CHAPTER 13: BARGAINING AND NEGOTIATION WITHIN POTENTIALLY COOPERATIVE GAMES

In this chapter we will set up a way of thinking about – a way of analyzing and justifying – a certain kind of decision in which bargaining and negotiations have important roles to play in arriving at a rational choice.

13.1 Bargaining and negotiations: basic concepts

“Bargaining” and “negotiation” are familiar terms, commonly used in fairly standard ways. Many people think of “negotiations” as a discussion process of making demands and compromises intended to settle a dispute, often between parties that have traditional disagreements such as labor and management concerning wages and benefits, or nations with a history of rivalry over territory, or divorcing couples over child custody and property. In this sense, “negotiation” means conflict resolution by a method of discussion that seeks to avoid increased hostility and the damage each side might do to the other.

As for “bargaining,” many people associate it with buying and selling, especially when there is no fixed price attached to an item. Without a set price, an interested buyer (who wants to buy low) is typically invited to make an offer, and the seller (who wants to sell high) is free to accept the offer or to make a counteroffer. This back-and-forth process of offer-counteroffer continues in the hopes of arriving at a mutually agreed exchange price. A person selling her car, or owners selling their house, and potential buyers making bids, are examples that easily come to people’s mind when they think of bargaining.

These are, of course, perfectly good uses of these two terms, but they are not quite the meanings that they will have in this chapter. We are going to consider bargaining and negotiation as methods of practical reasoning used to frame and solve a certain kind of decision problem – called a bargaining problem – methods that ideally not only arrive at a rational choice but also a choice resulting in an outcome that agents should accept as fair.

Here is an example of the kind decision problem we'll be working with in this chapter. A widow and widower, both elderly, have been long time friends. They decided a few years ago to live together for the convenience and company each could provide the other, and especially for the savings each would gain as housemates, for each had only a very modest retirement income. He sold his small house and moved into her small house. For several years, things went as planned; together they managed to save money, to purchase several antiques, to add an exercise room to the house, and even to buy a new car. But now things have gone downhill. Each has become increasingly disagreeable and difficult to live with. They've reached the point of ending their arrangement. Instead of fighting and perhaps turning to lawyers (which would very likely mean an end to their friendship even if they could afford lawyers) how might they cooperatively agree on their mutual gains and fairly divide them up so that they can separate yet remain friends?

If you think back to the kind of decision problems we looked at in the previous two chapters, potentially cooperative games, they will serve as our starting point and as a context for getting a fix on the ideas of bargaining and negotiation. Provisionally, then, we will think of **negotiation** as practical reasoning that results in a decision to cooperate in a potentially cooperative game, and we'll think of **bargaining** as practical reasoning that results in selecting an equilibrium outcome. The possibility of a better outcome for agents through cooperation, and a worse outcome without cooperation, sets up the conditions in which negotiations can (and should) take place. The possibility that the better outcome can give different agents different degrees or amounts of goal achievement sets up the conditions in which bargaining can (and should) take place. Negotiations, then, will be linked to agents coming together to gain the goal jointly, to pool their outcomes (that is: to maximize joint utility), while bargaining will be linked to dividing the goal up among the agents who decided to cooperate.

Let's consider the matrix of a potentially cooperative game that we'll be able to use as an example of a bargaining problem to see how the methods of bargaining and negotiations can be applied. (Note: while it is convenient to think of Row and Col as 2 people, it is important to remember that they could be nations, businesses, clubs, etc. What we cover below applies to such agents as much as it does to personal agents.) Suppose Row and Col must make interdependent decisions that looks like this (we'll skip the story and move directly to the matrix, but if you feel that this is a bit too abstract for you at this point, then think of these two agents as each having to make, say, a financial decision and each outcome unit of utility equals $1000; so, e.g. -5 means $5000 lost, and 2 means $2000 gained):

		Col	
		C	D
Row	C	10, 2	-4, 3
	D	5, 2	-5, -1

<u>Example 1</u>

As you can see, Row can choose the cooperative option for two possible outcomes, 10 or -4, or choose to defect for two possible outcomes, 5 or -5, in each case the outcome depending on Col's decision. Likewise, Col can choose to cooperate for equal outcomes 2, or defect for a payoff of 3 or -1 depending on Row's decision. If each agent makes a decision individually (as if they could not communicate), what principles of practical reasoning should each use to make a rational choice? Each sees that Row C and Col D yield a Nash equilibrium outcome (-4, 3). Also, both Row and Col see that Row has a dominant option C, and so Col (reasoning that Row, being rational, will drop D as an option and choose C), will choose D for a payoff of 3. By equilibrium and by domination, then, this game has a rational choice solution (C, D). Col will do well with a payoff of 3, but notice that Row will do pretty poorly gaining a -4 outcome.

Now suppose Row is very unhappy, perhaps even frightened, that she will have to suffer major loss with a -4 outcome. She thinks how she might avoid -4, how she might do better in this decision problem. She uses the opportunity to communicate with Col and approaches him about the possibility of a mutually cooperative solution. Row offers Col a “deal.” Here is how the conversation might go (check back to the matrix while reading it):

Row: “Col, before we make our separate individual decisions I’d like you to consider something: if you choose to cooperate instead of defecting, I will guarantee that you get 3 instead of 2 as an outcome. I will choose to cooperate, of course, and then I will subtract 1 from my gain of 10 and add it to your outcome 2, and so you’ll have your 3.”

Col: “Interesting offer Row, but why should I go along for just 1? I’m guaranteed my 3 without you. If you want me to switch from D to C you’ll have to do better than that; after all, by switching I’m in effect allowing you to go from -4 to 10, a huge gain of 14. Even if you guarantee me 1 of your units of goal achievement, do you think it’s fair that you gain 13 while I stay the same at 3 for having switched?”

Row: “Why should you care what I get? As long as you remain with 3 at no further cost to you, you should be indifferent between choosing option C or D. But if you are going to be like that, then I’ll guarantee you 2 of my 10 units of goal achievement for choosing to cooperate. 4 is better than 3 and so now you can’t be indifferent, your practical reasoning requires you to accept this offer, for now your option C dominates your option D.”

Col: “Quite the contrary! Reason tells me not to accept. If we each choose option C, then we get a total of 12 (your 10 plus my 2). But I’m the one with the power to switch or not, not you Row. So, I should get at least half and reason tells me I should get even more than half. But as a favor to you, I’ll settle for an even split, 6 each. Why should I accept your offer of 4, when 6 is a reasonable bottom line to expect for switching? And not just reasonable; you should realize how generous I am. After all, by switching from D to C I’m saving you from -4; in a way, I should get all 12 for switching and you still gain by going from -4 to 0.”

Row: “You have me in a bind, Col. Ok, I’ll guarantee you an even split, 6 each.”

Col: “Well, before we make our joint decisions I need to be sure that you’ll follow through. You can’t expect me just to trust your word. You’re in a bind and so might say anything to get me to go along. How about this: you give me 4 now, before I switch. Then we'll make our joint decision to cooperate and I'll get my outcome 2, walking away with a total of 6. You'll have your outcome 10, less the 4 you give me now, and that means you'll equally walk away with a total of 6."

Row: "Are you kidding! First, I just don't have 4 up-front to give you. But even if I did, you can't expect me just to trust your word that you'll switch. With 4 in hand, you might be tempted to choose D and get away with 7, leaving me -8! No, I won't agree to such a pre-decision payoff."

Col: "Too bad. Fine, don't trust me. Let's just forget the whole idea of mutual cooperation and make our decision as separate individual rational agents."

Row: "Ok, but be advised that, due to our little discussion here in which you seem so uncooperative, I will now individually choose to defect. It may be an irrational choice, but it's one that you should think carefully about."

Col: "Hmmm…I see what you mean, Row. About your offer of a guaranteed even split of 6 each for my switching from D to C: Is it still on the table? I hope so, for if not then you should be advised that I will choose D no matter what!"

In this example game, and in the in the hypothetical discussion between Row and Col, we can see that two quite different but closely related things are going on. First, we have the attraction of the mutual cooperation outcome (C, C). It is unequal; one agent gains a large part and the other a much smaller part of the goal. This means that one agent is at a disadvantage in a game like this; the agent who has the greater goal achievement potential has *more* to loose if mutual cooperation doesn't happen. In this example, Row needs Col's decision to cooperate more than Col needs Row's. This game is **asymmetrical**. (We will also be analyzing decisions in which agents are **symmetrical**, that is: equal in their outcomes.) In connection with the attraction of the mutual cooperation outcome, we have the effort of these agents to get each other to "join together" and, rather than make decisions as individual agents, make a "joint decision" to cooperate. There are several ways we might describe this (using some of the ideas that have been introduced in earlier chapters):

(a) going from the standpoint of individual decisions to the standpoint of a social choice, or

(b) going from 2 agents who have an interest in common (the goal) to 2 agents sharing a common interest (the joint benefit of mutual cooperation), or

(c) transforming a potentially cooperative game into a bargaining problem.

This aspect of the example 1, specifically the pattern of reasoning involve in going from an individual decision (C, D) to a joint decision (C, C), is called **negotiations**.

The second thing we see going on in example 1 is that these two agents can look at their decision problem as if the mutual cooperation outcome had already been achieved, and from this point of view they must now decide how to divide it up between themselves. In other words, they leave the standpoint of a joint decision, the standpoint of two agents sharing a common interest, and take the standpoint of individual agents, each desiring as much of the goal as the other will agree to give up. This aspect of example 1, specifically the pattern of reasoning involved in going from the total gain resulting from a joint cooperative decision (C, C) to a division of this gain back to the individuals, is called **bargaining**.

We can abstractly picture these two practical reasoning processes like this:

	Going from		To
Negotiations:	individual rational choice	→	joint decision to cooperate
Bargaining:	cooperation outcome	→	individual decisions about division

You can see from this example that bargaining and negotiation are two closely intertwined processes of reasoning. On the one hand, the result of bargaining (how the goal is to be apportioned to the agents who have decided to cooperate) ideally will present each agent with an attractive enough expectation of goal achievement to justify entering into, and continuing the effort at, negotiations. And on the other hand, it is the success of negotiations that allows the bargaining solution to take effect. Without successful negotiations, bargaining can't take place; but as we will see, bargaining might fail even though negotiations have been successful. As intertwined as these two patterns of practical reasoning are, however (and in realistic bargaining and negotiations situations, they are often so complexly intertwined that great efforts of analyses are required to tease them apart), it is conceptually important to keep them distinct, each having their own principles and methods of practical reasoning.

We will think of a **bargaining problem**, then, on two levels. *Broadly,* we will take a bargaining problem to be any potentially cooperative interdependent decision in which it is rational for agents to enter the processes of bargaining and negotiations. *Narrowly*, we will define a bargaining problem to be a decision problem, resulting from successful negotiations, of sharing or dividing the goal. As mentioned above, we will consider bargaining problems within the framework of potentially cooperative games. This will give us use of the familiar terms and principles we set up in the last chapter to apply in this chapter. Let's look at negotiations first.

