5

Simultaneous-Move Games: Continuous Strategies, 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.

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 the 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, which has two restaurants, Xavier's Tapas Bar and Yvonne's Bistro. To keep the story simple, we assume 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 prices to maximize profit, the payoff in this game. We suppose that they must get their menus printed separately without knowing the other's prices, 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 the game and to determine equilibrium prices. Let us call Xavier's price P_{x} and Yvonne's price P_{y}.

[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 10.

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:

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

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:

$$\begin{aligned}\Pi_x &= -8(44 + P_y) + (16 + 44 + P_y)P_x - 2(P_x)^2\\ &= -8(44 + P_y) + (60 + P_y)P_x - 2(P_x)^2.\end{aligned}$$

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 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.

[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).

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]

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

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.

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

[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.

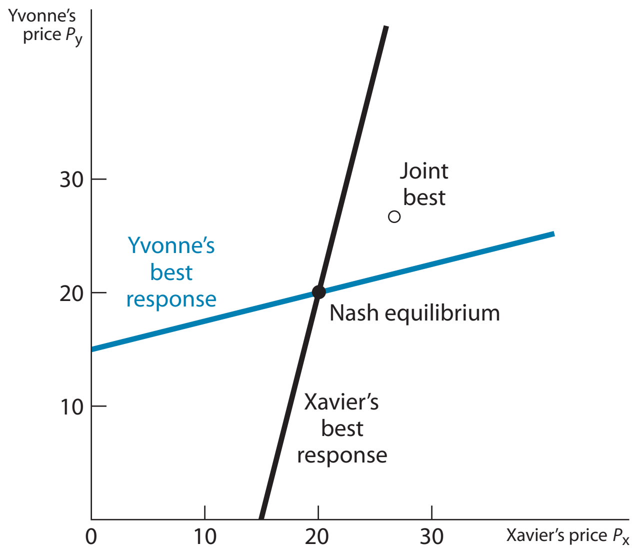

FIGURE 5.1 Best-Response Curves and Equilibrium in the Restaurant Pricing Game

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

[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.

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 firm's 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 profit from these new customers by raising its price part of the way. Thus, a restaurant that raises its price is also helping to increase its rival'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 of their profits? Yes. Suppose the two restaurants charged \$24 each. Then each would make a profit of \$16 on each of the 2,000 customers [$2{,}000 = (44 - 2 \times 24 + 24)$ hundred] that it would serve each month, for a total profit of \$32,000 per month.

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 $15 + 0.25 \times 24 = 21$. Then Yvonne will come back with her best response to that: $15 + 0.25 \times 21 = 20.25$. 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 10.[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 would accrue to each of them if they both lowered their prices. 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 ads that emphasize the bad things about the opponent. To keep matters simple, suppose the voters start out entirely ignorant and unconcerned and form opinions solely as a result of 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

[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).

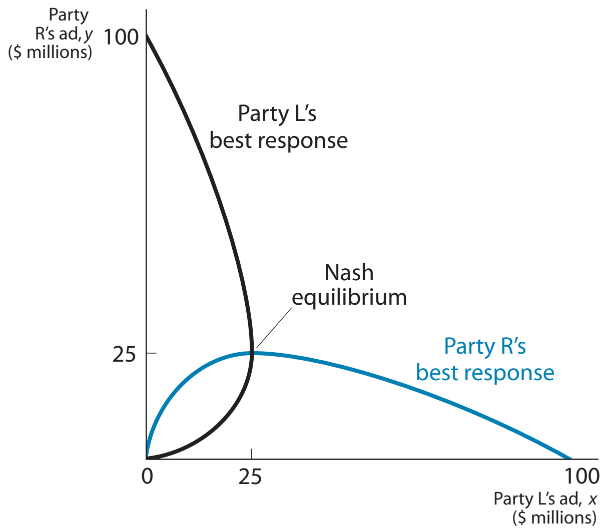

FIGURE 5.2 Best Responses and Nash Equilibrium in the Campaign Advertising Game

voters in detail in Chapter 15.) Even more simply, suppose the vote share of a party equals its share of the total campaign advertising that is done. Call the 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 in words its general meaning intuitively. 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 0. 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 party's expenditures by spending more oneself to compete harder. But if the other party already spends a great deal on ads, then one's own ads get only a small return in relation to their cost, so it is better to respond to the other party's increase in spending 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 (which keeps prices high and hurts consumers), a politicians' cartel to advertise less would probably benefit voters and society, like a producers' cartel for complements would lead to lower prices and benefit consumers. 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 vote share 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 ends up advertising 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 that 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

