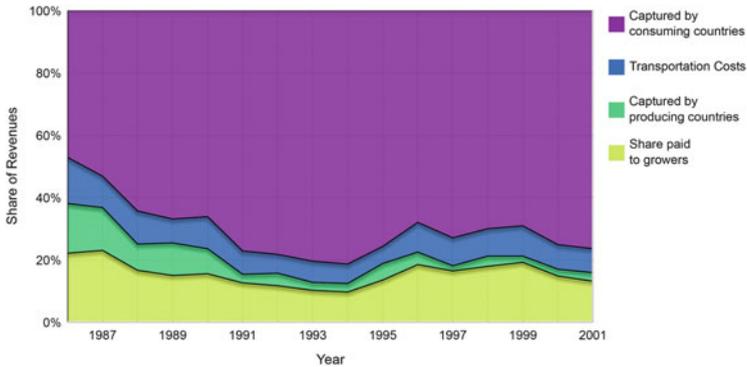


In Chap. 5 we visited the troubled world of coffee, where growers suffered from declining incomes and the economies of whole countries were at risk. We saw that prices fell after a supply management system was ended in 1989, and that subsequent increases in production capacity, combined with inelastic demand, caused the bottom to fall out of the market. This gave us an example of the power of supply and demand analysis to explain important developments in economic and social life.

But there is more to the story than this. Look at Fig. 14.1 on the next page, which shows the percentage of the coffee dollar going to each of four groups:

- the growers in coffee countries who actually produce the coffee beans
- other groups, particularly government agencies, in the coffee-producing countries
- exporters and shippers, in the form of fees, transportation costs and shrinkage
- companies in the coffee-consuming countries, especially the well-known coffee importing and retailing corporations

Several facts stand out. First, growers get a small share of the money generated in the coffee industry. It was never much above 20 %, and it has fallen much lower at times. It would be fair to say that the desperation among these coffee farmers in the first years of the new century was a product of both the declining price of the bean *and* their declining share of that price. The bulk of the proceeds, on the other hand, go to a small number of multinational coffee corporations, who account for most of the blue area in the graph. Shipping costs, which include the reduction in the measured quantity of coffee as the beans “settle” in transit, has been a constant factor over the years. The biggest reduction has been experienced by the governments in the coffee-growing countries. During the heyday of the Coffee Agreement, they levied export taxes to support marketing boards, technical services and other programs for the coffee industry. (Some of this money was no doubt siphoned off by well-placed officials, although this varied according to how well-governed the country was.) With the advent of a competitive free-for-all in coffee production, the ability of governments to dip into the revenue stream has been sharply limited, since high taxes would price their growers’ goods out of the market.



**Fig. 14.1** The division of coffee revenues among major players. The share of coffee revenues, adding up to 100 %, taken at each of four stages: coffee growers, governments of coffee-producing countries, shipping and roasting/wholesaling/retailing in the consuming countries

Figure 14.1 clearly shows that it would be naive to try to help farmers just by raising prices on the retail end. If the shares remain unchanged, little of this extra money would make its way to the growers themselves. One response by some consumers in Europe and the United States has been to purchase “Fair Trade” coffee. This is coffee which has been purchased directly from growers (usually organized cooperatively) at a price premium by certified Fair Trade companies. Because the price is higher *and* the growers get a larger share, this is a beneficial arrangement for them. On the other hand, the demand for this more expensive coffee has not been able to keep pace with the number of growers who would like to supply it. The system depends on consumers being willing to pay more for their cup out of solidarity with distant farmers, but most still search out the lowest prices. Unfortunately, some growers who would otherwise switch to other crops have tried to stay with coffee, hoping to break into the Fair Trade market, and this may have kept supply a little higher than it would have been otherwise. (Go back to the analysis in Chap. 5 to see what effect this may have on the coffee market.) The result is that, while the Fair Trade system has helped some farmers, many more are still on the brink.

What can economics offer in this situation? It can provide tools to help analyze *why* the coffee dollar is divided the way it is—what underlying forces are responsible for a world in which those who grow the coffee receive perhaps a fourth of the revenue captured by the companies that market it to consumers. Our approach will be to examine the nature of **bargaining power** in economic life.

#### Box 14.1: Five Kinds of Power

People often use the term “power” loosely, but those who have studied it have generally come to the conclusion that there are several different kinds of

(continued)

**Box 14.1** (continued)

power, and that the distinctions between them are important. In this book we have already introduced two such categories, market power and bargaining power. What is meant by power in general, and what other kinds of power are there?

These questions will not find a consensus answer among social scientists and philosophers. What is offered here is one possible mainstream approach, not too far from what most students of this subject would say.

In general, power is the ability to get someone else to do something you want them to do, which they wouldn't do otherwise. Thus the concept of power is inseparable from that of interest—what people want. It is only because there are conflicts of interest that power is a meaningful notion. If everyone's interest were always in harmony with everyone else's, no one would have any use for power over others.

