KELLOGG SCHOOL OF MANAGEMENT

Strategic Decision-Making

DECS-452 Professor Bob Weber
Week #6

1. Please carry out the Riverside / DEC negotiation exercise. Submit your results using the form on the class webpage no later than 11 PM Thursday (day section) / 11 PM Wednesday (evening section).

2. Read the *Managing Negotiations* handout (Chapters 6 and 8) and the attached material.

3. (Carryover from Week 4: This time for sure!) Write up your strategy for problem (3) in the "Common Knowledge" notes. On **Tuesday, May 10** (day section) / **Thursday, May 12** (evening section), I will collect the strategies, shuffle them, and then hand them out again, so be sure that yours is written up clearly enough for a classmate to be able to carry it out.

4. If you'd like to try a few informal exercises:

   (a) Analyze the "dissolving a partnership" example (in the attached section of the "Negotiation and Arbitration" readings) using the approach we took to the first-week bidding experiment

   (b) Show that, for either of the "mechanisms" illustrated by the two figures (in the attached section of the "Negotiation and Arbitration" readings), the best you can do is to truthfully report your own valuation of the item up for trade.

5. On Friday, May 13 (day section) / Thursday, May 12 (evening section), I'll distribute the final project materials. You'll have three weeks to work on them.

The "Salty Dog": Distributive Bargaining

Buyer and Seller Types

```
      $1000        weak        strong        $4000
(Seller) [-------------|--------------|---------------]
                                    [---------------|--------------|-------------] (Buyer)
                                  $3000          strong        weak       $6000
```

			seller	
	Pairings		$1,900	$2,900
			weak	**strong**
buyer	$4,100	**strong**	3,7,11,…	1,5,9,…
	$5,100	**weak**	4,8,12,…	2,6,10,…

The Salty Dog: Summary and Notes

Summary of Final Agreements, 1987-2016:

	\$1900	\$2900	
\$4100	\$3,330	\$3,502	strong buyer
\$5100	\$3,468	\$3,723	weak buyer
	weak seller	strong seller	

Likelihood of Agreement:

	\$1900	\$2900	
\$4100	98.1%	97.6%	strong buyer
\$5100	98.1%	97.2%	weak buyer
	weak seller	strong seller	

In the role-play, each side is given no guarantee that mutual gain is available, although in fact such gain *is* available to all pairs. On each side, a party can be in either a strong position (a relatively attractive alternative to a negotiated agreement is available), or a weak position (the alternative is relatively unattractive); hence there are actually four different types of pairings: strong-strong (the narrowest range of mutually-advantageous agreements), strong-weak, weak-strong, and weak-weak (the widest range). The strong and weak positions are set so that every party knows (privately) that a mutually-advantageous agreement exists, but doesn't know whether the other party also knows this.

Note:

(1) Just about every pair reaches agreement.

(2) The actual agreements vary by pair, but, on average, strong parties do better (i.e., reach agreements closer to their ends of the distributive axis) than weak parties against either strong or weak partners.

On the surface, (1) and (2) seem reasonable. **But (1) and (2) together are inconsistent with effective negotiating behavior!** If every pair "should" reach agreement, weak parties should act as if they are in strong positions. They will still reach agreements, and the agreements will be more favorable to them: They should do as well as the strong parties. Alternatively, if strong parties "should" do better, it must be because weak parties can't afford to mimic their behavior: That "strong" behavior must involve taking risks that weak parties can't justify, in which case, strong parties negotiating with other strong parties must sometimes fail to reach agreement.

Strategic Decision-Making (DECS-452 Course Outline)

The Salty Dog case: Distributive bargaining

1. Opening offers

    a) If you move first
        1) Too aggressive: Antagonistic
        2) Too generous: Costly - signals either weakness or softness
        3) Just right (as tough as possible without being ridiculous): Fine - gives mixed signal, either strong or tough, without revealing information; forces response. Can anchor further negotiations in desirable region, and increase opposing concern about reaching *any* agreement.
        4) Fuzzy: Difficult to manage, but can provide a way to force opponent to make first "real" offer
        5) Moral - If you have good information about your competitor's best alternative, start with a tough offer. If you know more about him than he thinks you do, start below his best alternative and let him drag you up.
    b) If he moves first
        1) With luck, he'll be too generous at the start
        2) Perhaps he'll betray a useful piece of information
        3) If he comes in tough, resist anchoring: Respond toughly.
        4) Moral - If you are not well informed, try to avoid the first move. If you are well-informed, but think he may be soft and feel you can undo his anchoring, make him move first.