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

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 0. 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 can 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 CRITICAL DISCUSSION OF THE NASH EQUILIBRIUM CONCEPT

Although Nash equilibrium is the primary solution concept for simultaneous games, it has been subject to several theoretical criticisms. In this section, we briefly review some of these criticisms and some rebuttals, in each case by using an example.[6] Some of the 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 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 led Nobel laureate Roger Myerson to rebut those criticisms of the Nash equilibrium that were based on the intuitive appeal of playing a different strategy. His rebuttal simply shifted the burden of proof onto the critic. "When asked why players in a game should behave as in some Nash equilibrium," he said, "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."[7]

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

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

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 might therefore choose those strategies. We offer two examples of this kind. The first comes from John Morgan, an economics professor at the University of California, Berkeley; Figure 5.3 shows the game table.

Best-response analysis quickly reveals that this game has a unique Nash equilibrium—namely, (A, A), yielding the payoffs (2, 2). But you may think, as 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 if 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 a 2 by playing C.

Myerson would respond, "Not so fast. If you really believe that the other player would think this way and play C, then you should play B to get the payoff 3. And if you think the other person would think this way and play B, then your best response to B should be A. And if you think the other person 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 comes from David Kreps, an economist at Stanford Business School, and 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 the 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

		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

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

FIGURE 5.4 Disastrous Nash Equilibrium?

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 payoff of 9.9 earned from 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, but not all, people assigned to be Player A would not 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 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 payoffs were as illustrated in Figure 5.4 but in reality B's payoffs were the reverse—the 9.9 payoff went with Left and the 10 payoff went with Right? What if the 9.9 payoff were only an approximation and the exact payoff was actually 10.1? What if B was a player with a substantially different value system or was not a truly rational player and 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 about 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 illustrate 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 (which we won't have the tools to discuss until Chapter 8). 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 what B's actual payoffs are, she must choose her strategy to be "best on 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 $10(1 - p) - 1{,}000p$. Therefore, it is better for A to choose Up if

$$9(1 - p) + 8p > 10(1 - p) - 1{,}000p, \quad \text{or} \quad p > 1/1{,}009.$$

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 8, we show how the concept can be used for understanding how people respond to the presence of risk in their lives—for example, by sharing the risk with others or by buying 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. Consider the following coordination game, played by students at Stanford University. One player was assigned the city of Boston and the other was assigned San Francisco. Each was then 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 those cities. The two chose simultaneously and independently. If and only if their choices divided up the nine cities completely and without any overlap between them, both got a prize. Despite the existence of 512 different Nash equilibria, when both players were Americans or long-time U.S. residents, more than 80% of the time they chose a unique equilibrium based on geography. The student assigned Boston chose all the cities east of the Mississippi, and the student assigned San Francisco chose all the cities 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 with much less coordination on the same dividing point.[8]

The features of the game itself, combined with shared cultural background, can help player expectations to converge. As another example of multiplicity of equilibria, consider a game where two players write down, simultaneously and independently, the share that each wants from a total prize of \$100. If the amounts that they write down add up to \$100 or less, each player receives 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

[8] See David Kreps, *A Course in Microeconomic Theory* (Princeton: Princeton University Press, 1990), pp. 392–93, 414–15.

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 is 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 will 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 8, we will 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 Chapter 4 in Section 4.7 and said that, by extending the concept of strategy to random mixtures, Nash equilibrium could be restored. In Chapter 7, we will 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?

		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

To begin to address this question, we consider the three-by-three game in Figure 5.5. Best-response analysis 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 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.

3 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 that we used in Section 2.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 3.A below. 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 3.B below 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.[9] 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?

[9] This example comes from Douglas Bernheim, "Rationalizable Strategic Behavior," *Econometrica*, vol. 52, no. 4 (July 1984), pp. 1007–1028, an article that originally developed the concept of rationalizability. See also Andreu Mas-Colell, Michael Whinston, and Jerry Green, *Microeconomic Theory* (New York: Oxford University Press, 1995), pp. 242–45.

		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

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.[10] 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 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.[11]

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.

[10] 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.