We can identify at least five forms that power can take:

1. **Power to withhold.** This is what monopoly power, the topic of the previous chapter, is all about. If I have something you want or need, and there is no one else you can get it from, you are in my power. By withholding, or threatening to withhold, this good I can compel you to do what I want. As we saw, there are two dimensions to this power, the degree of monopoly (how much do you depend on me alone?) and the elasticity of demand (how inflexible is your need for what I have?). What the monopolist compels the consumer to do is pay a higher price than would otherwise be set in a competitive market.
2. **Bargaining power.** This is the subject of the current chapter, so it is not a good idea to give away all the key points just yet. But we can say this: bargaining power is a function of the degree to which the parties need an agreement; the party with the least need has the most power.
3. **Coercive power.** This was discussed in the appendix to Chap. 6. Coercion arises when one side is in a position to make threats against the other, and the threatened side is unable to escape the interaction.
4. **Institutional power.** The societies we all live in are organized through institutions, such as governments, corporations and other well-established structures. The educational institution you are (probably) ensconced in is a good example: it has decision-making hierarchies, and it puts teachers in a position of authority over students. (They give grades and assign credit.) In normal times these institutions are largely beyond challenge; we must live within their rules whether or not we like them. These rules give power to some over others. Teachers, for instance, have the power to make students do assignments they might not choose to do if there were no inequalities of power at work.
5. **Cultural power.** The deepest source of power is located in the systems of thought and language that govern how we view the world. To be able to

(continued)

**Box 14.1** (continued)

persuade others to do what you want, especially if they are unaware that this is taking place, is the ideal solution. This is the power of propagandists and advertisers, but even more it is a power that permeates a culture, often behind the backs of both those who are manipulated and those who benefit from it. Indeed, sometimes the resort to more visible forms of power, such as coercion, is an indication that cultural structures are no longer doing their job. Teachers, for instance, often have a fair amount of cultural power that comes from the prestige of their position and the expectations that society has built up around their role in the educational process. Students may be following their lead without even being aware of it. If teachers use this influence to advance their own interests (psychological, economic, political or otherwise), the outcome is an exercise of power.

## 14.1 Where Does Bargaining Power Appear?

When we introduced the supply and demand apparatus in Chap. 5, we were careful to note that it depends on a large number of assumptions—simplifications that might be good enough in some situations but would probably never be entirely correct. Now we are in a position to see what will happen when some of them are changed.

Perhaps the most important of them was the assumption of anonymity. In the supply and demand world, nobody knows who anyone else is. Each buyer and seller confronts a faceless market, where there is nothing to be done but to accept or reject the going price. In the real world, however, participants in the economy often know the identities of the people they are dealing with. In that case, they do not have to take an offer as given; they can *negotiate*.

This type of recognition can arise when markets are “thin”: when there are relatively few buyers or sellers, as in a small town with just a few auto repair shops and perhaps a few hundred cars. Specialty markets often have this characteristic too. Often there are just a few suppliers of a highly specialized input into a finished product, like certain types of computer chips or high-performance industrial ceramics. The companies on both ends of these supply relationships are likely to be aware of one another.

A second crucial assumption of the standard supply and demand approach is that each transaction is a one-time-only affair. Prices are quoted, an exchange takes place, and that’s the end of the story. There is no sequel, at least none that any parties to the transaction need to take into consideration when they decide what to buy or sell today. In the real world, however, many economic relationships occur over and over, and this leads to knowledge of who’s who in the business, and also to more strategic attitudes about what offers to make or accept.

As in most aspects of life that repeat themselves (and this includes political life, spiritual life, romance, etc.), structures emerge. Rather than reinvent the wheel every day, people fall into routines which may be informal or may take the form of organizations with explicit rules and expectations. So we have internet dating systems, churches and other religious institutions, political parties, and so on. In the realm of the economy we have relatively stable networks that make it easier for people to take care of their ongoing needs. These include business associations, labor unions and **supply chains**. This last is somewhat unfamiliar but important: it is the name given to the linkages that connect businesses engaged in different stages in the production of some good or service. A contractor overseeing the building of a house might need the services of an electrician or a mason, and all of them may be working under a builder or real estate developer. This describes a supply chain in the construction sector. Making a computer requires purchasing a large number of components, such as motherboards and computer screens, and the final product may then be delivered to a retailer who markets it to the ultimate consumer. This too is a supply chain.

The general point is this: wherever economic actors have repeated relationships with one another, formal or informal institutions are likely to arise. The economy can be viewed as crisscrossed by supply chains and other structures that break down the abstract, anonymous universe of supply and demand. When this happens, the result is likely to be bargaining, where each side tries to reach an agreement that satisfies not only its current needs, but also strengthens its ability to safeguard its future interests.

