48

RAIFFA

4

Analytical Models and Empirical Results

With Elmtree House as a basis, we can now simplify and abstract. Later we will begin building up the complexities.

Consider the case in which two bargainers must jointly decide on a determinate value of some continuous variable (like money) that they can mutually adjust. One bargainer wants the value to be high —the higher the better—whereas the other bargainer wants the value to be low—the lower the better. We could label these agents "high aspirer" and "low aspirer," but for our purposes "buyer" and "seller" will be sufficient, even though the context we'll be dealing with is much broader than that consisting of simple business transactions in which there is an actual seller and buyer.

To simplify matters, let's assume that each bargaining agent is monolithic: he or she does not have to convince the members of some constituency that they should ratify the agreement. Let's also assume that the bargaining agents are primarily concerned about this deal only, that linkages to similar problems over repetitive plays, or linkages to other outstanding problems, are minimal—or, better yet, are nil. Setting precedents, cashing credits for past favors, and log-rolling between problems are not appropriate concerns. Time is a more troublesome matter. We shall try at first to deemphasize the role of time, or at most to keep it only informally in mind.

The two agents come together to bargain. The setting, the language, the costumes are all irrelevancies for us. We'll assume that the bargainers are honorable people—at least according to the code of ethics of our time—and we shall also assume that contracts made are enforceable and inviolable. No neutral third-party intervenors

are present to assist the bargainers. We'll also assume a single-threat environment: at most, any party can threaten only to break off negotiations and revert to the status quo before bargaining. The bargaining milieu can be classified as nonstrident.

Taking our cues from the Elmtree House illustration, we shall assume that each bargaining party has reflected on the decision problem he or she faces if no contract is made. Each has tried to determine his BATNA, or best alternative to a negotiated agreement.[1] We shall assume that by analyzing the consequences of no agreement, each bargainer establishes the threshold value that he or she needs. The seller has a reservation price, s, that represents the very minimum he will settle for; any final-contract value, x^*, that is less than s represents a situation for the seller that is worse than no agreement. If x^* is greater than s, then we can think of $x^* - s$ as the seller's surplus. The seller wants to maximize his surplus.[2] The buyer has some reservation price, b, that represents the very maximum she will settle for; any final-contract price, x^*, that is greater than b represents a situation for the buyer that is worse than no agreement. If x^* is less than b, then we can think of $b - x^*$ as the buyer's surplus[3] (see Figure 2).

If $b < s$—that is, if the maximum price the buyer will settle for is lower than the minimum price the seller will settle for—there is no possible zone of agreement. However, if $s < b$, then the *zone of agreement* (for the final contract x^*) is the interval from s to b. Suppose that the final agreement is some value x^* where x^* is between s and b; the buyer's surplus value is then $b - x^*$ and the seller's surplus value is $x^* - s$. The sum of the surplus values is $b - s$, which is independent of the intervening x^* value. In this sense, the "game"—if we think of the bargaining problem as a game—appears to be constant-sum (in surplus values). But not quite, because if $s < b$ (where a potential zone of agreement exists), the parties still might not come to an agreement—they might not agree to settle for a mutually acceptable x^* in the zone of agreement. So at most we can only think of this as a quasi-constant-sum game. To make it

1. I am indebted to Fisher and Ury (1981) for this term.
2. In the Elmtree House case, Steve's reservation price, s, was \$220,000. The bargainers settled at $x^* =$ \$325,000, so Steve, as the seller, had a surplus of \$105,000.
3. In the Elmtree House case we were not privy to Wilson's reservation price. Let us suppose that it was \$400,000. Then $b =$ \$400,000 and the buyer's surplus would have been \$75,000.

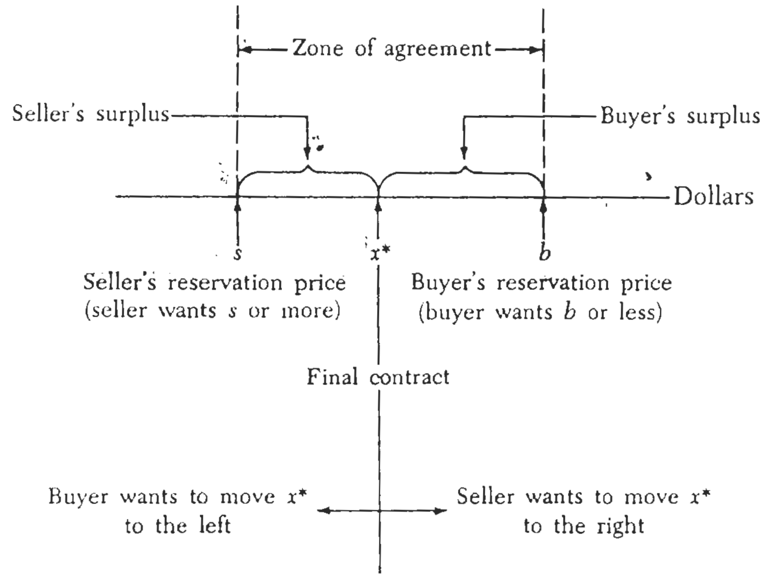

Figure 2. The geometry of distributive bargaining. (Note: If $b < s$, there is no zone of agreement.)

even more "quasi," the players generally do not know the size of the pie, $b - s$, that they have to divide.

In the abstraction we shall develop, each bargainer knows his or her reservation price, but has only probabilistic information about the other party's reservation price. Very often in practice the parties have but an imprecise feel for their own reservation price and make no formal attempt to assess a probability distribution of the other party's reservation price.

