Bargaining

A NEWLY ELECTED trade union leader went to his first tough bargaining session in the company boardroom. Nervous and intimidated by the setting, he blurted out his demand: “We want ten dollars an hour or else.”

“Or else what?” challenged the boss.

The union leader replied, “Nine fifty.”

Few union leaders are so quick to back down, and bosses need the threat of Chinese competition, not their own power, to secure wage concessions. But the situation poses several important questions about the bargaining process. Will there be an agreement? Will it occur amicably, or only after a strike? Who will concede and when? Who will get how much of the pie that is the object of the haggling?

In chapter 2, we sketched a simple story of the ultimatum game. The example illustrated the strategic principle of looking ahead and reasoning back. Many realities of the bargaining process were sacrificed in order to make that principle stand out. This chapter uses the same principle, but with more attention to issues that arise during bargaining in business, politics, and elsewhere.

We begin by recapitulating the basic idea in the context of union-management negotiation over wages. To look forward and reason backward, it helps to start at a fixed point in the future, so let us think of an enterprise with a natural conclusion, such as a hotel in a summer resort. The season lasts 101 days. Each day the hotel operates, it makes a profit of $1,000. At the beginning of the season, the employees’ union confronts the management over wages. The union presents its demand. The management either accepts this or rejects it and

returns the next day with a counteroffer. The hotel can open only after an agreement is reached.

First suppose bargaining has gone on for so long that, even if the next round leads to an agreement, the hotel can open for only the last day of the season. In theory, bargaining will not go on that long, but because of the logic of looking ahead and reasoning back, what actually happens is governed by a thought process that starts at this extreme. Suppose it is the union's turn to present its demand. At this point the management should accept anything as being better than nothing. So the union can get away with the whole $1,000.*

Now look at the day before the last day of the season, when it is the management's turn to make an offer. It knows that the union can always reject this, let the process go on to the last day, and get $1,000. Therefore the management cannot offer any less. And the union cannot do any better than get $1,000 on the last day, so the management need not offer any more on the day before.† Therefore the management's offer at this stage is clear: of the $2,000 profit over the last two days, it asks half. Each side gets $500 per day.

Next let the reasoning move back one more day. By the same logic, the union will offer the management $1,000 and ask for $2,000; this gives the union $667 per day and the management $333. We show the full process in the following table:

Successive rounds of wage bargaining

		Union's share		Management's share	
Days to go	Offer by	Total	Per day	Total	Per day
1	Union	$1,000	$1,000	$0	$0
2	Management	$1,000	$500	$1,000	$500
3	Union	$2,000	$667	$1,000	$333
4	Management	$2,000	$500	$2,000	$500
5	Union	$3,000	$600	$2,000	$400
...					
100	Management	$50,000	$500	$50,000	$500
101	Union	$51,000	$505	$50,000	$495

Each time the union makes an offer, it has an advantage, which stems from its ability to make the last all-or-nothing offer. But the advantage gets smaller as the number of rounds increases. At the start of a season 101 days long, the two

sides' positions are almost identical: $505 vs. $495. Almost the same division would emerge if the management were to make the last offer, or indeed if there were no rigid rules like one offer a day, alternating offers, *etc.*[1]

The appendix to this chapter shows how this framework generalizes to include negotiations in which there is no predetermined last period. Our restrictions to alternating offers and a known finite horizon were simply devices to help us look ahead. They become innocuous when the time between offers is short and the bargaining horizon is long—in these cases, looking ahead and reasoning backward leads to a simple and appealing rule: split the total down the middle.

There is a second prediction of the theory: the agreement will occur on the first day of the negotiation process. Because the two sides look ahead to predict the same outcome, there is no reason why they should fail to agree and jointly lose $1,000 a day. Not all instances of union-management bargaining have such a happy beginning. Breakdowns in negotiations do occur, strikes or lockouts happen, and settlements favor one side or the other. By refining our example and changing some of the premises, we can explain these facts.

THE HANDICAP SYSTEM IN NEGOTIATIONS

One important element that determines how the pie will be split is each side's cost of waiting. Although both sides may lose an equal amount of profits, one party may have other alternatives that help partially recapture this loss. Suppose that the members of the union can earn $300 a day in outside activities while negotiations with the hotel management go on. Now each time the management's turn comes, it must offer the union not only what the union could get a day later but also at least $300 for the current day. The entries in our table shift in the union's favor; we show this in a new table. Once again the agreement occurs at the season opening and without any strike, but the union does much better.

