17

Bargaining

PEOPLE ENGAGE IN BARGAINING throughout their lives. Children start by negotiating to share toys and to play games with other children. Couples bargain about matters of housing, child rearing, and the adjustments that each must make for the other's career. Buyers and sellers bargain over price, workers and bosses over wages. Countries bargain over policies of mutual trade liberalization; superpowers negotiate mutual arms reduction. And the two original authors of this book had to bargain with one another—generally very amicably—about what to include or exclude, how to structure the exposition, and so forth. To get a good result from such bargaining, the participants must devise good strategies. In this chapter, we raise and explicate some of these basic ideas and strategies.

All bargaining situations have two things in common. First, the total payoff that the parties to the negotiation are capable of creating and enjoying as a result of reaching an agreement should be greater than the sum of the individual payoffs that they could achieve separately—the whole must be greater than the sum of the parts. Without the possibility of this excess value, or "surplus," the negotiation would be pointless. If two children considering whether to play together cannot see a net gain from having access to a larger total stock of toys or from one another's company in play, then it is better for each to "take his toys and play by himself." The world is full of uncertainty, and the expected benefits may not materialize. But when engaged in bargaining, the parties must at least perceive some gain therefrom: when he agreed to sell his soul to the Devil, Faust thought the benefits of knowledge and power that he gained were worth the price that he would eventually have to pay.

The second important general point about bargaining follows from the first: it is not a zero-sum game. When a surplus exists, the negotiation is about how to divide it up. Each bargainer tries to get more for himself and leave less for the others. This may appear to be zero sum, but behind it lies the danger that if the agreement is not reached, no one will get any surplus at all. This mutually harmful alternative, as well as *both* parties' desire to avoid it, is what creates the potential for the threats—explicit and implicit—that make bargaining such a strategic matter.

Before the advent of game theory, one-on-one bargaining was generally thought to be a difficult and even indeterminate problem. Observation of widely different outcomes in otherwise similar-looking situations lent support to this view. Theorists were not able to achieve any systematic understanding of why one party gets more than another and attributed this result to vague and inexplicable differences in "bargaining power."

Even the simple theory of Nash equilibrium does not take us any further. Suppose two people are to split $1. Let us construct a game in which each is asked to announce what he would want. The moves are simultaneous. If their announcements x and y add up to 1 or less, each gets what he announced. If they add up to more than 1, neither gets anything. Then *any* pair (x, y) adding to 1 constitutes a Nash equilibrium in this game: *given* the announcement of the other, each player cannot do better than to stick to his own announcement.[1]

Further advances in game theory have brought progress along two quite different lines, each using a distinct mode of game-theoretic reasoning. In Chapter 2, we distinguished between cooperative-game theory, in which the players decide and implement their actions jointly, and noncooperative-game theory, in which the players decide and take their actions separately. Each of the two lines of advance in bargaining theory uses one of these two approaches. One approach views bargaining as a *cooperative* game, in which the parties find and implement a solution jointly, perhaps using a neutral third party such as an arbitrator for enforcement. The other approach views bargaining as a *noncooperative* game, in which the parties choose strategies separately and we look for an equilibrium. However, unlike our earlier simple game of simultaneous announcements, whose equilibrium was indeterminate, here we impose more structure and specify a sequential-move game

[1] As we saw in Chapter 5, Section 3.B, this type of game can be used as an example to bolster the critique that the Nash-equilibrium concept is too imprecise. In the bargaining context, we might say that the multiplicity of equilibria is just a formal way of showing the indeterminacy that previous analysts had claimed.

of offers and counteroffers, which leads to a determinate equilibrium. As in Chapter 2, we emphasize that the labels "cooperative" and "noncooperative" refer to joint versus separate actions, not to nice versus nasty behavior or to compromise versus breakdown. The equilibria of noncooperative bargaining games can entail a lot of compromise.

1 NASH'S COOPERATIVE SOLUTION

In this section, we present Nash's cooperative-game approach to bargaining. First we present the idea in a simple numerical example; then we develop the more general algebra.[2]

A. Numerical Example

Imagine two Silicon Valley entrepreneurs, Andy and Bill. Andy produces a microchip set that he can sell to any computer manufacturer for $900. Bill has a software package that can retail for $100. The two meet and realize that their products are ideally suited to each other and that, with a bit of trivial tinkering, they can produce a combined system of hardware and software worth $3,000 in each computer. Thus, together they can produce an extra value of $2,000 per unit, and they expect to sell millions of these units each year. The only obstacle that remains on this path to fortune is to agree to a division of the spoils. Of the $3,000 revenue from each unit, how much should go to Andy and how much to Bill?

Bill's starting position is that without his software, Andy's chip set is just so much metal and sand, so Andy should get only the $900 and Bill himself should get $2,100. Andy counters that without his hardware, Bill's programs are just symbols on paper or magnetic signals on a diskette, so Bill should get only $100, and $2,900 should go to him, Andy.

Watching them argue, you might suggest they "split the difference." But that is not an unambiguous recipe for agreement. Bill might offer to split the profit on each unit equally with Andy. Under this scheme, each will get a profit of $1,000, meaning that $1,100 of the revenue goes to Bill and $1,900 to Andy. Andy's response might be that they should have an equal percentage of profit on their contribution to the joint enterprise. Thus, Andy should get $2,700 and Bill $300.

The final agreement depends on their stubbornness or patience if they negotiate directly with one another. If they try to have the dispute arbitrated by a third party, the arbitrator's decision depends on his sense of the relative value

[2] John F. Nash Jr., "The Bargaining Problem," *Econometrica*, vol. 18, no. 2 (1950), pp. 155–62.

of hardware and software and on the rhetorical skills of the two principals as they present their arguments before the arbitrator. For the sake of definiteness, suppose the arbitrator decides that the division of the profit should be 4:1 in favor of Andy; that is, Andy should get four-fifths of the surplus while Bill gets one-fifth, or Andy should get four times as much as Bill. What is the actual division of revenue under this scheme? Suppose Andy gets a total of x and Bill gets a total of y; thus Andy's profit is $(x - 900)$ and Bill's is $(y - 100)$. The arbitrator's decision implies that Andy's profit should be four times as large as Bill's; so $x - 900 = 4(y - 100)$, or $x = 4y + 500$. The total revenue available to both is \$3,000; so it must also be true that $x + y = 3{,}000$, or $x = 3{,}000 - y$. Then $x = 4y + 500 = 3{,}000 - y$, or $5y = 2{,}500$, or $y = 500$, and thus $x = 2{,}500$. This division mechanism leaves Andy with a profit of $2{,}500 - 900 = \$1{,}600$ and Bill with $500 - 100 = \$400$, which is the 4:1 split in favor of Andy that the arbitrator wants.

We now develop this simple data into a general algebraic formula that you will find useful in many practical applications. Then we go on to examine more specifics of what determines the ratio in which the profits in a bargaining game get split.

B. General Theory

Suppose two bargainers, A and B, seek to split a total value v, which they can achieve if and only if they agree on a specific division. If no agreement is reached, A will get a and B will get b, each by acting alone or in some other way acting outside of this relationship. Call these their *backstop* payoffs or, in the jargon of the Harvard Negotiation Project, their **BATNAs (best alternative to a negotiated agreement)**.[3] Often a and b are both zero, but, more generally, we only need to assume that $a + b < v$, so that there is a positive **surplus** $(v - a - b)$ from agreement; if this were not the case, the whole bargaining would be moot because each side would just take up its outside opportunity and get its BATNA.

Consider the following rule: each player is to be given his BATNA plus a share of the surplus, a fraction h of the surplus for A and a fraction k for B, such that $h + k = 1$. Writing x for the amount that A finally ends up with, and similarly y for B, we translate these statements as

$$x = a + h(v - a - b) = a(1 - h) + h(v - b)$$
$$x - a = h(v - a - b)$$

and

$$y = b + k(v - a - b) = b(1 - k) + k(v - a)$$
$$y - b = k(v - a - b).$$

[3] See Roger Fisher and William Ury, *Getting to Yes*, 2nd ed. (New York: Houghton Mifflin, 1991).

We call these expressions the Nash formulas. Another way of looking at them is to say that the surplus $(v - a - b)$ gets divided between the two bargainers in the proportions of $h\colon k$, or

$$\frac{y - b}{x - a} = \frac{k}{h}$$

or, in slope-intercept form,

$$y = b + \frac{k}{h}(x - a) = \left(b - \frac{ak}{h}\right) + \frac{k}{h}x.$$

To use up the whole surplus, x and y must also satisfy $x + y = v$. The Nash formulas for x and y are actually the solutions to these last two simultaneous equations.

A geometric representation of the **Nash cooperative solution** is shown in Figure 17.1. The backstop, or BATNA, is the point P, with coordinates (a, b). All points (x, y) that divide the gains in proportions $h\colon k$ between the two players lie along the straight line passing through P and having slope k/h; this slope is just the line $y = b + (k/h)(x - a)$ that we derived earlier. All points (x, y) that use up the whole surplus lie along the straight line joining $(v, 0)$ and $(0, v)$; this line is the second equation that we derived—namely, $x + y = v$. The Nash solution is at the intersection of the lines, at the point Q. The coordinates of this point are the parties' payoffs after the agreement.