If we take the asymmetric point of view of one of the bargainers —say, the seller—the seller would be well advised before the negotiations start to ascertain s and to probabilistically assess $\tilde{b}$.[4] During the negotiation, the seller wants to periodically reassess $\tilde{b}$, at least informally; but he also wants to lead the buyer to think that $\tilde{s}$ is

4. I use the convention of a tilde to denote an uncertain quantity, or random variable. Thus, the seller knows s but assesses a distribution for $\tilde{b}$; the buyer assesses $\tilde{s}$ but knows b. In the Elmtree House case the seller, Steve, knows that s = \$220,000 and his assessment for the uncertain buyer's reservation, $\tilde{b}$, was depicted in Figure 1. Wilson, the buyer, would know b and assess $\tilde{s}$.

higher than it really is. The seller should also be aware that the buyer may be analogously motivated—that is, the buyer wants to make the seller think that b is lower than it really is. To what lengths a player might be willing to go to mislead his or her quasi-adversary (I say "quasi" since we are not discussing a strictly constant-sum game) depends on the culture. In some cultures, it is acceptable to marshal forcefully, but truthfully, all the arguments for one's own side and to avoid giving gratuitous help to the other side. In other cultures it is acceptable to exaggerate or even to bend the truth here and there—but not too much. In still other cultures a really big whopper, if accomplished with flair and humor, is something to brag about and not to hide after the fact, especially if it is successful.

A simple laboratory bargaining problem can be introduced with less than one page of confidential instructions to the seller and buyer.[5] The context is the sale of a used car, the Streaker, and the setting is dated to justify a seller's reservation price of \$300 and a buyer's reservation price of \$550. The instructions to each give only the vaguest of hints about the other person's *RP*. The challenge for a buyer is to get a good deal for herself, and she will be judged in terms of how well she has done in comparison to other buyers in an identical situation; the seller is judged similarly, in comparison to other sellers. This is like a duplicate bridge scoring system.

Players who put themselves in the role of one or the other of these negotiators will naturally ask a number of questions. What analyses should be done? What bargaining ploys seem to work? Should I open first with an offer? If I open first, how extreme should I be? Am I better off giving a reasonable value that would yield me a respectable surplus and remaining firm, or should I start with a more extreme value and pace my concessions with those of the other party? What is a reasonable pattern of concessions? Our data indicate that in this situation most pairs of negotiators come to an agreement.

A typical pattern of concessions is depicted in Figure 3, where s_1, b_1, s_2, b_2, and so on represent the prices ~~successively proposed~~ by the seller and buyer. I call this pattern "the negotiation dance." The seller might open with a value of \$700 (s_1 in the figure); the

5. I am indebted to John Hammond for this example.

50

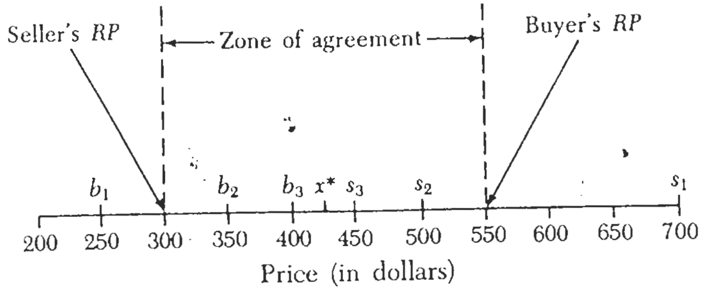

Figure 3. The negotiation dance (x^* = final-contract price).

buyer retorts with b_1 = \$250; then in succession come s_2 = \$500 (breaking the buyer's *RP*), b_2 = \$300 (breaking the seller's *RP*), s_3 = \$450, b_3 = \$400, and a final-contract price of x^* = \$425. Would x^* be higher if s_1 were \$900 instead of \$700? If so, why not make s_1 = \$2,000?

racheting

Our data yielded a number of interesting findings. First, the final contracts ranged over the entire zone of agreement, from \$300 to \$550. A sprinkling (less than 1 percent) of cases were settled out of the zone of agreement for a value less than \$300 or more than \$550; the subjects in these cases misinterpreted the directions. In some cases, but surprisingly few (around 3 percent), agreement was never achieved.

Second, the average of the final contracts was \$415 with a standard deviation of 52, indicating a surprising spread of outcomes. The average opening offer of the sellers was \$525 (standard deviation of 116); the average opening offer of the buyers was \$261 (standard deviation 112).

Third, the Boulware[6] strategy of making a reasonable opening and remaining firm works sometimes, but more often than not it antagonizes the other party, and many of the no-agreements resulted from this strategy. Advice: don't embarrass your bargaining partner by forcing him or her to make all the concessions.

Too firm on opening

Fourth, once two offers are on the table (s_1 and b_1), the best prediction of the final contract is the midpoint, $(s_1 + b_1)/2$—provided that the midpoint falls within the zone of agreement. If the midpoint falls outside this zone, then it's hard to predict where the final contract will fall. It is not true that x^* will be near the reservation price that is closer to the midpoint. The reason is that the concessions will have to be lopsided, and it's hard to predict the consequences. Thus, if b_1 = \$250 and s_1 = \$2,000, with the midpoint being \$1,125, the seller is going to be forced to make huge concessions and x^* might end up closer to \$300 than to \$550.

6. Lemuel Boulware, former vice-president of the General Electric Company, rarely made concessions in wage negotiations; he started with what he deemed to be a fair opening offer and held firm. This is commonly referred to as Boulwarism.

