3.6) as Figure 6.12 here and remind you of the rollback equilibrium.

The equilibrium strategy of the first mover (Emily) is simply a move, "Don't contribute." The second mover chooses from among four possible strategies (choice of two responses at each of two nodes) and chooses the strategy "Don't contribute (D) if Emily has chosen her Contribute, and Contribute (C) if Emily has chosen her Don't contribute," or, more simply, "D if C, C if D," or even more simply "DC." Talia has 16 available strategies (choice of two responses at each of four nodes), and her equilibrium strategy is "D following Emily's C and Nina's C, C following their CD, C following their DC, and D following their DD," or "DCCD" for short.

Remember, too, the reason for these choices. The first mover has the opportunity to choose Don't, knowing that the other two will recognize that the nice garden won't be forthcoming unless they contribute and that they like the nice garden sufficiently strongly that they will contribute.

Now we change the rules of the game to make it a simultaneous-move game. (In Chapter 4, we solved a simultaneous-move version with somewhat different payoffs; here we keep the payoffs the same as in Chapter 3.) The payoff matrix is in Figure 6.13. Cell-by-cell inspection shows very easily that there are four Nash equilibria.

In three of the Nash equilibria of the simultaneous-move game, two players contribute, while the third does not. These equilibria are similar to the rollback equilibrium of the sequential-move game. In fact, each one corresponds to the rollback equilibrium of the sequential game with a particular order of play. Further, any given order of play in the sequential-move version of this game leads to the same simultaneous-move payoff table.

But there is also a fourth Nash equilibrium here, where no one contributes. *Given* the specified strategies of the other two—namely, Don't contribute—any one player is powerless to bring about the nice garden and therefore chooses

TALIA chooses:

Contribute

EMILY \ NINA	Contribute	Don't
Contribute	3, 3, 3	3, 4, 3
Don't	4, 3, 3	2, 2, 1

Don't Contribute

EMILY \ NINA	Contribute	Don't
Contribute	3, 3, 4	1, 2, 2
Don't	2, 1, 2	2, 2, 2

FIGURE 6.13 The Street-Garden Game with Simultaneous Moves

not to contribute as well. Thus, in the change from sequential to simultaneous moves, the first-mover advantage has been lost. Multiple equilibria arise, only one of which retains the original first mover's high payoff.

Next we return to the sequential-move version—Emily first, Nina second, and Talia third—but show the game in its normal or strategic form. In the sequential-move game, Emily has 2 pure strategies, Nina has 4, and Talia has 16; so this means constructing a payoff table that is 2 by 4 by 16. With the use of the same conventions as we used for three-player tables in Chapter 4, this particular game would require a table with 16 "pages" of two-by-four payoff tables. That would look too messy; so we opt instead for a reshuffling of the players. Let Talia be the row player, Nina be the column player, and Emily be the page