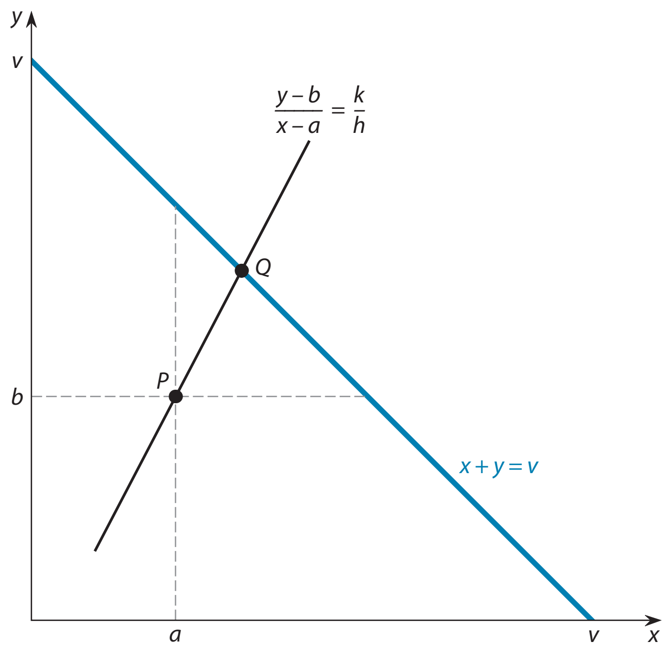

FIGURE 17.1 The Nash Bargaining Solution in the Simplest Case

The Nash formula says nothing about how or why such a solution might come about. And this vagueness is its merit—it can be used to encapsulate the results of many different theories taking many different perspectives.

At the simplest, you might think of the Nash formula as a shorthand description of the outcome of a bargaining process that we have not specified in detail. Then *h* and *k* can stand for the two parties' relative bargaining strengths. This shorthand description is a cop-out; a more complete theory should explain where these bargaining strengths come from and why one party might have more than the other. We do so in a particular context later in the chapter. In the meantime, by summarizing any and all of the sources of bargaining strength in these numbers *h* and *k*, the formula has given us a good tool.

Nash's own approach was quite different—and indeed different from the whole approach to game theory that we have taken thus far in this book. Therefore, it deserves more careful explanation. In all the games that we have studied so far, the players chose and played their strategies separately from each other. We have looked for equilibria in which each player's strategy was in his own best interest, given the strategies of the others. Some such outcomes were very bad for some or even all of the players, the prisoners' dilemma being the most prominent example. In such situations, there was scope for the players to get together and agree that all would follow some particular strategy. But in our framework, there was no way in which they could be sure that the agreement would hold. After reaching an agreement, the players would disperse, and, when it was each player's turn to act, he would actually take the action that served his own best interest. The agreement on joint action would unravel in the face of such separate temptations. True, in considering repeated games in Chapter 10, we found that the implicit threat of the collapse of an ongoing relationship might sustain an agreement, and, in Chapter 8, we did allow for communication by signals. But individual action was of the essence, and any mutual benefit could be achieved only if it did not fall prey to the selfishness of separate individual actions. In Chapter 2, we called this approach to game theory *noncooperative*, emphasizing that the term signified how actions are taken, not whether outcomes are jointly good. The important point, again, is that any joint good has to be an equilibrium outcome of separate action in such games.

What if joint action *is* possible? For example, the players might take all their actions immediately after the agreement is reached, in one another's presence. Or they might delegate the implementation of their joint agreement to a neutral third party, or to an arbitrator. In other words, the game might be *cooperative* (again in the sense of joint action). Nash modeled bargaining as a cooperative game.

The thinking of a collective group that is going to implement a joint agreement by joint action can be quite different from that of a set of individual people who know that they are *interacting* strategically but are *acting* noncooperatively. Whereas the latter set will think in terms of an equilibrium and then delight or grieve, depending on whether they like the results, the former can think first of what is a good outcome and then see how to implement it. In other words, the theory defines the outcome of a cooperative game in terms of some general principles or properties that seem reasonable to the theorist.

Nash formulated a set of such principles for bargaining and proved that they implied a unique outcome. His principles are roughly as follows: (1) the outcome should be invariant if the scale in which the payoffs are measured changes linearly; (2) the outcome should be **efficient**; and (3) if the set of possibilities is reduced by removing some that are irrelevant in the sense that they would not be chosen anyway, then the outcome should not be affected.

The first of these principles conforms to the theory of expected utility, which we discussed briefly in the appendix to Chapter 8. We saw there that a nonlinear rescaling of payoffs represents a change in a player's attitude toward risk and a real change in behavior; a concave rescaling implies risk aversion, and a convex rescaling implies risk preference. A linear rescaling, being the intermediate case between these two, represents no change in the attitude toward risk. Therefore, it should have no effect on expected payoff calculations and no effect on outcomes.

The second principle simply means that no available mutual gain should go unexploited. In our simple example of A and B splitting a total value of v, it would mean that x and y has to sum to the full amount of v available, and not to any smaller amount; in other words, the solution has to lie on the $x + y = v$ line in Figure 17.1. More generally, the complete set of logically conceivable agreements to a bargaining game, when plotted on a graph as in Figure 17.1, will be bounded above and to the right by the subset of agreements that leave no mutual gain unexploited. This subset need not lie along a straight line such as $x + y = v$ (or $y = v - x$); it could lie along any curve of the form $y = f(x)$.

In Figure 17.2, all of the points on and below (that is, "south" and to the "west" of) the thick blue curve labeled $y = f(x)$ constitute the complete set of conceivable outcomes. The curve itself consists of the efficient outcomes; there are no conceivable outcomes that include more of both x and y than the outcomes on $y = f(x)$; so there are no unexploited mutual gains left. Therefore, we call the curve $y = f(x)$ the **efficient frontier** of the bargaining problem.

We can illustrate a curved efficient frontier using the example of efficient risk allocation from Chapter 8, Section1.A. Two farmers, each with a square root utility function, face the risk that equally likely good or bad weather would make their incomes either \$160,000 or \$40,000, yielding each an expected utility of

$$1/2 \times \sqrt{160{,}000} + 1/2 \times \sqrt{40{,}000} = 1/2 \times 400 + 1/2 \times 200 = 300.$$

But their risks are perfectly negatively correlated. One gets good weather only when the other gets bad, so their combined income is \$200,000 no matter which of them gets the good weather. If they negotiate so that the first of them

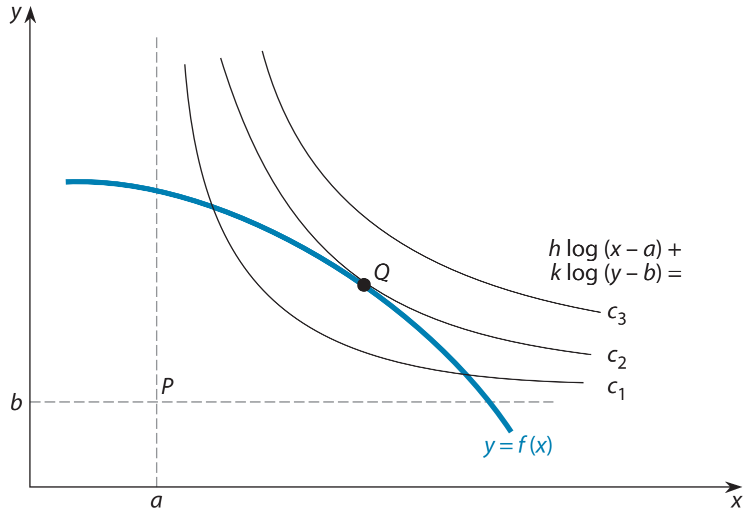

FIGURE 17.2 The General Form of the Nash Bargaining Solution

gets z of the combined income and the other gets the remaining $(200{,}000 - z)$, their respective utilities x and y will be

$$x = \sqrt{z} \quad \text{and} \quad y = \sqrt{200{,}000 - z}.$$

Therefore, we can describe the set of possible risk-sharing outcomes by the equation

$$x^2 + y^2 = z + (200{,}000 - z) = 200{,}000.$$

This equation defines a quarter-circle in the positive quadrant and represents the efficient frontier of the farmers' bargaining problem. The BATNA of each farmer is the expected utility 300 he would get if the two are not able to come to any risk-sharing agreement. Substituting this value into the equation above yields $300^2 + 300^2 = 90{,}000 + 90{,}000 = 180{,}000 < 200{,}000$. So the farmers' BATNA point lies inside the quarter-circle efficient frontier.

The third principle also seems appealing. If an outcome that a bargainer wouldn't have chosen anyway drops out of the picture, what should it matter? This assumption is closely connected to the "independence of irrelevant alternatives" assumption of Arrow's impossibility theorem, which we met in Chapter 15, Section 3, but we must leave the development of this connection to more advanced treatments of the subject.