To investigate economic transactions between parties who know who each other are and who are thinking strategically, and not just reacting to current conditions, we obviously need a different type of model. Ideally, it would tell us how economic outcomes reflect what each side brings to the table, and it would give us a language to analyze strategic decision-making. What we are looking for can be found, in an imperfect but still evolving form, in game theory—in the branch of game theory that studies the bargaining process. We will begin with an extremely simple version of such a model and then consider how it could be elaborated to make it more applicable to real-world questions.

---

## 14.2 An Elementary Model of Bargaining

Just as we did with supply and demand, we will make several assumptions that, despite their unreality, serve to isolate an important aspect of the economy, as a starting point for more complex applications. In this case we will begin with two individuals who can either come to an agreement with one another or else go their separate ways. Agreeing results in the creation of something of economic value which is divided up between the two bargainers in the agreed-upon way. In the event of disagreement they each have another outside option to fall back on. These options also can be measured in economic terms and compared to the returns each side would get from a particular agreement. Finally, as is often the case when

simple economic models are being constructed, we assume that they each have only one interest, acquiring as much benefit for themselves as possible. (All of these assumptions will be lifted shortly.)

In addition, in order to make the model as simple as possible, we temporarily make one further assumption: the effect of the agreement, if they can reach it, is to divide a fixed sum of money. This sometimes occurs in the real world—for example, in bargaining over wages or prices. A simple case occurs when one person offers to sell a possession like a used car, and another person offers to buy. Any additional money paid by the buyer is received by the seller. Economists call such a situation a **zero-sum game**, since the gains and losses from bargaining exactly cancel each other out. For now, let's assume that this is the case.

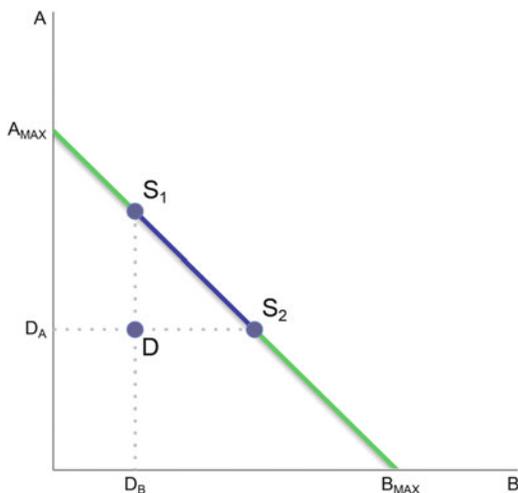
We can see this in Fig. 14.2 on the following page, which identifies the two bargainers as A and B. Let's call them Abe and Bev. The line going from the upper left of the diagram to the lower right is an **agreement curve**; it indicates the benefits for both Abe and Bev of each possible agreement. Since a fixed sum of money is at stake, the line has a slope of  $-1$ : any vertical change (for Abe) results in the exact opposite horizontal change (for Bev).

Identifying the disagreement option—the possibility that results from not agreeing—sheds further light on the situation. In Fig. 14.2 this is represented by the point D. By not agreeing, Abe would be settling for the benefit  $D_A$ ; similarly, Bev would be settling for  $D_B$ . Immediately we can see that this rules out at least some of the potential bargaining outcomes. Abe would never willingly agree to any bargain to the southeast of  $S_2$ , since that would result in even less utility than opting out of any agreement at all. By the same token, Bev would not accept a deal to the northwest of  $S_1$ . Thus, only these two points and the line segment between them constitute a potential zone of agreement.

So far, so good—but what would determine where in this zone the two bargainers would end up? There is no conclusive way to answer this question, but the famed mathematician/economist John Nash proposed one solution that has been influential ever since. Nash enumerated a set of restrictions on what ought to constitute a solution and then proved that there would be a single outcome that could meet all of them simultaneously. For our purposes, one important restriction is that the solution be symmetrical, in the sense that the identities of the bargainers shouldn't matter; any two individuals placed in the role of Abe or Bev, with the same disagreement options, should arrive at the same results. When you think about it, this could be controversial, since one side might have many social or personal advantages, such as prestige or skill, compared to the other; we will show later how Nash acknowledged this possibility.

Without going into the full proof, it is enough for now to say that Nash's solution turns out to be simple and elegant: it is the point on the agreement curve which maximizes the product of each side's net benefit. What does this mean? The net benefit is the value of an agreement over and above the value of disagreeing; the product is what results from multiplying one net benefit by the other. This may still sound mysterious, but in our simple zero-sum case it has an obvious implication: the parties will agree to split the difference. Figure 14.3 tells the story.