[11] When one allows for mixed strategies, as we will do in Chapter 7, there arises the possibility of a pure strategy being dominated by a mixture of other pure strategies. With such an expanded definition of a dominated strategy, iterated elimination of strictly dominated strategies turns out to be equivalent to rationalizability. The details are best left for a more advanced course in game theory.

Thus, R4 and C4 can never figure in the set of rationalizable strategies. 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 2.C above. 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 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 to this chapter 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 0, which gives

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

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

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

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

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

[12] 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.

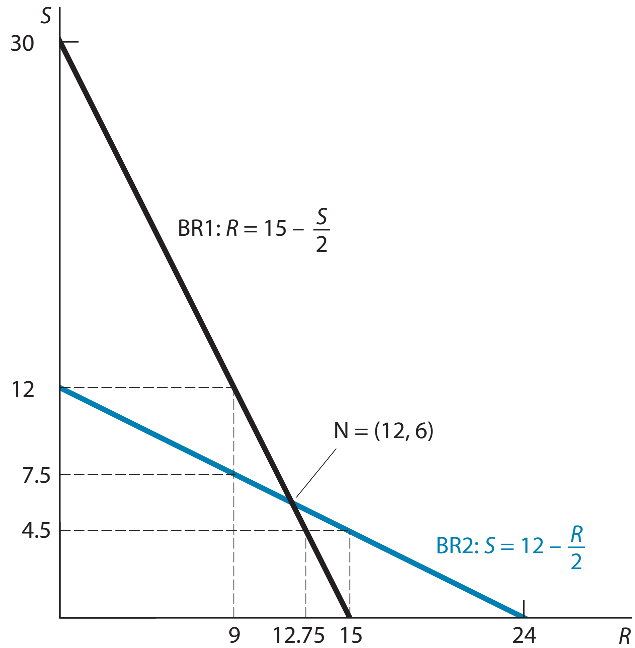

FIGURE 5.7 Nash Equilibrium through Rationalizability

$S = 0$ 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.[13] 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;

[13] 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.

it will work for any game that has a unique Nash equilibrium at the intersection of downward-sloping best-response curves.[14]

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.

4 EMPIRICAL EVIDENCE CONCERNING NASH EQUILIBRIUM

In Chapter 3, when we considered empirical evidence on sequential-move games and rollback, we presented empirical evidence from observations on games actually played in the world, as well as games deliberately constructed for testing the theory in the laboratory. There we pointed out the different merits and drawbacks of the two methods for assessing the validity of rollback equilibrium predictions. Similar issues arise in securing and interpreting the evidence on Nash equilibrium play in simultaneous-move games.

Real-world games are played for substantial stakes, often by experienced players who have the knowledge and the incentives to employ good strategies. But these situations include many factors beyond those considered in the theory. In particular, in real-life situations, it is difficult to observe the quantitative

[14] 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.

payoffs that players would have earned for all possible combinations of strategies. Therefore, if their behavior does not bear out the predictions of the theory, we cannot tell whether the theory is wrong or whether some other factors overwhelm the strategic considerations.

Laboratory experiments attempt to control for other factors in an attempt to provide cleaner tests of the theory. But they bring in inexperienced players and provide them with little time and relatively weak incentives to learn the game and play it well. Confronted with a new game, most of us would initially 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 player would learn to play. Experiments often 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 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, such research asks, "Do participants play their Nash equilibrium strategies?" Reviewing this work, Douglas Davis and Charles Holt conclude that, in relatively simple single-move games with a unique Nash equilibrium, the equilibrium "has considerable drawing power . . . after some repetitions with different partners."[15] But the theory's success is more mixed in more complex situations, such as when multiple Nash equilibria exist, when emotional factors modify payoffs beyond the stated cash amounts, when the calculations for finding a Nash equilibrium are more complex, or when the game is repeated with fixed partners. We will briefly consider the performance of Nash equilibrium in several of these circumstances.

I. CHOOSING AMONG MULTIPLE EQUILIBRIA In Section 2.B above, we presented examples demonstrating that focal points sometimes emerge to help players choose among multiple Nash equilibria. Players may not manage to coordinate 100% of the time, but circumstances often enable players to achieve much more coordination than would be experienced by random choices across possible equilibrium strategies. Here we present a coordination game designed with an interesting trade-off: the equilibrium with the highest payoff to all players also happens to be the riskiest one to play, in the sense of Section 2.A above.