Fifth, from a linear regression analysis it appears that if the buyer's opening bid is held constant, then on the average adding \$100 to the opening bid of the seller nets an increase of about \$28 to the final contract. If the seller's opening bid is held constant, then on the average subtracting \$100 from the opening bid of the buyer nets a decrease of about \$15 from the final contract.

With one group of 70 subjects I ran a variation of the Streaker experiments with some fascinating but inconclusive results. In the variation, the instructions to the buyers were the same: as in the original experiments, they still had a reservation price of \$550. But the instructions to the sellers were altered: they still had to get at least \$300, but they were told not to try to get as much as possible because of the desirability of later amicable relationships with the buyers. The sellers were told that they would receive a maximum score if they could sell the car for \$500 and that every dollar above \$500 would detract from their score; a sale of x dollars above \$500 would yield them the same satisfaction as x dollars below \$500. Thus, for example, a score of \$525 would be equivalent to a score of \$475. Of course, the buyers were not aware of these confidential instructions to the sellers.

Surprisingly, the sellers did better playing this variation with benevolent intentions toward the buyer than they did with aggressive intentions to squeeze out as much as possible. In the variation, the average price for the car was \$457 instead of \$415. One reason for this might have been that in the original version, the sellers were told only to get more than \$300 and they did not have any target figure. In the variation, they were told that the best achievable value was \$500 and this became a target value. Indeed, the sellers' opening offers averaged higher in the variation than in the original exercise (\$592 versus \$525). In the variation, the sellers came down faster from high values (above \$500) but they became more reluctant to reduce their prices as they pierced their \$500 aspiration level, thus making it seem to the buyers that they were approaching their reservation values.

Some sellers said that they felt some qualms when they let themselves be bargained back from $600 to $500, knowing that this was the direction in which they wanted to move. Some sellers told the buyers that they thought $500 was the fair price and that they did not want to get a higher value; but the buyers they were bargaining with tended not to believe them, and these sellers on the average hurt themselves.

Analytical elaboration. It would be interesting to run some additional variations, such as the following:

1. Give the seller a specified reservation value of $300. Hint at a "fair" or "reasonable" value of $500, but suggest that getting more would be still better. Let the buyer remain with a reservation value of $550.
2. Go back to the first variation in which the seller needs $300 and wants $500, and in which getting x dollars above $500 is like getting x dollars below $500. Push the buyer's reservation value below $500—to, say, $450. It is likely that some sellers will get confused between what they absolutely need ($300) and what they aspire to ($500).
3. Make the seller's reservation value of $300 more vague. Tell the seller, for example, that if the negotiations fall through he will have to sell the car to a dealer, who will offer him one of the three equally likely values: $200, $300, or $400. Since $300 is the expected value of the alternative, it should serve as the effective reservation value for the present negotiation; but in this case the seller might bargain more aggressively for values over $300.

In distributive bargaining, successive offers by the seller are usually monotonely decreasing, whereas those by the buyer are monotonely increasing. Indeed, one of the principles of good-faith bargaining is that once a concession is made, it is not reversed. The following anecdote depicts an amusing counterexample.[7]

7. I am indebted to Philburn Ratoosh for this anecdote.

5

Larry M. gazed somewhat disinterestedly at a briefcase displayed in the window of a luggage store in Mexico City. The proprietor, who spoke English, approached him outside the store and said, "Are you interested in that briefcase?"

"No, I'm just window shopping," Larry replied.

"You can have it for $15. That's a good buy."

Larry had a perfectly acceptable briefcase and said that he was not interested.

"All right, you can have it for $14." Declined.

"How about $13? That's a fantastic buy." Declined.

At this point, Larry became interested. He didn't want the briefcase, but he was curious about how far the shopkeeper would lower his price. So he stayed around saying nothing.

"I'll sell it for $12. You can't get anything like this at that price in the States." Declined.

"All right, since you're obviously a tourist with a limited budget, just for you I'll give it to you for $11." Declined.

"My final offer: if you promise not to tell anybody, I'll sell it to you for $12."

"Hey, wait a second," interrupted Larry. "You just offered it to me for $11."

"Did I do that? I made a terrible mistake. I shouldn't have done that. But even a mistake must be honored, so for you and only for you I'll sell it for $11."

Larry bought the briefcase for $11.

Now let's employ the typical mathematician's device: pushing to extreme cases. It might seem that we've already reached the simplest level, but we haven't. Consider the following three special cases.

EACH PARTY KNOWS THE OTHER'S RESERVATION PRICE

Suppose that the seller and buyer each know their own and their adversary's reservation price. If $b < s$, then there is no zone of agreement: no deal is possible and the parties know it. If $b > s$, then a zone of agreement exists and the parties have a potential gain of $b - s$ to share. Of course, they get nothing if they can't agree on a sharing rule. Instead of carrying around excess symbols, suppose that s = \$400, b = \$600, and $b - s$ = \$200. How should they share that \$200 surplus? The obvious focal point would be in the middle (\$100 to each), and that's what happens overwhelmingly in experimental negotiations—provided that some care is taken to balance the environment.

In one interesting experiment conducted by Richard Zeckhauser, many pairs of subjects were each asked to divide \$2 between themselves; no agreement meant no money. In the symmetrical version, practically all settled on the \$1 focal point. In some pairs, one party was secretly prompted to hold out for \$1.20 and to hold firm; as expected, the reactions of the opposing parties were also firm—they would rather take nothing than 80 cents. Would this be your preference, too, if you had to share \$200 and someone demanded \$120?