Nash proved that the cooperative outcome that satisfied all three of these assumptions could be characterized by the mathematical maximization problem: choose x and y to

$$\text{maximize } (x - a)^h (y - b)^k \quad \text{subject to } y = f(x).$$

Here x and y are the outcomes, a and b the backstops, and h and k two positive numbers summing to 1, which are like the bargaining strengths of the Nash formula. The values for h and k cannot be determined by Nash's three assumptions alone; thus they leave a degree of freedom in the theory and in the outcome. Nash actually imposed a fourth assumption on the problem—that of symmetry between the two players; this additional assumption led to the outcome $h = k = 1/2$ and fixed a unique solution. We have given the more general formulation that subsequently became common in game theory and economics.

Figure 17.2 also gives a geometric representation of the objective of the maximization. The black curves labeled c_1, c_2, and c_3 are the level curves, or contours, of the function being maximized; along each such curve, $(x - a)^h(y - b)^k$ is constant and equals c_1, c_2, or c_3 (with $c_1 < c_2 < c_3$) as indicated. The whole space could be filled with such curves, each with its own value of the constant, and curves farther to the "northeast" would have higher values of the constant.

It is immediately apparent that the highest possible value of the function is at that point of tangency, Q, between the efficient frontier and one of the level curves.[4] The location of Q is defined by the property that the contour passing through Q is tangent to the efficient frontier. This tangency is the usual way to illustrate the Nash cooperative solution geometrically.[5]

In our example of Figure 17.1, we can also derive the Nash solution mathematically; to do so requires calculus, but the ends here are more important—at least to the study of games of strategy—than the means. For the solution, it helps to write $X = x - a$ and $Y = y - b$. Thus, X is the amount of the surplus that goes to A, and Y is the amount of the surplus that goes to B. The efficiency of the outcome guarantees that $X + Y = x + y - a - b = v - a - b$, which is just the total surplus and which we will write as S. Then $Y = S - X$, and

$$(x - a)^h(y - b)^k = X^hY^k = X^h(S - X)^k.$$

In the Nash solution, X takes on the value that maximizes this function. Elementary calculus tells us that the way to find X is to take the derivative of this expression with respect to X and set it equal to zero. Using the rules of calculus for taking the derivatives of powers of X and of the product of two functions of X, we get

$$hX^{h-1}(S - X)^k - X^hk(S - X)^{k-1} = 0.$$

[4] One and only one of the (convex) level curves can be tangential to the (concave) efficient frontier; in Figure 17.2, this level curve is labeled c_2. All lower-level curves (such as c_1) cut the frontier in two points; all higher-level curves (such as c_3) do not meet the frontier at all.

[5] If you have taken an elementary microeconomics course, you will have encountered the concept of social optimality, illustrated graphically by the tangent point between the production possibility frontier of an economy and a social indifference curve. Our Figure 17.2 is similar in spirit; the efficient frontier in bargaining is like the production possibility frontier, and the level curves of the objective in cooperative bargaining are like social indifference curves.

When we cancel the common factor $X^{h-1}(S - X)^{k-1}$, this equation becomes

$$h(S - X) - kX = 0$$
$$hY - kX = 0$$
$$kX = hY$$
$$\frac{X}{h} = \frac{Y}{k}.$$

Finally, expressing the equation in terms of the original variables x and y, we have $(x - a)/h = (y - b)/k$, which is just the Nash formula. The punch line: Nash's three conditions lead to the formula we originally stated as a simple way of splitting the bargaining surplus.

The three principles, or desired properties, that determine the Nash cooperative-bargaining solution are simple and even appealing. But in the absence of a good mechanism to make sure that the parties take the actions stipulated by the agreement, these principles may come to nothing. A player who can do better by strategizing on his own than by using the Nash solution may simply reject the principles. If an arbitrator can enforce a solution, the player may simply refuse to go to arbitration. Therefore, Nash's cooperative solution will seem more compelling if it can be given an alternative interpretation—namely, as the Nash equilibrium of a noncooperative game played by the bargainers. This can indeed be done, and we will develop an important special case of it in Section 5.

2 VARIABLE-THREAT BARGAINING

In this section, we embed the Nash cooperative solution within a specific game—namely, as the second stage of a sequential-play game. We assumed in Section 1 that the players' backstops (BATNAs) a and b were fixed. But suppose there is a first stage to the bargaining game in which the players can make strategic moves to manipulate their BATNAs within certain limits. After they have done so, the Nash cooperative outcome starting from those BATNAs will emerge in a second stage of the game. This type of game is called **variable-threat bargaining**. What kind of manipulation of the BATNAs is in a player's interest in this type of game?

We show the possible outcomes from a process of manipulating BATNAs in Figure 17.3. The originally given backstops (a and b) are the coordinates for the game's backstop point P; the Nash solution to a bargaining game with these backstops is at the outcome Q. If player A can increase his BATNA to move the game's backstop point to P_1, then the Nash solution starting there leads to the outcome Q', which is better for player A (and worse for B). Thus, a strategic move that improves one's own BATNA is desirable. For example, if you have a

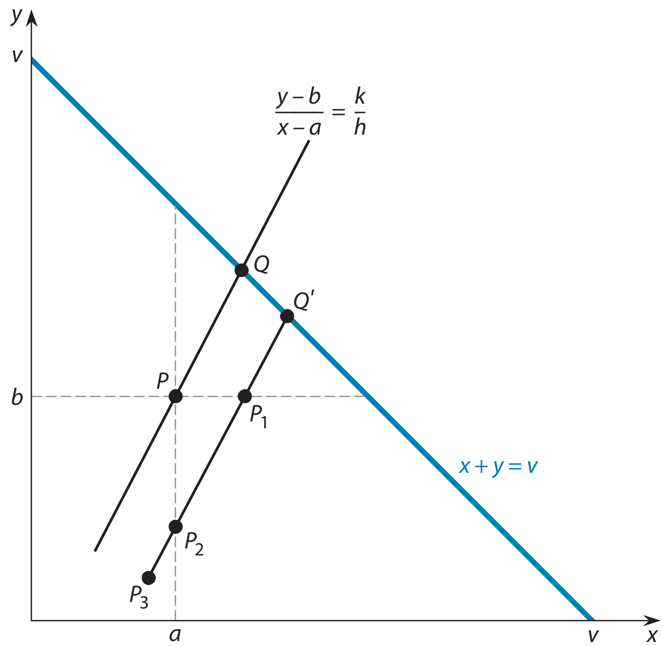

FIGURE 17.3 Bargaining Game of Manipulating BATNAs

good job offer in your pocket—a higher BATNA—when you go for an interview at another company, you are likely to get a better offer from that employer than you would if you did not have the first alternative.

The result that improving your own BATNA can improve your ultimate outcome is quite obvious, but the next step in the analysis is less so. It turns out that if player A can make a strategic move that *reduces* player B's BATNA and moves the game's backstop point to P_2, the Nash solution starting there leads to the *same* outcome Q' that was achieved after A increased his own BATNA to get to the backstop point P_1. Therefore, this alternative kind of manipulation is equally in player A's interest. As an example of decreasing your opponent's BATNA, think of a situation in which you are already working and want to get a raise. Your chances are better if you can make yourself indispensable to your employer so that without you his business has much worse prospects; his low outcome in the absence of an agreement—not offering you a raise and your leaving the firm—may make him more likely to accede to your wishes.

Finally and even more dramatically, if player A can make a strategic move that lowers *both players' BATNAs* so that the game's backstop point moves to P_3, that again has the same result as each of the preceding manipulations. This particular move is like using a threat that says, "This will hurt you more than it hurts me."

In general, the key for player A is to shift the game's BATNA point to somewhere below the line PQ. The farther southeast the BATNA point is moved, the better it is for player A in the eventual outcome. As is usual with threats, the idea is not actually to suffer the low payoff but merely to use its prospect as a lever to get a better outcome.

The possibility of manipulating BATNAs in this way depends on the context. We offer just one illustration. In 1980 there was a baseball players' strike. It took a very complicated form. The players struck in spring training, then resumed working (playing, really) when the regular season began in April, and went on strike again starting on Memorial Day. A strike is costly to both sides, employers and employees, but the costs differ. During spring training the players do not have salaries, but the owners make some money from vacationing spectators. At the start of the regular season, in April and May, the players get salaries but the weather is cold and the season is not yet exciting; therefore the crowds are small, and so the cost of a strike to the owners is low. The crowds start to build up from Memorial Day onward, which raises the cost of a strike to the owners, but the salaries that the players stand to lose stay the same. So we see that the two-piece strike was very cleverly designed to lower the BATNA of the owners *relative to* that of the players as much as possible.[6]

One puzzle remains: Why did the strike occur at all? According to the theory, everyone should have seen what was coming; a settlement more favorable to the players should have been reached so that the strike would have been unnecessary. A strike that actually happens is a threat that has "gone wrong." Some kind of uncertainty—asymmetric information or brinkmanship—must be responsible.

3 ALTERNATING-OFFERS MODEL I: TOTAL VALUE DECAYS

Here we move back to the more realistic noncooperative-game theory and think about the process of individual strategizing that may produce an equilibrium in a bargaining game. Our standard picture of this process is one of **alternating offers**. One player—say, A—makes an offer. The other—say, B—either accepts it or makes a counteroffer. If he does the latter, then A can either accept it or come back with another offer of his own. And so on. Thus, we have a sequential-move game and look for its rollback equilibrium.