John Van Huyck, Raymond Battalio, and Richard Beil describe a 16-player game in which each player simultaneously chooses an "effort" level between 1

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

and 7. Individual payoffs depend on group "output," a function of the minimum effort level chosen by anyone in the group, minus the cost of one's individual effort. The game has exactly seven Nash equilibria in pure strategies; any outcome in which all players choose the same effort level is an equilibrium. The highest possible payoff ($1.30 per player) occurs when all subjects choose an effort level of 7, while the lowest equilibrium payoff ($0.70 per player) occurs when all subjects choose an effort level of 1. The highest-payoff equilibrium is a natural candidate for a focal point, but in this case there is a risk to choosing the highest effort; if just one other player chooses a lower effort level than you, then your extra effort is wasted. For example, if you play 7 and at least one other person chooses 1, you get a payoff of just $0.10, far worse than the worst equilibrium payoff of $0.70. This makes players nervous about whether others will choose maximum effort, and as a result, large groups typically fail to coordinate on the best equilibrium. A few players inevitably choose lower than the maximum effort, and in repeated rounds play converges toward the lowest-effort equilibrium.[16]

II. EMOTIONS AND SOCIAL NORMS In Chapter 3, we saw several examples in sequential-move games where players were more generous to each other than Nash equilibrium would predict. Similar observations occur in simultaneous-move games such as the prisoners' dilemma game. One reason may be that the players' payoffs are different from those assumed by the experimenter: in addition to cash, their payoffs may also include the experience of emotions such as empathy, anger, or guilt. In other words, 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.[17] Seen through this lens, these observations do not show any deficiency of the Nash equilibrium concept itself, but they do warn us against using the concept under naive or mistaken assumptions about people's payoffs.

[16] See John B. Van Huyck, Raymond C. Battalio, and Richard O. Beil, "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure," *American Economic Review*, vol. 80, no. 1 (March 1990), pp. 234–48. Subsequent research has suggested methods that can promote coordination on the best equilibrium. Subhasish Dugar, "Non-monetary Sanction and Behavior in an Experimental Coordination Game," *Journal of Economic Behavior & Organization*, vol. 73, no. 3 (March 2010), pp. 377–86, shows that players gradually manage to coordinate on the highest-payoff outcome merely by allowing players, between rounds, to express the numeric strength of their disapproval for each other player's decision. Roberto A. Weber, "Managing Growth to Achieve Efficient Coordination in Large Groups," *American Economic Review*, vol. 96, no. 1 (March 2006), pp. 114–26, shows that starting with a small group and slowly adding additional players can sustain the highest-payoff equilibrium, suggesting that a firm may do well to expand slowly and make sure that employees understand the corporate culture of cooperation.

[17] 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, UK: Palgrave MacMillan, 2004), pp. 85–104.

It might be a mistake, for example, to assume that players are always driven by the selfish pursuit of money.

III. COGNITIVE ERRORS As we saw in the experimental evidence on rollback equilibrium in Chapter 3, players do not always fully think through the entire game before playing, nor do they always expect other players to do so. Behavior in a game known as the travelers' dilemma illustrates a similar limitation of Nash equilibrium in simultaneous-move games. In this game, two travelers purchase identical souvenirs while on vacation, and the airline loses both of their bags on the return trip. The airline announces to the two players that it intends to reimburse them for their losses, but it does not know the exact amount to reimburse. It knows the correct amount is between $80 and $200 per person, so it designs a game as follows. Each player may submit a claim between $80 and $200. The airline will reimburse both players at an amount equal to the lower of the two claims submitted. In addition, if the two claims differ, the airline will pay a reward of $5 to the person making the smaller claim and deduct a penalty of $5 from the reimbursement of the person making the larger claim.

With these rules, irrespective of the actual value of the lost luggage, each player has an incentive to undercut the other's claim. In fact, it turns out that the only Nash equilibrium, and indeed the only rationalizable outcome, is for both players to report the minimum number of $80. However, in the laboratory, players rarely claim $80; instead they claim amounts much closer to $200. (Real payoff amounts in the laboratory are typically in cents rather than in dollars.) Interestingly, if the penalty/reward parameter is increased by a factor of 10, from $5 to $50, behavior conforms much more closely to the Nash equilibrium, with reported amounts generally near $80. Thus, behavior in this experiment varies tremendously with a parameter that does not affect the Nash equilibrium at all; the unique equilibrium is $80, regardless of the size of the penalty/reward amount.