Successive rounds of wage bargaining (with outside activities)

		Union's share		Management's share	
Days to go	Offer by	Total	Per day	Total	Per day
1	Union	$1,000	$1,000	$0	$0
2	Management	$1,300	$650	$700	$350
3	Union	$2,300	$767	$700	$233
4	Management	$2,600	$650	$1,400	$350
5	Union	$3,600	$720	$1,400	$280
...					
100	Management	$65,000	$650	$35,000	$350
101	Union	$66,000	$653	$35,000	$347

This result can be seen as a natural modification of the principle of equal division, allowing for the possibility that the parties start the process with different "handicaps," as in golf. The union starts at $300, the sum its members could earn on the outside. This leaves $700 to be negotiated, and the principle is to split it evenly, $350 for each side. The union gets $650 and the management only $350.

In other circumstances the management could have an advantage. For example, it might be able to operate the hotel using scabs while the negotiations with the union go on. But because those workers are less efficient or must be paid more, or because some guests are reluctant to cross the union's picket lines, the management's profit from such operation will be only $500 a day. Suppose the union members have no outside income possibilities. Once again there will be an immediate settlement with the union without an actual strike. But the prospect of the scab operation will give the management an advantage in the negotiation, and it will get $750 a day while the union gets $250.

If the union members have an outside income possibility of $300 *and* the management can operate the hotel with a profit of $500 during negotiations, then only $200 remains free to be bargained over. They split that $200 evenly so that the management gets $600 and the union gets $400. The general idea is that the better a party can do by itself in the absence of an agreement, the larger its share of the bargaining pie will be.

MEASURING THE PIE

The first step in any negotiation is to measure the pie correctly. In the example just above, the two sides are not really negotiating over $1,000. If they

reach an agreement, they can split $1,000 per day. But if they don't reach an agreement, then the union has a fallback of $300 and the management has a fallback of $500. Thus an agreement only brings them an additional $200. In this case, the best way to think about the size of the pie is that it is $200. More generally, the size of the pie is measured by how much value is created when the two sides reach an agreement compared to when they don't.

In the lingo of bargaining, the fallback numbers of $300 for the union and $500 for management are called BATNAs, a term coined by Roger Fisher and William Ury. It stands for Best Alternative to a Negotiated Agreement.[2] (You can also think that it stands for Best Alternative to No Agreement.) It is the best you can get if you don't reach an agreement with this party.

Since everyone can get their BATNA without having to negotiate, the whole point of the negotiation is how much value can be created above and beyond the sum of their BATNAs. The best way to think about the pie is how much more value can be created beyond giving everyone his or her BATNA. This idea is both profound and deceptively simple. To see how easy it is to lose sight of BATNAs, consider the following bargaining problem adapted from a real-world case.

Two companies, one in Dallas and one is San Francisco, were using the same New York–based lawyer. As a result of coordinating their schedules, the lawyer was able to fly NY–Houston–SF–NY, a triangle route, rather than make two separate trips.

The one-way airfares were:

NY–Houston	$666
Houston–SF	$909
SF–NY	$1,243
Total:	$2,818

The total cost of the trip was $2,818. Had the lawyer done each of the trips separately, the round-trip fares would have been just double the one-way fares (as there was no time to book the trip in advance).

Our question considers how the two companies might negotiate the division of the airfare. We realize that the stakes are small here, but it is the principle we are looking for. The simplest approach would be to split the airfare in two: $1,409 to each of Houston and San Francisco.* In response to such a proposal you might well hear from Houston: we have a problem. It would have been cheaper for Houston to have paid for the round trip to Houston all by itself. That

fare is only twice $666, or $1,332. Houston would never agree to such a split.

Another approach is to have Houston pay for the NY–Houston leg, to have SF pay for the SF–NY leg, and for the two to split the Houston–SF leg. Under that approach, SF would pay $1,697.50 and Houston would pay $1,120.50.

The two companies could also agree to split the total costs proportionately, using the same ratio as their two round-trip fares. Under this plan, SF would pay $1,835, about twice as much as Houston, who would pay $983.

When faced with such a question, we tend to come up with ad hoc proposals, some of which are more reasonable than others. Our preferred approach is to start with the BATNA perspective and measure the pie. What will happen if the two companies can't agree? The fallback is that the lawyer would make two separate trips. In that case, the cost would be $1,332 to Houston and $2,486 to SF, for a total of $3,818. Recall that the triangle route cost only $2,818. This is the key point: the extra cost of doing the two round-trips over the triangle route is $1,000. That is the pie.