**Fig. 14.2** Negotiating possibilities with default options in a zero-sum bargaining game. A and B can agree to create an economic value and divide it up. If A gets the entire benefit he receives  $A_{MAX}$ ; if B gets it all she gets  $B_{MAX}$ . If they fail to agree they fall back on the disagreement option D, leaving  $D_A$  for A and  $D_B$  for B. A would not accept any offer worse for him than  $S_2$  while B would not accept any worse for her than  $S_1$



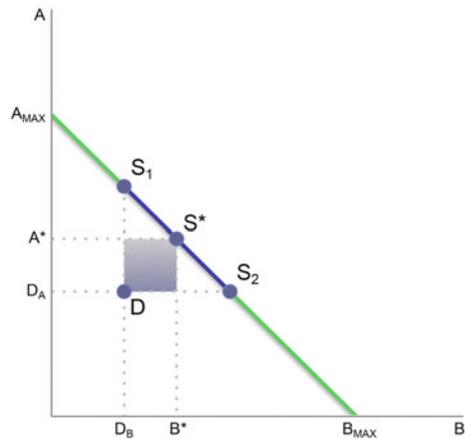
Consider any possible solution point  $S$ . Abe’s net benefit is given by the vertical distance from  $D_A$  to height of  $S$ , and Bev’s by the distance from  $D_B$  to the horizontal value of  $S$ . These two form the length and width of a rectangle whose area is the product of the net benefits. It should be clear that the sum of these two net benefits must be equal, since if, for example,  $S$  slides down to the right in Bev’s direction, any increase in Bev’s net benefit is exactly offset by a decrease in Abe’s. So what  $S$  gives us the rectangle with the largest area? We know from elementary geometry that this has to be a square, with equal length and width. In other words, the agreement that satisfy’s Nash’s criteria is  $S^*$ , halfway between  $S_1$  and  $S_2$ .

This seems to be a wonderful outcome, as fair as anyone could want. Indeed, it harkens back to an ancient debate in economics over the meaning of a **just price**. Back in the European Middle Ages, philosophers and theologians argued over how prices ought to be set in negotiations: what approach to economics would be regarded as “spiritually correct”? The most influential answer was provided by Thomas Aquinas, who was canonized (made a saint) by the Roman Catholic church. He said that the just price would be the one that equalized the benefits of an exchange between the parties to it, measured as the gain from having the exchange as against not having it. St. Thomas would no doubt be pleased by Fig. 14.3 and would see it as a vindication of his argument.

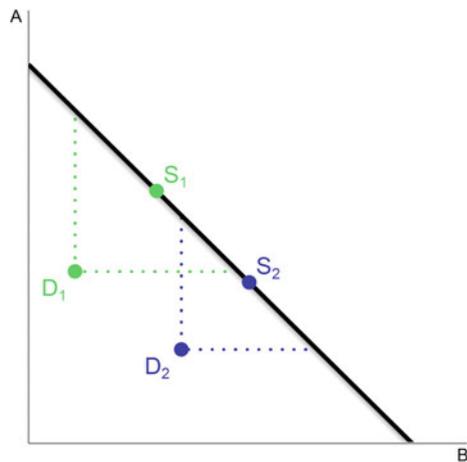
Nevertheless, it is a fair outcome only to the extent that the disagreement options are fair. To see this, look at Fig. 14.4, which shows the effect of a change in disagreement options on Nash’s solution. The first disagreement point,  $D_1$ , favors Abe and results in an agreement tilted in his direction. The second,  $D_2$ , favors Bev and has the opposite effect. The implication is that, while net benefits remain equal, this only means that those who start out from the best position end up in the best position.

In the simplest possible way, this diagram captures the essence of any bargaining theory useful to economists: it shows that agreements will tend to favor the parties

**Fig. 14.3** Nash’s bargaining solution in a zero-sum bargaining game. The solution point is  $S^*$ , where the product of Abe’s net benefit ( $A^* - D_A$ ) and Bev’s net benefit ( $B^* - D_B$ ) is maximized. This is the area of the shaded square



**Fig. 14.4** The effect of a change in the disagreement point on Nash’s solution. When the disagreement point favors Abe, as does  $D_1$ , so does solution  $S_1$ . When the disagreement point  $D_2$  favors Bev, so does the corresponding solution  $S_2$

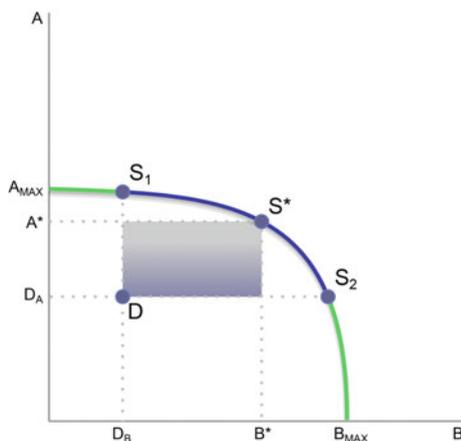


who are best off *without* them. Having a better outside option is the soul of bargaining power.