2. The middle game

    a) Don't let the desire to reach agreement overshadow your search for a good agreement, i.e., don't become more risk-averse than you truly are. A useful approach is to formulate an aspiration level for yourself, and to update this as information comes forth.
    b) Recognize the time already spent in discussion as a sunk cost.
    c) Carefully listen to him, and continually update your perception of his likely best alternative.
    d) Watch for impatience (usually a sign of weakness), and don't appear impatient yourself.
    e) Moral – Be aware of your own psychological state, as well as his.

3. The close

    a) Lock him into a commitment to agreement.
    b) Congratulate him on his toughness, and both of you on your mutual success.
    c) If it seems useful, pull out the "quivering quill".
    d) If you were prepared to walk, and he calls your threat, walk!

His ideas were thought reactionary at the time, but Lemuel Boulware has lived to see many of them accepted.

A look back at "Boulwarism"

A chat with Lemuel R. Boulware
by Peter Brimelow (*Forbes* - May 29, 1989, pp. 246-248)

As part of his long-ago effort to cope with the power of trade unions, Lemuel R. Boulware in the late 1950s hired a movie star called Ronald Reagan, who bolstered the viewership of General Electric's TV show. In carrying out his assignment, Reagan visited GE plants where his political views were deeply influenced by the discovery that "we didn't chain the workers to the machines." Ronald Reagan went on to use his convictions and his powers of persuasion to help restore America's faith in the free enterprise system.

Now 94 and retired for nearly 30 years, Boulware is restricted by recent illness to a wheelchair in his Delray Beach, Fla. oceanfront home. But this spring he talked at length with *Forbes* about many things, including the past and the future of labor relations in the U.S.

In the 1950s "Boulwarism" was a household word. It was coined to describe Lemuel Boulware's seemingly rigid style in opening wage negotiations as head of employee relations for General Electric. He would not bargain, labor leaders complained. He would simply make whatever offer he had determined was in "the balanced best interest" of company, work force and consumer — and thereafter refuse to budge. Boulwarism was widely denounced as arrogant paternalism. But in fact it was rather more subtle. As Boulware pointed out in his book *The Truth About Boulwarism*, written after he retired in 1961, it could work only where the company was offering a pay scale that was acceptable by prevailing standards — otherwise resistance would be too intense.

Because Boulware recognized the need to pay competitively, his wage offer was never simply a stonewall refusal to consider an increase. And, contrary to legend, he says that he always undertook to negotiate further "on getting any old or new information proving that change would be in the balanced best interest of all." He reports that in the 14 years he handled labor negotiations for General Electric, only one of his initial offers was actually accepted in the end without amendment.

But from the labor union leaders' point of view, Boulwarism had a major disadvantage: It made it difficult for them to go to their members and claim all the credit for the settlement. They went all out to discredit his ideas and succeeded in making Boulwarism a dirty word to many people.

Because of this essentially political need on the part of the labor leaders, many negotiations had tended in his time to become "amateur theatricals," as Boulware puts it. "For instance, if everything pointed to a five-cent increase being about right, there was a strong tendency among employers in those days to offer nothing at first. Then, under public strike threat pressure, about half would be offered. Then, after all the union representatives had been called in from the plants and the resulting vote for a strike had been well-aired in a receptive press, management would 'capitulate' by upping the offer to the full five cents an hour."

He recollects, "Time and again I was told in private — and even occasionally before mediators and 20 or 30 others at the bargaining table — that there was nothing wrong with the offer except it was ours and not theirs, and that they had to justify themselves with their members by showing they could force something more out of US."

This was exactly what Boulware wanted to prevent. He explains: "We all had 'quite a shock when, at plant after plant in 1946 [when General Electric was idled by a seven-week strike] union officials had proved they had a 'push-button' control over employees and could cause them not only to strike but also, in too many cases, to do senseless and frightful damage to the investment in their places of work."