13.2 Negotiations within potentially cooperative games: the decision to cooperate

The first potentially cooperative game we looked at was harmony. In harmony, you will recall, agents rationally choose to cooperate as separate individuals; they need not communicate with each other or even know about each other. As individuals separately pursuing a goal, the very nature of harmony is a decision problem that automatically brings rational agents separately to choose the cooperation option, and for each agent the cooperation option yields the maximum utility outcome. In harmony, then, there is no need, there is no place, for negotiations because it is not possible to improve goal achievement.

There are many potentially cooperative games, however, in which agents can gain more of the goal than the individual rational choice of each gains them. This, recall, is the sub-optimal outcome problem we encountered in the last chapter in connection with the stag hunt and the prisoner's dilemma. Games in which it is possible to improve goal achievement (that is: games in which there is a sub-optimal outcome problem) are just the games in which negotiations are needed to gain agents a better outcome than individual rational choice gains them. In order to improve goal achievement in such games the agents, through a process of negotiations, must transform the original game into a different kind decision problem – a bargaining problem. It is important to note, and we will repeat this point, that negotiations does not solve the sub-optimal outcome problem, but rather avoids it by changing one kind of decision problem into another kind that doesn't contain the problem. (By analogy: if your car has a flat tire, you have a problem. If you fix the flat you solve the problem; but if you don't fix it and instead switch to another car, you avoid the problem but haven't solved it.)

Negotiations must bring about 3 central points of agreement between the agents in order to achieve this transformation of a potentially cooperative game containing a sub-optimal outcome problem into a bargaining problem: (1) individual security levels, (2) mutual cooperation payoff, and (3) guaranteed protection. Let's consider each.

(1) The agents must agree on exactly what can be gained of the goal by each if there is a failure to cooperate. This involves determining, for each agent, an **individual security level**. What outcome utility (or disutility) can each agent guarantee herself as a payoff, if each makes a separate individual rational choice? This is an agent's maximin outcome. An agent's security level is not (necessarily) the agent's worse outcome. Also, it is not an outcome that requires the other agent to make a rational choice. It is the maximum goal achievement (or minimum goal loss) an agent can guarantee herself in the given potentially cooperative game, if the agent were to make a rational choice alone. For example, take the game we used above (maximin values are underlined):

		Col: C	Col: D	
Row	C	10, 2	-4, 3	-4
	D	5, 2	-5, -1	-5
		2	-1	

Example 1

By choosing C, Row can assure himself at least a –4 outcome disutility, no matter how Col chooses. This is Row's security level, and it is Row's maximin outcome. Col can guarantee herself at least a payoff utility of 2 by choosing C. This is her security level, and her maximin outcome. Of course, if Col were sure that Row would make a rational choice, Col would choose D for a payoff of 3. But this is not guaranteed, for Row just might be irrational and choose D giving Col a disutility of -1. But Col can be assured of avoiding -1 and gaining at least an outcome 2 by choosing C, no matter what option Row chooses.

For another example, look at this particular (asymmetrical) prisoner's dilemma:

		Col		
		C	D	
Row	C	5, 7	-10, 9	-10
	D	8, -10	2, -3	<u>2</u>
		-10	<u>-3</u>	

<u>Example 2</u>

Row's security level is outcome utility 2 (choosing D), and Col's security level outcome disutility is –3 (choosing D).

In order for agents to agree on each other's security levels, note that they must (a) communicate, (b) agree on how to frame (analyze) the particular decision problem they face, and (c) agree on the utility scales and particular values that are used to represent degrees of goal achievement or loss. It goes without saying that a breakdown in communications makes negotiations impossible, so keeping lines of communication open is a fundamental condition for all phases of the negotiations process. In addition, knowing how to analyze and represent interdependent decision problems <u>objectively</u> and <u>honestly</u> is obviously an important part of successful negotiations, for if one or both of the parties don't understand how the problem is to be framed, or have poor practical reasoning skills, or believe the other side is not negotiating in good faith but rather as a stalling tactic, it is hard to see how they will make any progress in this part of the negotiations.

Utility values can be a particularly troubling sticking point in negotiations. It is important, in this regard, to recall how "utility" has been defined (see Chapter 3.1.2) and that it is not the particular utility scale (the range) being used that matters, but rather the intervals – the gaps – that represent the information agents need to have successful negotiations.

The mutually agreed security level of each agent is represented in the following way:

For example 1: U(R) = -4 U(C) = 2 This reads: In this potentially cooperative game, the outcome utility (disutility) Row is sure of by maximin reasoning is –4, and the outcome utility Col is certain of gaining by maximin reasoning is 2.

For example 2, the security level of each agent is represented in the following way:

U(R) = 2 U(C) = -3 This reads: In this potentially cooperative game, the amount of goal achievement (outcome utility) Row is assured by a maximin choice is 2, and the goal achievement (disutility) Col is certain of by a maximin choice is -3.

(2) The second point of agreement that negotiations must bring about is: exactly what is gained if there is mutual cooperation. The agents must agree on what the total goal achievement is that they can expect from cooperation, over-and-above each agent's security level. What is meant by this is not what each agent individually gains if each separately chooses to cooperate (as if the agents were in a game of harmony). Rather, it is the total *mutual* gain from cooperation that must be agreed on. Here are some examples.

Imagine two college students: one has to buy books for her classes that will cost $300 used at the university bookstore. The other has already taken those classes, has the books, and he is about to sell them back to the university bookstore used for $75. If they can find out about each other and cooperate (she buys the textbooks from him and he sells them to her), then their mutual gain is $225 – the difference between the two university bookstore prices. (The bargaining problem, if they can negotiate a deal, will be how to split the $225 gain. He wants more than the bookstore offer of $75, but how much more should he get? She wants to give less than the bookstore price of $300, but how much less should she give him?)

Suppose two neighboring nations each claim as part of its territory the same square mile of land. Each alone can secure its borders if this disputed square mile in not included. But if they cooperate, then there is a joint gain of 1 square mile over-and-above each nation's security level. If they can't cooperate, neither gets it; it remains disputed. (The bargaining problem, if they negotiate successfully, will be how to divide or share this square mile and its resources.)

Suppose a person is sure she can sell her used car to her local used car dealer for no more than $2000 and suppose another person is sure he can buy this used car from his local used car dealer for no less than $3000. If this seller and buyer can find each other and cooperate (that is: she sells to him, he buys from her), then their joint gain is $1000 – the difference between her sure $2000 sale and his sure $3000 buy price. (The bargaining problem will be how to split this $1000).

Suppose two police departments each separately gather data on suspected criminals. On the basis of the data it gathers, one department is able to make on average 10 criminal arrests per year. On the basis of the data gathered by the other department, it makes on average 25 criminal arrests per year. If these two departments cooperate by pooling their intelligence gathering technology, however, they can together arrest on average 40 criminals per year, a mutual cooperation gain of 5 additional criminal arrests. (If they can negotiate successfully and cooperate, the bargaining problem will be how to share the credit, or reward, or grant money, or new equipment received, etc that this increase of 5 arrests will earn.)

With these examples in mind, let's continue to look at this second agreement point. In this part of the negotiations process, the agents join together into a unit – they form a temporary mini-society – and pool their expected cooperative outcome utilities. This part of the negotiations, naturally, builds on the decision problem analysis and utility assignments that the agents agreed to and already used to discover each one's individual security level. (Note: for convenience at this point, we will represent the expected mutual cooperative outcome utilities by adding the separate cooperative outcomes, but the 4 examples you just read directly above should alert you to the fact that this will need to be revised.) The mutual cooperation gain is represented as follows. For example 1: U(R and C) = 10+2 = 12 This reads: Row and Col together expect a joint cooperative outcome utility 12 in this potentially cooperative game. For example 2 : U(R and C) = 5+7 = 12 This reads: Row and Col jointly expect a mutual cooperation outcome utility 12 in this potentially cooperative game.

We now put these two negotiation results (the security level and the joint cooperative outcome) together. In the theory of rational choice, this is called **the characteristic function form** (which we will shorten to **characteristic form**). For example 1 the characteristic form is:

U(R) = -4 U(C) = 2

U(R and C) = 12

For example 2 the characteristic form is: U(R) = 2 U(C) = -3

U(R and C) = 12

Once the characteristic form of a potentially cooperative game is agreed to, it is easy for agents to see and agree on what will be lost by failing to cooperate: they <u>drop</u> in goal achievement from their expected joint payoff to at least their assured individual maximin security levels.

(3) The third point of agreement between agents that must be negotiated is the hardest, messiest, and contains the greatest danger of failed negotiations. Each agent must somehow be guaranteed that cooperation will not be exploited by the other. If there is no guaranteed protection from the sucker's payoff, no rational agent will continue negotiations. This is especially crucial for the more vulnerable agent in asymmetric games, a worry that the other side should appreciate and acknowledge. There are two parts to assuring agents that their cooperation will not set them up for exploitation: (a) one is the question of how much protection it is rational to demand, and (b) the other is the problem of enforcement. We'll take these in turn.

(a) Take another look at example 2: the sucker's outcome for each is pretty bad; much worse than each agent's security level. No rational agent will accept, or risk, a possible amount of goal achievement (loss!) by cooperation that is less than what the agent can be assured of by an individual maximin choice. Protection from exploitation, in this example, will need to be guaranteed to both agents. Negotiations must in effect void the two exploitation cell possibilities of the matrix, and transform the decision problem:

From this potentially cooperative game:

		Col C	Col D
Row	C	5, 7	-10, 9
	D	8, -10	2, -3

$\longrightarrow$

To this bargaining problem:

		Col C	Col D
Row	C	U(R and C) = 12	~~-10, 9~~
	D	~~8, -10~~	U(R) = 2 U(C) = -3

How much of a guarantee is needed? How much assurance that the other side will not defect is it rational for each agent to demand in this part of the negotiations? It depends on two factors. First, on how tempting the free ride outcome is. We need to know how much more goal achievement the free ride outcome gains an agent over and above what would be gained by that agent if each agent individually chose to cooperate. In example 1, a free ride gains Row 3 utility values (payoff 8 rather than payoff 5). A free ride lets Col go from 7 to 9, a gain of 2. This means that Row is slightly more tempted to exploit Col's cooperation by defecting, than Col is tempted to exploit Row. So, Col should demand a slightly stronger guarantee that Row would not try to gain at Col's expense, than Row should demand that Col will not take advantage of Row by defecting. So, with regard to the free ride temptation, whatever level of guarantee it is rational for Row to accept concerning Col, Col has practical reason not to accept concerning Row, and instead require a <u>slightly</u> stronger guarantee. In this particular game, this fact will be common knowledge and so there should be complete agreement about Col's need for more protection than Row's need.

The second factor is how bad the sucker's payoff is. We need to know how much goal loss an agent must face in going from the agent's security level to the sucker's payoff. As you can easily see, Row has more to fear (going from 2 to -10 = a loss of 12) than Col (going from -3 to -10 = a loss of 7) concerning the possibility of exploitation. So, with regard to the sucker's payoff, whatever level of assurance is rational for Col to demand that Row won't exploit Col's cooperation by defecting, Row has practical reason not to accept concerning Col, and instead require roughly 1/5th more assurance. Again, in this particular game, this fact will be common knowledge and so mutually agreed to.