To find a rollback equilibrium, we must start at the end and work backward. But where is the end point? Why should the process of offers and counteroffers ever terminate? Perhaps more drastic, why would it ever start? Why would

[6] See Larry DeBrock and Alvin Roth, "Strike Two: Labor-Management Negotiations in Major League Baseball," *Bell Journal of Economics*, vol. 12, no. 2 (Autumn 1981), pp. 413–25.

the two bargainers not stick to their original positions and refuse to budge? It is costly to both if they fail to agree at all, but the benefit of an agreement is likely to be smaller to the one who makes the first or the larger concession. The reason that anyone concedes must be that continuing to stand firm would cause an even greater loss of benefit. This loss takes one of two broad forms. The available pie, or surplus, may **decay** (shrink) with each offer, a possibility that we consider in this section. The alternative possibility is that time has value and **impatience** is important, and so a delayed agreement is worth less; we examine this possibility in Section 5.

Consider the following story of bargaining over a shrinking pie. A fan arrives at a professional football (or basketball) game without a ticket. He is willing to pay as much as $25 to watch each quarter of the game. He finds a scalper who states a price. If the fan is not willing to pay this price, he goes to a nearby bar to watch the first quarter on the big-screen TV. At the end of the quarter, he comes out, finds the scalper still there, and makes a counteroffer for the ticket. If the scalper does not agree, the fan goes back to the bar. He comes out again at the end of the second quarter, when the scalper makes him yet another offer. If that offer is not acceptable to the fan, he goes back into the bar, emerging at the end of the third quarter to make yet another counteroffer. The value of watching the rest of the game is declining as the quarters go by.[7]

Rollback analysis enables us to predict the outcome of this alternating-offers bargaining process. At the end of the third quarter, the fan knows that, if he does not buy the ticket then, the scalper will be left with a small piece of paper of no value. So the fan will be able to make a very small offer that, for the scalper, will still be better than nothing. Thus, on his last offer, the fan can get the ticket almost for free. Backing up one period, we see that, at the end of the second quarter, the scalper has the initiative in making the offer. But he must look ahead and recognize that he cannot hope to extract the whole of the remaining two quarters' value from the fan. If the scalper asks for more than $25—the value of the *third* quarter to the fan—the fan will turn down the offer because he knows that he can get the fourth quarter later for almost nothing, so the scalper can ask for $25 at most. Now consider the situation at the end of the first quarter. The fan knows that if he does not buy the ticket now, the scalper can expect to get only $25 later, and so $25 is all that the fan needs to offer now to secure the ticket. Finally, before the game even begins, the scalper can look ahead and ask for $50; this $50 includes the $25 value of the *first* quarter to the fan plus the $25 for which the fan can get the remaining three quarters' worth. Thus, the two will strike an immediate

[7] Just to keep the argument simple, we imagine this process as one-on-one bargaining. Actually, there may be several fans and several scalpers, turning the situation into a *market*. You can access our supplemental chapter on interactions in markets on the textbook Web site.

agreement, and the ticket will change hands for $50, but the price is determined by the full forward-looking rollback reasoning.[8]

This story can be easily turned into a more general argument for two bargainers, A and B. Suppose A makes the first offer to split the total surplus, which we call v (in some currency—say, dollars). If B refuses the offer, the total available drops by x_1 to $(v - x_1)$; B offers a split of this amount. If A refuses B's offer, the total drops by a further amount x_2 to $(v - x_1 - x_2)$; A offers a split of this amount. This offer and counteroffer process continues until finally, say, after 10 rounds, $v - x_1 - x_2 - \cdots - x_{10} = 0$, so the game ends. As usual with sequential-play games, we begin our analysis at the end.

If the game has gone to the point where only x_{10} is left, B can make a final offer whereby he gets to keep "almost all" of the surplus, leaving a measly cent or so to A. Left with the choice of that or absolutely nothing, A should accept the offer. To avoid the finicky complexity of keeping track of tiny cents, let us call this outcome "x_{10} to B, 0 to A." We will do the same in the other (earlier) rounds.

Knowing what is going to happen in round 10, we turn to round 9. Here A is to make the offer, and $(x_9 + x_{10})$ is left. A knows that he must offer at least x_{10} to B or else B will refuse the offer and take the game to round 10, where he can get that much. Bargainer A does not want to offer any more to B. So, on round 9, A will offer a split where he keeps x_9 and leaves x_{10} to B.

Then on the round before, when $x_8 + x_9 + x_{10}$ is left, B will offer a split where he gives x_9 to A and keeps $(x_8 + x_{10})$. Working backward, on the very first round, A will offer a split where he keeps $(x_1 + x_3 + x_5 + x_7 + x_9)$ and gives $(x_2 + x_4 + x_6 + x_8 + x_{10})$ to B. This offer will be accepted.

You can remember these formulas by means of a simple trick. *Hypothesize* a sequence in which all offers are refused. (This sequence is *not* what actually happens.) Then add up the amounts that would be destroyed by the refusals of one player. This total is what the other player gets in the actual equilibrium. For example, when B refused A's first offer, the total available surplus dropped by x_1, and x_1 became part of what went to A in the equilibrium of the game.

If each player has a positive BATNA, the analysis must be modified somewhat to take them into account. At the last round, B must offer A at least the BATNA a. If x_{10} is greater than a, B is left with $(x_{10} - a)$; if not, the game must terminate before this round is reached. Now at round 9, A must offer B the larger of the two amounts—the $(x_{10} - a)$ that B can get in round 10 or the BATNA b that B can get outside this agreement. The analysis can proceed all the way back to round 1 in this way; we leave it to you to complete the rollback reasoning for this case.

[8] To keep the analysis simple, we omitted the possibility that the game might get exciting, and so the value of the ticket might actually increase as the quarters go by. The uncertainty makes the problem much more complex but also more interesting. The ability to deal with such possibilities should inspire you to go beyond this book or course to study more advanced game theory.

We have found the rollback equilibrium of the alternating-offers bargaining game, and in the process of deriving the outcome, we have also described the full strategies (complete contingent plans of action) behind the equilibrium—namely, what each player *would* do if the game reached some later stage. In fact, actual agreement is immediate on the very first offer. The later stages are not reached; they are off-equilibrium nodes and paths. But as usual with rollback reasoning, the foresight about what rational players would do at those nodes if they were reached is what informs the initial action.

The other important point to note is that *gradual decay* (several potential rounds of offers) leads to a more even or fairer split of the total than does *sudden decay* (only one round of bargaining permitted). In the latter, no agreement would result if B turned down A's very first offer; then, in a rollback equilibrium, A would get to keep (almost) the whole surplus, giving B an "ultimatum" to accept a measly cent or else get nothing at all. The subsequent rounds give B the credible ability to refuse a very uneven first offer.

4 EXPERIMENTAL EVIDENCE

The theory of this particular type of bargaining process is fairly simple, and many people have staged laboratory or classroom experiments that create such conditions of decaying totals, to observe what the experimental subjects actually do. We mentioned some of them briefly in Chapter 3 when considering the validity of rollback reasoning; now we examine them in more detail in the context of bargaining.[9]

The simplest bargaining experiment is the **ultimatum game**, in which there is only one round: player A makes an offer and, if B does not accept it, the bargaining ends and both get nothing. The general structure of these games is as follows. A pool of players is brought together, either in the same room or at computer terminals in a network. They are paired; one person in the pair is designated to be the *proposer* (the one who makes the offer or is the seller who posts a price) and the other to be the *chooser* (the one who accepts or refuses the offer or is the customer who decides whether to buy at that price). The pair is given a fixed surplus, usually $1 or some other sum of money, to split.

Rollback reasoning suggests that A should offer B the minimal unit—say, 1 cent out of a dollar—and that B should accept such an offer. Actual results are dramatically different. In the case in which the subjects are together in a

[9] For more details, see Douglas D. Davis and Charles A. Holt, *Experimental Economics* (Princeton: Princeton University Press, 1993), pp. 263–69, and *The Handbook of Experimental Economics*, ed. John H. Kagel and Alvin E. Roth (Princeton: Princeton University Press, 1995), pp. 255–74.

room and the assignment of the role of proposer is made randomly, the most common offer is a 50 : 50 split. Very few offers worse than 75 : 25 are made (with the proposer to keep 75% and the chooser to get 25%), and if made, they are often rejected.

This finding can be interpreted in one of two ways. Either the players cannot or do not perform the calculation required for rollback or the payoffs of the players include something other than what they get out of this round of bargaining. Surely the calculation in the ultimatum game is simple enough that anyone should be able to do it, and the subjects in most of these experiments are college students. A more likely explanation is the one that we put forth in Chapter 3, Section 6, and Chapter 5, Section 3—that the theory, which assumed payoffs to consist only of the sum earned in this one round of bargaining, is too simplistic.

Participants can have payoffs that include other things. They may have self-esteem or pride that prevents them from accepting a very unequal split. Even if the proposer A does not include this consideration in his own payoff, if he thinks that B might, then it is a good strategy for A to offer enough to make it likely that B will accept. Proposer A balances his higher payoff with a smaller offer to B against the risk of getting nothing if B rejects an offer deemed too unequal.