The value of reaching an agreement is that it creates $1,000 in savings that is otherwise lost. Each of the two companies is equally valuable in reaching that agreement. Thus, to the extent that they are equally patient in the negotiations, we would expect them to split this amount evenly. Each party saves $500 over the round-trip fare: Houston pays $832 and SF pays $1,986.

You can see that this is a much lower number for Houston than any of the other approaches. It suggests that the division between two parties should not be based on the mileage or the relative airfares. Although the airfare to Houston is smaller, that doesn't mean they should end up with less of the savings. Remember, if they don't agree to the deal the whole $1,000 is lost. We would like to think that this is a case where you might have started off with one of the alternative answers, but, having seen how to apply BATNAs and thereby measure the pie correctly, you are persuaded that the new answer is the most equitable outcome. If you started right away with Houston paying $832 and SF paying $1,986, hats off to you. It turns out that this approach to dividing costs can be traced back to the Talmud's principle of the divided cloth.[3]

In the negotiations we've looked at, the BATNAs were fixed. The union was able to get $300 and management, $500. The round-trip airfares for NY–Houston and NY–SF were given exogenously. In other cases, the BATNAs are not fixed. That opens up the strategy of influencing the BATNAs. Generally speaking, you will want to raise your BATNA and lower the BATNA of the other side. Sometimes these two objectives will be in conflict. We now turn to this subject.

THIS WILL HURT YOU MORE THAN IT HURTS ME

When a strategic bargainer observes that a better outside opportunity translates into a better share in a bargain, he will look for strategic moves that improve his outside opportunities. Moreover, he will notice that what matters is his outside opportunity *relative* to that of his rival. He will do better in the bargaining even if he makes a commitment or a threat that lowers both parties' outside opportunities, so long as that of the rival is damaged more severely.

In our example, when the union members could earn $300 a day on the outside while the management could make a profit of $500 a day using scab labor, the result of the bargaining was $400 for the union and $600 for the management. Now suppose the union members give up $100 a day of outside income to intensify their picketing, and this reduces the management's profit by $200 a day. Then the bargaining process gives the union a starting point of $200 ($300 minus $100) and the management $300 ($500 minus $200). The two starting points add up to $500, and the remaining $500 of daily profit from regular operation of the hotel is split equally between them. Therefore the union gets $450 and the management gets $550. The union's threat of hurting both (but hurting the management more) has earned it an extra $50.

Major League Baseball players employed just such a tactic in their wage negotiations in 1980. They went on strike during the exhibition season, returned to work at the start of the regular season, and threatened to strike again starting on Memorial Day weekend. To see how this "hurt the team owners more," note that during the exhibition season the players got no salaries, while the owners earned revenue from vacationers and locals. During the regular season the players got the same salary each week. For the owners, the gate and television revenues are low initially and rise substantially during and after the Memorial Day weekend. Therefore the loss of the owners *relative* to that of the players was highest during the exhibition season and again starting Memorial Day weekend. It seems the players knew the right strategy.[4]

The owners gave in just before the second half of the threatened strike. But the first half actually occurred. Our theory of looking ahead and reasoning back is clearly incomplete. Why is it that agreements are not always reached before any damage is done—why are there strikes?

BRINKMANSHIP AND STRIKES

Before an old contract expires, the union and the firm begin the negotiations

for a new labor contract. But there is no sense of urgency during this period. Work goes on, no output is sacrificed, and there is no apparent advantage to achieving an agreement sooner rather than later. It would seem that each party should wait until the last moment and state its demand just as the old contract is about to expire and a strike looms. That does happen sometimes, but often an agreement is reached much sooner.

In fact, delaying agreement can be costly even during the tranquil phase when the old contract still operates. The process of negotiation has its own risk. There can be misperception of the other side's impatience or outside opportunities, tension, personality clashes, and suspicion that the other side is not bargaining in good faith. The process may break down despite the fact that both parties want it to succeed.

Although both sides may want the agreement to succeed, they may have different ideas about what constitutes success. The two parties do not always look forward and see the same end. They may not have the same information or share the same perspective, so they see things differently. Each side must make a guess about the other's cost of waiting. Since a side with a low waiting cost does better, it is to each side's advantage to claim its cost is low. But these statements will not be taken at face value; they have to be proven. The way to prove one's waiting costs are low is to begin incurring the costs and then show you can hold out longer, or to take a greater risk of incurring the costs—lower costs make higher risks acceptable. It is the lack of a common view about where the negotiations will end that leads to the beginning of a strike.