The subjects were next told to share \$2 but they were each penalized 5 cents for every minute it took them to decide on their sharing rule. They quickly jumped to the \$1 focal point. Then came an interesting variation: Party A was penalized 5 cents per minute of negotiations, whereas Party B was penalized 10 cents per minute. Clearly Party A had a strategic advantage. But what had become the natural focal point? The surprising thing is that empirically, averaging over many subjects, Party A (the stronger party) in this variation did worse—not better, as might have been expected. Once symmetry was destroyed it invited power confrontations, and the seemingly advantaged Party A ended up, on the average, worse off than he had been in the symmetric case.

There is a famous example used by game theorists: How should a rich man and a poor man agree to share \$200? The rich man could argue for a \$150-to-\$50 split in his favor because it would grieve the poor man more to lose \$50 than the rich man to lose \$150. Of course, an arbitrator, keeping in mind the needs of the rich man and the poor man, might suggest the reverse apportionment. The rich man could also argue for an even split on the grounds that it would be wrong to mix business and charity: "Why should I be asked to give charity to this poor man? I would rather get my fair share of \$100 and give charity to a much poorer person."

Instead of dividing up \$200, let's introduce another asymmetry by having two bargainers divide up 200 poker chips; as before, no

agreement means no chips to either. Suppose further that Player A can convert the chips to dollars in equal amounts—one chip equals one dollar—but that Player B is given a complicated nonlinear schedule for converting chips to dollars. Figure 4 depicts one possible case. If A gets x chips, then B can cash in the remaining $(200 - x)$ chips for an amount in dollars equivalent to the vertical distance above x from the horizontal axis to the negotiation curve. If Player A argues that the game is symmetric in chips and that each should get 100 chips, Player B would receive \$45. If Player B argues that the real currency is dollars, not chips, the symmetric solution would give \$58 to each: A would get 58 chips, and B would get 142 chips that are convertible to \$58. This is analogous to the rich man's claiming that the real currency involved in his negotiation with the poor man should be after-tax dollars; and because he is in a higher tax bracket than the poor man, he should get more than \$100 in a "symmetric" split of \$200.

Another way of disturbing an apparently symmetric strategic situ-

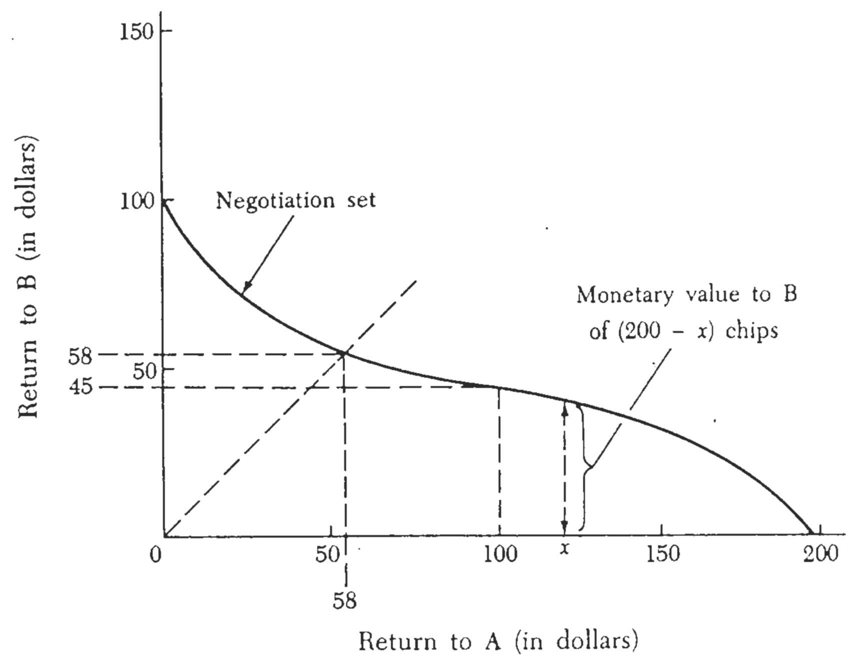

Figure 4. Example of a negotiation set with symmetry in chips but not in conversion to money.

ation is to have different numbers of people on each side of the bargaining situation. A simple case might be one in which Party A and Party B have to divide \$200. No agreement means no money. But now let Party A comprise two people (A′ and A″) who have agreed to their share, and let B represent one person. At one focal point, \$100 could go to Party A and \$100 to Party B; A's \$100 could then be split, \$50 to A′ and \$50 to A″. At another focal point, each of the three could get \$66.66; each one, after all, has full veto power.

This compendium of possible asymmetries is far from complete, but the examples it presents are instructive: differences in initial endowments or wealth, differences in time-related costs, differences in perceived determination or aggressiveness, differences in marginal valuations (as in tax brackets), differences in needs, and differences in the number of people comprising each side. There are, of course, many others.

The notion of symmetry and focal points is often associated by bargainers with their notion of "fairness." But one person's symmetry is frequently another's asymmetry, and the discussion of what is symmetric can be divisive. Even in the extremely simple case of two-party distributive bargaining, in which each side knows the other's reservation price and in which a zone of agreement is known to exist, there is a possibility that the players might not agree to an apportionment of the potential surplus $b - s$.

ONE PARTY KNOWS THE ADVERSARY'S RESERVATION PRICE

Suppose that the buyer knows the seller's reservation price (s) as well as his own (b); the seller knows s but has only a probability distribution for b. To be less general, assume that in a laboratory situation s is set at \$10 and each party knows this. Next, let b be chosen from a rectangular distribution from \$0 to \$30—that is, all values in the interval from \$0 to \$30 are equally likely.[8] Suppose that the chosen value of b is \$25. How might the players negotiate?