A second possibility is that, when the participants in the experiment are gathered in a room, the anonymity of pairing cannot be guaranteed. If the participants come from a group such as college students or villagers who have ongoing relationships outside this game, they may value those relationships. Then the proposers fear that, if they offer too unequal a split in this game, those relationships may suffer. Therefore, they would be more generous in their offers than the simplistic theory suggests. If this is the explanation, then ensuring greater anonymity should enable the proposers to make more unequal offers, and experiments do find this to be the case.

Finally, people may have a sense of fairness drilled into them during their nurture and education. This sense of fairness may have evolutionary value for society as a whole and may therefore have become a social norm. Whatever its origin, it may lead the proposers to be relatively generous in their offers, quite irrespective of the fear of rejection. One of us (Skeath) has conducted classroom experiments of several different ultimatum games. Students who had partners previously known to them with whom to bargain were noticeably "fairer" in their split of the pie. In addition, several students cited specific cultural backgrounds as explanations for behavior that was inconsistent with theoretical predictions.

Experimenters have tried variants of the basic game to differentiate between these explanations. The point about ongoing relationships can be handled by stricter procedures that visibly guarantee anonymity. Doing so by itself has some effect on the outcomes but still does not produce offers as extreme as those predicted by the purely selfish rollback argument of the theory. The remaining explanations—namely, "fear of rejection" and the "ingrained sense of fairness"—remain to be sorted out.

The fear of rejection can be removed by considering a variant called the *dictator game*. Again, the participants are matched in pairs. One person (say, A) is designated to determine the split, and the other (say, B) is simply a passive recipient of what A decides. Now the split becomes decidedly more uneven, but even here a majority of As choose to keep no more than 70%. This result suggests a role for an ingrained sense of fairness.

But such a sense has its limits. In some experiments, a sense of fairness was created simply when the experimenter randomly assigned roles of proposer and chooser. In one variant, the participants were given a simple quiz, and those who performed best were made proposers. This created a sense that the role of proposer had been earned, and the outcomes did show more unequal splits. When the dictator game was played with earned rights and with stricter anonymity conditions, most As kept everything, but some (about 5%) still offered a 50 : 50 split.

One of us (Dixit) carried out a classroom experiment in which students in groups of 20 were gathered together in a computer cluster. They were matched randomly and anonymously in pairs, and each pair tried to agree on how to split 100 points. Roles of proposer and chooser were not assigned; either could make an offer or accept the other's offer. Offers could be made and changed at any time. The pairs could exchange messages instantly with their matched opponent on their computer screens. The bargaining round ended at a random time between 3 and 5 minutes; if agreement was not reached in time by a pair, both got zero. There were 10 such rounds with different random opponents each time. Thus, the game itself offered no scope for cooperation through repetition. In a classroom context, the students had ongoing relationships outside the game, but they did not generally know or guess with whom they were playing in any round, even though no great attempt was made to enforce anonymity. Each student's score for the whole game was the sum of his point score for the 10 rounds. The stakes were quite high, because the score accounted for 5% of the course grade!

The highest total of points achieved was 515. Those who quickly agreed on 50 : 50 splits did the best, and those who tried to hold out for very uneven scores or who refused to split a difference of 10 points or so between the offers and ran the risk of time running out on them did poorly.[10] It seems that moderation and fairness do get rewarded, even as measured in terms of one's own payoff.

[10] Those who were best at the mathematical aspects of game theory, such as problem sets, did a little worse than the average, probably because they tried too hard to eke out an extra advantage and met resistance. And women did slightly better than men.

5 ALTERNATING-OFFERS MODEL II: IMPATIENCE

Now we consider a different kind of cost of delay in reaching an agreement. Suppose the actual monetary value of the total available for splitting does not decay, but players have a "time value of money" and therefore prefer early agreement to later agreement. They make offers alternately as described in Section 3, but their time preferences are such that money now is better than money later. For concreteness, we will say that both bargainers believe that having only 95 cents right now is as good as having \$1 one round later.

A player who prefers having something right away to having the same thing later is impatient; he attaches less importance to the future relative to the present. We came across this idea in Chapter 10, Section 2, and saw two reasons for it. First, player A may be able to invest his money—say, \$1—now and get his principal back along with interest and capital gains at a rate of return r, for a total of $(1 + r)$ in the next period (tomorrow, next week, next year, or whatever is the length of the period). Second, there may be some risk that the game will end between now and the next offer (as in the sudden end at a time between 3 and 5 minutes in the classroom game described earlier). If p is the probability that the game continues, then the chance of getting a dollar next period has an expected value of only p now.

Suppose we consider a bargaining process between two players with zero BATNAs. Let us start the process with one of the two bargainers—say, A—making an offer to split \$1. If the other player, B, rejects A's offer, then B will have an opportunity to make his own offer one round later. The two bargainers are in identical situations when each makes his offer, because the amount to be split is always \$1. Thus, in equilibrium the amount that goes to the person currently in charge of making the offer (call it x) is the same, regardless of whether that person is A or B. We can use rollback reasoning to find an equation that can be solved for x.

Suppose A starts the alternating offer process. He knows that B can get x in the next round when it is B's turn to make the offer. Therefore, A must give B at least an amount that is equivalent, in B's eyes, to getting x in the next round; A must give B at least $0.95x$ now. (Remember that, for B, 95 cents received now is equivalent to \$1 received in the next round; so $0.95x$ now is as good as x in the next round.) Player A will not give B any more than is required to induce B's acceptance. Thus, A offers B exactly $0.95x$ and is left with $(1 - 0.95x)$. But the amount that A gets when making the offer is just what we called x. Therefore, $x = 1 - 0.95x$, or $(1 + 0.95)x = 1$, or $x = 1/1.95 = 0.512$.

Two things about this calculation should be noted. First, even though the process allows for an unlimited sequence of alternating offers and counteroffers, in the equilibrium the very first offer A makes gets accepted and the bargaining ends. Because time has value, this outcome is efficient. The cost of delay governs how much A must offer B to induce acceptance; it thus affects A's rollback reasoning. Second, the player who makes the first offer gets more than half of the pie—namely, 0.512 rather than 0.488. Thus, each player gets more when he makes the first offer than when the other player makes the first offer. But this advantage is far smaller than that in an ultimatum game with no future rounds of counteroffers.

Now suppose the two players are not equally patient (or impatient, as the case may be). Player B still regards \$1 in the next round as being equivalent to 95 cents now, but A regards it as being equivalent to only 90 cents now. Thus, A is willing to accept a smaller amount to get it sooner; in other words, A is more impatient. This inequality in rates of impatience can translate into unequal equilibrium payoffs from the bargaining process. To find the equilibrium for this example, we write x for the amount that A gets when he starts the process and y for what B gets when he starts the process.

Player A knows that he must give B at least $0.95y$; otherwise B will reject the offer in favor of the y that he knows he can get when it becomes his turn to make the offer. Thus, the amount that A gets, x, must be $1 - 0.95y$; $x = 1 - 0.95y$. Similarly, when B starts the process, he knows that he must offer A at least $0.90x$, and then $y = 1 - 0.90x$. These two equations can be solved for x and y:

$$\begin{aligned} x &= 1 - 0.95(1 - 0.9x) \\ [1 - 0.95(0.9)]x &= 1 - 0.95 \\ 0.145x &= 0.05 \\ x &= 0.345 \end{aligned} \qquad \text{and} \qquad \begin{aligned} y &= 1 - 0.9(1 - 0.95y) \\ [1 - 0.9(0.95)]y &= 1 - 0.9 \\ 0.145y &= 0.10 \\ y &= 0.690 \end{aligned}$$

Note that x and y do not add up to 1, because each of these amounts is the payoff to a given player when he makes the first offer. Thus, when A makes the first offer, A gets 0.345 and B gets 0.655; when B makes the first offer, B gets 0.69 and A gets 0.31. Once again, each player does better when he makes the first offer than when the other player makes the first offer, and once again the difference is small.

The outcome of this case with unequal rates of impatience differs from that of the preceding case with equal rates of impatience in a major way. With unequal rates of impatience, the more impatient player, A, gets a lot less than B even when he is able to make the first offer. We expect that the person who is willing to accept less to get it sooner ends up getting less, but the difference is very dramatic. The proportion of A's shares and B's shares is almost 1 : 2.

As usual, we can now build on these examples to develop the more general algebra. Suppose A regards \$1 immediately as being equivalent to \$$(1 + r)$ one offer later or, equivalently, A regards \$$1/(1 + r)$ immediately as being equivalent to \$1 one offer later. For brevity, we substitute a for $1/(1 + r)$ in the calculations that follow. Likewise, suppose player B regards \$1 today as being equivalent to \$$(1 + s)$ one offer later; we use b for $1/(1 + s)$. If r is high (or equivalently, if a is low), then player A is very impatient. Similarly, B is impatient if s is high (or if b is low).

Here we look at bargaining that takes place in alternating rounds, with a total of \$1 to be divided between two players, both of whom have zero BATNAs. (You can do the even more general case easily once you understand this one.) What is the rollback equilibrium?

We can find the payoffs in such an equilibrium by extending the simple argument used earlier. Suppose A's payoff in the rollback equilibrium is x when he makes the first offer; B's payoff in the rollback equilibrium is y when he makes the first offer. We look for a pair of equations linking the values x and y and then solve these equations to determine the equilibrium payoffs.[11]

When A is making the offer, he knows that he must give B an amount that B regards as being equivalent to y one period later. This amount is $by = y/(1 + s)$. Then, after making the offer to B, A can keep only what is left: $x = 1 - by$.