Assigned to labor relations after the strike, Boulware concluded that union propaganda had persuaded General Electric's employees that their welfare did not depend on the economic health of their company but on whatever muscle they could apply to their employers. Managements that allowed themselves to be cast as villains in any such union—scripted charades would merely increase the labor leaders' power — particularly if, as with General Electric when Boulware took over, the unions were the only side that would talk to the press. The end result: major corporations becoming "the slaves of unions" — which is exactly what happened, says Boulware, in the American automobile industry.

Depriving unions of credit for pay increases that were really the result of market forces was only part of Boulware's counteroffensive. He also launched a company-wide "corrective education" program. A regular bombardment of short, punchy

commentaries and cartoons on elementary economics (see below) began in employee newspapers, and were even included with mailings to shareholders and distributed to community leaders. (“The clergymen were the worst,” says Boulware. “They were always against us.”)

In an uncanny foreshadowing of T-groups and Japanese corporate communalism, General Electric supervisors held discussion groups for employees in their homes, focusing on a specially prepared economics primer that dealt with such intriguing questions as “Where does the government get its money?” Significantly, Boulware says, the most interested participants were the wives.

Boulware's broadsides may have been simple to read, but they were in fact subtly sophisticated. For example, he never blamed unions for inflation — high wages could force companies out of business and kill jobs, but an increase in the overall price level would result only if the government tried to counter this effect by printing money. Thus wage and price controls were irrelevant to inflation — a point not understood even by subsequent Republican Presidents.

General Electric had no major strikes during Boulware's tenure. But it has been argued that his policy as executed by his successors eventually resulted in a protracted strike in 1969, after which it was abandoned. Boulware, a loyal company man, says only: “They thought the job was done.”

Boulware in 1950 and illustration from G.E. employee publication

More than capital and labor were at the bargaining table.

Much of Boulware's experience is now history. Even heavy industrial unionism is a shadow of its former self, with General Electric's work force sharply deunionized.

What does he think of the current scene? Boulware likes George Bush but says he's surrounded by “go-along-to-get-along types.” Boulware still thinks that the answer to the nation's problems is civic virtue, which can be injected only by the efforts of the minority of enlightened businessmen. Alone in his big house since the recent death of his wife, attended by a servant couple and encroached upon by luxury condominiums, he still receives a procession of emissaries from struggling free market think tanks and magazines and dispenses advice and money to them, an investment in his country's future designed to pay off slowly and in the long term.

Of the future for free enterprise, he says: “You have to believe.” He knows that the victory isn't won, but he has lived long enough to see labor and capital retreat considerably from their formerly adversarial positions and to have watched as Ronald Reagan went on to apply many of Boulware's basic ideas to the national scene.

Fish or Shark?

The Ecology of Strategic Behavior

The difficulty in teaching a course on strategic behavior is in distinguishing between what lessons I **would** teach if everyone in the world would follow my teachings religiously, what I **could** teach if everyone would listen to me but felt free to choose whether to follow my teachings or not, and what I feel I **should** teach to a group of managers who will subsequently be dealing with others who haven't heard me.

The finitely-repeated Prisoners' Dilemma illustrates my quandary. If everyone were required to obey my teachings, the solution would be simple: I **would** tell you to always "cooperate" (even in the one-stage game). The world would be a better place for us all. (The problem, of course, is that I can't **force** even **one** person to follow my advice.) If instead, everyone in the world were to listen to me but each had the freedom subsequently to decide for himself whether or not to follow my teachings, the only thing I **could** teach is to always "defect." If I taught anything else, and you (an individual) thought that everyone else would follow my teachings, you would find yourself preferring to disobey me. But in the real world, "Tit-for-Tat" makes sense, and (subject to minor modifications) is what I believe I **should** teach. As long as there is some chance that you are facing someone who will cooperate to the bitter end (or as long as you think that they think you may be such a person, or ...), it is in your interest to cooperate through the early stages, defecting only very near the end.

More picturesquely, think of those who blindly cooperate as "fish," and those who act totally aggressively as "sharks." In an ocean without sharks (i.e., in a naively Utopian society), the fish thrive. In an ocean of sharks alone, everyone starves (at least, in relative terms). God forbid that my goal in life be to make everyone a shark! Instead, my goal is to teach you to thrive in an ocean containing both fish and sharks (i.e., in the "real ocean"), by being a cross between a spiny fish and a dull-toothed shark. Tit-for-Tat starts off fishy, but bites back.