Think of a strike as an example of signaling. While anyone can say that he or she has a low cost of going on strike or taking on a strike, to actually do so is the best proof possible. As always, actions speak louder than words. And, as always, conveying information by a signal entails a cost, or sacrifice of efficiency. Both the firm and the workers would like to be able to prove their low costs without having to create all the losses associated with a work disruption.

The situation is tailor-made for the exercise of brinkmanship. The union could threaten an immediate breakdown of talks followed by a strike, but strikes are costly to union members as well. While time for continued negotiation remains, such a dire threat lacks credibility. But a smaller threat can remain credible: tempers and tensions are gradually rising, and a breakdown may occur even though the union doesn't really want it to. If this bothers the management more than it bothers the union, it is a good strategy from the union's perspective. The argument works the other way around, too; the strategy of brinkmanship is a weapon for the stronger of the two parties—namely, the one that fears a breakdown less.

Sometimes wage negotiations go on after the old contract has expired but without a strike, and work continues under the terms of the old contract. This might seem to be a better arrangement, because the machinery and the workers are not idle and output is not lost. But one of the parties, usually the union, is seeking a revision of the terms of the contract in its favor, and for it the arrangement is singularly disadvantageous.* Why should the management concede? Why should it not let the negotiations spin on forever while the old contract remains in force de facto?

Again the threat in the situation is the probability that the process may break down and a strike may ensue. The union practices brinkmanship, but now it does so after the old contract has expired. The time for routine negotiations is past. Continued work under an expired contract while negotiations go on is widely regarded as a sign of union weakness. There must be some chance of a strike to motivate the firm to meet the union's demands.

When the strike does happen, what keeps it going? The key to commitment is to reduce the threat in order to make it credible. Brinkmanship carries the strike along on a day-by-day basis. The threat never to return to work would not be credible, especially if the management comes close to meeting the union's demands. But waiting one more day or week is a credible threat. The losses to the workers are smaller than their potential gains. Provided they believe they will win (and soon), it is worth their while to wait. If the workers are correct in their beliefs, management will find it cheaper to give in and in fact should do so immediately. Hence the workers' threat would cost them nothing. The problem is that the firm may not perceive the situation the same way. If it believes the workers are about to concede, then losing just one more day's or week's profits is worth getting a more favorable contract. In this way, both sides continue to hold out, and the strike continues.

Earlier, we talked about the risk of brinkmanship as the chance that both sides would fall together down the slippery slope. As the conflict continues, both sides risk a large loss with a small but increasing probability. It is this increasing exposure to risk that induces one side to back down. Brinkmanship in the form of a strike imposes costs differently, but the effect is the same. Instead of a small chance of a large loss, there is a large chance, even certainty, of a small loss when a strike begins. As the strike continues unresolved, the small loss grows, just as the chance of falling off the brink increases. The way to prove determination is to accept more risk or watch strike losses escalate. Only when one side discovers that the other is truly the stronger does it decide to back down. Strength can take many forms. One side may suffer less from waiting, perhaps because it has valuable alternatives; winning may be very important,

perhaps because of negotiations with other unions; losing may be very costly, so that the strike losses look smaller.

Brinkmanship applies to the bargaining between nations as well as that between firms. When the United States tries to get its allies to pay a greater share of the defense costs, it suffers from the weakness of negotiating while working under an expired contract. The old arrangement in which the Americans bear the brunt of the burden continues in the meantime, and the U.S. allies are happy to let the negotiations drag on. Can—and should—the United States resort to brinkmanship?

Risk and brinkmanship change the process of bargaining in a fundamental way. In the earlier accounts of sequences of offers, the prospect of what would come later induced the parties to reach an agreement on the first round. An integral aspect of brinkmanship is that sometimes the parties do go over the brink. Breakdowns and strikes can occur. They may be genuinely regretted by both parties but may acquire a momentum of their own and last surprisingly long.

SIMULTANEOUS BARGAINING OVER MANY ISSUES

Our account of bargaining has so far focused on just one dimension, namely the total sum of money and its split between the two sides. In fact, there are many dimensions to bargaining: the union and management care not just about wages but health benefits, pension plans, conditions of work, and so on. The United States and its trading partners care not just about total CO_2 emissions but how they are allocated. In principle, many of these are reducible to equivalent sums of money, but with one important difference: each side may value the items differently.