To explain these results from their laboratory, Monica Capra and her coauthors employed a theoretical model called **quantal-response equilibrium (QRE)**, originally proposed by Richard McKelvey and Thomas Palfrey. This model's mathematics are beyond the scope of this text, but its main contribution is that it allows for the possibility that players make errors, with the probability of a given error being much smaller for costly mistakes than for mistakes that reduce one's payoff by very little. Furthermore, the model incorporates players who expect each other to make errors in this way. It turns out that quantal-response analysis can explain the data quite well. Reporting a high claim is not very costly when the penalty is only $5, so players are more willing to report values near $200—especially knowing that their rivals are likely to behave similarly, so the payoff to reporting a high number can be quite large. However, with a penalty/reward of $50 instead of $5, reporting a high claim becomes quite costly, so players are very unlikely to expect each other to make such a mistake. This expectation pushes behavior toward the Nash equilibrium claim of \$80. Building on this success, quantal-response equilibrium has become a very active area of game-theoretic research.[18]

IV. COMMON KNOWLEDGE OF RATIONALITY We just saw that to better explain experimental results, QRE allows for the possibility that players may not believe that others are perfectly rational. Another way to explain data from experiments is to allow for the possibility that different players engage in different levels of reasoning. A strategic guessing game that is 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 two-thirds—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 67, so for each player, any choice above 67 is dominated by 67.[19] But all players should rationally figure this out, so the average can never exceed 67 and two-thirds of it can never exceed 44, and so any choice above 44 is dominated by 44. The iterated elimination of dominated strategies goes on until only 0 is left.

However, when a group actually plays this game for the first time, the winner is not a person who plays 0. Typically, the winning number is somewhere around 15 or 20. The most commonly observed choices are 33 and 22, suggesting that a large number of players perform one or two rounds of iterated dominance without going further. That is, "level-1" players imagine that all other players will choose randomly, with an average of 50, so they best-respond with a choice of two-thirds of this amount, or 33. Similarly, "level-2" players imagine that everyone else will be a "level-1" player, and so they best-respond by playing two-thirds of 33, or 22. Note that all of these choices are far from the Nash equilibrium of 0. It appears that many players follow a limited number of steps of

[18] See Kaushik Basu, "The Traveler's Dilemma," *Scientific American*, vol. 296, no. 6 (June 2007), pp. 90–95. The experiments and modeling can be found in C. Monica Capra, Jacob K. Goeree, Rosario Gomez, and Charles A. Holt, "Anomalous Behavior in a Traveler's Dilemma?" *American Economic Review*, vol. 89, no. 3 (June 1999), pp. 678–90. Quantal-response equilibrium (QRE) was first proposed by Richard D. McKelvey and Thomas R. Palfrey, "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, vol. 10, no. 1 (July 1995), pp. 6–38.

[19] 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 two-thirds of this, so $x = (2/3)[x + (N - 1)100]/N$, or $x = 100(2N - 2)/(3N - 2)$. If $N = 10$, then $x = (18/28)/100 = 64$ (approximately). So any choice above 64 is dominated by 64. The same reasoning applies to the successive rounds.

iterated elimination of dominated strategies, in some cases because they expect others to be limited in their number of rounds of thinking.[20]

V. LEARNING AND MOVING TOWARD EQUILIBRIUM What happens when the strategic guessing game is repeated with the same group of players? In classroom experiments, we find that the winning number can easily drop 50% in each subsequent round, as the students expect all their classmates to play numbers as low as the previous round's winning number or lower. By the third round of play, winning numbers tend to be as low as 5 or less.

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 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 tend to agree with this latter viewpoint.

Interestingly, we have found that people learn somewhat faster by observing others play a game than by playing 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.

[20] You will analyze similar games in Exercises S12 and U11. For a summary of results from large-scale experiments run in European newspapers with thousands of players, see Rosemarie Nagel, Antoni Bosch-Domènech, Albert Satorra, and Juan Garcia-Montalvo, "One, Two, (Three), Infinity: Newspaper and Lab Beauty-Contest Experiments," *American Economic Review*, vol. 92, no. 5 (December 2002), pp. 1687–1701.

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

While the field does not allow for as much direct observation as the laboratory does, observations outside the laboratory can also provide valuable evidence about the relevance of Nash equilibrium. Conversely, Nash equilibrium often provides a valuable starting point for social scientists to make sense of the real world.