I have no qualms concerning the message that came from the bidding experiment. In most economic situations, people **do** worry too much about getting something, and not enough about how much they get. (Similarly, in emotionally-charged situations with more than one possible resolution, people frequently worry too much about the precise form of **what** they get — Fisher and Ury, in *Getting to Yes*, properly emphasize the importance of focusing on issues, rather than on positions.) Optimizing against your best model of opposing behavior (i.e., being sharkish) is, in many cases, clearly the best way to deal with general economic confrontations.

But perhaps there are exceptions. If you are negotiating with someone who is obviously a fish, and indeed, such a fish that you could swallow him whole, should you do so? Society at times says, "No!" (For example, a divorce judge is supposed to refuse to ratify a mutually-accepted settlement that he deems clearly inequitable, i.e., unconscionable. The legal presumption is that such a contract was agreed to under duress.) The line is certainly drawn somewhere short of "conning" the fish.

In a world of pure fish, disputes would be easy to resolve. Have all parties honestly reveal all of their concerns, all of their alternative opportunities, all of their preferences over tradeoffs between issues (in multi-issue disputes), and then determine the agreement which yields the greatest aggregate gain and divide that gain equally among the parties. (The question of interpersonal utility comparisons — "This means more to me than it does to you" — need not come up: The Nash bargaining solution, or any of several alternative schemes, provides a reasonable "split-the-difference" procedure.)

Unfortunately, there **are** sharks in these here waters (not necessarily taught by me — if they didn't exist, God would evolve them). And in consequence, I cannot counsel totally fish-like behavior. The Salty Dog case provides a good example. If as a buyer you truthfully and convincingly reveal your reservation price at the beginning of the negotiations, a seller with any reservation price below yours faces an overwhelming temptation to exaggerate his own reservation price and offer to split the difference. For in this case, as in many in the business and personal worlds, the exaggeration will never be subsequently detected. (The shark need only wear fish clothing.)

What counsel can I give? If I taught you precisely how to formulate an opening offer on the basis of your reservation price, but never to walk away when a profitable deal is on the table, others could "invert" my teachings, accurately infer your reservation price from your opening offer, and squeeze you as before.

Fortunately, there **is** something that I can teach, if you must, as a buyer, make the first offer. Make an offer (even an invertible one), and make it absolutely clear that any higher counter-offer has some probability of being rejected, *even if your best alternative is worse than the counter-offer*. If your statement (threat?) is believed, then the best the seller can do is maximize his expected gain; if your stated probability of accepting an offer very close to your reservation price is quite low, his optimal counter-offer will probably offer you gains of substance. (If he squeezes **too** hard, he is very likely to gain **nothing**.) Of course, the same teachings apply to the seller, if he must move first.

In the real world, it is not socially graceful to be so blunt. We make a starting offer, and signal our walk-away likelihood through our reluctance to move, or through the size of our concessions. Being able to read another's signals accurately, and to correctly give the proper signals ourselves, is a skill which comes only through experience. (Role-playing exercises provide a riskless way to gain some such experience. I don't have time in this course to do an adequate number of skill-building role-plays, but then, mine is not the only course at Kellogg which deals with negotiations.)

The bottom line is that, in order to be at least a spiny fish, you **must** be able to convince others that you are willing, at times, to walk away from an agreement that offers you gain. And unless you are an accomplished hypnotist, the only way to be convincing is to **be** someone who sometimes walks away.

What is a shark? Someone who can convince others that he is prepared to **permanently** walk away, and still come back gracefully if they don't budge. A typical tactic is to let an affiliated party "override" your decision: A lawyer, having walked, can later return saying, "My client wants me to try again." An arms negotiator can be "sent back to the table" by the President. A car buyer can be "overruled" by a spouse. Alternatively, you can come back yourself, saying, "I just realized that there is another way we can structure the deal that might be acceptable." (For example: "Oh, the car comes with tires? Then maybe I can raise my offer a bit.") Still, Lincoln's caution is of relevance: "You can fool all of the people some of the time, and some of the people all of the time. But you can't fool all of the people, all of the time." If you **never** truly walk, you're definitely exploitable.