Similarly, when B is making the offer, he must give A the equivalent of x one period later—namely, ax. Therefore $y = 1 - ax$. Solving these two equations is now a simple matter. We have $x = 1 - b(1 - ax)$, or $(1 - ab)x = 1 - b$. Expressed in terms of r and s, this equation becomes

$$x = \frac{1-b}{1-ab} = \frac{s+rs}{r+s+rs}.$$

Similarly, $y = 1 - a(1 - by)$, or $(1 - ab)y = 1 - a$. This equation becomes

$$y = \frac{1-a}{1-ab} = \frac{r+rs}{r+s+rs}.$$

Although this quick solution might seem a sleight of hand, it follows the same steps used earlier, and we soon give a different reasoning yielding exactly the same answer. First, let us examine some features of the answer.

First note that, as in our simple unequal-impatience example, the two magnitudes x and y add up to more than 1:

$$x + y = \frac{r+s+2rs}{r+s+rs} > 1.$$

Remember that x is what A gets when he has the right to make the first offer, and y is what B gets when he has the right to make the first offer. When A makes the first offer, B gets $(1 - x)$, which is less than y; this just shows A's advantage from

[11] We are taking a shortcut; we have simply assumed that such an equilibrium exists and that the payoffs are uniquely determined. More rigorous theory proves these conditions. For a step in this direction, see John Sutton, "Non-Cooperative Bargaining: An Introduction," *Review of Economic Studies*, vol. 53, no. 5 (October 1986), pp. 709–24. The fully rigorous (and quite difficult) theory is given in Ariel Rubinstein, "Perfect Equilibrium in a Bargaining Model," *Econometrica*, vol. 50, no. 1 (January 1982), pp. 97–109.

being the first proposer. Similarly, when B makes the first offer, B gets y and A gets $(1 - y)$, which is less than x.

However, usually r and s are small numbers. When offers can be made at short intervals such as a week or a day or an hour, the interest that your money can earn from one offer to the next or the probability that the game ends precisely within the next interval is quite small. For example, if r is 1% (0.01) and s is 2% (0.02), then the formulas yield $x = 0.668$ and $y = 0.337$; so the advantage of making the first offer is only 0.005. (A gets 0.668 when making the first offer, but $1 - 0.337 = 0.663$ when B makes the first offer; the difference is only 0.005.) More formally, when r and s are each small compared with 1, then their product rs is very small indeed; thus we can ignore rs to write an approximate solution for the split that does not depend on which player makes the first offer:

$$x = \frac{s}{r+s} \qquad \text{and} \qquad y = \frac{r}{r+s}.$$

Now $x + y$ is approximately equal to 1.

Most important, x and y in the approximate solution are the shares of the surplus that go to the two players, and $y/x = r/s$; that is, the shares of the players are inversely proportional to their rates of impatience as measured by r and s. If B is twice as impatient as A, then A gets twice as much as B; so the shares are 1/3 and 2/3, or 0.333 and 0.667, respectively. Thus, we see that patience is an important advantage in bargaining. Our formal analysis supports the intuition that, if you are very impatient, the other player can offer you a quick but poor deal, knowing that you will accept it.

This effect of impatience hurts the United States in numerous negotiations that our government agencies and diplomats conduct with other countries. The American political process puts a great premium on speed. The media, interest groups, and rival politicians all demand results and are quick to criticize the administration or the diplomats for any delay. Under this pressure to deliver, the negotiators are always tempted to come back with results of any kind. Such results are often poor from the long-run U.S. perspective; the other countries' concessions often have loopholes, and their promises are less than credible. The U.S. administration hails the deals as great victories, but they usually unravel after a few years. The financial crisis of 2008 offers another and dramatic example. When the housing boom collapsed, some major financial institutions that held mortgage-backed assets faced bankruptcy. That led them to curtail credit, which in turn threatened to throw the U.S. economy into a severe recession. The crisis exploded in September, in the midst of a presidential election campaign. The Treasury, the Federal Reserve, and political leaders in Congress all wanted to act fast. This impatience led them to offer far more generous terms of rescue to many financial insitutions, when a slower process would have yielded an outcome that cost the taxpayers much less and offered them much better prospects of sharing in future gains on the assets being rescued.

Individuals who suffer losses are in a much weaker position when they negotiate with insurance companies on coverage. The companies often make lowball offers of settlement to people who have suffered a major loss, knowing that their clients urgently want to make a fresh start and are therefore very impatient.

As a conceptual matter, the formula $y/x = r/s$ ties our noncooperative game approach to bargaining to the cooperative approach of the Nash solution discussed in Section 1. The formula for shares of the available surplus that we derived there becomes, with zero BATNAs, $y/x = k/h$. In the cooperative approach, the shares of the two players stood in the same proportions as their bargaining strengths, but these strengths were assumed to be given somehow from the outside. Now we have an explanation for the bargaining strengths in terms of some more basic characteristics for the players—h and k are inversely proportional to the players' rates of impatience r and s. In other words, Nash's cooperative solution can also be given an alternative and perhaps more satisfactory interpretation as the rollback equilibrium of a noncooperative game of offers and counteroffers, if we interpret the abstract bargaining-strength parameters in the cooperative solution correctly in terms of the players' characteristics, such as impatience.

Finally, note that agreement is once again immediate—the very first offer is accepted. As usual, the whole rollback analysis disciplines by making the first proposer recognize that the other would credibly reject a less adequate offer.

To conclude this section, we offer an alternative derivation of the same (precise) formula for the equilibrium offers that we derived earlier. Suppose this time that there are 100 rounds; A is the first proposer and B the last. Start the backward induction in the 100th round; B will keep the whole dollar. Therefore in the 99th round, A will have to offer B the equivalent of \$1 in the 100th round—namely, b, and A will keep $(1 - b)$. Then proceed backward:

In round 98, B offers $a(1 - b)$ to A and keeps
$1 - a(1 - b) = 1 - a + ab$.
In round 97, A offers $b(1 - a + ab)$ to B and keeps
$1 - b(1 - a + ab) = 1 - b + ab - ab^2$.
In round 96, B offers $a(1 - b + ab - ab^2)$ to A and keeps
$1 - a + ab - a^2b + a^2b^2$.
In round 95, A offers $b(1 - a + ab - a^2b + a^2b^2)$ to B and keeps
$1 - b + ab - ab^2 + a^2b^2 - a^2b^3$.

Proceeding in this way and following the established pattern, we see that, in round 1, A gets to keep

$$\begin{aligned} &1 - b + ab - ab^2 + a^2b^2 - a^2b^3 + \cdots + a^{49}b^{49} - a^{49}b^{50} \\ &= (1 - b)[1 + ab + (ab)^2 + \cdots + (ab)^{49}] \end{aligned}$$

The consequence of allowing more and more rounds is now clear. We just get more and more of these terms, growing geometrically by the factor *ab* for every two offers. To find A's payoff when he is the first proposer in an infinitely long sequence of offers and counteroffers, we have to find the limit of the infinite geometric sum. In the appendix to Chapter 10 we saw how to sum such series. Using the formula obtained there, we get the answer

$$(1 - b)\,[1 + ab + (ab) + (ab)^2 + \cdots + (ab)^{49} + \cdots] = \frac{1 - b}{1 - ab}.$$

This is exactly the solution for x that we obtained before. By a similar argument, you can find B's payoff when he is the proposer and, in doing so, improve your understanding and technical skills at the same time.

6 MANIPULATING INFORMATION IN BARGAINING

We have seen that the outcomes of a bargain depend crucially on various characteristics of the parties to the bargain, most important their BATNAs and their impatience. We have proceeded thus far by assuming that the players knew each other's characteristics as well as their own. In fact, we have assumed that each player knew that the other knew, and so on; that is, the characteristics were common knowledge. In reality, we often engage in bargaining without knowing the other side's BATNA or degree of impatience; sometimes we do not even know our own BATNA very precisely.

As we saw in Chapter 8, a game with such uncertainty or informational asymmetry has associated with it an important game of signaling and screening of strategies for manipulating information. Bargaining is replete with such strategies. A player with a good BATNA or a high degree of patience wants to signal this fact to the other. However, because someone without these good attributes will want to imitate them, the other party will be skeptical and will examine the signals critically for their credibility. And each side will also try screening, by using strategies that induce the other to take actions that will reveal its characteristics truthfully.

In this section, we look at some such strategies used by buyers and sellers in the housing market. Most Americans are active in the housing market several times in their lives, and many people are professional real-estate agents or brokers who have even more extensive experience in the matter. Moreover, housing is one of the few markets in the United States where haggling or bargaining over price is accepted and even expected. Therefore, considerable experience of strategies is available. We draw on this experience for many of our examples and interpret it in the light of our game-theoretic ideas and insights.[12]

When you contemplate buying a house in a new neighborhood, you are unlikely to know the general range of prices for the particular type of house in which you are interested. Your first step should be to find this range so that you can then determine your BATNA. And that does not mean looking at newspaper ads or realtors' listings, which indicate only asking prices. Local newspapers and some Internet sites list recent actual transactions and the actual prices; you should check them against the asking prices of the same houses to get an idea of the state of the market and the range of bargaining that might be possible.