Such differences open up new possibilities for mutually acceptable bargains. Suppose the company is able to secure group health coverage on better terms than the individual workers would obtain on their own—say, $1,000 per year instead of $2,000 per year for a family of four. The workers would rather have health coverage than an extra $1,500 a year in wages, and the company would rather offer health coverage than an extra $1,500 in wages, too. It would seem that the negotiators should throw all the issues of mutual interest into a common bargaining pot, and exploit the difference in their relative valuations to achieve outcomes that are better for everyone. This works in some instances; for example, broad negotiations toward trade liberalization in the General Agreement on Tariffs and Trade (GATT) and its successor, the World Trade

Organization (WTO), have had better success than ones narrowly focused on particular sectors or commodities.

But joining issues together opens up the possibility of using one bargaining game to generate threats in another. For example, the United States may have more success in extracting concessions in negotiations to open up the Japanese market to its exports if it threatened a breakdown of the military relationship, thereby exposing Japan to a risk of Korean or Chinese aggression. The United States has no interest in actually having this happen; it would be merely a threat that would induce Japan to make the economic concession. Therefore, Japan would insist that the economic and military issues be negotiated separately.[5]

THE VIRTUES OF A VIRTUAL STRIKE

Our discussion of negotiation has also left out the effect on all the players who aren't a party to the deal. When UPS workers go on strike, customers end up without packages. When Air France baggage handlers go on strike, holidays are ruined. A strike hurts more than the two parties negotiating. A lack of agreement on global warming and CO_2 emissions could prove devastating to all future generations (who don't get a seat at the table).

But the parties negotiating have to be willing to walk away in order to demonstrate the strength of their BATNA or to hurt the other side more. Even for an ordinary strike, the collateral damage can easily eclipse the size of the dispute. Until President Bush stepped in on October 3, 2002, invoking the Taft-Hartley Act, the ten-day dock-worker lockout disrupted the U.S. economy to the tune of more than $10 billion. The conflict was over $20 million of productivity enhancements. The collateral damage was *500 times* larger than the amounts that the workers and managers were squabbling over.

Is there some way that the two parties can resolve their differences without imposing such large costs on the rest of us? It turns out that for more than fifty years there has been a clever idea to virtually eliminate all of the waste of strikes and lockouts without altering the relative bargaining power of labor and management.[6] Instead of a traditional strike, the idea is to have a *virtual* strike (or virtual lockout), in which the workers keep working as normal and the firm keeps producing as normal. The trick is that during the virtual strike neither side gets paid.

In a regular strike, workers lose their wages and an employer loses its profits. So during a virtual strike, the workers would work for nothing and the employer would give up all of its profits. Profits might be too hard to measure

and short-term profits might also understate the true cost to the firm. Instead, we have the firm give up all of its revenue. As to where the money would go, the revenue could go to Uncle Sam or a charity. Or, the product could be free so that the revenues would be given to customers. During a virtual strike, there is no disruption to the rest of economy. The consumer is not left stranded without service. Management and labor feel the pain and thus have an incentive to settle, but the government, charities, or customers get a windfall.

An actual strike (or a lockout that the management initiates to preempt a strike) can permanently destroy consumer demand and risk the future of the whole enterprise. The National Hockey League imposed a lockout in response to a threatened strike during the 2004–5 season. The whole season was lost, there was no Stanley Cup, and it took a long while for attendances to recover after the dispute was finally settled.

The virtual strike is not just a wild idea waiting to be tested. During World War II, the navy used a virtual strike to settle a labor dispute at the Jenkins Company valve plant in Bridgeport, Connecticut. A virtual strike arrangement was also used in a 1960 Miami bus strike. Here, the customers got a free ride, literally.

In 1999, Meridiana Airline's pilots and flight attendants staged Italy's first virtual strike. The employees worked as usual but without being paid, while Meridiana donated the receipts from its flights to charities. The virtual strike worked just as predicted. The flights that were virtually struck were not disrupted. Other Italian transport strikes have followed the Meridiana lead. In 2000, Italy's Transport Union forfeited 100 million lire from a virtual strike carried out by 300 of its pilots. The virtual pilots' strike provided a public relations opportunity, as the strike payments were used to buy a fancy medical device for a children's hospital. Instead of destroying consumer demand, as in the 2004–5 NHL lockout, the virtual strike windfall provides an opportunity to increase the brand's reputation.

Somewhat perversely, the public relations benefit of virtual strikes may make them harder to implement. Indeed, a strike is often designed to inconvenience consumers so that they put pressure on management to settle. Thus asking an employer to forfeit its profits may not replicate the true costs of a traditional strike. It is notable that in all four historical examples, management agreed to forfeit more than its profits—and instead forfeited its entire gross revenue on all sales during the duration of the strike.