It is overly simplistic to say that concerns about long-term relationships or reputation should restrain a shark. It's worth noting here that "walking away" can happen in the middle of a continuing relationship: Rather than settling a disagreement quickly and amicably, a husband and wife might subject each other to a day or two of sulking silence. Or a union might strike. But the relationship endures. The occasional failure-to-settle maintains a healthy atmosphere of mutual respect.

Of course, many (nay, **most**) people indeed never walk away from gainful agreements. And others **do** exploit this. And you, in turn, can exploit **them.** An effective salesman will expect you (a "typical"

purchaser) to walk only if the price on the table is above your reservation price. He'll start well above and slowly come down, reading you all the way. As soon as he thinks you're no longer willing to walk (i.e., as soon as he thinks he's hit your reservation price), he'll freeze. Feeling that **he's** the shark and **you're** the fish, he'll come all the way to **his** reservation price only if he feels that your reservation price is very near his. So you, a spiny fish, wear your pure-fish clothes. You do your research beforehand, deciding what you want and what reservation price the salesman is likely to have. You approach him, and indicate a desire to spend some amount **less** than this. You let him boost you to more expensive models by showing disinterest for those listed near your "expressed" budget, but you never admit the ability to pay any more. (He'll assume you can be dragged up somewhat.) When he finally shows you what you want, the game begins. He makes some concessions, you come up in a few tiny steps, then freeze just above his reservation price. If necessary, you begin to walk. Don't worry: If you did your research well, and froze **anywhere** above his reservation price, he (not knowing that you're spiny) will mistakenly conclude that he's hit **your** reservation price, and will close the deal.

Must you lie about (or at least, "strategically misrepresent") your position? Not necessarily. But in practice most people back themselves into situations where they must. The salesman asks, "What is the most you can afford?" An indignant "If I told you that, you'd drag me up to it!" simply alerts him to your spinyness. Without prior thought (and practice!) you'll probably cite a falsely-low amount. (He expects this: **Most** people will. In that case, are you really lying?) A more principled approach might be to turn the question around: "Well... How far down can **you** come?" But the turn-around itself is a bit of a tipoff that you know what you're doing.

Fish working with other fish build excellent agreements, and share equally in the joint gains. But fish working with disguised sharks (and most true sharks **do** wear fish clothing) build pretty decent agreements, and then the sharks take most of the gains for themselves. If the sea contains both fish and sharks (as, in fact, it assuredly does), the fish frequently lose out. But sharks against sharks get nothing.

A shark wants it all; a spiny fish recognizes the needs of other parties. Spiny fish build decent agreements with other spiny fish, and, after dancing the strategic dance, share in the gains. Spiny fish are relatively safe from sharks (and can sometimes turn the tables).

The biggest problem a spiny fish (i.e., an effective negotiator) faces is in deciding how to deal with pure fish (i.e., those less-skilled). Does he mimic a pure fish (sacrificing potential gain to protect the other party), or a shark in pure-fish clothing? I wish I had an easy answer. But it seems clear that, once we abandon the unrealistic dream of an all-pure-fish ocean, the best we can work for is a sea in which most of the fish have spines. (Perhaps we can even starve out the sharks.) Wearing our own spines proudly, and nipping at the pure fish without devouring them, leaves us in excellent shape and encourages them to grow their own spines. And maybe, just maybe, the right ecological balance can be achieved.

Example 8 (dissolving a partnership). Two individuals jointly own a piece of property. They have decided to sever their relationship, and for one of the two to buy the land from the other. Each knows how valuable the land is to him, but is unsure of its worth to the other. They agree that each will write down a bid; the high bidder will keep the land, and pay the amount of his bid to the other.

Assume that each is equally likely to value the land at any level between $0 and $1200, and that both know this. Then the unique Bayesian equilibrium point of the bidding game is for each to bid one-third of his own valuation. If, for example, one of them values the land at $300 and believes the other to be following the indicated equilibrium strategy, then by bidding $100 he has an expected payoff of $1/4 \cdot \$200 + 3/4 \cdot \$250$; he expects to win with probability 1/4, and when he loses, he expects the other's (winning) bid to be between $100 and $400. This private strategy is optimal for him, given his belief about the other's behavior. (Given his belief that his partner will bid a third of the partner's valuation, his own expected payoff, when his valuation is v and he bids b, is

$$(3b/1200) \cdot (v-b) + (1 - 3b/1200) \cdot (b+400)/2.$$

In general, this is maximized by taking $b = v/3$.)