Next comes finding out (screening) the other side's BATNA and level of impatience. If you are a buyer, you can find out why the house is being sold and how long it has been on the market. If it is empty, why? And how long has it been that way? If the owners are getting divorced or have moved elsewhere and are financing another house on an expensive bridge loan, it is likely that they have a low BATNA or are rather impatient.

You should also find out other relevant things about the other side's preferences, even though these preferences may seem irrational to you. For example, some people consider an offer too far below the asking price an insult and will not sell at any price to someone who makes such an offer. Norms of this kind vary across regions and times. It pays to find out what the common practices are.

Most important, the *acceptance* of an offer more accurately reveals a player's true willingness to pay than anything else and therefore is open to exploitation by the other player. A brilliant game-theorist friend of ours tried just such a ploy. He was bargaining for a floor lamp. Starting with the seller's asking price of $100, the negotiation proceeded to a point where our friend made an offer to buy the lamp for $60. The seller said yes, at which point our friend thought: "This guy is willing to sell it for $60, so his true rock-bottom price must be even lower. Let me try to find out whether it is." So our friend said, "How about $55?" The seller got very upset, refused to sell for any price, and asked our friend to leave the store and never come back.

The seller's behavior conformed to the norm that it is utmost bad faith in bargaining to renege on an offer once it is accepted. It makes good sense as a norm in the whole context of all bargaining games that take place in society. If an offer on the table cannot be accepted in good faith by the other player without fear of the kind of exploitation attempted by our friend, then each bargainer will wait to get the other to accept an offer, thereby revealing the limit of his true rock-bottom acceptance level, and the whole process of bargains will grind to a

[12] We have taken the insights of practitioners from Andrée Brooks, "Honing Haggling Skills," *New York Times*, December 5, 1993.

halt. Therefore, such behavior has to be disallowed. Making it a social norm to which people adhere instinctively, as the seller in the example did, is a good way for society to achieve this aim.

The offer may explicitly say that it is open only for a specified and limited time; this stipulation can be part of the offer itself. Job offers usually specify a deadline for acceptance; stores have sales for limited periods. But in that case the offer is truly a *package* of price and time, and reneging on either dimension provokes a similar instinctive anger. For example, customers get quite angry if they arrive at a store in the sale period and find an advertised item unavailable. The store must offer a rain check, which allows the customer to buy the item at its sale price when next available at the regular price; even this offer causes some inconvenience to the customer and risks some loss of goodwill. The store can specify "limited quantities, no rain checks" very clearly in its advertising of the sale; even then, many customers get upset if they find that the store has run out of the item.

Next on our list of strategies to use in one-on-one bargaining, as in the housing market, comes signaling your own high BATNA or patience. The best way to signal patience is to *be* patient. Do not come back with counteroffers too quickly, "let the sellers think they've lost you." This signal is credible because someone not genuinely patient would find it too costly to mimic the leisurely approach. Similarly, you can signal a high BATNA by starting to walk away, a tactic that is common in negotiations at bazaars in other countries and some flea markets and tag sales in the United States.

Even if your BATNA is low, you may commit yourself to not accepting an offer below a certain level. This constraint acts just like a high BATNA, because the other side cannot hope to get you to accept anything less. In the housing context, you can claim your inability to concede any further by inventing (or creating) a tightwad parent who is providing the down payment or a spouse who does not really like the house and will not let you offer any more. Sellers can try similar tactics. A parallel in wage negotiations is the *mandate.* A meeting is convened of all the workers who pass a resolution—the mandate—authorizing the union leaders to represent them at the negotiation but with the constraint that the negotiators must not accept an offer below a certain level specified in the resolution. Then, at the meeting with the management, the union leaders can say that their hands are tied; there is no time to go back to the membership to get their approval for any lower offer.

Most of these strategies entail some risk. While you are signaling patience by waiting, the seller of the house may find another willing buyer. As employer and union wait for one another to concede, tensions may mount so high that a strike that is costly to both sides nevertheless cannot be prevented. In other words, many strategies of information manipulation are instances of brinkmanship. We saw in Chapter 14 how such games can have an outcome that is bad for both parties. The same is true in bargaining. *Threats* of breakdown of negotiations or of strikes are strategic moves intended to achieve quicker agreement or a better deal for the player making the move; however, an *actual* breakdown or strike is an instance of the threat "gone wrong." The player making the threat—initiating the brinkmanship—must assess the risk and the potential rewards when deciding whether and how far to proceed down this path.

7 BARGAINING WITH MANY PARTIES AND ISSUES

Our discussion thus far has been confined to the classic situation where two parties are bargaining about the split of a given total surplus. But many real-life negotiations include several parties or several issues simultaneously. Although the games get more complicated, often the enlargement of the group or the set of issues actually makes it easier to arrive at a mutually satisfactory agreement. In this section, we take a brief look at such matters.[13]

A. Multi-Issue Bargaining

In a sense, we have already considered multi-issue bargaining. The negotiation over price between a seller and a buyer always comprises *two* things: (1) the object offered for sale or considered for purchase and (2) money. The potential for mutual benefit arises when the buyer values the object more than the seller does—that is, when the buyer is willing to give up more money in return for getting the object than the seller is willing to accept in return for giving up the object. Both players can be better off as a result of their bargaining agreement.

The same principle applies more generally. International trade is the classic example. Consider two hypothetical countries, Freedonia and Ilyria. If Freedonia can convert 1 loaf of bread into 2 bottles of wine (by using less of its resources such as labor and land in the production of bread and using them to produce more wine instead) and Ilyria can convert 1 bottle of wine into 1 loaf of bread (by switching its resources in the opposite direction), then between them they can create more goods "out of nothing." For example, suppose that Freedonia can produce 200 more bottles of wine if it produces 100 fewer loaves of bread and that Ilyria can produce 150 more loaves of bread if it produces 150 fewer bottles of wine. These switches in resource utilization create an extra 50 loaves of bread and 50 bottles of wine relative to what the two countries produced originally. This extra bread and wine is the "surplus" that they can create if they can agree

[13] For a more thorough treatment, see Howard Raiffa, *The Art and Science of Negotiation* (Cambridge, Mass.: Harvard University Press, 1982), parts III and IV.

on how to divide it between them. For example, suppose Freedonia gives 175 bottles of wine to Ilyria and gets 125 loaves of bread. Then each country will have 25 more loaves of bread and 25 more bottles of wine than it did before. But there is a whole range of possible exchanges corresponding to different divisions of the gain. At one extreme, Freedonia may give up all the 200 extra bottles of wine that it has produced in exchange for 101 loaves of bread from Ilyria, in which case Ilyria gets almost all the gain from trade. At the other extreme, Freedonia may give up only 151 bottles of wine in exchange for 150 loaves of bread from Ilyria, and so Freedonia gets almost all the gain from trade.[14] Between these limits lies the frontier where the two can bargain over the division of the gains from trade.

The general principle should now be clear. When two or more issues are on the bargaining table at the same time and the two parties are willing to trade more of one against less of the other at different rates, then a mutually beneficial deal exists. The mutual benefit can be realized by trading at a rate somewhere between the two parties' different rates of willingness to trade. The division of gains depends on the choice of the rate of trade. The closer it is to one side's willingness ratio, the less that side gains from the deal.

Now you can also see how the possibilities for mutually beneficial deals can be expanded by bringing more issues to the table at the same time. With more issues, you are more likely to find divergences in the ratios of valuation between the two parties and are thereby more likely to locate possibilities for mutual gain. In regard to a house, for example, many of the fittings or furnishings may be of little use to the seller in the new house to which he is moving, but they may be of sufficiently good fit and taste that the buyer values having them. Then if the seller cannot be induced to lower the price, he may be amenable to including these items in the original price to close the deal.

However, the expansion of issues is not an unmixed blessing. If you value something greatly, you may fear putting it on the bargaining table; you may worry that the other side will extract big concessions from you, knowing that you want to protect that one item of great value. At the worst, a new issue on the table may make it possible for one side to deploy threats that lower the other side's BATNA. For example, a country engaged in diplomatic negotiations may be vulnerable to an economic embargo; then it would much prefer to keep the political and economic issues distinct.

[14] Economics uses the concept *ratio of exchange*, or price, which here is expressed as the number of bottles of wine that trade for each loaf of bread. The crucial point is that the possibility of gain for both countries exists with any ratio that lies between the 2 : 1 at which Freedonia can just convert bread into wine and the 1 : 1 at which Ilyria can do so. At a ratio close to 2 : 1, Freedonia gives up almost all of its 200 extra bottles of wine and gets little more than the 100 loaves of bread that it sacrificed to produce the extra wine; thus Ilyria has almost all of the gain. Conversely, at a ratio close to 1 : 1, Freedonia has almost all of the gain. The issue in the bargaining is the division of gain and therefore the ratio or the price at which the two should trade.