Why would workers ever agree to work for nothing? For the same reason that workers are willing to strike now—to impose pain on management and to prove that they have a low cost of waiting. Indeed, during a virtual strike, we

might expect to see labor work *harder* because every additional sale represents additional pain to the manufacturer, who has to forfeit the entire revenue on the sale.

Our point is to replicate the costs and benefits of the negotiation to the parties involved while at the same time leaving everyone else unharmed. So long as the two sides have the same BATNAs in the virtual strike as they do in the real one, they have no advantage in employing the real strike over a virtual one. The right time to go virtual is when the two sides are still talking. Rather than wait until the strike is real, labor and management might agree in advance to employ a virtual strike in the event their next contract negotiations fail. The potential gains from eliminating the entire inefficiency of traditional strikes and lockouts justify efforts to experiment with this new vision for managing labor conflict.

CASE STUDY: 'TIS BETTER TO GIVE THAN TO RECEIVE?

Recall our bargaining problem in which a hotel's management and its labor were negotiating over how to divide the season's profits. Now, instead of labor and management alternating offers, imagine that *only* the management gets to make offers, and labor can only accept or reject.

As before, the season lasts 101 days. Each day the hotel operates, it makes a profit of $1,000. Negotiations start at the beginning of the season. Each day, the management presents its offer, which is either accepted or rejected by labor. If accepted, the hotel opens and begins making money, and the remaining profits are split according to the agreement. If rejected, the negotiations continue until either an offer is accepted or the season ends and the entire profits are lost.

The following table illustrates the declining potential profits as the season progresses. If both labor and management's only concern is to maximize its own payoff, what do you expect will happen (and when)? If you were labor, what would you do to improve your position?

Wage bargaining—Management makes all offers

Days to go	Offer by	Total profits to divide	Amount offered to labor
1	Management	$1,000	?
2	Management	$2,000	?
3	Management	$3,000	?
4	Management	$4,000	?
5	Management	$5,000	?
…			
100	Management	$100,000	?
101	Management	$101,000	?

Case Discussion

In this case, we expect the outcome to differ substantially from 50:50. Because management has the sole power to propose, it is in the stronger bargaining position. Management should be able to get close to the entire amount and reach agreement on the first day.

To predict the bargaining outcome, we start at the end and work backward. On the last day there is no value in continuing, so labor should be willing to accept any positive amount, say $1. On the penultimate day, labor recognizes that rejecting today's offer will bring only $1 tomorrow; hence they prefer to accept $2 today. The argument continues right up to the first day of the season. Management proposes to give labor $101, and labor, seeing no better alternative in the future, accepts. This suggests that in the case of making offers, 'tis better to give than to receive.

This analysis clearly exaggerates management's true bargaining power. Postponing agreement, even by one day, costs management $999 and labor only $1. To the extent that labor cares not only about its payments but also how these payments compare to management's, this type of radically unequal division will not be possible. But that does not mean we must return to an even split. Management still has more bargaining power. Its goal should be to find the minimally acceptable amount to give to labor so that labor prefers getting that amount over nothing, even though management may get more. For example, in the last period, labor might be willing to accept $200 while management gets $800 if labor's alternative is zero. If so, management can perpetuate a 4:1 split throughout each of the 101 days and capture 80 percent of the total profit.

Wage bargaining—Management makes all offers

Days to go	Offer by	Total profits to divide	Amount offered to labor
1	Management	$1,000	$1
2	Management	$2,000	$2
3	Management	$3,000	$3
4	Management	$4,000	$4
5	Management	$5,000	$5
…			
100	Management	$100,000	$100
101	Management	$101,000	$101

The value of this technique for solving bargaining problems is that it suggests some of the different sources of bargaining power. Splitting the difference or even division is a common but not universal solution to a bargaining problem. Look forward and reason backward provides a reason why we might expect to find unequal division. Yet there is reason to be suspicious of the look forward and reason backward conclusion. What if you try it and it doesn't work? Then what should you do?

The possibility that the other side could prove your analysis wrong makes this repeated version of the game different from the one-shot version. In the one-shot version of divide the $100, you can assume that the receiver will find it enough in his interest to accept $20 so that you can get $80. If you end up wrong in this assumption, the game is over and it is too late to change your strategy. Thus the other side doesn't have an opportunity to teach you a lesson with the hope of changing your future strategy. In contrast, when you play 101 iterations of the ultimatum game, the side receiving the offer might have an incentive to play tough at first and thereby establish that he is perhaps irrational (or at least has a strong conviction for the 50:50 norm).*

What should you do if you offer an 80:20 split on day one and the other side says no? This question is easiest to answer in the case where there are only two days total so that the next iteration will be the last. Do you now think that the person is the type who will reject anything other than 50:50? Or do you think that was just a ruse to get you to offer 50:50 in the final round?