It is important to remind ourselves, here, that the free ride temptation – the danger of having one's cooperation exploited – is not due to bad, or devious, or greedy agents. Not at all! It is practical reasoning, nothing less, that "drives" a rational agent to make a rational choice, given the goal and the available options. Negotiations, then, are not about protecting oneself from "the other side" as if "your side" represents good and the "other side" represents evil. Rather, it is because each agent assumes the other is <u>rational</u>, and <u>equally</u> skillful at practical reasoning, and it is because this is <u>common knowledge</u>, that negotiations must include provisions against exploitation in the form of secure, reliable guarantees.

What the agents must agree to, then, is a way to devalue the free ride outcome, to make it less rationally attractive than the individual cooperative outcome. This will assure each agent that the other has no rational justification to "stick" her with the sucker's payoff. If the goal were money, the agents might agree on paying a fine for defecting. The amount of the fine would have to be sufficient to bring the value of the free ride below that of the cooperative outcome (e.g. a bigger fine for richer agents, a smaller fine for poorer agents, due to the relative value of money). Or, the agents might agree on forfeiting part of the goal, namely, that part gained by the free ride outcome. Or, perhaps the agents could agree to turn over something important, for instance a future opportunity to cooperate for an even bigger goal or the right to enter future negotiations, to a third party with the power to withhold it should an agent defect.

The general principle of practical reasoning that is being applied here is that negotiations must alter the outcome utilities of the original potentially cooperative game (in the case of example 2, a prisoner's dilemma) enough to transform it into another decision problem: a bargaining problem that is represented by a characteristic function form, not by a 2 x 2 matrix.

Let's work through this third part of negotiations again using example 1. (As above, if you are not comfortable working with just the abstract matrix, then imagine a story – say two college students are considering merging their individual small computer repair businesses and each utility point equals $1000 gained or lost.) We immediately see that the decision situation is quite different from example 2. First of all, example 1 is not a prisoner's dilemma, so we expect the agents not to have the same worries we found in example 2. Note that Row's sucker's payoff is worse than Col's (-4 as opposed to 2), but that these values are the agent's security levels; that is, neither agent need worry that the other's defection will bring goal achievement below what each can already guarantee herself. Nevertheless, Row is much more vulnerable (threatened) when it comes to being exploited by Col than Col is by Row. This is not because Col is powerfully tempted by the free ride payoff (Col would only gain 1 utility value by singly defecting), but rather because Row would experience a large loss of goal achievement (going from 10 to -4) if Col defected. In negotiations here, Row has practical reason to demand a guarantee that Col won't defect, but Col need not worry that Row will defect (Row will loose by Col defecting while Col remains with the same outcome 2 should Row defect). Col can exploit Row, but Row can't exploit Col, and so Col would not be justified making an equal demand for protection (to do so would be an irrational worry).

Negotiations, then, must work to cancel the possibility that Col's will exploit Row; Row needs Col's guarantee. Successful negotiations will transform the decision problem:

From this potentially cooperative game:

Row \ Col	C	D
C	10, 2	-4, 3
D	5, 2	-5, -1

→

To this bargaining problem:

Row \ Col	C	D
C	U(R and C) = 12	U(R) = -4, ~~3~~
D	~~5,~~ U(C) = 2	~~-5, -1~~

(b) What makes this 3rd part of negotiations especially hard, messy, and libel to fail is the **problem of enforcement**. Even if the agents agree on the levels of guaranteed protection against exploitation each rationally requires in order to cooperate, and on what kinds of adjustments will devalue (and thus cancel) the temptation to free ride, they still must rely on an enforcement mechanism or agency strong enough to back up and assure the guarantee. Here is a vivid way to illustrate this problem.

> Suppose a person tries to negotiate with himself about quitting smoking. He says to himself as he finishes a cigarette: "I promise to myself this is my last cigarette. And just to be sure, in case I find myself about to give in to the temptation to have a cigarette, I will enforce my pledge to give up smoking with a threat: I will depriving myself of something I really enjoy. Let's see, it will have to be a future pleasure greater than the pleasure I will get by giving in to a future temptation to have a cigarette." (Let's say that this person loves to see old movies.) He says to himself, "6 months of seeing old movies is much more rewarding to me than having a cigarette, so if I smoke again, then I will not let myself see an old movie for 6 months!"

You can easily see why this smoker's negotiation agreement (with himself) is doomed to fail. Because he believes he might prove too weak to keep his pledge to give up smoking when the temptation to smoke arises, our smoker looks for a way to counterbalance the appeal of smoking, a way to subtract from the enjoyment of a cigarette. To make the counterbalancing threat "real", that is: to make it do its intended job, it must be enforced. But if this person does not feel strong enough to keep his pledge not to smoke any more, why should he feel strong enough to enforce the threat; the latter will take even greater strength, for it must deprive him of something he values <u>more</u> than smoking or it can't work as a threat in the first place. Where is this extra inner strength to come from?

Now apply this to the third part of negotiations. Agents come together to form a unit that has a common interest: goal achievement from joint cooperative. Each agent rationally will demand a certain level of protection against being exploited. This level of protection must be guaranteed with assured enforcement or rational agents will not respect it; without backing, a verbal guarantee might be only so many empty words that, for all the agent knows, could be setting the agent up to be taken advantage of if accepted. The agents can't turn to each other for enforcement, any more than the smoker in our story can turn to the "old-movie-loving side of himself" to enforce the pledge to quit made by the "smoking side of himself". Where can they turn for enforcement?

Here, then, is where negotiations become extremely difficult. "Internal" enforcement mechanisms like honor, moral principles, religious beliefs, duty, an agent's good word, honesty, the agent's sense of fair play, the agent's good character, conscience, …, all such "forces" seem to suffer from <u>insufficient power</u> to provide enforcement. They are, to be blunt, not threatening enough, not binding enough, for they are not sufficiently independent of the agents that are negotiating. Practical reason tends to doubt such internal enforcers, the more so the more important the goal is for the well-being of the agents, for too much is made to depend on trust and faith in the other agent's (flexible!) goodness and not enough to depend on an inflexible process.

If internal enforcement systems are too plastic, what about turning to an outside power? "External" enforcement mechanisms, third parties that are (hopefully!) neutral like government agencies, police, armies, local, national, and international judicial powers,…, are typically <u>too powerful</u>, too binding, and pose too threatening an enforcing agency for agents willingly to submit to; they are too independent of the agent's who are negotiating. The danger here is that such external powers might not only provide enforcement of the negotiated agreement to devalue the free ride payoff or counterbalance it with penalties; they might impose a decision on the agents that is not a rational choice, or provide no way out for an agent who has a legitimate change of heart. Practical reason tends to doubt such external enforcers, the more so the more important the goal is for the well-being of the agents, for too much is made to depend on "blind" inflexible processes and not enough to depend on human flexibility. If such external enforcers, say, grew impatient at the pace of negotiations, agents have every right to fear what they might do. Think about it: would you want the courts to force you to accept a settlement if, for example, you were going through a divorce and in negotiations your disagreement about the joint value of your property or the proper penalty for violating visiting rights was thought to be holding up progress or had reached an inflexible deadline?

The problem of enforcement, as you can appreciate, is a major hurdle in negotiations with no single rational solution. The best we can do is to make some general recommendations based on the way negotiating (in our sense of this term) agents actually seem to deal with this problem. The rules-of-thumb seem to be something alone these lines: when negotiations are informal, agents tend to rely on the <u>stronger</u> of the internal enforcements mechanisms. For example, agents will require each other to make explicit promises, or warn one another of the psychological damage that will result from cheating, or make explicit to each other how a breach of trust will ruin future relations. Agents might appeal to a code of honor by explicitly "shaking hands", or "sacrificing" something valuable like opening a rare bottle of wine and "drinking on it" to seal the agreement. If an agent knows the other is very religious, or very superstitious, then the agent will typically require the other to connect the agreement to these "enforcing" beliefs (swearing on the Bible, for example, or swearing an oath on the health or life of one's loved ones).

When negotiations are formal, agents tend to rely on the <u>weaker</u> of the external enforcement mechanisms. For example, agents will require witnesses, official signatures to document stating that they agree to abide by the terms of the negotiations, or perhaps officially involve a neutral third party with limited authority to penalize cheating. In formal negotiations, then, agents try to make their agreements either public (so that appropriate levels of social norms and pressures can be brought to bear against anyone trying to dodge the agreement), or legal (so that appropriate levels of official force and political coercion serve as enforcement mechanisms), and then build in "wiggle room" and "a way out" in the form of loopholes and limits on the authority's enforcement power.

Negotiations, then, is a process of practical reasoning that transforms a potentially cooperative game containing a sub-optimal outcome problem into a bargaining problem. You can well imagine how many different potentially cooperative games there are that contain the sub-optimal outcome problem. We have worked through the 3-part negotiations agreements for only two (asymmetrical) examples, but the following methods of practical reasoning and the principles of negotiations will generally be the same for every case:

(1) The decision problem must be framed (analyzed) to show agents, and have them agree, that it is a potentially cooperative game containing a sub-optimal outcome problem.
(2) Individual security levels must be clearly represented and agreed on.
(3) The joint mutual cooperation outcome must be clearly represented and agreed on.
(4) Acceptable guarantees, with acceptable enforcement, that there is no longer the temptation to free ride (exploit cooperation by defecting).
(5) Combine (2) and (3) to form the decision problem's characteristic form.

Achieving these 5 negotiation steps will transform the original potentially cooperative game into a bargaining problem.

While the terminology is not 100% fixed, agents who have made progress coming to agreement especially on the 3rd part of negotiations (guaranteed protection against exploitation) are sometimes said to be **negotiating in good faith.** And agents who have made significant progress on all 3 points of agreement are often said to be negotiating seriously or to have entered **serious negotiations**.

EXERCISE:

1) Transform the following potentially cooperative games into bargaining problems. In which of the original games are the agents in a symmetrical or an asymmetrical decision problem? Are the original games stag hunts or prisoner's dilemmas? How might the agents solve the problem of enforcement?

a) 2 gangs want control of the same city territory. Gang R is larger and more powerful than gang C. If they fight for control both fail to gain it, but gang C comes out worse (some of its members leave and join R). If they form a pact to share control of the territory they gain money, members, and power worth U(20) to R and U(10) to C on an interval scale: (-50…0…50).
But if R can use the pact to put C off guard and gain full control of the territory, R can gain money, members, and power worth U(25). If C can use the pact to put R off guard and take over the territory for itself, C gains money, members, and power worth U(25). The worse case, being suckered by the pact, would mean the end of the gang, all its members would join the other gang, worth U(-50). How should these gangs negotiate?

b) 2 neighboring countries have been engaged in a joint effort to search for oil deposits along their shared border. It is an increasingly expensive undertaking, but there is a very good chance of finding significant oil reserves that they would shared (worth U(15) for each). However, the government of each county is going through unstable times, causing secret doubts on each one's part about the other's ability to continue the search for oil. Meanwhile, each country could put its efforts and resources into tourist industry development, not as much of a national asset as oil, but much better than continuing the search for oil alone, an effort that can't succeed with the resources of just one side. Of course, if both countries develop their tourist industry potential, neither will do as well as one would if it could be the only tourist destination in the region. If the goal of each country is economic prosperity, how might negotiations help these countries achieve their goal?