Once the buyer shows an interest in negotiating, the seller can update his knowledge about the unknown b. He knows that b is not less than 10. The final determination will depend not only on the bargaining skills of the two contenders but on their obstinacy levels. The buyer should be able to push the seller down to a value close to \$10. The buyer could act as if b were on the order of 14, rather than 25.

In these simple negotiations, in which only a single number b is unknown to the seller, the behavior of the bargainers will depend critically on whether b will become known to the seller after the negotiations are completed. In most real negotiations a reservation price is not just handed to the players: they have to analyze what their alternatives might be if there is no agreement, and uncertainties are usually involved. When inconvenience, transaction costs, and risk aversion are taken into account, it might never be possible, even after the negotiations, for one party to determine the reservation price of the other. Laboratory results depend to a crucial extent on whether true reservation prices are revealed after the termination of the bargain.

Imagine a case in which a business is acquired for a price of \$7.2 million. A couple of months after the transaction is completed the seller asks the buyer, "What was the very maximum amount you would have been willing to pay for my firm?" The buyer's reservation price was \$12 million, but if she reveals that high number she might make the seller feel miserable and she might tarnish her reputation. Of course, there are those who might gleefully and boastfully admit to \$12 million. More likely the response of the buyer might be, "You did quite well—I might have gone up to \$8 million, but I'm not sure." That's not a truthful response, but it's a kind one. The misrepresentation is not offered for the purpose of squeezing out a few extra dollars—at least not immediately—but in a self-serving way it does enhance the reputation of the buyer. The best alternative is probably a truthful but evasive answer: "Sorry, that's a number that just should not be disclosed." Of course, an analytically minded seller might then muse, "Hmm—she wouldn't use that ploy unless she'd really gotten the better of me."

Suppose that the buyer's reservation price happens to be extremely low, either by chance drawing in a laboratory setting, or in a real-world setting because of unexpected exogenous factors. If the buyer reveals this true reservation price—and it may be in her in-

8. This procedure is implemented in experiments by taking thirty-one blank cards; labeling them 0, 1, and so on up to 30; shuffling them; and letting the buyer choose a card at random from the deck. Once the experimenter has shown the card to the buyer, he returns it to the deck.

54

terest to do so—the seller might suspect that this is merely a ploy. The buyer might be better served if she refrained from making such truthful pronouncements, especially if her *RP* appears to be self-servingly low: the buyer can actually lose credibility by being honest. In one experiment involving successive bargaining rounds with different, independent, randomly drawn reservation prices for each round, a perspicacious buyer who drew an extremely low reservation price in one round decided to make believe that his *RP* was higher than it actually was; he announced a b' that was *higher* than his observed b. He was willing to lose money in that round in order not to jeopardize his credibility for further rounds of repeated negotiations. lying

EACH PARTY HAS PROBABILISTIC INFORMATION ABOUT THE OTHER'S RESERVATION PRICE

The following highly structured bargaining problem might be called the *canonical case of distributive bargaining*. Those who know game theory will recognize it as a formulation based on the work of Harsanyi (1965).

A seller and a buyer each have a probability distribution, one for the seller's *RP* and one for the buyer's *RP*. Both distributions are known to both parties. A random drawing is made to establish the buyer's *RP* and is shown only to the buyer; a second random drawing is made to establish the seller's *RP* and is shown only to the seller. The seller and the buyer then negotiate, face to face, and the payoffs are the surplus values that the parties can achieve. If the random values for b and s are such that $b < s$, there is no zone of agreement; if $b > s$, there is a zone of agreement and the bargainers have to share the excess, $b - s$. They do not know before they start bargaining whether there is an excess and, if so, how large it is. Since each bargainer knows only his or her own reservation price, each has a different probability assessment of the amount of excess to be shared.

To be specific and to keep the probabilistic elements simple, let s be drawn from a rectangular distribution from 50 to 150 and let b be drawn from a rectangular distribution from 100 to 200 (see Figure 5). All values between 50 and 150 are equally likely for s; all values between 100 and 200 are equally likely for b.

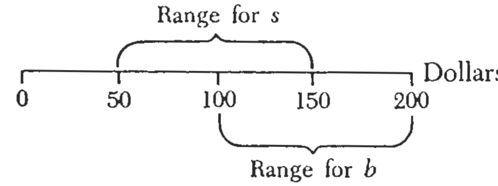

Figure 5. Distribution of reservation prices for the canonical case.

We will assume that the drawings are independent[9]—that the seller's knowledge of the outcome of s does not affect his probabilistic assessments for b, and vice versa. A particular joint drawing can be represented by a point (s, b) in the square shown in Figure 6. All points in that square are equally likely outcomes. There is a one-eighth chance that s will be greater than b and that no zone of agreement will exist; there is a seven-eighths chance that a zone of agreement *will* exist.

Subjects are assigned roles and each is given a randomly drawn *RP*. They negotiate outside any experimental setting and follow no structured rules. They can negotiate face to face or over the phone or write notes to each other. They can make up their own rules but they *cannot* show their confidential *RP*s to each other. They are given ample time to negotiate—roughly twenty-four hours, during which they may meet several times, for as little as a few minutes each meeting. They must turn in their negotiation forms at a specified time.