If the other party says yes, he will get 200 for both days, for a total of 400. Even a cold, calculating machine would say no to 80:20 if he thought that doing

so would get him an even split in the last period, or 500. But if this is just a bluff, you can stick with 80:20 in the final round and be confident that it will be accepted.

The analysis gets more complicated if your initial offer was 67:33 and that gets turned down. Had the receiver said yes, he would have ended up with a total of 333 for two days, or 666. But now that he's said no, the best he can reasonably hope for is a 50:50 split in the final round, or 500. Even if he gets his way, he will end up worse off. At this point, you have some evidence that this isn't a bluff. Now it might well make sense to offer 50:50 in the final round.

In sum, what makes a multiround game different from the one-shot version, even if only one side is making all the offers, is that the receiving side has an opportunity to show you that your theory isn't working as predicted. At that point, do you stick with the theory or change your strategy? The paradox is that the other side will often gain by appearing to be irrational, so you can't simply accept irrationality at face value. But they might be able to do so much damage to themselves (and to you along the way) that a bluff wouldn't help them. In that case, you might very well want to reassess the objectives of the other party.

APPENDIX: RUBINSTEIN BARGAINING

You might think that it is impossible to solve the bargaining problem if there is no end date to the game. But through an ingenious approach developed by Ariel Rubinstein, it is possible to find an answer.[7]

In Rubinstein's bargaining game, the two sides alternate making offers. Each offer is a proposal for how to divide the pie. For simplicity, we assume that the pie is of size 1. A proposal is something like (X, 1–X). The proposal describes who gets what; thus if X = 3/4, that means 3/4 for me, 1/4 for you. As soon as one side accepts the other's proposal, the game is over. Until then, the offers alternate back and forth. Turning down an offer is expensive, as this leads to a delay in reaching an agreement. Any agreement that the parties reach tomorrow would be more valuable if reached today. An immediate settlement is in their joint best interest.

Time is money in many different ways. Most simply, a dollar received earlier is worth more than the same dollar received later, because it can be invested and earn interest or dividends in the meantime. If the rate of return on investments is 10 percent a year, then a dollar received right now is worth $1.10 received a year later. The same idea applies to union and management, but there are some additional features that may add to the impatience factor. Each week the

agreement is delayed, there is a risk that old, loyal customers will develop long-term relationships with other suppliers, and the firm will be threatened with permanent closure. The workers and the managers will then have to move to other jobs that don't pay as well, the union leaders' reputation will suffer, and the management's stock options will become worthless. The extent to which an immediate agreement is better than one a week later is the probability that this will come to pass in the course of the week.

Just as with the ultimatum game, the person whose turn it is to make the proposal has an advantage. The size of the advantage depends on the degree of impatience. We measure impatience by how much is left if one does the deal in the next round rather than today. Take the case where there is an offer each week. If a dollar next week is worth 99¢ today, then 99 percent of the value remains (99¢ in the hand is worth a $1 in next week's bush). We represent the cost of waiting by the variable δ. In this example, $\delta = 0.99$. When δ is close to one, such as 0.99, then people are patient; if δ is small, say 1/3, then waiting is costly and the bargainers are impatient. Indeed, with $\delta = 1/3$, two-thirds the value is lost each week.

The degree of impatience will generally depend on how much time elapses between bargaining rounds. If it takes a week to make a counteroffer, then perhaps $\delta = 0.99$. If it only takes a minute, then $\delta = 0.999999$, and almost nothing is lost.

Once we know the degree of impatience, we can find the bargaining division by considering the least one might ever accept and the most one will ever be offered. Is it possible that the least amount you would ever accept is zero? No. Say it were and the other side offers you zero. Then you know that if you were to turn down the zero today and it comes time to make your counteroffer tomorrow, you can offer the other side δ and he will accept. He will accept because he would rather get δ tomorrow than have to wait one more period to get 1. (He would only get 1 in his best-case scenario that you accept 0 in two periods.) So once you know that he will surely take δ tomorrow, that means you can count on $1–\delta$ tomorrow, and so you should never accept anything less than $\delta(1–\delta)$ today. Hence neither today nor in two periods should you accept zero.*

The argument wasn't fully consistent, in that we found the minimal amount you would accept assuming that you would take zero in two periods. What we really want to find is the minimal amount you would accept where that number holds steady over time. What we are looking for is the number such that when everyone understands this is the least you will ever accept, it leads you to a position where you should accept nothing less.