	EMILY							
	Contribute				Don't			
	NINA				NINA			
TALIA	CC	CD	DC	DD	CC	CD	DC	DD
CCCC	3, 3, 3	3, 3, 3	3, 4, 3	3, 4, 3	3, 3, 4	1, 2, 2	3, 3, 4	1, 2, 2
CCCD	3, 3, 3	3, 3, 3	3, 4, 3	3, 4, 3	3, 3, 4	2, 2, 2	3, 3, 4	2, 2, 2
CCDC	3, 3, 3	3, 3, 3	3, 4, 3	3, 4, 3	2, 1, 2	1, 2, 2	2, 1, 2	1, 2, 2
CDCC	3, 3, 3	3, 3, 3	2, 2, 1	2, 2, 1	3, 3, 4	1, 2, 2	3, 3, 4	1, 2, 2
DCCC	4, 3, 3	4, 3, 3	3, 4, 3	3, 4, 3	3, 3, 4	1, 2, 2	3, 3, 4	1, 2, 2
CCDD	3, 3, 3	3, 3, 3	3, 4, 3	3, 4, 3	2, 1, 2	2, 2, 2	2, 1, 2	2, 2, 2
CDDC	3, 3, 3	3, 3, 3	2, 2, 1	2, 2, 1	2, 1, 2	1, 2, 2	2, 1, 2	1, 2, 2
DDCC	4, 3, 3	4, 3, 3	2, 2, 1	2, 2, 1	3, 3, 4	1, 2, 2	3, 3, 4	1, 2, 2
CDCD	3, 3, 3	3, 3, 3	2, 2, 1	2, 2, 1	3, 3, 4	2, 2, 2	3, 3, 4	2, 2, 2
DCDC	4, 3, 3	4, 3, 3	3, 4, 3	3, 4, 3	2, 1, 2	1, 2, 2	2, 1, 2	1, 2, 2
DCCD	4, 3, 3	4, 3, 3	3, 4, 3	3, 4, 3	3, 3, 4	2, 2, 2	3, 3, 4	2, 2, 2
CDDD	3, 3, 3	3, 3, 3	2, 2, 1	2, 2, 1	2, 1, 2	2, 2, 2	2, 1, 2	2, 2, 2
DCDD	4, 3, 3	4, 3, 3	3, 4, 3	3, 4, 3	2, 1, 2	2, 2, 2	2, 1, 2	2, 2, 2
DDCD	4, 3, 3	4, 3, 3	2, 2, 1	2, 2, 1	3, 3, 4	2, 2, 2	3, 3, 4	2, 2, 2
DDDC	4, 3, 3	4, 3, 3	2, 2, 1	2, 2, 1	2, 1, 2	1, 2, 2	2, 1, 2	1, 2, 2
DDDD	4, 3, 3	4, 3, 3	2, 2, 1	2, 2, 1	2, 1, 2	2, 2, 2	2, 1, 2	2, 2, 2

FIGURE 6.14 Street-Garden Game in Strategic Form

player. Then "all" that is required to illustrate this game is the 16 by 4 by 2 game table shown in Figure 6.14. The order of payoffs still corresponds to our earlier convention in that they are listed row, column, page player; in our example, that means the payoffs are now listed in the order Talia, Nina, and Emily.

As in the monetary–fiscal policy game between the Fed and the Congress, there are multiple Nash equilibria in the simultaneous street-garden game. (In Exercise S8, we ask you to find them all.) But there is only one subgame-perfect equilibrium, corresponding to the rollback equilibrium found in Figure 6.13. Although cell-by-cell inspection finds all of the Nash equilibria, iterated elimination of dominated strategies can reduce the number of reasonable equilibria for us here. This process works because elimination identifies those strategies that include noncredible components (such as "High always" for the Fed in Section 3.B). As it turns out, such elimination can take us all the way to the unique subgame-perfect equilibrium.

In Figure 6.14, we start with Talia and eliminate all of her (weakly) dominated strategies. This step eliminates all but the strategy listed in the eleventh row of the table, DCCD, which we have already identified as Talia's rollback equilibrium strategy. Elimination can continue with Nina, for whom we must compare outcomes from strategies across both pages of the table. To compare her CC to CD, for example, we look at the payoffs associated with CC in *both pages* of the table and compare these payoffs with the similarly identified payoffs for CD. For Nina, the elimination process leaves only her strategy DC; again, this is the rollback equilibrium strategy found for her above. Finally, Emily has only to compare the two remaining cells associated with her choice of Don't and Contribute; she gets the highest payoff when she chooses Don't and so makes that choice. As before, we have identified her rollback equilibrium strategy.

The unique subgame-perfect outcome in the game table in Figure 6.14 thus corresponds to the cell associated with the rollback equilibrium strategies for each player. Note that the process of iterated elimination that leads us to this subgame-perfect equilibrium is carried out by considering the players in reverse order of the actual play of the game. This order conforms to the order in which player actions are considered in rollback analysis and therefore allows us to eliminate exactly those strategies, for each player, that are not consistent with rollback. In so doing, we eliminate all of the Nash equilibria that are not subgame-perfect.

SUMMARY

Many games include multiple components, some of which entail simultaneous play and others of which entail sequential play. In *two-stage* (and multistage) games, a "tree house" can be used to illustrate the game; this construction allows

the identification of the different stages of play and the ways in which those stages are linked together. Full-fledged games that arise in later stages of play are called *subgames* of the full game.