The number of actual agreements reached was surprisingly large. One might think that if there were a small zone of agreement—for example, if $s = 110$ and $b = 115$—the parties often would not be able to agree on a final price. Not so. It is true that the smaller the zone, the longer it may take for the parties to locate it, but they almost always come to agreements when agreements are possible. Inefficiencies occur only when there is a zone of agreement and the parties do not come to an agreement. Informal bargaining, without

9. The laboratory procedure can be implemented as follows. The seller has a deck of 101 cards labeled 50, 51, and so on to 150; one of these cards is drawn at random, shown to the seller and the experimenter, and returned to the deck. The buyer has a deck of 101 cards labeled 100, 101, and so on to 200; one of these cards is drawn at random, shown to the buyer and the experimenter, and returned to the deck. The payoffs to the buyer and seller are made in a confidential manner so that each player never knows the other's *RP*.

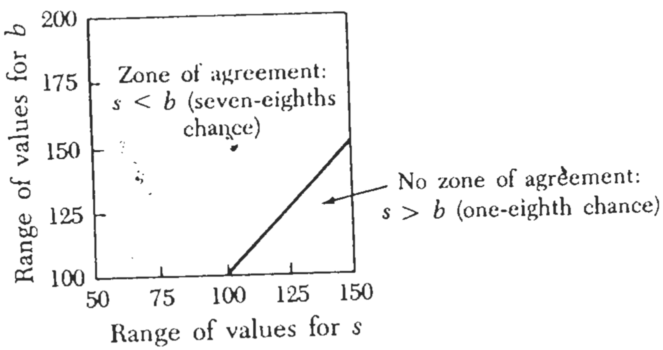

Figure 6. Joint representation of equally likely outcomes.

any imposed structure for negotiations and without tight time constraints, leads to more efficient outcomes than do most formal methods. One proposed structured alternative to informal bargaining is the procedure by which both parties reveal their reservation prices at the same time. This alternative, though appealing, does not work very well, as we will see below.

SIMULTANEOUS-REVELATION RESOLUTION

In any negotiation experiment there will usually be some bargaining pair who decide to devise rules of their own.[10] A seller says to his adversary, "Let's not waste time. My reservation price is $300. What's yours?" What a temptation to a competitive buyer! Let's assume that her reservation price is $550. Should she be honest and honesty say so? Is the seller trying to take advantage of her? Perhaps the true reservation price of the seller is really $200. According to a commonly proposed symmetric resolution procedure, the parties simultaneously disclose their reservation prices: "I'll write down my reservation price if you'll write down yours at the same time. If we're compatible, we'll split." Let these disclosed values be s' (not necessarily the true value s) for the seller, and b' (not necessarily the true b) for the buyer. If $b' < s'$, then negotiations are broken off; if $s' < b'$, the final contract will be $x^* = (b' + s')/2$, the midpoint between b' and s'. (See Figure 7.)

10. This section and the following one are fairly technical and can be skipped by nonmathematically inclined readers.

55

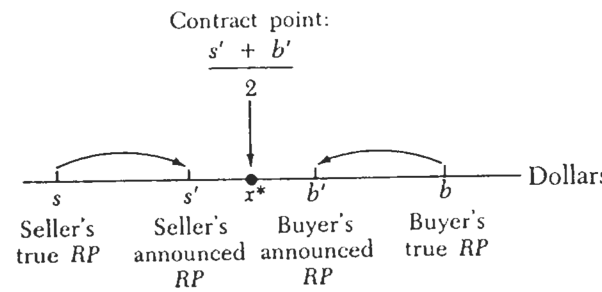

Figure 7. The simultaneous-revelation procedure. (Note: An inefficiency would result if $s' > b'$.)

When this simultaneous-revelation procedure was tried, most parties gave truthful revelations: s' equaled s, and b' equaled b. However, in some cases s' was greater than s, and b' less than b; indeed, in some of these cases, there was in fact a zone of agreement (s was less than b) but the parties did not detect it (s' was greater than b') and an inefficiency resulted.

Suppose that a seller draws a very low s value—say, 60. Should his announced value s' be 60, or a higher value such as 110? Remember that as long as the announced b' is higher than his s', the final-contract price will be midway between these announced values.

In a nonlaboratory, real-world setting a bargainer may have no way of ever ascertaining the other party's true reservation price. In an experimental setting, on the other hand, it's difficult to keep these true reservation prices secret after the fact. Is it "ethically correct" for someone to lie about his or her reservation price when the parties agree to reveal their values simultaneously?[11] Some would say that this behavior was absolutely inappropriate, but others would claim that the purpose of laboratory exercises is to provide vicarious experiences: "In real-world settings most people don't even have firm reservation prices. Besides, it's culturally acceptable to exaggerate a bit in your own favor. What's wrong with that? If my adversary did it to me, I wouldn't be angry. I do to others as I

11. In a debriefing session following one laboratory exercise, a buyer defended her behavior as follows: "My confidential reservation price, b, was 170 and my announced bid, b', was 130. I didn't think of b' as a distortion of the truth but as a strategic bid, not unlike any sealed bid for a contract."

expect others to do to me." We'll look closely at this philosophy later.

Here is a simple exercise. Suppose that the subject playing the role of seller receives a value of s drawn from the interval \$50 to \$150, and that the subject playing the role of buyer receives a value of b drawn from the interval \$100 to \$250. All values within these intervals are equally likely. What strategies can the seller devise to determine his value of s' as a function of s (for $\$50 < s < \$150$)? Figure 8 depicts three such strategies: (1) a representative strategy where, for example, the seller would say \$112 if his actual *RP* were \$75; (2) a strategy of truthful revelation, where $s' = s$ for all s; and (3) a strategy of truncated truthful revelation, where $s' = \$100$ for all $s < \$100$ and $s' = s$ for all $s > \$100$.