Nash’s analysis applies in similar fashion to bargaining situations in which tradeoffs are more complicated. In many contexts it is reasonable to suppose that, the more the parties share with one another, the greater their combined benefits will be—a **positive sum game**. This is reflected in Fig. 14.5, where the agreement curve bulges out in the middle and flattens at either end:

In this particular case the agreement curve is flatter for some distance to the right of  $S_1$ , indicating that over this range small concessions by Abe yield larger returns to Bev. This has the effect of pushing the solution further toward Bev’s side; at  $S^*$ , rather than at some point equidistant between  $S_1$  and  $S_2$ , the area of the shaded rectangle is maximized. A different way of saying this is that bargains that favor

**Fig. 14.5** Nash's bargaining solution when the bargaining is positive-sum. When the agreement curve allows for less than a one-to-one tradeoff between benefits to A and B, the Nash solution still calls for the rectangle of net benefits whose area is the largest



Bev offer greater gains to cooperation, and this is reflected in Bev being able to command a larger share of the enlarged pie.

One crucial assumption underlying Nash's solution is symmetry, that it doesn't matter what Abe and Bev bring to the table other than the simple facts resulting from the disagreement point and the shape of the agreement curve. Everything else—their histories, psychology, bargaining skills, etc.—is irrelevant. This is why there is a tendency toward outcomes that equalize net benefits or respect differences in tradeoffs. Nash had a simple suggestion, however, for incorporating personal differences: if one side is stronger in some sense than the other, we should convert this into a pair of weights and adjust the solution accordingly. For instance, if Bev is stronger than Abe, her weight might be 1 and Abe's only  $\frac{1}{2}$ . Then, instead of an equal division of net benefits in the zero-sum game, Bev would get  $\frac{2}{3}$  and Abe just  $\frac{1}{3}$ . A weighted solution could also be applied to the positive-sum case of Fig. 14.5.

The account of Nash's bargaining solution provided above may seem a bit arbitrary. This is because the amount of space required to present all of his mathematical reasoning would be too great in an introductory textbook—but also because it *is* somewhat arbitrary. Many game theorists have criticized it for making too many assumptions that have little empirical basis. Other solution concepts have been put forward and criticized as well. The reason we have focused on Nash is that his is the most popular approach; if you read articles in economic or business journals about bargaining, it is likely that they will refer to the Nash solution at some point. Nevertheless, it doesn't matter too much for our purposes what solution process we endorse. Any set of rules for predicting the result of a bargaining situation, as long as they are applied consistently, will give us the relationship between the disagreement point and the final outcome that we saw in Fig. 14.4, and *that* is the most important thing game theory has to offer: bargaining power is based on having favorable outside options.

**Box 14.2: John Forbes Nash**

John Nash was born in Bluefield, West Virginia in 1928. By the time he was in high school he showed signs of an extraordinary aptitude for math and science. He received a scholarship to attend the Carnegie Institute of Technology (now called Carnegie Mellon University) in Pittsburgh and then went to Princeton University for graduate study in mathematics.

As is often the case with mathematicians, his first work was his most influential. His Ph.D dissertation was published in the form of three articles in the early 1950s, establishing an equilibrium outcome for positive-sum non-cooperative games and proposing his solution for a two-person bargaining game (which we use in this chapter). He also published important papers in the theory of algebra, and within a few years he had been awarded tenure at the Massachusetts Institute of Technology.

It was at this moment of maximum success that Nash developed the symptoms of severe mental illness that would plague him thereafter. In 1959 he was placed in a mental hospital with a diagnosis of schizophrenia, and he lost his job at MIT. For more than a decade he wandered Europe and North America, in and out of mental institutions. Eventually he settled at Princeton University, where he earned enough money to live on, and where he would spend his days alone, filling notebooks and blackboards with equations.

Eventually his skills began to come back to him, and he was able to do more advanced theoretical work in game theory and other branches of mathematics. In 1994 he was awarded a Nobel Prize in economics in conjunction with two other game theorists, John Harsanyi and Reinhard Selten. The Nobel committee singled out for recognition the path-breaking papers Nash had published from his student work at Princeton.

Nash's life was the subject of a best-selling book, *A Beautiful Mind*, by Sylvia Nasar and a film, (very) loosely based on his biography, directed by Ron Howard.

A little reflection should lead to the realization that bargaining power of the sort we have analyzed can be found almost everywhere in the economy. It explains why, in most cases, employers have more bargaining power than employees: there are usually more workers seeking jobs in our economy than jobs seeking workers. The consequences of a failure to agree (on employment) usually leaves the worker in a worse situation than the employer. (But we will also see exceptions to this rule in Chap. 16.) A large manufacturer usually has bargaining power relative to a small supplier, particularly when many suppliers are available to fill an order. Men often have more bargaining power than women in family situations, since the result of splitting up would typically be that the woman would bear the greater burden of caring for children. (This would be intensified if the couple were currently sharing their income, but if the man's job pays more than the woman's.)