13.3 Bargaining and the bargaining problem within potentially cooperative games

A bargaining problem is a decision problem about the best way the jointly achievable goal is to be shared by or divided up between agents. The agents who have successfully negotiated will each want "my fair share", nothing less – indeed, that's why they negotiated to transform a potentially cooperative game containing a sub-optimal outcome problem into a bargaining problem in the first place. The method of practical reasoning called bargaining will ideally achieve this "fair share" division of the goal (or that part of the goal achieved by mutual cooperation).

A bargaining problem is represented by a characteristic form:

$$U(R) = x \qquad\qquad U(C) = y$$

$$U(R \text{ and } C) = z$$

where x, y, and z are 3 outcome utility (or disutility) values, and "R" and "C" represent any two individual agents (humans, organizations, nations, etc.). This reads: x is agent R's security level, and y is agent C's security level, in this bargaining problem, and the best joint cooperative outcome R and C can achieve together in this bargaining problem is z. What we want to do now is imagine that agents R and C have achieved z (better: we want to have agents R and C imagine that they have achieved z), and are looking for the best way to divide z between them. Depending on the goal, "divide" might mean breaking z into parts and each agent taking a part, as two nations might do with territory or two businesses might do with profits. Or it might mean sharing time with z, as a divorcing couple might do with a vacation house or children, or two drivers must do sharing the same car. Either way, the agents must decide together, that is: they must come to an agreement, how best to distribute z to each as individuals. This is their bargaining problem, and bargaining is the process of practical reasoning that will result in a solution.

We'll first look at a convenient way a bargaining problem, in general, is set up and within this framework define a few key terms in the theory of bargaining. After this, we'll turn to the general question of what a rational choice solution to a bargaining problem would be. We will then, in section 13.3, apply this bit of theory to several examples.

Suppose that successful negotiations has transformed

<u>this potentially cooperative game:</u>

		Col	
		C	D
Row	C	10, 10	-10, 5
	D	5, -10	0, 0

<u>to this bargaining problem:</u>

		Col	
		C	D
Row	C	U(R and C) = 20	~~-10, 5~~
	D	~~5, -10~~	U(R) = 0 U(C) = 0

You will recognize that this potentially cooperative game is a stag hunt and that it contains a sub-optimal outcome problem. (As above, if you are uncomfortable not having a concrete decision problem described, feel free to go back to Chapter 12 and use any of the stag hunt examples as the problem this matrix represents.) Also note that the agents are symmetrical; if we could "switch" them without their knowing it, they would not be able to tell that they had been switched. They are indifferent as to who is in the Row spot and who is in the Col spot. The characteristic form is: U(R) = 0 U(C) = 0

U(R and C) = 20

A convenient way to visualize this information is to set it up as a graph (a coordinate space). Let's place Row's utility values along a line (it is standard practice to make it the horizontal or x axis), and Col's utility values along another line (the vertical or y axis). Because the security level of each agent is 0 and their joint cooperation outcome is 20, we will need an interval utility scale at least from (0 – 20); make it (0 – 25). And we want to enclose the space that represents every possible division between Row and Col of the joint outcome 20. Our graph looks like this:

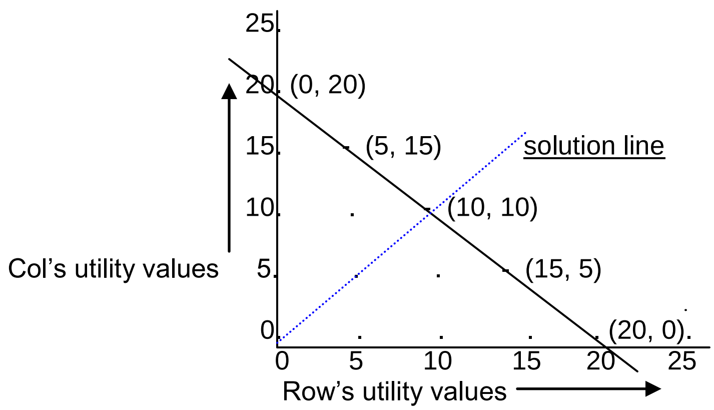

The area within the triangle from the point (0,0) up to the point (0,20) down to point (20,0) and back to point (0,0) is called the **bargaining space**, or **bargaining room**, or **bargaining region**. Note that every point within this space is associated with two numbers (x,y); x = Row's outcome utility at that point, and y = Col's outcome utility at that point. For example, the point (5,10) means that Row gains 5 of the joint cooperative outcome 20, and Col gains 10 (twice as much as Row) of the joint cooperative outcome. The point (7,3) means that Row's share of 20 is 7, and Col's share is 3. The point (0,20) means that the result of bargaining gives Row nothing and Col all of the joint cooperative outcome. The general idea is that bargaining must result in a division of the joint cooperative outcome represented by a point within the bargaining region.

The side of the bargaining space triangle from point (0,20) to point (20,0) is special. It contains all the possible divisions of the joint cooperative outcome that sum to the full value of that outcome – they all sum to 20 in this example. This line is called the **bargaining line**, and the possible divisions of the goal that lie on this line are called the **bargaining set** or the **negotiation set**. We can easily see that any split of the joint cooperative outcome that is on the bargaining line form an equilibrium pair; that is: no agent can singly improve her share by switching to another point, given that the other agent stands firm. The bargaining set, then, contains multiple equilibrium bargaining outcomes.

The dotted line running from the point (0,0) toward the upper-right is called the **solution line**. Row wants the tip of the solution line to extend out toward the bargaining line and wants to pull the line down toward Row's utility values, the best division for Row being the point (20,0). Col, naturally, also wants the tip of the solution line touching the bargaining line, but pulled up toward Col's utility values, the best split for Col being point (0,20). The ability of an agent in a bargaining problem to move the solution line in a favorable direction is called **bargaining ability** or **bargaining power**.

The flip side of bargaining power is the idea of **concession**. Whatever point in the bargaining space that ends up being the solution to the bargaining problem, at least one agent, but typically both, will end up with a share that is less than the share gained in the agent's most favorable division. Both agents can't end up with their most favorable division, or there wouldn't be a bargaining problem in the first place. A concession, then, is the share an agent loses to the other agent's bargaining power; that is: an agent concedes the amount of the joint cooperative outcome between the agent's most favorable share and the share received in the solution. So, for example, if Row and Col end up with a (8,9) split, then Row has conceded 12 to Col's bargaining power (12 = the difference between Row's most favorable point (20,0) and the accepted (8,9) division. Col has conceded 11 to Row's bargaining power (11 = the difference between Col's most favorable point (0,20) and the accepted (8,9) division. But if, for example, Row and Col settle on a (0,20) division, then clearly Row has conceded everything of the joint cooperative outcome to Col's bargaining power and accepts nothing more than her security level.

Finally, let's examine the important point (0,0). This point represents several ideas that are fundamental to the theory of bargaining. First, if the agents cannot reach a solution to the bargaining problem, if they cannot settle on any division of the joint cooperative outcome, this point represents **bargaining failure**. In effect, the bargaining problem reverts back to the original potentially cooperative game containing a sub-optimal problem, and agents make their separate individual rational choices. In the example we are working with, maximin equilibrium makes (D,D) the unhappy rational choice. So, the triangle's origin point (in this example it's (0,0)) indicates the outcome that will result in case the agents fail to agree on a solution to their bargaining problem; this is a bargaining agent's "cushion".

There is a second central idea that (0,0) represents: **fundamental bargaining power**. In order for there to be a solution to a bargaining problem, agents must bargain; that is: agents must apply practical reasoning in a systematic way in an effort to discover a rational choice solution to the problem of a fair division of the benefits of cooperation. If one agent backs out, the whole bargaining problem falls apart. One agent alone can't create or sustain, much less solve, a bargaining problem. This means that each agent has the power to deprive the other of any share of the benefits of cooperation over-and-above that agent's security level. The ability *not-to-cooperate* is a basic bargaining power that each agent must respect in the other, if a bargaining problem is to reach a solution. This ability to deprive, this power to "punish", the other agent simply by "walking away" and exiting the game is a very valuable threat, and keeps an agent from accepting too-readily less than a fair share. In a bargaining problem, a **threat** is a communication from one agent to another that the agent has a course of action that will be more costly to the other agent than it is to the one agent, and that the agent intends to take this course of action if the other agent causes a certain condition to arise. The course of action each agent has is: cease bargaining. For example, in the bargaining problem we are working with suppose that Row demands a (19,1) split, and refuses to consider any counter-offer. On what basis can Col refuse? 1 is surely better than 0, Col's security level, so it seems that Col should accept the offer – unless, that is, Col has the bargaining ability to pull the solution line in a more favorable direction. And clearly Col has this power; Col can not only refuse to accept the offer, Col can threaten to "leave the table" thereby plunging Row to his security level. Note that the threat, if carried out, hurts Col as well, but not nearly so much as bargaining failure costs Row – if, that is, Row sincerely hoped for a (19,1) split. Given Col's fundamental bargaining power and willingness to use it as a threat, Row must concede more than 1 or accept bargaining failure.

There is a third key idea in bargaining theory that the point (0,0) represents, and it is the most important: it is a measure of independent goal achievement each agent "brings to the bargaining table". It represents each agent's initial bargaining status. It is a measure, so to speak, of each agent's basic "worth" relative to the bargaining problem. In this respect, it serves as a reference point, a fixed base, a guaranteed fall-back position, for agents to claim a share of the joint cooperative outcome. The point (0,0) is referred to as the **status quo** point to highlight this particular idea. There are two important possibilities concerning the status quo point. (1) When x and y are equal (as they are in the example we are working with, namely $x = 0$ and $y = 0$), the agents are in a symmetrical bargaining problem, and should regard each other as equals.
(2) When x and y are unequal (as they are in the characteristic forms for examples 1 and 2 in section 13.1 above on negotiations), the agents are in an asymmetrical bargaining problem, and should regard each other as unequal; that is, they both regard one agent "worth" more than the other relative to the bargaining problem.