Changing the rules of a game to alter the timing of moves may or may not alter the equilibrium outcome of a game. Simultaneous-move games that are changed to make moves sequential may have the same outcome (if both players have dominant strategies), may have a first-mover or second-mover advantage, or may lead to an outcome in which both players are better off. The sequential version of a simultaneous game will generally have a unique rollback equilibrium even if the simultaneous version has no equilibrium or multiple equilibria. Similarly, a sequential-move game that has a unique rollback equilibrium may have several Nash equilibria when the rules are changed to make the game a simultaneous-move game.

Simultaneous-move games can be illustrated in a game tree by collecting decision nodes in *information sets* when players make decisions without knowing at which specific node they find themselves. Similarly, sequential-move games can be illustrated by using a game table; in this case, each player's full set of strategies must be carefully identified. Solving a sequential-move game from its strategic form may lead to many possible Nash equilibria. The number of potential equilibria can be reduced by using the criteria of *credibility* to eliminate some strategies as possible equilibrium strategies. This process leads to the *subgame-perfect equilibrium (SPE)* of the sequential-move game. These solution processes also work for games with additional players.

KEY TERMS

continuation (197)
credibility (198)
information set (194)
off-equilibrium path (197)
off-equilibrium subgame (197)
subgame (182)
subgame-perfect equilibrium (SPE) (198)

SOLVED EXERCISES

S1. Consider the simultaneous-move game with two players that has no Nash equilibrium in pure strategies, illustrated in Figure 4.15. If the game were transformed into a sequential-move game, would you expect that game to exhibit a first-mover advantage, a second-mover advantage, or neither? Explain your reasoning.

S2. Consider the game represented by the game tree below. The first mover, Player 1, may move either Up or Down, after which Player 2 may move either Left or Right. Payoffs for the possible outcomes appear below. Reexpress this game in strategic (table) form. Find all of the pure-strategy Nash equilibria in the game. If there are multiple equilibria, indicate which one is subgame-perfect. For those equilibria that are not subgame-perfect, identify the reason (the source of the lack of credibility).

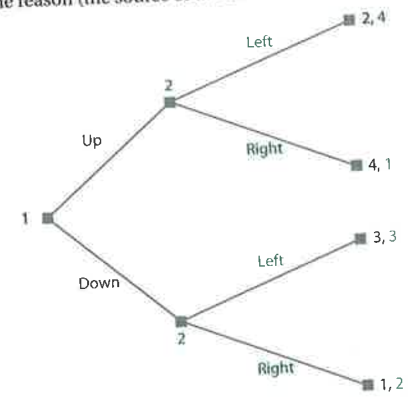

S3. Consider the Airbus–Boeing game in Exercise S4 in Chapter 3. Show that game in strategic form and locate all of the Nash equilibria. Which one of the equilibria is subgame-perfect? For those equilibria that are not subgame-perfect, identify the source of the lack of credibility.

S4. Return to the two-player game tree in part (a) of Exercise S2 in Chapter 3.
(a) Write the game in strategic form, making Scarecrow the row player and Tinman the column player.
(b) Find the Nash equilibrium.

S5. Return to the two-player game tree in part (b) of Exercise S2 in Chapter 3.
(a) Write the game in strategic form. (Hint: Refer to your answer to Exercise S2 of Chapter 3.) Find all of the Nash equilibria. There will be many.
(b) For the equilibria that you found in part (a) that are not subgame-perfect, identify the credibility problems.

S6. Return to the three-player game tree in part (c) of Exercise S2 in Chapter 3.
(a) Draw the game table. Make Scarecrow the row player, Tinman the column player, and Lion the page player. (Hint: Refer to your answer to Exercise S2 of Chapter 3.) Find all of the Nash equilibria. There will be many.

(b) For the equilibria that you found in part (a) that are not subgame-perfect, identify the credibility problems.