I. APPLICATIONS OF NASH EQUILIBRIUM One of the earliest applications of the Nash equilibrium concept to real-world behavior 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 negotiations 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.[21]

As we saw earlier in this chapter, price competition is one important application of Nash equilibrium. Other strategic choices by firms include product quality, investment, R&D, and so on. The theory has also 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. Game-theoretic models, based on the Nash

[21] For those who would like to see more applications, here are some suggested sources. Thomas Schelling's *The Strategy of Conflict* (Cambridge, Mass.: Harvard University Press, 1960) and *Arms and Influence* (New Haven: Yale University Press, 1966) are still required for all students of game theory. The classic textbook on game-theoretic treatment of industries is Jean Tirole, *The Theory of Industrial Organization* (Cambridge, Mass.: MIT Press, 1988). In political science, an early classic is William H. Riker, *Liberalism Against Populism* (San Francisco: W. H. Freeman, 1982). For advanced-level surveys of research, see several articles in *The Handbook of Game Theory with Economic Applications*, ed. Robert J. Aumann and Sergiu Hart (Amsterdam: North-Holland/Elsevier Science B. V., 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," both of which are in volume 3.

equilibrium concept and its dynamic generalizations, fit the data for many major industries, such as automobile manufacturers, reasonably well. They also give us a better understanding of the determinants of competition than the older models, which assumed perfect competition and estimated supply and demand curves.[22]

Pankaj Ghemawat, a professor at the IESE Business School in Barcelona, has developed a number of case studies of individual firms or industries, supported by statistical analysis of the data. His game-theoretic models are remarkably successful in improving our understanding of several initially puzzling observed business decisions on pricing, capacity, innovation, and so on. For example, DuPont constructed an enormous amount of manufacturing capacity for titanium dioxide in the 1970s. It added capacity in excess of the projected growth in worldwide demand over the next decade. At first glance, this choice looked like a terrible strategy because the excess capacity could lead to lower market prices for this commodity. However, DuPont successfully foresaw that, by having excess capacity in reserve, it could punish competitors who cut prices by increasing its production and driving prices even lower in the future. This ability made it a price leader in the industry, and it enjoyed high profit margins. The strategy worked quite well, with DuPont continuing to be a worldwide leader in titanium dioxide 40 years later.[23]

More recently, game theory has become the tool of choice for the study of political systems and institutions. As we shall see in Chapter 15, game theory has shown how voting and agenda setting in committees and elections can be strategically manipulated in pursuit of one's ultimate objectives. Part Four of this book will develop other applications of Nash equilibrium in auctions, voting, and bargaining. We also develop our own case study of the Cuban missile crisis in Chapter 14.

Some critics remain unpersuaded of the value of Nash equilibrium, claiming that the same understanding of these phenomena can be obtained using previously known general principles of economics, political science, and so on. In one sense they are right. A few of these analyses existed before Nash equilibrium 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 100 years. One can think of Nash equilibrium as just a general formulation of that equilibrium concept for all games. Some theories

[22] For simultaneous-move models of price competition, see Timothy F. Bresnahan, "Empirical Studies of Industries with Market Power," in *Handbook of Industrial Organization*, vol. 2, ed. Richard L. Schmalensee and Robert D. Willig (Amsterdam: North-Holland/Elsevier, 1989), pp. 1011–57. For models of entry, see Steven Berry and Peter Reiss, "Empirical Models of Entry and Market Structure," in *Handbook of Industrial Organization*, vol. 3, ed. Mark Armstrong and Robert Porter (Amsterdam: North-Holland/Elsevier, 2007), pp. 1845–86.

[23] Pankaj Ghemawat, "Capacity Expansion in the Titanium Dioxide Industry," *Journal of Industrial Economics*, vol. 33, no. 2 (December 1894), pp. 145–63. For more examples, see Pankaj Ghemawat, *Games Businesses Play: Cases and Models* (Cambridge, Mass.: MIT Press, 1997).

of strategic voting date to the eighteenth century, and some notions of credibility can be found in history as far back as Thucydides' *Peloponnesian War*. However, what Nash equilibrium does is to unify all these applications and thereby facilitate the development of new ones.