Each seller must submit a seller strategy and each buyer must submit a buyer strategy. Each seller is then "scored" by pitting his or her strategy against each buyer's strategy in turn; the seller's score is then his average return—averaged over all s values and over all buyer-adversaries. Buyers are scored analogously.[12]

If s and b are the actual *RPs*, and if s' and b' are the revealed values, the payoffs can be formulated as follows:

$$\text{to the seller:} \quad \begin{cases} (s' + b')/2 - s & \text{if } s' < b' \\ 0 & \text{if } s' > b'; \end{cases}$$

$$\text{to the buyer:} \quad \begin{cases} b - (s' + b')/2 & \text{if } s' < b' \\ 0 & \text{if } s' > b'. \end{cases}$$

The difference between s and s' can be said to be the amount of exaggeration (or distortion) at s. Subjects in general—even those students who helped me design the game—played it very badly: they exaggerated too much. When truthful revelation strategies, or even truncated-truthful revelation strategies (see Figure 8) are pitted against each other, the probability of getting an (s, b) pair with no zone of agreement is .125 (see Figure 6). But averaging over all subject strategy responses and over all (s, b) pairs yielded an extremely large probability, .46, that no zone of agreement (in re-

12. This game has been extensively analyzed by Chatterjee and Samuelson (1981).

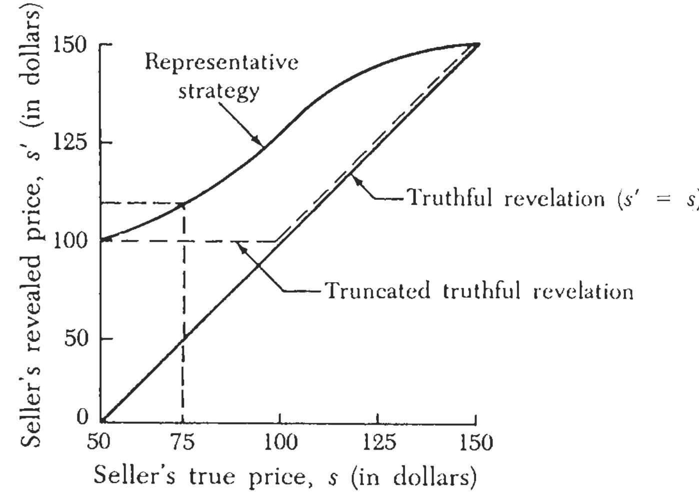

Figure 8. Strategies for the seller in the simultaneous-revelation resolution procedure.

vealed values) would exist! Thus, over one-third of simulated trials resulted in no agreement when in fact a zone of agreement did exist. Not very efficient. This happened because there was so much exaggeration—so much, in fact, that those subjects who used a truncated truthful strategy did exceptionally well comparatively. They found that a good retort against an extreme exaggeration is (truncated) truth telling. If both parties exaggerate a lot, then the chances for an agreement are very poor (see Figure 9).

Thus, although the simultaneous-revelation resolution procedure was devised to eliminate the need for haggling, it is obviously not a very good substitute.

Figures 10 and 11 depict a pair of equilibrium strategies: one for

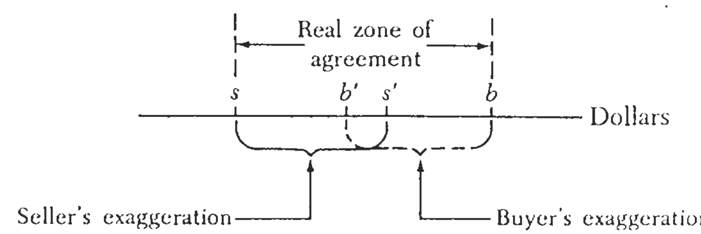

Figure 9. Case in which there is a zone of agreement in real but not in revealed values.

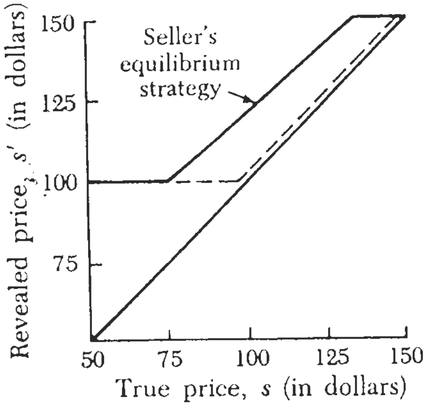

Figure 10. Seller's equilibrium strategy for the simultaneous-revelation resolution procedure.

the seller and one for the buyer. As long as one party adopts his part of the equilibrium strategy, the other would find it to his advantage to do likewise. But the equilibrium pair is not efficient, because for many (s, b) pairs where there is a zone of agreement, the revealed (s', b') pairs yield no agreement. With a pair of equilibrium strategies in contention, 38 percent of all (s, b) pairs result in no agreement. (The empirical percentage of no-agreements was 46.) Two truth tellers do better than two equilibrium strategists.

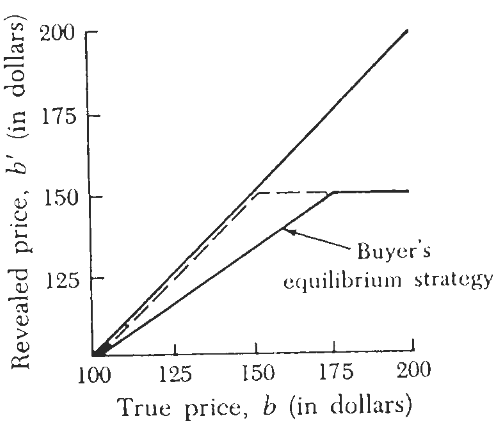

Figure 11. Buyer's equilibrium strategy for the simultaneous-revelation resolution procedure.