Here is how we solve that circular reasoning. Assume that the worst (or

lowest) division you will ever accept gives you L, where L stands for lowest. To figure out what L must be, let's imagine that you decide to turn down today's offer in order to make a counteroffer. As you contemplate possible counteroffers, you can anticipate that the other side can never hope for more than 1–L when it is their turn again. (They know you won't accept less than L, and so they cannot get more than 1–L.) Since that is the best they can do two periods later, they should accept δ(1–L) tomorrow.

Thus today when you are contemplating accepting their offer, you can be confident that, were you to reject their offer today and counter with δ(1–L) tomorrow, they would accept. Now we are almost done. Once you know that you can always get them to accept δ(1–L) tomorrow, that leaves you 1–δ(1–L) tomorrow for sure.

Therefore, you should never take anything today less than

$$\delta(1-\delta(1-L)).$$

That gives us a minimum value for L:

$$L > \delta(1-\delta(1-L))$$

or

$$L \geq \frac{\delta(1-\delta)}{(1-\delta^2)} = \frac{\delta}{(1+\delta)}$$

You should never accept anything less than δ/(1 + δ), because you can get more by waiting and making a counteroffer that the other side is sure to accept. What is true for you is also true for the other side. By the same logic, the other side will also never accept less than δ/(1 + δ). That tells us what the most you can ever hope for is.

Using M for most, let's look for a number that is so large you should never turn it down. Since you know that the other side will never accept less than δ/(1 + δ) next period, in the best possible case you can get at most 1–δ/(1 + δ) = 1/(1 + δ) next period. If that is the best you can do next period, then today you should always accept δ(1/(1 + δ)) = δ/(1 + δ).

So we have

$$L \geq \frac{\delta}{(1+\delta)}$$

and

$$M \leq \frac{\delta}{(1+\delta)}.$$

That means that the least you will ever accept is $\delta/(1 + \delta)$ and that you will always accept anything at or above $\delta/(1 + \delta)$. Since these two amounts are exactly the same, that is what you will get. The other side won't offer less, as you will turn it down. They won't offer you more, as you will surely accept $\delta/(1 + \delta)$.

The division makes sense. As the time period between offers and counteroffers shrinks, it is reasonable to say that participants are less impatient; or, mathematically, δ gets close to 1. Look at the extreme case, where $\delta = 1$. Then the proposed division is

$$\frac{\delta}{(1+\delta)} = \frac{1}{2}.$$

The pie is split evenly between the two sides. If waiting a turn is essentially costless, then the person who goes first doesn't have any advantage, and so the division is 50:50.

At the other extreme, imagine that the pie all disappears if the offer isn't accepted. This is the ultimatum game. If the value of an agreement tomorrow is essentially zero, then $\delta = 0$, and the split is (0, 1), just as in the ultimatum game (with all the caveats, too).

To take an intermediate case, imagine that time is of the essence so that each delay loses half the pie, $\delta = 1/2$. Now the division is

$$\frac{\delta}{(1+\delta)} = \frac{\frac{1}{2}}{(1+\frac{1}{2})} = \frac{1}{3}.$$

Think of it this way. The person making me an offer has a claim to all of the pie that will be lost if I say no. That gives him 1/2 right there. Of the half that remains, you can get half of that or 1/4 total, as this amount would be lost if he doesn't accept your offer. Now after two rounds, he will have collected 1/2 and you will have 1/4 and we are back to where we started. Thus in each pair of offers, he can collect twice as much as you, leading to the 2:1 division.

The way we solved the game, the two sides are equally patient. You can use this same approach to find a solution when the two parties have differing costs of waiting. As you might expect, the side that is more patient gets a bigger slice of the pie. Indeed, as the time period between offers gets shorter, the pie is split in the ratio of waiting costs. Thus if one side is twice as impatient as the other, it gets one-third of the pie, or half as much as the other.*

The fact that the greater share in bargaining agreements goes to the more patient side is unfortunate for the United States. Our system of government, and its coverage in the media, fosters impatience. When negotiations with other nations on military and economic matters are making slow progress, interested lobbyists seek support from congressmen, senators, and the media, who pressure the administration for quicker results. Our rival nations in the negotiations know this very well and are able to secure greater concessions from us.