Furthermore, the development of game theory has also led directly to a wealth of new ideas and applications that did not exist before—for example, how the existence of a second-strike capability reduces the fear of surprise attack, how different auction rules affect bidding behavior and seller revenues, how governments can successfully manipulate fiscal and monetary policies to achieve reelection even when voters are sophisticated and aware of such attempts, and so on. If these examples had all been amenable to previously known approaches, they would have been discovered long ago.

II. REAL-WORLD EXAMPLES OF LEARNING We conclude by offering an interesting illustration of equilibrium and the learning process in the real-world game of major-league baseball. In this game, the stakes are high and players play more than 100 games per year, creating strong motivation and good opportunities to learn. Stephen Jay Gould discovered this beautiful example.[24] The best batting averages recorded in a baseball season declined over most of the twentieth century. 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 have kept 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]

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

[24] Stephen Jay Gould, "Losing the Edge," in *The Flamingo's Smile: Reflections in Natural History* (New York: W. W. Norton & Company, 1985), pp. 215–29.

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.

Michael Lewis's 2003 book *Moneyball* (later made into a movie starring Brad Pitt) describes a related example of movement toward equilibrium in baseball. Instead of focusing on the strategies of individual players, it focuses on the teams' back-office strategies of which players to hire. The book documents Oakland A's general manager Billy Beane's decision to use "sabermetrics" in hiring decisions—that is, paying close attention to baseball statistics based on the theory of maximizing runs scored and minimizing runs given up to opponents. These decisions involved paying more attention to attributes undervalued by the market, such as a player's documented ability to earn walks. Such decisions arguably led to the A's becoming a very strong team, going to the playoffs in five out of seven seasons, despite having less than half the payroll of larger-market teams such as the New York Yankees. Beane's innovative payroll strategies have subsequently been adopted by other teams, such as the Boston Red Sox, who, under general manager Theo Epstein, managed to break the "curse of the Bambino" in 2004 and win their first World Series in 86 years. Over the course of a decade, nearly a dozen teams decided to hire full-time sabermetricians, with Beane noting in September 2011 that he was once again "fighting uphill" against larger teams that had learned to best-respond to his strategies. Real-world games often involve innovation followed by gradual convergence to equilibrium; the two examples from baseball both give evidence of moving toward equilibrium, although full convergence may sometimes take years or even decades to complete.[25]

We take up additional evidence about other game-theoretic predictions 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, especially as a first approach. On the whole, we believe you should have considerable confidence in using the Nash equilibrium

[25] Susan Slusser, "Michael Lewis on A's 'Moneyball' Legacy," *San Francisco Chronicle*, September 18, 2011, p. B-1. The original book is Michael Lewis, *Moneyball: The Art of Winning an Unfair Game* (New York: W. W. Norton & Company, 2003).

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.[26] Often the reason will be your misunderstanding of the players' objectives, not the players' failure to play the game correctly giving their true objectives.

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.*

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.

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

[26] In an article probing the weaknesses of Nash equilibrium in experimental data and proposing QRE-style alternative models for dealing with them, two prominent researchers write, "we will be the first to admit that we begin the analysis of a new strategic problem by considering the equilibria derived from standard game theory before considering" other possibilities. Jacob K. Goeree and Charles A. Holt, "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions," *American Economic Review*, vol. 91, no. 5 (December 2001), pp. 1402–22.

of game theory have helped economists and political scientists, in particular, to understand important consumer, firm, voter, legislative, and government behaviors.

KEY TERMS

best-response curve (137)
best-response rule (134)
continuous strategy (133)
never a best response (150)
quantal-response equilibrium (QRE) (158)
rationalizability (150)
rationalizable (150)
refinement (148)

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 an advertising 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 $x = 10\sqrt{y} - y$ 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 restaurant pricing game illustrated in Figure 5.1 defines customer 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 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 3.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 (each cart incurs 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 $Q = q_{\text{Korea}} + q_{\text{Japan}}$. (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, $c_{\text{Korea}} = c_{\text{Japan}} = 30$, 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 $c_{\text{Korea}} = 20$ and $c_{\text{Japan}} = 40$, what is the market share of each country (that is, the percentage of ships that each country sells relative to the total number of ships 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 $Q = q_{\text{Korea}} + q_{\text{Japan}} + q_{\text{China}}$. Assume that all three countries have a per-ship cost of \$30 million: $c_{\text{Korea}} = c_{\text{Japan}} = c_{\text{China}} = 30$.