As we will see, bargaining power is not everything, but it is an important factor in many economic contexts. To apply it, however, we will need a more elaborate theory than the bare-bones model we have just outlined.

---

### 14.3 Extensions to the Basic Bargaining Model

Here we will take up, one at a time, complications that students of bargaining have examined in great detail.

1. The role of time. One of the least satisfactory aspects of the simple theory is that there is no bargaining process to speak of. That is, Abe and Bev are not doing the sorts of things bargainers do (making offers and counteroffers), nor is the passage of time, an important aspect of any economic relationship, represented in any way. Indeed, the notion that the outcome that results will be the one that two people would agree on if they had an infinite amount of time and could try out any proposal, not stopping until they had found the very best, is far removed from reality.

In the last 20 years an alternative approach to the study of bargaining has been developed that has time at its core. The basic idea is this: Abe and Bev would once again negotiate, but using a procedure that describes how they make offers to each other. Specifically, one side (say Bev) begins by proposing the terms of an agreement. The other side, Abe, can either accept or propose an alternative. This round of bargaining is thought of as taking a certain amount of time, which we could measure in minutes, days, or months. Now it is Bev's turn again: she can either accept Abe's counteroffer or make one of her own. And so it goes, back and forth, until a final agreement is reached or the two parties simply give up.

The first thing to notice about this new wrinkle is that it adds a second cost to delaying agreement or failing to agree. In the simple model we began with, the cost was that the parties would be brought back to their disagreement options, assumed to be less desirable than the potential agreements available to them. Now, in addition, by rejecting offers and extending the process, they put off enjoying these benefits, and postponing a benefit is a type of cost. (Hunger due to a meal that has been put off for too many hours is a clear example.) Thus there is an interest on both sides in speeding up the negotiation. This takes the form not only of a greater willingness to accept a less-than-perfect proposal, but also (and symmetrically) a greater willingness to make proposals that will be seen as attractive by the other side. In the simplest case, when both parties have complete information about all the costs and benefits of each possible proposal, as well as the desires of both themselves and their negotiating partner, and if each puts the same value on coming to an agreement quickly, the forces cancel out, and the result is the same division of benefits that Nash predicts.

Of course, for every assumption there is another possibility. In particular, one important difference between the instantaneous process we began with and the time-consuming process we are considering now is that differences in the value of

time can lead to changes in the distribution of bargaining power. If Abe is in a hurry and needs an agreement *now*, while Bev is willing to wait a bit longer, this shifts power from Abe to Bev. This distinction has obvious economic significance. Individuals, companies and others with secure sources of income and assets to fall back on can afford to take a more relaxed attitude toward time. Those facing a lack of funds to provide for essential needs, or who are barely one step ahead of their creditors, or who live from paycheck to paycheck with little in the way of safety net—they are the ones whose bargaining power is likely to fall.

2. Long-term relationships. The simple approaches we have considered assume that each episode of bargaining is free-standing: the two parties come into it without any history (or at least none that's relevant), and they do not expect to meet again in the future. All that matters is what they can get out of the process *now*. Sometimes this is a reasonably accurate description of what happens. For instance, in many countries it is expected that buyers and sellers will haggle at the open-air markets that are common places for trade. If you are not a regular customer—if you are a tourist, for example—you will find yourself in a negotiating situation not all that different from the simple Nash model of Fig. 14.4. But in all likelihood this is the exception and not the rule. In most real-world contexts we bargain over and over with the same people: landlords, suppliers, workers or employers, etc. What difference does this make?

The answer is, quite a lot, but it is difficult to generalize. One way to think about the problem is that, in a long-term context, each bargaining proposal (or response to someone else's proposal) serves two purposes at once: it has an effect on the current negotiation, and it also sends a signal about one's intentions in future negotiations. We could imagine two different signals parties might want to send:

- Where there are large gains to cooperation, it may pay to moderate one's bargaining in the present in order to signal greater willingness to cooperate in the future. A familiar example is a close relationship between two people, such as a marriage. Each partner's well-being depends greatly on the ability of both to work together with as little friction as possible. Of course, from time to time issues come up over which there is a strong difference of opinion or interest. When deciding how forcefully to make its case, each partner, if it is acting rationally, should consider the message their behavior sends about how cooperative they intend to be in the future. The result will normally be less aggressive bargaining than would be expected between two strangers.
- Where differences in fundamental interests are likely to continue into the future, parties will be tempted to bargain even more aggressively than otherwise, in order to send the message that they are strong and determined and should not be challenged. This can sometimes be seen in negotiations between unions and managers. Each side knows there will be future negotiations, and they see an interest not only in getting a favorable agreement in the present, but also presenting a general image of toughness. This attitude, when adopted on both sides, can sometimes lead to escalating displays of aggression and ultimately conflicts that would otherwise not be in anyone's interest.