57

Analytical elaboration. What can bargainers do when they know about equilibrium strategies but do not have the analytical skills necessary to compute these equilibrium strategies, or do not have the time to devote to such intricate analyses? Let's take the vantage point of the seller. One simple analysis is to boldly hypothesize a reasonable strategy for the buyer and by trial-and-error figure out a reasonable counterresponse for selected values of s—say, for $s = 60, 80, 100, 140$; these can be compared with a curve for interpolated values of s by inspection. A second simple analysis seeks the best retort against a truncated truthful revelation strategy; this retort distorts the truth more than the equilibrium strategy. Next, one can seek the best retort against the best-retort-against-the-truncated-truthful-revelation-strategy; this retort distorts the truth *less* than the equilibrium strategy. It can be proved that successive iterations—that is, the best against the best-against-the-best and so on, and finally against the truncated truth—yield a sequence of strategies that converge to the equilibrium strategy, and that these strategies oscillate ever closer and closer around the equilibrium strategy. Two or three stages in that sequence already give a practical approximation of the equilibrium strategy.

The simultaneous-revelation resolution procedure is inefficient because it encourages exaggerations; but it's fast and uncomplicated. If time is at a premium or if one is engaged in many such bargaining problems, then this resolution procedure still has merit —especially if the parties can refrain from undue exaggeration.

A MODIFICATION THAT INDUCES TRUTHFULNESS

The simultaneous-revelation resolution procedure can be altered ~~in such a way as to engender truthful~~ness.[13] This modified form exists in theory, but no one has yet discovered how to apply it to real-world situations; it would be wonderful if someone could.

It's important to keep in mind that the distributive bargaining

13. The material in this section is based on research done by Chatterjee (1979) and by Pratt and Zeckhauser (1979).

problem being modified is in canonical form: private reservation prices are drawn from commonly known probability distributions. Furthermore, the parties must agree to the modified payoff procedure before drawing their reservation prices. If these assumptions are violated, the modified procedure will not be strictly truth-generating, but it still will encourage less exaggeration.

Suppose that there is a seller (let's call him Jim), a buyer (Jane), and a rules manipulator (George). Imagine that George can induce Jane to make honest revelations: her declared price, b', is the same as her real price, b. How can George get Jim to be equally honest? If Jim's real price is s and if he announces s' while Jane announces b', then let Jim's payoff be $\{[(s' + b')/2] - s\}$ if $s' < b'$ and 0 otherwise (the formula given earlier), *plus* an adjusted amount that Jane will pay him that depends solely on the s' he announces (see Figure 12). Notice that the higher Jim's s' the lower the adjusted payment he will receive from Jane. Hence, with the adjustment there is less incentive for him to exaggerate as much. He will want to decrease s', and now the trick is to manipulate the adjustment function so that if Jane tells the truth by announcing $b' = b$, then Jim's best overall response is also to tell the truth—that is, to announce $s' = s$. Of course, the adjustment function may go too far: it may be so steep that Jim may want to select s' below s. The idea is to adjust it in a way that causes him to announce $s' = s$ for all s. All this assumes that Jim is trying to maximize his expected overall monetary return.

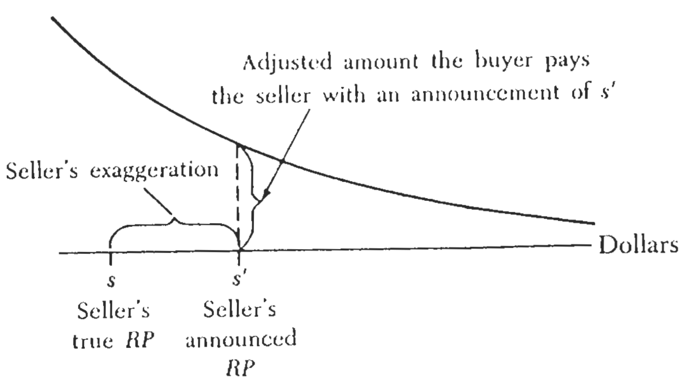

Figure 12. The extra payoff the seller receives as a function of s'

Now, assuming that Jim agrees to always announce $s' = s$, how can we determine the adjustment that Jane will have to pay him? George will induce her to pay this adjustment by reversing the roles and making Jim pay her an adjustment value that depends solely on her announced value b'; he'll manipulate her adjustment function so as to make it most profitable for her (on an expected-value basis) to announce b. Then he'll raise or lower the adjustment function so that Jim's net expected side payment (the amount he receives from Jane: a function of s', less the amount he pays to Jane, which is a function of b') is zero. Her expected net side payments also will be zero.

Can all this be done? Yes, say the experts. But in order to implement this scheme, the seller and the buyer have to agree to it before the seller knows s and the buyer knows b; and in order for the rules manipulator to calculate appropriate adjustment functions, he needs to know the probability distributions that underlie the drawings of s and b. Those are rather restrictive assumptions. But the result is so appealing that it should not be lightly dismissed. With suitable adjustment functions, honest revelations are in equilibrium: each party should tell the truth if the other does. Furthermore, because the equilibrium strategies for the unadjusted game are not jointly efficient, the equilibrium expected payoff from the adjusted game is higher for each than the equilibrium expected payoff of the original game.