S7. Consider a simplified baseball game played between a pitcher and a batter. The pitcher chooses between throwing a fastball or a curve, while the batter chooses which pitch to anticipate. The batter has an advantage if he correctly anticipates the type of pitch. In this constant-sum game, the batter's payoff is the probability that the batter will get a base hit. The pitcher's payoff is the probability that the batter fails to get a base hit, which is simply one minus the payoff of the batter. There are four potential outcomes:
(i) If a pitcher throws a fastball, and the batter guesses fastball, the probability of a hit is 0.300.
(ii) If the pitcher throws a fastball, and the batter guesses curve, the probability of a hit is 0.200.
(iii) If the pitcher throws a curve, and the batter guesses curve, the probability of a hit is 0.350.
(iv) If the pitcher throws a curve, and the batter guesses fastball, the probability of a hit is 0.150.

Suppose that the pitcher is "tipping" his pitches. This means that the pitcher is holding the ball, positioning his body, or doing something else in a way that reveals to the batter which pitch he is going to throw. For our purposes, this means that the pitcher-batter game is a sequential game in which the pitcher announces his pitch choice before the batter has to choose his strategy.
(a) Draw this situation, using a game tree.
(b) Suppose that the pitcher knows he is tipping his pitches but can't stop himself from doing so. Thus, the pitcher and batter are playing the game you just drew. Find the rollback equilibrium of this game.
(c) Now change the timing of the game, so that the batter has to reveal his action (perhaps by altering his batting stance) before the pitcher chooses which pitch to throw. Draw the game tree for this situation, and find the rollback equilibrium.

Now assume that the tips of each player occur so quickly that neither opponent can react to them, so that the game is in fact simultaneous.
(d) Draw a game tree to represent this simultaneous game, indicating information sets where appropriate.
(e) Draw the game table for the simultaneous game. Is there a Nash equilibrium in pure strategies? If so, what is it?

S8. The street-garden game analyzed in Section 4 of this chapter has a 16-by-4-by-2 game table when the sequential-move version of the game is expressed in strategic form, as in Figure 6.14. There are *many* Nash equilibria to be found in this table.

(a) Use best-response analysis to find all of the Nash equilibria in the table in Figure 6.14.
(b) Identify the subgame-perfect equilibrium from among your set of all Nash equilibria. Other equilibrium outcomes look identical to the subgame-perfect one—they entail the same payoffs for each of the three players—but they arise after different combinations of strategies. Explain how this can happen. Describe the credibility problems that arise in the nonsubgame-perfect equilibria.

S9. As it appears in the text, Figure 6.1 represents the two-stage game between CrossTalk and GlobalDialog with a combination of tables and trees. Instead, represent the entire two-stage game in a single, very large game tree. Be careful to label which player makes the decision at each node, and remember to draw information sets between nodes where necessary.

S10. Recall the mall location game in Exercise S8 in Chapter 3. That three-player sequential game has a game tree that is similar to the one for the street-garden game, shown in Figure 6.12.
(a) Draw the tree for the mall location game. How many strategies does each store have?
(b) Illustrate the game in strategic form and find all of the pure-strategy Nash equilibria in the game.
(c) Use iterated dominance to find the subgame-perfect equilibrium. (Hint: Reread the last two paragraphs of Section 4.)

S11. The rules of the mall location game, analyzed in Exercise S10 above, specify that when all three stores request space in Urban Mall, the two bigger (more prestigious) stores get the available spaces. The original version of the game also specifies that the firms move sequentially in requesting mall space.
(a) Suppose that the three firms make their location requests simultaneously. Draw the payoff table for this version of the game and find all of the Nash equilibria. Which one of these equilibria do you think is most likely to be played in practice? Explain.

Now suppose that when all three stores simultaneously request Urban Mall, the two spaces are allocated by lottery, giving each store an equal chance of getting into Urban Mall. With such a system, each would have a two-thirds probability (or a 66.67% chance) of getting into Urban Mall when all three had requested space there, and a one-third probability (33.33% chance) of being alone in the Rural Mall.
(b) Draw the game table for this new version of the simultaneous-play mall location game. Find all of the Nash equilibria of the game. Which one of these equilibria do you think is most likely to be played in practice? Explain.

(c) Compare and contrast the equilibria found in part (b) with the equilibria found in part (a). Do you get the same Nash equilibria? Why or why not?