3. Behavioral complications. The bargaining models we have examined are all based on the assumption of pure self-interest: each bargainer has no objective other than getting as much benefit for him- or herself as possible, whether this is a one-time benefit (as in Nash) or a flow of benefits over time from an ongoing relationship. Assuming pure self-interest makes it easier to analyze the bargaining process (and most other aspects of economics as well), but it is not an accurate representation of how real human beings think and act.

Research into real-world bargaining behavior has grown by leaps and bounds during the past two decades, and it is hardly possible to summarize all the interesting findings. Here are just a few:

- (a) Perceptions of *fairness* play a crucial role. In many situations people will not press a bargaining advantage because the result would appear unfair. They would feel ashamed to be benefitting from the weakness of others, or from advantages they do not feel they deserve. Laboratory evidence conclusively demonstrates this in the context of the “ultimatum game”, for example. In this game, two players get to split a given amount of money if they can agree on how to divide it, but they get nothing if they can’t. The bargaining process is simple and dramatic: player 1 gets to propose a split, and player 2 can either accept or reject it. That’s all: no discussion or counter-offers. This gives player 1 a huge advantage, since he or she can propose a highly favorable split knowing that, as long as some tiny amount is left over for #2, accepting is more beneficial than rejecting. Yet this is rarely what happens, and a large percentage of offers are for an even division of the money—the triumph of pure fairness over pure self-interest.
- (b) People often act on the basis of *reciprocity* even when it does not appear to be in their direct self-interest. Reciprocity can take two forms, making gifts (in the form of favorable offers) to those who have previously given to you, and imposing punishments (such as a refusal to agree) to those seen as having violated the rules of proper behavior. Both can be costly, since they both forego potential gains that a strictly self-interested negotiator might capture.
- (c) People often adhere to *social norms* rather than push their own advantage as far as it will go. Over time, communities evolve agreed-upon solutions to many bargaining situations. Consider restaurant tipping, for instance. From a pure bargaining point of view, nearly all the advantage lies with the customer, particularly if this is not an encounter likely to be repeated. In fact, it is so one-sided that it is hard to call it bargaining at all. As a matter of self-interest, the diner, particularly one from another city, should get up from the table and leave nothing; there would be little the server could do about this. Nevertheless, each culture has a tipping norm, and most people obey it. (One of the first questions a traveler is likely to ask is, how much do people tip around here?)

(d) There are large differences in behavior across individuals. Some people appear to place a high value on fairness, others less. Some are far more reciprocal than others, or more reciprocal in one direction (such as giving) than another (punishing). There are also gender differences that have been identified in some cultures: women on average appear to be less aggressive in pursuing self-interest in bargaining situations and are more prone to reciprocating gifts than men. Why these differences exist—between individuals and across social groups—and how particular behaviors can be nurtured or discouraged is largely beyond our current knowledge, but despite (or because of) this it is one of the hottest areas for debate.

The field of behavioral economics, particularly as it illuminates bargaining, is in the midst of a revolution, and it is likely that a chapter on this topic written 5 years from now will look quite different from this one.

---

## 14.4 Bargaining Power in Action

The theories described in this chapter do not provide a magic decoding device for figuring out how bargaining works in real life, but they can be helpful in suggesting where to look for patterns. To illustrate, here are two important instances where bargaining power has changed in broad daylight, with consequences for large numbers of people.

1. Wal-Mart and its suppliers. Once upon a time in countries like the United States, manufacturers were large and retail stores were small. A few giant firms accounted for most of the production of consumer items, and many individual stores or small chains relied on them for merchandise. This gave the producers plenty of bargaining power: if one store wouldn't accept their terms, another would, but stores had few options for stocking their shelves. The result was much higher profits for the concentrated manufacturing sectors, and a scramble for much smaller profits in retail.

Not any more. Stores have gotten bigger and the chains more extensive, and no one illustrates this trend more than Wal-Mart in the US. Wal-Mart actually takes in more than 50 % of all retail dollars spent by US consumers; it is the indispensable connection between any manufacturer that wants to reach a mass market and the consumers themselves. For most manufacturers Wal-Mart accounts for such a large percentage of their sales that they simply *have* to agree to supply it. For their part, however, Wal-Mart has several manufacturers it can turn to, and it can also create its own house brand if that promises greater profits. In other words, the disagreement option for most manufacturers is much worse than it is for Wal-Mart. This gives the retail chain an enormous advantage in negotiating supply contracts.

The results are dramatic. Wal-Mart forces its suppliers to sell at lower prices and accept lower profit margins. It often withholds payment until *after* consumers have rung up their purchases. It can dictate packaging and shipping methods according to

its own convenience. It can even pressure a producer to change its product line, for instance by introducing a stripped-down version of its flagship products even when this has the potential to undermine the supplier's brand strategy. Some of the outsourcing of production to lower-cost regions like China that has characterized the US economy is the result of price pressure exerted by Wal-Mart on its supply chain.