Observe that this bidding arrangement always yields a Pareto-efficient result, i.e., the individual who values the land more highly always ends up in possession of it. Hence, the appropriate choice of a dispute-resolution procedure can, at times, circumvent inefficiencies of the type which arise in the lemon problem.

Note also that an intervenor could suggest the use of this procedure, if the parties found themselves unable to work out an agreement on their own.

4.3 **The revelation principle**

There are, of course, numerous other procedures that an intervenor could suggest in order to resolve the dispute in Example 8. Let us consider the (seemingly appalling) general question of what outcomes can result, at equilibrium, from *any* procedure which might be used to resolve a given dispute.

A simple, yet conceptually deep, type of analysis has become standard. Consider any equilibrium pair of strategies in a particular game. Each party's strategy can be viewed as a book, with each chapter detailing the private strategy of one of that party's types. Given the two actual types, a pairing of the private strategies in the two appropriate chapters will lead to an outcome of the game.

Next, step back from this setting, and picture the two parties in separate rooms, each instructing an agent on how to act on his behalf. Each agent holds in hand the strategy book of his side; all he must be told is which chapter to use. From this new perspective, the two parties can be thought of as playing an "agent-instruction" game, in which the strategy books are prespecified and each must merely tell his agent his type (or, equivalently, point to a chapter in his strategy book). An equilibrium point in this new "chapter-selection" game is for each to tell the truth to his agent. Otherwise, the original strategies could not have been in equilibrium in the original game.

Consequently, *anything* which can be accomplished at equilibrium through the use of *any* particular dispute-resolution procedure, can also be accomplished through the use of some other procedure in which the only actions available to the parties are to state their (respective) types, and in which it is in equilibrium for each to truthfully reveal his type.

Example 9. Two parties must share the cost of a public works project (e.g., the planting of a tree on the boundary line separating their homes); if they cannot agree, the project will not be carried out. The project will cost \$100. Both parties are risk-neutral, and it is known by both that Party A will derive \$90 in benefit from the project. Party B knows the benefit he will receive, but all that is known to A is that there is a 50% chance it is worth \$90, and a 50% chance it is worth only \$30, to B. What possible agreements could they reach?

The revelation principle tells us that any agreement which could be arrived at through any negotiation procedure will be an agreement which could also be achieved in a formally-structured game in which each simply names his type, and each has no incentive to lie. (Since A's type is known to both, only B will actually have a move in this game.)

An outcome of the revelation game will be, most generally, a probability that the project will be carried out, and a sharing of the \$100 cost if it is indeed carried out. Since a different outcome might result from each of the two type-declarations B might make, the full spectrum of possible agreements can be characterized by four parameters: p_H and p_L, the probabilities of project commencement given that B announces his type to be "high" (\$90) or "low" (\$30), and e_H and e_L, the payments to be made by B given his announcement and that the project is carried out.

In order for truth-telling to be optimal (i.e., a best response to A's null action) for B, these parameters must satisfy two *incentive constraints*:

$(90 - e_H)\cdot p_H \geq (90 - e_L)\cdot p_L$ (the \$90-type must prefer announcing "H" over "L")

$(30 - e_H)\cdot p_H \leq (30 - e_L)\cdot p_L$ (the \$30-type must prefer announcing "L" over "H")

Furthermore, in order for A, and for both types of B, to be willing to agree to the procedure, it must satisfy the following *participation constraints*:

$1/2 \cdot p_H \cdot (e_H - 10) + 1/2 \cdot p_L \cdot (e_L - 10) \geq 0$ (for A to participate)

$e_H \leq 90$, $e_L \leq 30$ (for both of B's types to participate)

It follows (algebraically) from all this that we must have $p_H \geq p_L$ and $e_H \geq e_L$; that is, when B reports himself to be the \$90-type (in practice, when he acts as if he is that type), the project is more likely to be carried out, and he will be charged a larger share of the cost.