To avoid possible error, let's consider carefully what "equal worth" and "unequal worth" mean in the context of a bargaining problem. The idea here is not that some agents – be they people, organizations, or nations – are better than others in some absolute sense (whether they are or not is certainly not for rational choice theory to decide). Rather, the idea is that an agent might merit or deserve a bigger share of the goal than another agent based purely on a restricted relevant measure or criterion. Such a measure or criterion must derive from the goal. Here is an example to consider. Suppose two friends, one tall and one short, pool their savings to buy a car; their goal is to own a car as a reliable means of transportation. The short friend contributes, say, \$5000, and the tall friend contributes, say, \$2000. They buy a \$7000 car, which is much more reliable than any \$5000 car or \$2000 car each could have bought on their own. Now they must agree how to share the use of the car. On what basis should they go about deciding on a "fair share"? With respect to height, they clearly treat each other as unequal, but this inequality is not relevant to the issue of car use. You can see how odd it would be for the tall friend to claim that she should get more use of the car than the short friend simply because she is taller. As friends, however, they treat each other as equals; for each friend, one is not any better than the other as a human being. But as with height, his should not enter into the picture either, for it is not relevant in figuring out how to divvy up use of the car (if it were, then every person whom these friends treat as equal human being should be given an equal share of the use of their car!). To agree to equally share car use because as human beings they are, on some fundamental level, equals would be to miss what is special about their particular situation. Clearly, we all see that (at least one) relevant measure or criterion of "worth" in this example is the amount contributed toward buying the car. Were it not for the ability to buy the car, there wouldn't be the use of the car to bargain over. In this narrow respect, then, the friends are seriously unequal, and both should realize this. In effect, the short friend "owns" 5/7 of the car and the tall friend only 2/7, based on the initial contribution toward purchase. So, out of every 7 days the short friend should get the car for 5 days and the tall friend 2 days, or perhaps out of every 7 hours the one should have the car for 5 hours and the other for just 2 hours. The tall friend might like the car more, or need the car more, or whatever…. But unless the short friend concedes some of his 5 days or 5 hours car use to his tall friend (for whatever reason: generosity, a trade for something else, payback for past favors, needs it less, …,), the 5-short – 2-tall split gives each friend just what is merited or deserved as narrowly measured by their "unequal worth"; that is: their initial financial contribution that made goal achievement (owning a reliable car) possible in the first place.

13.2.1 <u>Arbitration: rational choice solution to a bargaining problem</u>

Now that we have set up a general way to represent a bargaining problem (that is: displaying its characteristic form as a graph whose lines and points represent key bargaining concepts), we'll look at a method of finding a rational choice solution. According to the common use of the word "bargaining", one might expect that the agents will now make offers and counter-offers to each other of division points within the bargaining space until they both agree on a single point, which then is their solution. Not at all! This would leave the bargaining process much too open to irrational choice. After all, two (real) agents might agree on a solution point out of pure exhaustion, or just to “get it over with”, or just to be able to “move on”, or out of confusion, or because of any number of psychological motives having little to do with a rational or a fair solution. Ideal rational agents will avoid such a hit-or-miss, back-and-forth offer-counteroffer process. Instead, as we have seen in earlier chapters, practical reasoning applies a set of principles to discover a rational choice, in this case principles that will pick out a point in the bargaining space as a rational choice solution to the bargaining problem. Such principles must embody both norms of practical rationality and norms of fairness, so that the bargaining point that satisfies these principles will be, relative to these principles, a rational and a fair choice. In the theory of rational choice, using such a set of principles is called **arbitration** or an **arbitration scheme**. So, the rational choice solution to a bargaining problem is the result of arbitration.

The arbitration method we will use is based (very!) loosely on the Nash arbitration scheme (due to John Nash, whom we met in Chapter 11 in connection with equilibrium points in potentially cooperative games). There are two principles that are taken to be conditions of rationality, and two that are conditions of fairness. We’ll use the general characteristic form above as our reference: U(R) = x U(C) = y

U(R and C) = z

where x, y, and z are 3 outcome utility (or disutility) values, and “R” and “C” represent any two individual agents (humans, organizations, nations, etc.). Recall that this reads: this is a bargaining problem, resulting from successful negotiations, in which x is agent R’s security level, and y is agent C’s security level, and the joint cooperative outcome R and C can achieve together is z.

(1) <u>A bargaining solution must not give an agent less than the agent’s security level</u>. In other words, R’s share of z must be more than, if not at least equal to, x. And C’s share of z must likewise be more than, if not at least equal to, y. This is a practical rationality principle saying, in effect, that no rational agent will agree to a division point that leaves the agent with less of the

goal than the agent can gain by not cooperating. To do so would be an irrational choice, given the goal.

(2) A bargaining solution must be a point on the bargaining line. This rationality principle says that a rational solution must be Pareto-optimal; that is: it gives agents shares that sum to z. In this way, all of z gets divided up with nothing wasted, nothing left over to tempt an agent to gain it behind the other's back or to give any agent even the suspicion that this might be possible, and nothing going to a 3rd party who was not involved in the negotiations (not a stakeholder). Another way to think of this principle is this: A bargaining solution must be a Nash equilibrium. If it weren't, an agent could gain a larger share by individually moving to another point in the bargaining region, provided the other agent accepted her original share, and so it would be an irrational choice to decide on a sub-optimal solution point.

(3) A bargaining solution must respect the proportion of the goal achievable at each agent's status quo point. This is a principle of fairness that has a long tradition in the field of Ethics and goes back to Aristotle who argued that a fair distribution will treat equals equally, and unequals unequally in proportion to their inequality. In terms of the bargaining concepts introduced above, this principle says that if the problem is symmetrical, the solution should be symmetrical; and if the problem is asymmetrical, the solution should be asymmetrical in the same proportion. What is key – one wants to say beautiful! – about this principle, is that it makes the agent's initial status quo point a relevant measure of "worth" or "deservingness". How odd, and how unfair, it would be if an agent's share of the joint cooperative outcome made the agent's individual goal achievement ability irrelevant. The "dark" side of this principle, as you might have guessed, is that it makes the first point of agreement in negotiations, fixing the agent's individual security levels, very contentious. Agents might try to introduce all kinds of self-serving "measures of worth" into the original potentially cooperative game in an effort to boost their initial bargaining status. This is why the next principle is needed.

(4) A bargaining solution must not be influenced by irrelevant norms, measures, options, or characteristics of the agents. This is a fairness principle that works back-to-back with the 3rd principle. It is intended to exclude things like: manipulation of the utility scales, appeals to goals and values that are not explicitly defined as part of the specific goal of the bargaining problem under consideration, appeals to unrealistic options, to the social or economic status of the agents, the whim of the ruler, the impatience of an enforcing agency, … etc. Principle 4 reminds us that a bargaining problem is a decision problem, and a decision problem is framed and analyzed according to strict guidelines for good reason. Once the goal (and its objectives) is defined, the options listed, and the problem put into the right model, the line between what is relevant for a rational choice solution (if there is one) and what is irrelevant becomes pretty clear. We can think of principle 4, then, as a kind of "anti-discrimination" principle. It says that once we account for any inequality falling under principle 3, in all other ways the arbitration scheme and the bargaining solution must treat the agents impartially as equals; it says that equality is the "presumption" unless inequality is justified by principle 3. This means, in effect, that once we subtract the relevant inequalities allowed under principle 3 from the agent's status quo point, the remainder must represent the agents as symmetrical.

Let's now apply this arbitration scheme (that is: these 4 principles of fair rational choice for bargaining problems) to the example above that we displayed as a graph. If we let the solution line move within the triangle according to these 4 principles, where will it stop?

(1) By principle 1, we eliminate all points containing 0, for (0,0) is what Row and Col can achieve without cooperating. Thus, (1,0) is out, (0,20) is out, (20,0) is out, (0,0) is out, (15,0) is out,…, etc.

(2) By principle 2, the solution must be on the bargaining line (that is: sum to 20), exclusive of the two end points of this line containing 0's.

(3) By principle 3, we see that the agents are symmetrical; that is, they have equal status or worth as determined by their security levels. So, they deserve equal shares of 20. The only point satisfying principle 3 that has also satisfied principles 1 and 2 is (10,10). They each must make equal concessions to the other's equal bargaining power.

(4) By principle 4, there is nothing else that requires either agent to concede any part of 10 to the bargaining power of the other. There is no relevant respect in which Row and Col are unequal in this particular bargaining problem.

By this arbitration scheme, then, (10,10) is the rational choice, a fair division of the benefits of mutual cooperation. You are probably thinking that the (10,10) solution was obvious all along. You are right, but the methods of practical reasoning and the principles of fair rational choice were not obvious, and these will serve us well for bargaining problems in which the fair rational choice isn't so obvious.

Let's note two features of the (10,10) solution that will help us set up and solve asymmetric bargaining problems. First, (10,10) is the midpoint of the bargaining line. Second, (10,10) obeys an interesting formula: subtract Row's security level from Row's share of the joint cooperative outcome, and then do the same for Col. (In this example we get 10 – 0 = 10 for both Row and Col). Now multiple these two remainders. What we end up with is the largest product of all the points on the bargaining line. So, 10 x 10 = 100, but the nearest point down (if we stick with whole numbers) is (11,9). (11 – 0) x (9 – 0) = 99. Nash proved that the point in the bargaining space that yields the largest product after subtracting the status quo values is the only one that satisfies his arbitration scheme. Isn't that lovely!

13.4 <u>Bargaining and negotiations: working through examples</u>

We will now practice the methods of practical reasoning involved in bargaining and negotiations on several examples. Here are the steps to solving a bargaining problem.

1) Negotiation steps:

   (1) Represent the decision problem as a potentially cooperative game containing a

sub-optimal outcome problem. (Chapter 12)

(2) Transform the potentially cooperative game into the characteristic form of a bargaining problem, by having agents agree on:

(a) their individual security levels

(b) the joint gain achieved by cooperating

(c) enforcement of the agreement not to exploit the other's cooperation

2) Bargaining steps:

(1) Represent the characteristic form as a graph containing:

(a) the bargaining region

(b) the status quo point

(c) the bargaining line

(2) Apply the arbitration scheme; that is: find the point on the bargaining line that satisfies the 4 principles of fair rational choice.

3) Rational choice solution for a bargaining problem: divide the goal between the agents according to the values represented by the point singled out by the arbitration scheme.

Example 1: Students merging

Rita and Carol are students starting their 2nd year of college. During their 1st year each worked part-time to help pay for college by offering house cleaning services to local residents. They had a goal in common: to help pay for college by providing a house cleaning services. Of course, they were competitors that 1st year and so each limited the goal achievement of the other. Over summer break, they each got the idea of talking to the other when school starts about the possibility of joining their efforts to enlarge their earnings; they could, each felt, do much better together than separately. But also each saw the potential to "make a little extra" on the side by doing house cleaning without letting the other know, once each had the cooperation of the other. Likewise, each saw the possibility that one could "play sick" and let the other do the work while the free-rider did more enjoyable activity. You know both students, and have told them about a practical reasoning course you took. They come to you to help negotiate the details of merging their separate house cleaning jobs to see if it would be worth it. What fair and reasonable solution to their bargaining problem can you suggest that would enable them to decide if they should merge?

Your first step is to find out what sort of game Rita and Carol are in. You find that they agree on a common goal: to provide a part-time house cleaning service to help pay for college. You also get them to define this goal in terms of 3 objectives that they evaluate equally: (1) attract local customers, (2) clean house well to assure repeat business, and (3) make money. Using these 3 objectives as criteria, you bring Rita and Carol together to evaluate their statuses, something that they must both agree on point by point. (Note that several steps in deriving criteria from the goal analysis are not being shown, in order to focus on the steps of negotiations. A review of Chapter 2 is advised, if you are unsure how this evaluation works.)