B. Multiparty Bargaining

Having many parties simultaneously engaged in bargaining also may facilitate agreement, because instead of having to look for pairwise deals, the parties can seek a circle of concessions. International trade is again the prime example. Suppose the United States can produce wheat very efficiently but is less productive in cars, Japan is very good at producing cars but has no oil, and Saudi Arabia has a lot of oil but cannot grow wheat. In pairs, they can achieve little, but the three together have the potential for a mutually beneficial deal.

As with multiple issues, expanding the bargaining to multiple parties is not simple. In our example, the deal would be that the United States would send an agreed amount of wheat to Saudi Arabia, which would send its agreed amount of oil to Japan, which would in turn ship its agreed number of cars to the United States. But suppose that Japan reneges on its part of the deal. Saudi Arabia cannot retaliate against the United States, because, in this scenario, it is not offering anything to the United States that it can potentially withhold. Saudi Arabia can only break its deal to send oil to Japan, an important party. Thus, enforcement of multilateral agreements may be problematic. The General Agreement on Tariffs and Trade (GATT) between 1946 and 1994, as well as the World Trade Organization (WTO) since then, have indeed found it difficult to enforce their agreements and to levy punishments on countries that violate the rules.

SUMMARY

Bargaining negotiations attempt to divide the *surplus* (excess value) that is available to the parties if an agreement can be reached. Bargaining can be analyzed as a *cooperative* game in which parties find and implement a solution jointly or as a (structured) *noncooperative* game in which parties choose strategies separately and attempt to reach an equilibrium.

Nash's cooperative solution is based on three principles of the outcomes' invariance to linear changes in the payoff scale, *efficiency*, and invariance to removal of irrelevant outcomes. The solution is a rule that states the proportions of division of surplus, beyond the backstop payoff levels (also called *BATNAs* or *best alternatives to a negotiated agreement*) available to each party, based on relative bargaining strengths. Strategic manipulation of the backstops can be used to increase a party's payoff.

In a noncooperative setting of *alternating offer and counteroffer*, rollback reasoning is used to find an equilibrium; this reasoning generally includes a first-round offer that is immediately accepted. If the surplus value *decays* with refusals, the sum of the (hypothetical) amounts destroyed owing to the refusals of a single player is the payoff to the other player in equilibrium. If delay in agreement is costly owing to *impatience*, the equilibrium offer shares the surplus roughly in inverse proportion to the parties' rates of *impatience*. Experimental evidence indicates that players often offer more than is necessary to reach an agreement in such games; this behavior is thought to be related to player anonymity as well as beliefs about fairness.

The presence of information asymmetries in bargaining games makes signaling and screening important. Some parties will wish to signal their high BATNA levels or extreme patience; others will want to screen to obtain truthful revelation of such characteristics. When more issues are on the table or more parties are participating, agreements may be easier to reach, but bargaining may be riskier or the agreements more difficult to enforce.

KEY TERMS

alternating offers (674)
best alternative to a negotiated agreement (BATNA) (666)
decay (675)
efficient frontier (669)
efficient outcome (669)
impatience (675)
Nash cooperative solution (667)
surplus (666)
ultimatum game (677)
variable-threat bargaining (672)

SOLVED EXERCISES

S1. Consider the bargaining situation between Compaq Computer Corporation and the California businessman who owned the Internet address www.altavista.com.[15] Compaq, which had recently taken over Digital Equipment Corporation, wanted to use this man's Web address for Digital's Internet search engine, which at that time had the address www.altavista.digital.com. Compaq and the businessman apparently negotiated long and hard during the summer of 1998 over a selling price for the latter's address.

Although the businessman was the "smaller" player in this game, the final agreement appeared to entail a \$3.35 *million* price tag for the Web address in question. Compaq confirmed the purchase in August and began using the address in September but refused to divulge any of the financial details of the settlement. Given this information, comment on

[15] Details regarding this bargaining game were reported in "A Web Site by Any Other Name Would Probably Be Cheaper," *Boston Globe*, July 29, 1998, and in Hiawatha Bray's "Compaq Acknowledges Purchase of Web Site," *Boston Globe*, August 12, 1998.

the likely values of the BATNAs for these two players, their bargaining strengths or levels of impatience, and whether a cooperative outcome appears to have been attained in this game.

S2. Ali and Baba are bargaining to split a total that starts out at $100. Ali makes the first offer, stating how the $100 will be divided between them. If Baba accepts this offer, the game is over. If Baba rejects it, a dollar is withdrawn from the total, so it is now only $99. Then Baba gets the second turn to make an offer of a division. The turns alternate in this way, a dollar being removed from the total after each rejection. Ali's BATNA is $2.25 and Baba's BATNA is $3.50. What is the rollback-equilibrium outcome of the game?

S3. Two hypothetical countries, Euphoria and Militia, are holding negotiations to settle a dispute. They meet once a month, starting in January, and take turns making offers. Suppose the total at stake is 100 points. The government of Euphoria is facing reelection in November. Unless the government produces an agreement at the October meeting, it will lose the election, which it regards as being just as bad as getting zero points from an agreement. The government of Militia does not really care about reaching an agreement; it is just as happy to prolong the negotiations or even to fight, because it would be settling for anything significantly less than 100.

(a) What will be the outcome of the negotiations? What difference will the identity of the first mover make?

(b) In light of your answer to part (a), discuss why actual negotiations often continue right down to the deadline.

UNSOLVED EXERCISES

U1. Recall the variant of the pizza pricing game in Exercise U2, part (b), in Chapter 10, in which one store (Donna's Deep Dish) was much larger than the other (Pierce's Pizza Pies). The payoff table for that version of the game is:

		PIERCE'S PIZZA PIES	
		High	Low
DONNA'S DEEP DISH	High	156, 60	132, 70
	Low	150, 36	130, 50

The noncooperative dominant-strategy equilibrium is (High, Low), yielding profits of 132 to Donna's and 70 to Pierce's, for a total of 202. If the two could achieve (High, High), their total profit would be 156 + 60 = 216, but Pierce's would not agree to this pricing.

Suppose the two stores can reach an enforceable agreement whereby both charge High and Donna's pays Pierce's a sum of money. The alternative to this agreement is simply the noncooperative dominant-strategy equilibrium. They bargain over this agreement, and Donna's has 2.5 times as much bargaining power as Pierce's. In the resulting agreement, what sum will Donna's pay to Pierce's?

U2. Consider two players who bargain over a surplus initially equal to a whole-number amount V, using alternating offers. That is, Player 1 makes an offer in round 1; if Player 2 rejects this offer, she makes an offer in round 2; if Player 1 rejects this offer, she makes an offer in round 3; and so on. Suppose that the available surplus decays by a constant value of $c = 1$ each period. For example, if the players reach agreement in round 2, they divide a surplus of $V - 1$; if they reach agreement in round 5, they divide a surplus of $V - 4$. This means that the game will be over after V rounds, because at that point there will be nothing left to bargain over. (For comparison, remember the football-ticket example, in which the value of the ticket to the fan started at $100 and declined by $25 per quarter over the four quarters of the game.) In this problem, we will first solve for the rollback equilibrium to this game, and then solve for the equilibrium to a generalized version of this game in which the two players can have BATNAs.

(a) Let's start with a simple version. What is the rollback equilibrium when $V = 4$? In which period will they reach agreement? What payoff x will Player 1 receive, and what payoff y will Player 2 receive?

(b) What is the rollback equilibrium when $V = 5$?

(c) What is the rollback equilibrium when $V = 10$?

(d) What is the rollback equilibrium when $V = 11$?

(e) Now we're ready to generalize. What is the rollback equilibrium for any whole-number value of V? (Hint: You may want to consider even values of V separately from odd values.)

Now consider BATNAs. Suppose that if no agreement is reached by the end of round V, Player A gets a payoff of a and Player B gets a payoff of b. Assume that a and b are whole numbers satisfying the inequality $a + b < V$, so that the players can get higher payoffs from reaching agreement than they can by not reaching agreement.

(f) Suppose that $V = 4$. What is the rollback equilibrium for any possible values of a and b? (Hint: You may need to write down more than one formula, just as you did in part (e). If you get stuck, try assuming specific values for a and b, and then change those values to see what happens. In order to roll back, you'll need to figure out the turn at which the value of V has declined to the point where a negotiated agreement would no longer be profitable for the two bargainers.)

(g) Suppose that $V = 5$. What is the rollback equilibrium for any possible values of a and b?

(h) For any whole-number values of a, b, and V, what is the rollback equilibrium?

(i) Relax the assumption that a, b, and V are whole numbers: let them be any nonnegative numbers such that $a + b < V$. Also relax the assumption that the value of V decays by exactly 1 each period: let the value decay each period by some constant amount $c > 0$. What is the rollback equilibrium to this general problem?

U3. Let x be the amount that player A asks for, and let y be the amount that B asks for, when making the first offer in an alternating-offers bargaining game with impatience. Their rates of impatience are r and s, respectively.

(a) If we use the approximate formulas $x = s/(r + s)$ for x and $y = r/(r + s)$ for y, and if B is twice as impatient as A, then A gets two-thirds of the surplus and B gets one-third. Verify that this result is correct.

(b) Let $r = 0.01$ and $s = 0.02$, and compare the x and y values found by using the approximation method with the more exact solutions for x and y found by using the formulas $x = (s + rs)/(r + s + rs)$ and $y = (r + rs)/(r + s + rs)$ derived in the chapter.