There *are* agreements which will lead to the project always being done, i.e., agreements with $p_H = p_L = 1$. However, such agreements must have $e_H = e_L \leq 30$, and hence A must bear at least 70% of the cost, independent of B's type. Any alternative agreement which lessens A's burden must have $p_L < 1$, and hence must require that the project is sometimes *not* carried out. For example, if gains from the project are to be split evenly between A and the announced type of B, then we must have $e_H = 50$ and $e_L = 20$; if the project is to be certainly carried out ($p_H = 1$) when the \$90-type is announced, then p_L can be at most 4/7.

One interpretation of this result is that efficiency and equity are, at times, at least partially incompatible. Only the threat by A of not doing the project can "separate" the two types of B, and this threat is only viable if, when B claims to be his \$30-type, A sometimes actually carries it through.

Example 10 (bilateral trade). One classical type of bargaining problem involves a seller and a buyer, each uncertain of how the other values an object currently held by the seller. Assume that each believes the other to be equally likely to value the object at any amount between \$0 and \$300; each, of course, knows his own valuation. According to the revelation principle, the possible agreements which can result from any choice of negotiation format can be characterized by a pair of functions $p(v_S,v_B)$ (the probability that trade takes place when the seller announces his valuation to be v_S and the buyer announces his to be v_B) and $e(v_S,v_B)$ (the amount to be paid by the buyer when these announcements are made and trade *does* take place). This pair of functions must satisfy a continuum of incentive constraints: Each buyer or seller type must prefer announcing truthfully to making any other announcement. Furthermore, the functions must satisfy a continuum of participation constraints: Every seller type must expect to be paid at least his valuation when trade takes place, and every buyer type must expect to pay no more than his valuation.

Consider one particular format for arranging a sale. Each writes down a price. If the seller writes a higher price than the buyer, no trade occurs; otherwise, the object is sold at the average of the two amounts. It is simple to show that it is not in equilibrium for both to write truthfully their valuations: If either is truthful, the other can gain by exaggeration. One natural equilibrium pair of strategies is for the seller to write down $2/3 \cdot t_S + 75$, where t_S is his actual valuation, and for the buyer to write down $2/3 \cdot t_B + 25$, where t_B is his valuation. Notice that when the buyer's valuation is only slightly greater than the seller's, trade does not take place; indeed, if the seller's type is greater than 225, there is never a trade.

Consider an alternative mechanism, wherein each writes down an amount (v_S and v_B, respectively), and trade takes place only if $v_S \geq v_B - 75$, at price of $(v_S+v_B)/3 + 50$. It can be verified that it *is* in equilibrium for both to tell the truth; furthermore, every pair of types faces (at equilibrium) the same outcome here as they did in the previous mechanism. This latter mechanism is, in fact, the "revelation game" derived from the former game and equilibrium point using the approach outlined at the beginning of this section. (It is known that the mechanism given here maximizes the traders' *ex ante* (before they learn their types) expected joint gains from trade; if one were organizing a market within which such traders would be forced to deal, this mechanism would be the natural choice.)

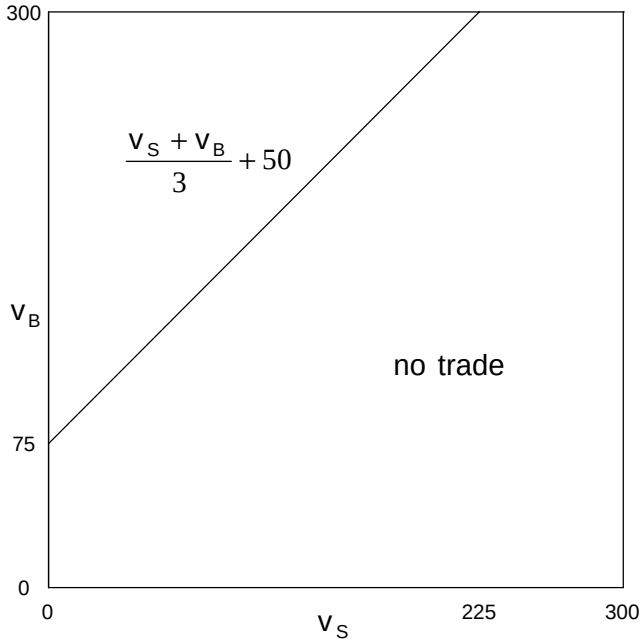

4.4 **Incentive-efficiency**

It would seem foolish for an arbitrator to suggest a mechanism for dispute resolution which left every type no better off, and some types worse off, than some other mechanism would have. In other words, the suggested mechanism should be *incentive-efficient*, i.e., efficient subject to the incentive constraints. A generalization of Nash's solution to the complete-information bargaining problem should therefore select a particular incentive-efficient mechanism.