As to (3) their current money situation, Rita admits that she's owed $300 from customers for cleaning services, and that Carol would have to take on this deficit if they merge. On a scale of (-10…0…30) they agree to give this -3. Carol, on the other hand, is neither owed back payment nor has any saved money, it all went toward college costs; for this they agree Carol gets 0.

Concerning (2) their house cleaning ability, both agree that Rita is a whiz and likes doing it, whereas Carol is only average at house cleaning and doesn't like the work all that much, especially cleaning bathrooms. They agree to rate Rita 10 on cleaning, and Carol only 3.

Finally, with regard to (1) attracting customers, Rita is not very good at bringing in business, whereas Carol has a better personality in this respect and finds it easier than Rita to get local residents to try her service; but because Carol is not into house cleaning, she tends not to get repeat business. They agree that Rita gets -4, and Carol gets 5 under this attribute.

You now total up Rita's utilities: -3 + 10 + -4 = 3. The same with Carol: 0 + 3 + 5 = 8. If they go it alone, don't merge and continue to compete for business, you have an idea how their strengths and weaknesses compare. Rita: 3, Carol: 8. You confirm this when they tell you that for the first year Rita earned around $2500 where as Carol earned close to $6000.

You can see that they would make a great pair if they merged, for the weakness of each matches well with the strength of the other. They agree to give the full 30 to this prospect of a joint outcome. But the worry that the agreement to cooperate will be taken advantage of behind the back of the other must be addressed openly. Rita and Carol aren't friends, they have come to college from different parts of the country and don't know one another very well. All they are interested in is looking into the possibility of a part-time partnership in a house cleaning service. Being made a sucker will not only take away customers, it will also deeply upset the victim. They agree to give the sucker's outcome -5, and 10 to the temptation to cheat, for there are certainly rewards in doing some house cleaning "behind the back" of the other if they were no longer competing for local customers, and each also realizes how tempting it is to let the other do the work so that the one who "called in sick" can enjoy other activities or get some extra study time.

Now you are ready to put this information into a game matrix:

		Carol	
		C1: merge	C2: go it alone
Rita	R1: merge	30	-5, 10
	R2: go it alone	10, -5	3, 8

You recognize that this is an asymmetrical stag hunt (assuming that the mutual cooperation cell splits the 30 outcome so the neither get less than 10). If Rita and Carol make individual rational choices, maximin equilibrium makes (R2, C2) the rational choice, clearly sub-optimal.

Now the messy work of transforming this potentially cooperative game with a sub-optimal outcome into a bargaining problem must begin. The security levels of each agent are clear: 3 for

Rita and 8 for Carol. But what will overcome their fears of being exploited. Many enjoyable campus activities arise that frequently make cleaning a customer's house an annoying task, so the worry is real. Because the free-ride and the sucker's outcomes are equal for each agent, and of equal distance from the cooperative outcome, you conclude that they will require equal assurance of non-exploitation. You get them to agree on the following enforcements (and hope for the best!):

1) They have mutual college friends who have said that they will quickly inform the other if they hear that either Rita or Carol is violating the partnership.

2) They make an explicit promise not to service any customers without letting the other know, and you are a witness to this promise.

3) For each job they do, each will get the signature of the customer on a piece of paper describing the cleaning done, the amount paid, any extra tips, etc. and both get to see these.

4) For any "covering" one does that gives the other time-off, the other must "cover" for the one to have equal time off – no questions asked, or excuses needed as to how the time-off gets used.

With this last step completed, you have transformed the original game into the following bargaining problem:

		Carol	
		C1: merge	C2: go it alone
Rita	R1: merge	30	~~-5, 10~~
	R2: go it alone	~~10, -5~~	3, 8

The characteristic form is: U(Rita) = 3 U(Carol) = 8
U(Rita and Carol) = 30

The negotiations are now complete, and the bargaining problem of a fair division of the expected gains from merging is ready to be addressed. You must first represent this characteristic form as a bargaining graph, and then find the rational choice point that satisfies the arbitration scheme. It will, then, be both rational and fair, according to the 4 principles of arbitration, if Rita and Carol decide to merge and to share the goal achievement resulting from their house cleaning services alone the lines this point represents.

Here is the bargaining graph (in solid lines) with the work filled in (in dotted lines):

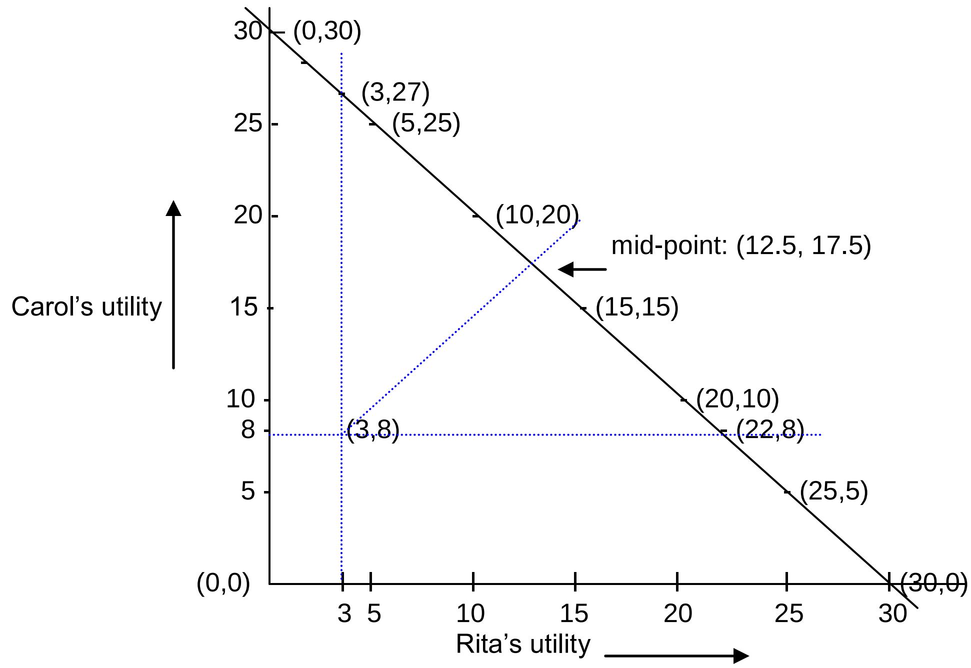

You first must depict the general bargaining space. Rita's (Row's) utility scale goes on the horizontal line (the x axis), and Carol's (Col's) is on the vertical line (the y axis). The bargaining space is the triangle: (0,0), (30,0), (0,30). Now you apply the 4 principles of the arbitration scheme.

1) By principle 1, you exclude all points in the bargaining space giving Rita less than 3, and all points giving Carol less than 8. You are left with a new smaller triangle within the original one, the triangle bounded by the points: (3,8), (3,27), (22,8).

2) By principle 2, the solution must be a point on the bargaining line between (3,27) and (22,8).

3) By principle 3, you must treat the point (3,8) as the new initial bargaining status quo point rather than the original point (0,0), for this respects the relevant inequality, the different worth or deservingness, between Rita and Carol in this bargaining problem. The solution line moves out from the point (3,8) instead of point (0,0), showing that this bargaining problem is asymmetrical

and the solution point treats the agents as unequal in proportion to their initial bargaining statuses.

4) By principle 4, once you have accounted for their relevant inequality by principle 3, you must now treat Rita and Carol as equals in all other ways. They must concede to each other equal amounts of goal achievement. The solution line must divide the shorter bargaining line equally, and it does this at the point (12.5, 17.5).

Point (12.5, 17.5) is the only point that satisfies these 4 principles. You now tell Rita and Carol that a rational and fair share of their joint goal achievement from merging their individual house cleaning services, according to the methods and principles of bargaining and negotiations you used, is the ratio: 12.5 for Rita and 17.5 for Carol. This split is the rational choice solution that is clearly based on the relative strengths and weaknesses that each agent brings to the proposed merger, by their own agreement, given the goal they are trying to achieve. For Rita, 12.5 is a better outcome than 3, and for Carol 17.5 is a better outcome than 8 (each better by an equal value!), so they should merge their house cleaning services on these terms rather than continue to go-it-alone.

An alternative way to find the point that satisfies the arbitration scheme is to use the Nash formula mentioned above: what two number x and y from the bargaining line has the largest product if you subtract 3 from x and 8 from y? You'll see that both 12.5 – 3 and 17.5 – 8 = 9.5 (subtracting the different "worths" leaves the agents equal), and that 9.5 x 9.5 = 90.25, a larger product than any other possible pair of utility values from the bargaining line will yield after the subtractions.

Pause:

Before turning to the next example, let's note some questions or objections someone might have – perhaps you yourself have them – with this bargaining solution. Addressing them will help deepen our understanding of the practical reasoning involved in bargaining and negotiations.

1) Not fair! If Rita and Carol do equal house cleaning work, why shouldn't they divide evenly any gain they achieve from cooperating? (We can think of this as the egalitarian objection. Egalitarians believe that people are fundamentally equal and therefore fairness requires equal sharing of any benefits or goods that result from people cooperating.)

Answer: The amount and the quality of work Rita and Carol do is only part (only 1/3) of the goal's objectives used to arrive at their individual security levels. If they don't merge and each continues their separate house cleaning service, they both believe and expect Carol to do better than Rita (by their own admission). But if, at this point, these agents believe that they didn't analyze the goal's objectives correctly (that is: the criteria of evaluation aren't weighted correctly), they are free to go back to this step in the negotiations and reconsider their agreement. A change in one or both agent's the initial bargaining status will certainly affect the bargaining solution.

2) Not fair! What if Rita needs more money than Carol because perhaps she has larger college expenses, or comes from a poorer family; shouldn't she be the one to get more out of a merger? (We can think of this as the social justice objection. Advocates of the social justice theory, Marxist for example, believe that the benefits of cooperation should be distributed in a way that most helps the most needy in society.)

Answer: If meeting larger college expenses or compensating for poverty matters to these agents, then these desires must be reflected in their goal. They were not part of the goal, and so must be treated in this bargaining problem as irrelevant. As we learned from Chapter 2 and 3, the criteria by which we assign utility values to outcomes come from the goal; once the set of criteria are formed, agents don't have to worry that "outside" influences will affect practical reasoning, for they literally "can't count", they have no weight. This is one of the "beauties" of the system. It would be very poor practical reasoning to overturn a rational choice on behalf of desires that were never entered as values into the reasoning, for this would make "rational choice" nothing but what agents desire! Likewise, it would be very poor practical reasoning if important desires were never entered as values in the first place, so that the “rational choice” solution resulted from an incomplete or distorted version of the decision problem. Of course, Rita and Carol are welcome to start over and to re-identify the goal so that things like college expenses and levels of poverty are brought into the picture, if they so desire. As a general rule: agents are always free (given the time and opportunity) – and indeed should try – to correct any step in solving a decision problem that they believe needs readjusting.