S12. Return to the game of Monica and Nancy in Exercise S10 of Chapter 5. Assume that Monica and Nancy choose their effort levels sequentially instead of simultaneously. Monica commits to her choice of effort first, and on observing this decision, Nancy commits to her own effort.
(a) What is the subgame-perfect equilibrium to the game where the joint profits are $4m + 4n + mn$, the effort costs to Monica and Nancy are m^2 and n^2 respectively, and Monica commits to an effort level first?
(b) Compare the payoffs of Monica and Nancy with those found in Exercise S10 of Chapter 5. Does this game have a first-mover or a second-mover advantage? Explain.

S13. Extending Exercise S12, Monica and Nancy need to decide which (if either) of them will commit to an effort level first. To do this, each of them simultaneously writes on a separate slip of paper whether or not she will commit first. If they both write "yes" or they both write "no," they choose effort levels simultaneously, as in Exercise S10 in Chapter 5. If Monica writes "yes" and Nancy writes "no," then Monica commits to her move first, as in Exercise S12. If Monica writes "no" and Nancy writes "yes," then Nancy commits to her move first.
(a) Use the payoffs to Monica and Nancy in Exercise S12 above as well as in Exercise S10 in Chapter 5 to construct the game table for the first-stage paper-slip decision game. (Hint: Note the symmetry of the game.)
(b) Find the pure-strategy Nash equilibria of this first-stage game.

UNSOLVED EXERCISES

U1. Consider a game in which there are two players, A and B. Player A moves first and chooses either Up or Down. If A chooses Up, the game is over, and each player gets a payoff of 2. If A moves Down, then B gets a turn and chooses between Left and Right. If B chooses Left, both players get 0; if B chooses Right, A gets 3 and B gets 1.
(a) Draw the tree for this game, and find the subgame-perfect equilibrium.
(b) Show this sequential-play game in strategic form, and find all of the Nash equilibria. Which is or are subgame-perfect? Which is or are not? If any are not, explain why.
(c) What method of solution could be used to find the subgame-perfect equilibrium from the strategic form of the game? (Hint: Refer to the last two paragraphs of Section 4.)

U2. Return to the two-player game tree in part (a) of Exercise U2 in Chapter 3.
- (a) Write the game in strategic form, making Albus the row player and Minerva the column player. Find all of the Nash equilibria.
- (b) For the equilibria you found in part (a) of this exercise that are not subgame-perfect, identify the credibility problems.

U3. Return to the two-player game tree in part (b) of Exercise U2 in Chapter 3.
- (a) Write the game in strategic form. Find all of the Nash equilibria.
- (b) For the equilibria you found in part (a) that are not subgame-perfect, identify the credibility problems.

U4. Return to the two-player game tree in part (c) of Exercise U2 in Chapter 3.
- (a) Draw the game table. Make Albus the row player, Minerva the column player, and Severus the page player. Find all of the Nash equilibria.
- (b) For the equilibria you found in part (a) that are not subgame-perfect, identify the credibility problems.

U5. Consider the cola industry, in which Coke and Pepsi are the two dominant firms. (To keep the analysis simple, just forget about all the others.) The market size is $8 billion. Each firm can choose whether to advertise. Advertising costs $1 billion for each firm that chooses it. If one firm advertises and the other doesn't, then the former captures the whole market. If both firms advertise, they split the market 50:50 and pay for the advertising. If neither advertises, they split the market 50:50 but without the expense of advertising.
- (a) Write the payoff table for this game, and find the equilibrium when the two firms move simultaneously.
- (b) Write the game tree for this game (assume that it is played sequentially), with Coke moving first and Pepsi following.
- (c) Is either equilibrium in parts (a) and (b) better from the joint perspective of Coke and Pepsi? How could the two firms do better?

U6. Along a stretch of a beach are 500 children in five clusters of 100 each. (Label the clusters A, B, C, D, and E in that order.) Two ice-cream vendors are deciding simultaneously where to locate. They must choose the exact location of one of the clusters.

If there is a vendor in a cluster, all 100 children in that cluster will buy an ice cream. For clusters without a vendor, 50 of the 100 children are willing to walk to a vendor who is one cluster away, only 20 are willing to walk to a vendor two clusters away, and no children are willing to walk the distance of three or more clusters. The ice cream melts quickly, so the walkers cannot buy for the nonwalkers.