(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 49 classmates. Each of them simultaneously and privately writes down a number between 0 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 one-half 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 if 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: $P = 120 - Q_{\text{DTC}}$. 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 $P = 120 - Q_{\text{DTC}} - Q_{\text{ADA}}$. 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, which shows first-run movies, and Sticky Shoe, which shows movies that have been out for a while at a cheaper price. The demand for movies at Modern Multiplex is given by $Q_{\text{MM}} = 14 - P_{\text{MM}} + P_{\text{SS}}$, while the demand for movies at Sticky Shoe is $Q_{\text{SS}} = 8 - 2P_{\text{SS}} + P_{\text{MM}}$, 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 in 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 S12 in Chapter 4?

U7. In the fishing-boat game of Section 3.B, 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 $Q = q_{\text{Intel}} + q_{\text{AMD}}$ 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 ($c_{\text{Korea}} = 20$, $c_{\text{Japan}} = 40$, and $c_{\text{China}} = 60$).

After it gains experience and adds production capacity, China's per-ship cost will decrease dramatically. Because labor is even cheaper in China than in 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 ($c_{\text{Korea}} = 20$, $c_{\text{Japan}} = 40$, and $c_{\text{China}} = 16$).

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 49 classmates (similar to the situation in Exercise S12). As before, each of the students simultaneously and privately writes down a number between 0 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.C 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 a share $y/(x+y)$. 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

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

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 (that is, 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 (that is, 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.[27]

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:

$$\begin{aligned} k &= [A + B(X + h) - C(X + h)^2] - (A + BX - CX^2) \\ &= Bh - C[(X + h)^2 - X^2] \\ &= Bh - C(X^2 + 2Xh + h^2 - X^2) \\ &= (B - 2CX)h - Ch^2. \end{aligned}$$

[27] 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.

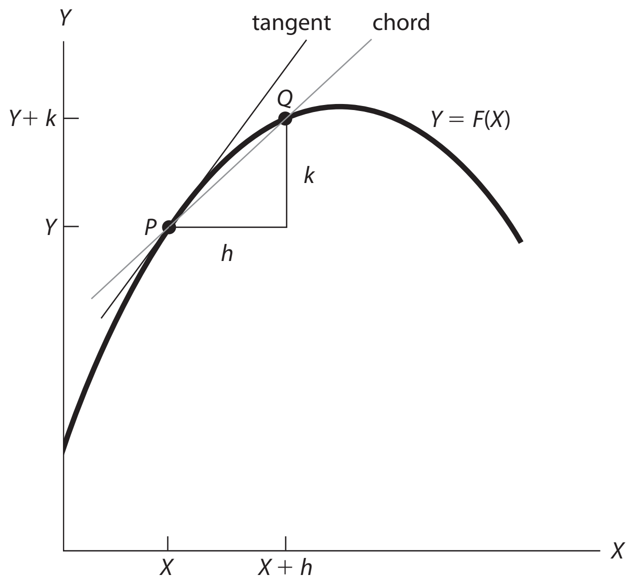

FIGURE 5A.1 Derivative of a Function Illustrated

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 0. 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 0, 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, should be negative at a maximum. Similar logic shows that the second

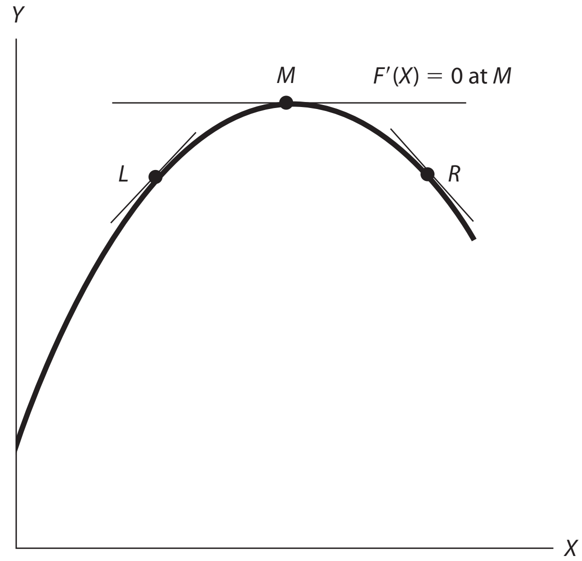

FIGURE 5A.2 Optimum of a Function

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 0. Setting it equal to 0 and solving for P_x gives the same equation as derived in Section 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 11 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

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

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.