This is not to say that the shift in bargaining power between supplier and retailer is the only source of Wal-Mart's success, but it has been one factor.

2. The global coffee market. We began this chapter with some information about the changing share of the coffee dollar going to producers, governments, exporters, and sellers. The multinational companies that brand and distribute the coffee have largely gained at the expense of the others, especially the farmers and government agencies, and this was an important aspect of the Coffee Crisis.

Simple bargaining theory has something to say about this. Two trends have altered the disagreement options facing different levels of the supply chain. First, the farmers are less organized than in the past. Before, they often sold their coffee through government-mandated marketing arrangements. In fact, it was often a marketing board representing thousands of producers, rather than each individual producer, that negotiated a price. Second, there has been increasing concentration on the part of the coffee multinationals. Over half the market in the US is accounted for by just four firms; in England the percentage is even higher, perhaps more than 90 %. These numbers alone suggest that there is an imbalance between a handful of corporations bargaining with thousands of small producers. The corporations can play one off against another, but any individual farmer has few options if no supply agreement is reached. In fact, the situation on the ground is often even more unequal, since farmers are often isolated, with little reliable information about global price trends and poor infrastructure for moving their goods to locations where they might command a higher price. There is nothing surprising, then, about the weakening position of small farmers in the coffee world; it is what we would expect given the changes that have been taking place. As this is being written, NGOs and the World Bank are promoting cooperative marketing arrangements between coffee growers to return to them some of the bargaining clout they lost when the national marketing boards were dismantled 20 years ago.

---

#### The Main Points

1. There are five kinds of power: the power to withhold, bargaining power, coercive power, institutional power and cultural power. They are not mutually exclusive.
2. Bargaining relationships are likely to arise when there are few participants to a market, when the identities of participants are known to each other, and when interaction is repeated.
3. The agreement curve is a technique for portraying the options open to two agents bargaining over the division of a good or resource. Each participant has a disagreement point, the fallback value that results from failing to agree with the other; no one would voluntarily agree to a bargain that leave him or her

worse off than the disagreement point. Any potential agreement that meets this constraint is a possible solution.

4. The Nash bargaining solution is the agreement that maximizes the product of its net benefit to the two parties, where the net benefit is the value of the agreement minus the value of the disagreement point. If the agreement curve is linear, this solution lies midway between the two disagreement values: each party gets the same net benefit. For any consistent solution process, however, it will be the case that the agreement will be more in favor of the party that has the least need for it. In other words, if someone's disagreement point falls, their bargaining power falls and they can expect a less favorable bargain.
5. A popular alternative model of the bargaining process presents it as a sequence of offers and counteroffers taking place through time. In this model, the party least able to wait (for whom the cost of delaying the agreement is greatest) has less bargaining power.
6. Repeated bargaining between the same parties is complicated; no general rule can be formulated to predict what will happen.
7. The purely rational, self-interested bargainer of standard game theory is not representative of real human beings. Most individuals are motivated by perceptions of fairness, demand reciprocity, and adhere to social norms governing bargaining situations. Moreover, there is wide variation across individuals: some attach more importance to fairness than others, are more reciprocating, etc.

► Terms to Define

Agreement curve  
 Bargaining power  
 Default option  
 Just price  
 Nash bargaining solution  
 Reciprocity (in bargaining)  
 Supply chain  
 Thin vs thick markets  
 Ultimatum game

**Questions to Consider**

1. What bargaining situations have you been in during the past month? Were these one-time situations, or were they repeating (part of a longer-term relationship)?
2. It is common for consumers and auto dealers to bargain over the price of new cars. It is not common for consumers and electronics retailers to bargain over the price of new computers. Why?
3. In one of the situations you listed in your answer to question 1, describe the disagreement options that you and your bargaining partner(s) faced. On whom did this confer the most bargaining power?

4. In the former Soviet Union there were often shortages of many consumer goods. In capitalist countries it is usually the other way around: firms produce a surplus of goods and have to convince consumers to buy them. Using the basic principles of bargaining theory, discuss who might gain and lose from each system.
5. For hundreds of years there have been regulations prohibiting grain merchants from driving up prices during times of famine. In many countries today there are still controls on the price of bread and other staples. What is the justification for this from the standpoint of bargaining theory? What disadvantages might there be with such a policy? Are there other ways to achieve the same objectives?
6. It has been said that wars could never arise unless at least one side miscalculates. Is this consistent with the bargaining theory presented in this chapter?
7. The example of Hurricane Charlie was brought up in Chap. 6. Reconsider it in light of the tension between self-interested and “fair” bargaining. Does this change your opinion? Explain.