If the two vendors choose the same cluster, each will get a 50% share of the total demand for ice cream. If they choose different clusters, then those children (locals or walkers) for whom one vendor is closer than the other will go to the closer one, and those for whom the two are equidistant will split 50% each. Each vendor seeks to maximize her sales.

- (a) Construct the five-by-five payoff table for the vendor location game; the entries stated here will give you a start and a check on your calculations:

  If both vendors choose to locate at A, each sells 85 units.

  If the first vendor chooses B and the second chooses C, the first sells 150 and the second sells 170.

  If the first vendor chooses E and the second chooses B, the first sells 150 and the second sells 200.
- (b) Eliminate dominated strategies as far as possible.
- (c) In the remaining table, locate all pure-strategy Nash equilibria.
- (d) If the game is altered to one with sequential moves, where the first vendor chooses her location first and the second vendor follows, what are the locations and the sales that result from the subgame-perfect equilibrium? How does the change in the timing of moves here help players resolve the coordination problem in part (c)?

U7. Return to the game among the three lions in the Roman Colosseum in Exercise S7 in Chapter 3.
- (a) Write out this game in strategic form. Make Lion 1 the row player, Lion 2 the column player, and Lion 3 the page player.
- (b) Find the Nash equilibria for the game. How many did you find?
- (c) You should have found Nash equilibria that are not subgame-perfect. For each of those equilibria, which lion is making a noncredible threat? Explain.

U8. Now assume that the mall location game (from Exercises S8 in Chapter 3 and S10 in this chapter) is played sequentially but with a different order of play: Big Giant, then Titan, then Frieda's.
- (a) Draw the new game tree.
- (b) What is the subgame-perfect equilibrium of the game? How does it compare to the subgame-perfect equilibrium for Exercise S8 in Chapter 3?
- (c) Now write the strategic form for this new version of the game.
- (d) Find all of the Nash equilibria of the game. How many are there? How does this compare with the number of equilibria from Exercise S10 in this chapter?

U9. Return to the game of Monica and Nancy in Exercise U10 of Chapter 5. Assume that Monica and Nancy choose their effort levels sequentially instead of simultaneously. Monica commits to her choice of effort first. On observing this decision, Nancy commits to her own effort.

(a) What is the subgame-perfect equilibrium to the game where the joint profits are $5m + 4n + mn$, the effort costs to Monica and Nancy are m^2 and n^2 respectively, and Monica commits to an effort level first?

(b) Compare the payoffs of Monica and Nancy with those found in Exercise U10 of Chapter 5. Does this game have a first-mover or second-mover advantage?

(c) Using the same joint profit function as in part (a), find the subgame-perfect equilibrium for the game where *Nancy* must commit first to an effort level.

U10. In an extension of Exercise U9, Monica and Nancy need to decide which (if either) of them will commit to an effort level first. To do this, each of them simultaneously writes on a separate slip of paper whether or not she will commit first. If they both write "yes" or they both write "no," they choose effort levels simultaneously, as in Exercise U10 in Chapter 5. If Monica writes "yes" and Nancy writes "no," they play the game in part (a) of Exercise U9, above. If Monica writes "no" and Nancy writes "yes," they play the game in part (c).

(a) Use the payoffs to Monica and Nancy in parts (b) and (c) in Exercise U9 above as well as those in Exercise U10 in Chapter 5 to construct the game table for the first-stage paper-slip decision game.

(b) Find the pure-strategy Nash equilibria of this first-stage game.

U11. In the faraway town of Saint James two firms, Bilge and Chem, compete in the soft-drink market (Coke and Pepsi aren't in this market yet). They sell identical products, and since their good is a liquid, they can easily choose to produce fractions of units. Since they are the only two firms in this market, the price of the good (in dollars), P, is determined by $P = (30 - Q_B - Q_C)$, where Q_B is the quantity produced by Bilge and Q_C is the quantity produced by Chem (each measured in liters). At this time both firms are considering whether to invest in new bottling equipment that will lower their variable costs.

(i) If firm j decides *not* to invest, its cost will be $C_j = Q_j^2 / 2$, where j stands for either B (Bilge) or C (Chem).