3) Not fair! What if one of these students deserves more money because her major is a more socially valuable field than the other’s major (like nursing or elementary school teaching compared to, say, a philosophy major!), or because she is a better and more serious student than the other? None of this seems to matter in a 12.5/Rita, 17.5/Carol split. (We can think of this as the utilitarian objection. Utilitarians believe that the benefits of cooperation should be distributed in a way the does the most good for the greatest number of people.)

Answer: This objection relates directly to the 2 fairness principles in the arbitration scheme. Rita and Carol, with regard to the bargaining problem they are in, are indeed unequal in what they “deserve” as the portion of goal achievement they should concede to each other. But, to repeat a important idea introduced above in connection with the third meaning of the status quo point, equality or inequality between agents is determined by very restricted measuring standards that come from the goal and, as part of the negotiations, are agreed to by the agents. To award Rita and Carol different amounts of money from their merged house cleaning service on the basis of differences between them that they have not agreed to, or that don’t come from their goal (like the social usefulness of each one’s major, or how serious a student each is), would be to violate the 2 fairness principles (especially principle 4) and thus be a form of unjustified discrimination. If, for example, the college were awarding scholarship money, then things like usefulness of major and student seriousness would probably count a great deal. But in Rita and Carol’s original stag hunt game and subsequent bargaining problem, such things are not allowed to sway the decision from the “outside”, given that they were never entered as criteria (i.e. attributes +value) that the fairness principles could apply to.

4) Not fair! Who are we to tell Rita and Carol whether they should merge their house cleaning services or not, or how they should share the benefits if they merge? They should be left alone to work it out for themselves, and whatever arrangements they voluntarily agree to are fair. It doesn’t matter if they arrived at their “contract” by a rational or an irrational process, so long as it’s voluntary and that they are consenting adults. (We can think of this as the libertarian objection. Libertarians believe in the basic freedom of individuals to decide things for themselves, and governments, religions, cultural traditions, or theories (like rational choice theory) shouldn’t be interfering or imposing schemes that limit individual liberty.)

Answer: This objection is interesting. On one level it simply does not apply; Rita and Carol (as we have set the problem up) have themselves turned to the theory of bargaining and negotiations for help and recommendations to solve their decision problem. Nothing forces them to accept what the theory presents as the rational choice solution. But on a deeper level, this objection seems to call into question one of the basic assumption about rationality that we set out in Chapter 1: that rationality is <u>normative</u>, and this means that people who reason automatically agree to be bound by standards of good reasoning and be critically judged by such standards as to how well or poorly they are reasoning.

Example 2: Two countries in conflict over water resources

Imagine two relatively poor countries, we’ll call them R-land and C-land, in a hot dry part of the world. They share a large border that includes a 2500 square mile body of water that they have fought over for the last half century. Historically, maps have shown the lake as part of one C-land’s territory, but R-land captured it in a war a century ago. After decades of chronic low-grade warfare during which possession of the lake has gone back-and-forth, the two countries have reached an impasse; each claims the lake as its territory on its official maps. R-land has a river flowing within its territory that empties into this 2500 square mile lake. C-land has very little water resources, only a small river coming from the lake that tends to dry up during the very hot season. R-land is now finishing a dam across this river and plans to “move” the lake from its present border location into its own territory behind the dam, further drying up any runoff from the lake into C-land. C-land, meanwhile, has built up military forces significantly stronger than R-land’s military forces and plans to use its army to destroy the dam and to secure full possession of the lake. The situation has reached the explosion point; there is real danger of an all-out war between the countries, which would result in wide-spread death, displacement of civilians, refugees – a real humanitarian crisis. Of course, R-land believes it can limit damage to the dam in a war and eventually repel C-land’s aggression. And C-land believes it can destroy the dam and capture the lake in the event of a war. As a last effort to avoid war, the two countries have agreed to let a United Nations negotiation team try to settle the conflict. But each side is naturally skeptical that the other side will actually stick to any UN agreement; each believes that the “enemy” will exploit any cooperation in order to gain a stronger position to capture the lake for itself, and of course each is tempted to do just that given the value of water in their part of the world. What bargain should the UN negotiation team present as a peaceful solution?

The UN team’s first step is to gather enough information to discover what sort of conflict of interest, what game, these countries are locked in. We can imagine the team going back-and-forth between government leaders, or splitting up into two sub-groups and each taking one country and then combining their data, or perhaps getting both side to sit down together to jointly provide the information. The team finds that each country has the same goal in common: to possess the lake as a valuable water resource.

This goal must now be analyzed into objectives that can be used as criteria for evaluating outcomes. Let's say, given the details of the narrative above, that the UN team gets the two governments to agree on these three points. (1) Because each country is culturally quite traditional, both believe in the importance of historical tradition in justifying a claim to possessing the lake. In this regard, C-land is better off than R-land. (2) Neither side can ignore military strength as "backing" the ability to take possession of the lake. In this respect, again C-land is better than R-land. (3) Finally, let suppose that the UN team gets these two governments to agree on a humanitarian point: the importance of the need for water in that part of the world. In this regard, both sides agree that C-land is worse off than R-land (R-land has a valuable water source without the lake, the river, whereas C-land doesn't). The UN team, using an interval utility scale (-10…0…20), get both governments to agree on the following values:

	R-land	C-land
(1) Historical tradition:	-5	7
(2) Military strength:	0	5
(3) Water resources without the lake:	0	-15
Total:	-5	-3

(Note that as in the above example we are skipping important details concerning the steps in deriving criteria from the goal analysis, evaluation of outcomes by such criteria, etc. in order to focus on the main steps in bargaining and negotiations. Refer back to Chapters 2 and 3, if you are unsure about these steps.)

The UN team continues to fill in the information that will identify what kind of game these countries are in. Both sides agree that R-land has two options – R1: leave the lake as is, or R2: dam the river and move the lake from the border into its own territory. Both sides further agree that C-land has two options – C1: don't attack the dam and capture the lake, or C2: use military force to attack the dam and capture the lake. For both sides, the cooperative option (R1, C1) means neither gain nor loss of the goal, and so both agree to rank this 0 on the interval utility scale (-10…0…20). Both governments also agree that (R2, C1) would mean a big gain for R-land and a big loss for C-land. They rank these outcomes (20, -10). They likewise agree with the UN team that (R1, C2) would be a big loss for R-land and a pretty big gain for C-land. They agree to rank these outcomes (-10, 15). Finally, as expected, there isn't much agreement on the outcome for the options (R2, C2). Each side believes a lot of damage would be done but that each country could "hold its own" against the aggression of the other. Fortunately, the UN team already has the (-5, -3) outcome evaluation of failed agreement, what both sides agree are the relative strengths and weaknesses in goal achievement if each side tries to achieve the goal on its own.

Here then is the game matrix for the conflict between R-land and C-land:

		C-land: C1: (cooperate)	C-land: C2: (defect)
R-land	R1: (cooperate)	0, 0	-10, 15
R-land	R2: (defect)	20, -10	-5, -3

The UN team sees that this is a classic one-time prisoner's dilemma. If R-land and C-land are individual rational agents, by dominance, by Nash equilibrium, and by maximin reasoning they will separately choose (R2, C2) and the result will be war. Both will have goal loss: historical tradition, military strength, and need for water on the side of C-land serves to keep R-land from having goal achievement (to a -5 degree), and the ability to dam the river on R-land's part serves to keep C-land from having goal achievement (to a -3 degree). The UN team also sees that this game is asymmetrical, so whatever arbitration is proposed, use of the lake will not be equal between these two conflicting nations.

The UN team must now start the negotiation process of transforming this potentially cooperative game having a sub-optimal outcome problem into a bargaining problem. The individual security levels are already done: U(R-land) = -5   U(C-land) = -3.   If the agents agree to re-adjust these values at a later point in the process, then goal identification and analysis will have to be revised.

Both governments have already agreed that possessing the lake is the goal; it should, then, receive the maximum value on the scale. So, if one side backs down to allow the other side possession of the lake, the outcome for the possessor is 20. U(R-land and C-land) = 20.

In a prisoner's dilemma, the free-ride and the sucker's outcomes are the main negotiation hurdles. These must be de-valued and enforcement that the other side will not cheat must be assured to each agent. Fortunately, the UN has forces available to monitor the activities of each country with respect to their negotiation agreements, providing the countries allow UN inspection teams to do their work. Let's suppose that the following three items have been agreed to by both governments:

(1) UN monitors have free access to observe and report to both countries the activities of each with respect to the lake conflict.

(2) Both countries have agreed to allow UN forces to intervene to stop any violations of the negotiations agreements that have been confirmed.

(3) The UN will send aid and funds from the international community into both countries, but will cut off all such assistance to the side that attempts to cheat on the cooperation of the other.

(Obviously, this is very unrealistic; two nations in a real conflict like this would never agree so easily to such terms, things would be far more messy and difficult than described here. The point of the example, however, is to gain practice applying certain methods and principles of practical reasoning, it is not, on this elementary level, to approach anything close to realism.)

Assuming that these three de-valuing and enforcing measures do the trick, the NU team has transformed the original prisoner's dilemma in to a bargaining problem:

		C-land	
		C1: (cooperate)	C2: (defect)
R-land	R1: (cooperate)	20	~~-10, 15~~
	R2: (defect)	~~20, -10~~	-5, -3

having the characterstic form: U(R-land) = -5 U(C-land) = -3

U(R-land and C-land) = 20

It now remains for the UN negotiating team to propose a formula for dividing the joint goal achievement (the lake) according to a rational and fair arbitration. Both countries, if they are rational agents, should see the solution as rational and fair, and whatever concessions each side is required to make, each should be much better off with the bargaining solution than with their original prisoner's dilemma solution.

Here is the bargaining graph. As we did with example 1, the work is filled in with dotted lines.

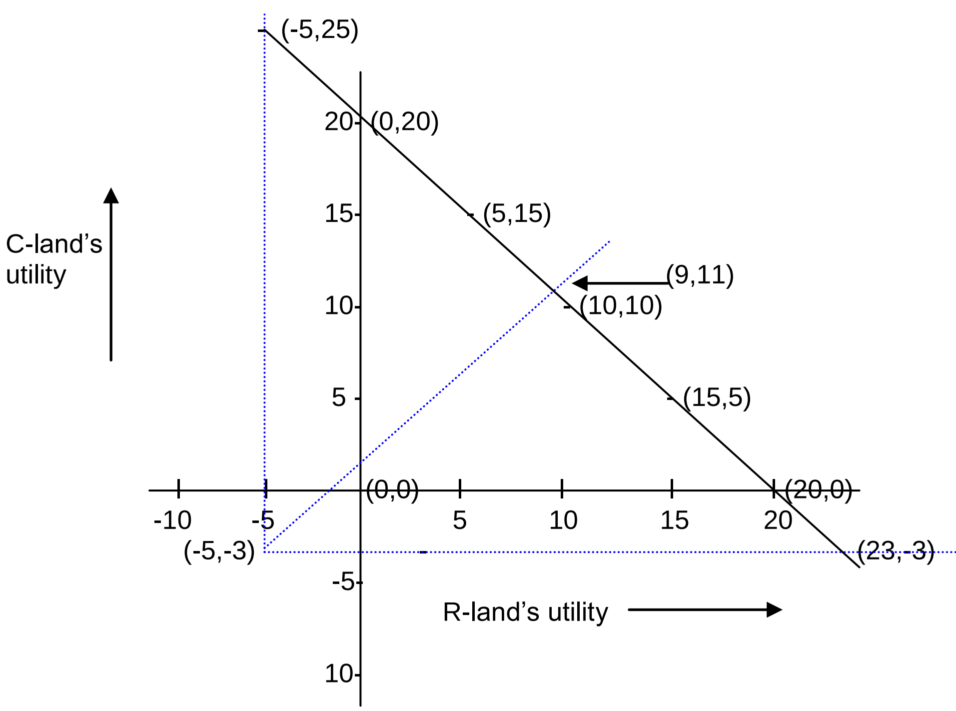

The basic bargaining region is pictured in the triangle within the points: (0,0), (0,20), (20,0).