In Example 9, the incentive-efficient agreements all have $p_H = 1$. Furthermore, if $e_L \geq 10$, then $p_L = (90\text{-}e_H) / (90\text{-}e_L)$.

In Example 10, the mechanism presented can be shown to be one of the incentive-efficient mechanisms. Furthermore, no incentive-efficient mechanism is *ex post* Pareto-efficient: It is impossible to arrange for advantageous trades to always take place.

4.5 **Equity and durability**

Although parties usually enter negotiations with a primary objective of reaching an agreement advantageous to themselves, much of the ensuing discussion between the parties concerns the "fairness" of different proposed agreements.

From where do the parties obtain their notions of what is fair? Certainly, there are commonly accepted principles which are culturally based. "The greatest good for the greatest number," "From each according to his ability; to each according to his needs," and "Whatever can be obtained from the sweat of the brow" are examples of such principles; clearly, they stand somewhat in contradiction to each other.

Sometimes, precedent plays a role in perceptions of fairness. Labor negotiations typically take the previous contract as a starting point, and each party will argue that a concession on one issue "should" be matched by an opposing concession on another. At other times, a neutral third party will be asked to resolve a dispute in terms of his external view of equity: The parties will submit their dispute to binding arbitration.

The Nash solution in settings of complete information was derived from a list of desired properties, at least two of which (individual rationality and symmetry) were directly concerned with equity. Furthermore, the use of threat-making to establish the original conflict outcome carries with it a notion of equity: Those who will suffer relatively more if agreement is *not* reached, receive relatively less from the agreement which *is* reached. (In Example 4, Alfred receives less than half of the monetary gains available from trade with Burton.)

Recently, Myerson has proposed a generalization of the Nash solution to bargaining games with incomplete information. His approach gives explicit regard to the inter-type competition we have previously discussed.

Example 10 (continued). Assume that the mechanism described earlier (each names a price; trade takes place if the named prices are compatible, at the average of the two named prices; each follows his specified equilibrium strategy) is proposed to two traders. At the time of the proposal, each trader of course knows his own type. If the seller's valuation is greater than \$225, he will naturally make an objection. For example, if his valuation is \$250, he expects no trade to take place if this mechanism is used. Instead, he could commit himself to a first-and-final offer of some higher amount, say, \$275. Although the buyer may be antagonized by this action, if *his* valuation is greater than \$275 and he believes the seller's commitment

then there is some chance that he will accept the offer. Similarly, if the buyer's valuation is less than $75, he will object to the use of this mechanism.

Even when the seller's valuation is less than $225, and he agrees to use the proposed mechanism, his mere agreement reveals something about him - namely, that his valuation is less than $225. (Otherwise, the buyer would expect him to object.) With this extra information, the buyer might choose to reject the proposed mechanism, and put additional pressure on the seller.

Consequently, this mechanism is not "durable," in the sense that either it will not be accepted by at least one party, or it will be accepted and the equilibrium will not persist, since the parties' beliefs about each other will be changed by their acceptances.

In response to this difficulty, Myerson has proposed a "neutral" bargaining solution, which takes into account the inter-type conflict which could upset a proposed mechanism. For the buyer-seller problem we are considering, the neutral bargaining solution is summarized by the figure below: Each trader names his valuation; trade takes place if the named valuations (v_S for the seller, and v_B for the buyer) lie in one of the two triangles, at the price indicated within the triangle. It can be shown that this mechanism is both incentive-compatible (it is optimal for each to tell the truth, given that the other does so) and incentive-efficient. Notice that trade takes place more often for the extreme types than under the previous mechanism, and slightly less often for the intermediate types. This must be so, in order for the extreme types to not wish to "upset" the mechanism.