(ii) If a firm decides to invest, its cost will be $C_j = 20 + Q_j^2 / 6$, where j stands for either B (Bilge) or C (Chem). This new cost function reflects the fixed cost of the new machines (20) as well as the lower variable costs.

The two firms make their investment choices simultaneously, but the payoffs in this investment game will depend on the subsequent duopoly games that arise. The game is thus really a two-stage game: decide to invest, and then play a duopoly game.

(a) Suppose both firms decide to invest. Write the profit functions in terms of Q_B and Q_C for the two firms. Use these to find the Nash equilibrium of the quantity-setting game. What are the equilibrium quantities and profits for both firms? What is the market price?

(b) Now suppose both firms decide not to invest. What are the equilibrium quantities and profits for both firms? What is the market price?

(c) Now suppose that Bilge decides to invest, and Chem decides not to invest. What are the equilibrium quantities and profits for both firms? What is the market price?

(d) Write out the two-by-two game table of the investment game between the two firms. Each firm has two strategies: Investment and No Investment. The payoffs are simply the profits found in parts (a), (b), and (c). (Hint: Note the symmetry of the game.)

(e) What is the subgame-perfect equilibrium of the overall two-stage game?

U12. Two French aristocrats, Chevalier Chagrin and Marquis de Renard, fight a duel. Each has a pistol loaded with one bullet. They start 10 steps apart and walk toward each other at the same pace, 1 step at a time. After each step, either may fire his gun. When one shoots, the probability of scoring a hit depends on the distance. After k steps it is $k/5$, and so it rises from 0.2 after the first step to 1 (certainty) after 5 steps, at which point they are right up against one another. If one player fires and misses while the other has yet to fire, the walk must continue even though the bulletless one now faces certain death; this rule is dictated by the code of the aristocracy. Each gets a payoff of 21 if he himself is killed and 1 if the other is killed. If neither or both are killed, each gets 0.

This is a game with five sequential steps and simultaneous moves (shoot or not shoot) at each step. Find the rollback (subgame-perfect) equilibrium of this game.

Hint: Begin at step 5, when the duelists are right up against one another. Set up the two-by-two table for the simultaneous-move game at this step, and find its Nash equilibrium. Now move back to step 4, where the probability of scoring a hit is 4/5, or 0.8, for each. Set up the two-by-two table for the simultaneous-move game at this step, correctly specifying in the appropriate cell what happens in the future. For example, if one shoots and misses, but the other does not shoot, then the other will wait until step 5 and score a sure hit. If neither shoots, then the game will go to the next step, for which you have already found the equilibrium. Using all this information, find the payoffs in the two-by-two table of step 4, and find the Nash equilibrium at this step. Work backward in the same way through the rest of the steps to find the Nash equilibrium strategies of the full game.

U13. Describe an example of business competition that is similar in structure to the duel in Exercise U12.

7

Simultaneous-Move Games with Mixed Strategies I: Two-by-Two Games

In our study of simultaneous-move games in Chapter 4, we came across a class of games that the solution methods described there could not solve; in fact, games in that class have no Nash equilibria in pure strategies. To predict outcomes for such games, we need an extension of our concepts of strategies and equilibria. This is to be found in the randomization of moves, which is the focus of this chapter and the next.

Consider the tennis-point game from the end of Chapter 4. This game is zero sum; the interests of the two tennis players are purely in mutual conflict. Evert wants to hit her passing shot to whichever side—down the line (DL) or crosscourt (CC)—is not covered by Navratilova, whereas Navratilova wants to cover the side to which Evert hits her shot. In Chapter 4, we pointed out that in such a situation, any systematic choice by Evert will be exploited by Navratilova to her own advantage and therefore to Evert's disadvantage. Conversely, Evert can exploit any systematic choice by Navratilova. To avoid being thus exploited, each player wants to keep the other guessing, which can be done by acting unsystematically or randomly.

However, randomness doesn't mean choosing each shot half the time, or alternating between the two. The latter would itself be a systematic action open to exploitation, and a 60-40 or 75-25 random mix may be better than 50-50 depending on the situation. In this chapter we develop methods for calculating the best mix and discuss how well this theory helps us understand actual play in such games.