Now the UN negotiation team must apply the 4 principles of arbitration.

1) By principle 1, R-land can't receive less than -5, and C-land can't receive less than -3. This means that the bargaining region must be enlarged, rather than decreased as in example 1. The new space of possible divisions is the dotted triangle within the points: (-5,-3), (-5,25), (23,-3).

2) By principle 2, the bargaining solution must be a point on the bargaining line running from the point (-5,25) to the point (23,-3), each point summing to the full value of the goal, and each point a Nash equilibrium. So far, each point is just as good a rational solution as any other. Principles 3 and 4 are needed to select a single special point from the bargaining line – one that represents a fair solution.

3) By principle 3, the point (-5,-3) becomes the new initial bargaining status quo point, representing the unequal bargaining power (in this case: bargaining weakness), the unequal worth, of R-land and C-land within this bargaining problem. Using (-5,-3) as the origin point for the solution line indicates that this is an asymmetrical bargaining problem, and depicts the exact proportion in which the agents should be treated unequally.

4) Once the relevant difference between R-land and C-land has been subtracted away, by principle 4 they must be counted as equals. What is the point (x,y) that, if you subtract -5 from x and -3 from y leaves the remainders equal? In other words, on the graph what point divides the bargaining line from (-5,25) to (23,-3) exactly in half? The UN negotiation team discovers that it's point (9,11). This is the point at which the agents concede equal lengths of the bargaining line to each other. 9 minus -5 = 14, and 11 minus -3 = 14. And 14 x 14 yields the largest product, in keeping with Nash's formula, than any other possible division of the goal.

So, the UN team proposes this split of the lake: out of every 20 units, 9 goes to R-land and 11 goes to C-land. These two countries, if they accept the solution, might position the official border so as to divide the 2500 square mile lake in a 9-to-11 proportion (1125 square miles will belong to R-land and 1375 square miles will belong to C-land). But this bargaining solution has another, perhaps a better, way of being implemented. Envision this: C-land allows R-land to build the dam and move the lake into its own territory. But for every 9 gallons of water held back by the dam, R-land must let 11 gallons flow into the old lake, which now completely becomes part of C-land's territory. They each end up with a smaller lake that are in 9 -11 proportion sizes, and C-land gains its historical territory back. (The potential problem here, however, is that R-land controls the flow of water, something that C-land ought to worry about, given this vital resource.)

Pause:

The last remark is an important general concern. You might remember it from the beginning of Chapter 2, but it's worth repeating. It is one thing to use practical reasoning to arrive at a rational choice, it is another thing for agents to act and actually do the rational thing, and not go back on agreements. Practical reasoning can (often!) bring us to the rational choice, but practical reasoning can't make an agent implement a decision, carry out an option, or stay the course if it's the right one. The latter issue belongs more to the field of psychology, perhaps, but it is not part of the study of practical reasoning and making rational choice.

Example 3: Two friends in conflict about seeing a movie

Rob and Cathy have been going together since high school. They are now students at the same university and are taking 2 of their classes together: a practical reasoning course and a film course. For the next assignment in the reasoning course, it is their turn to present and solve a bargaining problem to the class. For the film course, the next assignment is to see a movie and turn in a paper analyzing and evaluating it. Rob is doing well in the reasoning course, but he is having a hard time in the film course and needs a good grade on the paper assignment. Cathy is doing well in both courses. They naturally want to see a movie together for the film assignment, but are having a conflict about what movie to see, and they can see only one movie before the assignment is due. There are 4 possible movies available, but Rob believes that the one that Cathy would like to see wouldn't make for a very good paper. The one he believes would be best for his paper grade, Cathy wouldn't enjoy very much. As they discuss their movie conflict, they realize that they are in a bargaining game – one that would be perfect to present in their practical reasoning class. To, as it were, "kill two birds with one stone", they decide to formally work up their movie conflict as a bargaining problem both to present to the reasoning class and to help them decide what movie they should see together for the film class assignment.

The first step is to clarify the goal: see a movie. They analyze this goal into 3 objectives: (1) go out together, (2) enjoy a good movie, and (3) use the movie for a film class paper assignment. They each assign weights according to how each values these objectives, use these as criteria to evaluate the 4 movie options which they designate A,B,C,D (using an interval scale:1 - 10), and use the (rounded) outcome utilities to construct a game matrix.

Rob:	(1) .4		(2) .2		(3) .4	= 1.0	Final outcome utilities
Movie A:	(10) 4	+	(8) 1.6	+	(7) 2.8	=	8.4 (rounded to 8)
B:	(8) 3.2	+	(5) 1	+	(6) 1.8	=	6
C:	(7) 2.8	+	(3) .6	+	(10) 4	=	7.4 (rounded to 7)
D:	(2) .8	+	(0) 0	+	(5) 2	=	2.8 (rounded to 3)

It is clear that Rob believes movie C is the best one for his paper topic, but A turns out to be the best overall movie to see according to these 3 criteria. For Rob: (A p C p B p D), where "p" means "preferred to".

Cathy:	(1) .5		(2) .3		(3) .2	= 1.0	Final outcome utilities
Movie A:	(2) 1	+	(0) 0	+	(5) 1	=	2
B:	(6) 3	+	(4) 1.2	+	(10) 2	=	6.2 (rounded to 6)
C:	(8) 4	+	(6) 1.8	+	(2) .4	=	6.2 (rounded to 6)
D:	(10) 5	+	(8) 2.4	+	(3) .6	=	8

Cathy values being together more than Rob does, believes movie B would make a good topic for her film class paper, but discovers that D is the best movie for her to see by these 3 criteria. Her preference order is: (D p C i B p A), where "i" means "indifferent between".

The decision problem can't end here, however, with Rob seeing movie A (alone) and Cathy seeing movie D (alone). This would lose each a large chunk of goal achievement. They can both do better if they cooperate.

They now use these outcome utilities to construct a game matrix, subtracting the (rounded) values for objective (1) from the (rounded) final outcome utilities for any movie attended alone. Here, then, is their matrix, simplified with regard to decimals:

		Cathy A	Cathy B	Cathy C	Cathy D
Rob	A	8, 2	4, 3	4, 2	4, 3
Rob	B	3, 1	6, 6	3, 2	3, 3
Rob	C	4, 1	4, 3	7, 6	4, 3
Rob	D	2, 1	2, 3	2, 2	3, 8

Rob and Cathy now investigate the solution to this potentially cooperative game, if each were to make their separate individual rational choices. For Rob, option D is dominated by the other options, so both agents drop it as an option. The same for option A in regard to Cathy, it gets dropped. In the remaining 3 x 3 matrix, C dominates A for Rob, and so Rob's A drops out of the menu of options. For Cathy, B now dominates D, so they drop Cathy's D option. There remains a 2 x 2 game with B and C as Rob's options, and B and C as Cathy's options. They see that there is a maximin strategy (C,B) for a (4, 3) outcome, but this is unstable; it is not an equilibrium outcome. The game has two Nash equilibrium pairs: (B,B) for a (6,6) outcome, and (C,C) for a (7,6) outcome. They recognize (being good practical reasoning students!) that this game is an asymmetrical clash of wills, and that one equilibrium outcome is sub-optimal, even if only by a little. Cathy can do Rob a favor and give in to his preferred movie C (which would be nice on Cathy's part, but not necessarily rational or fair). Or Cathy can claim that the slight "edge" Rob seems to have is too small to matter, and they can try to "battle it out" between movie B or C to see if one can overpower the other to give in, risking deadlock and perhaps even severed relations. Instead, Rob and Cathy turn to bargaining and negotiation to arbitrate their conflict and justify their bargaining solution as both rational and fair.

The negotiation part is simple in the case of Rob and Cathy's bargaining problem. Each readily agrees not to defect; that is, see a movie without the other (their word and close relationship is sufficient, for the game isn't a prisoner's dilemma or a stag hunt having a strong free ride benefit or a big sucker's loss). The individual security limits are clear: U(Rob) = 4 U(Cathy) = 3. For the joint cooperative gain, they must use the largest sum of the 4 options:

U(Rob and Cathy) = 13.

Their bargaining problem, with all the non-cooperative outcomes cancelled, now looks like this:

Cathy

Rob		A	B	C	D
	A	8, 2			
	B		6, 6		
	C			7, 6	
	D				3, 8

or:

	A	B	C	D
Rob:	8	6	7	3
Cathy:	2	6	6	8

Rob and Cathy are ready to picture their bargaining problem as a graph. As above, the work is done in dotted lines.

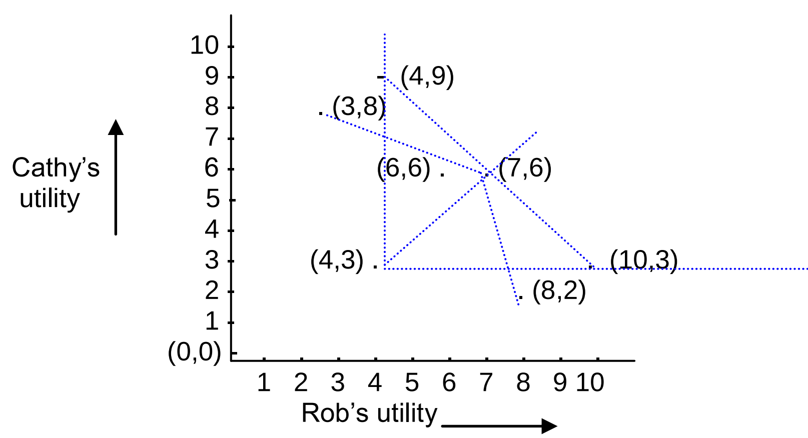

The bargaining space, before applying the arbitration principles, is the area bounded by solid lines: (0,0), (0,10), (10,0) (the line running from (0,10) to (10,0) has been omitted so as not to clutter the graph). Rob and Cathy now apply the principles.

(1) By principle 1, Rob's goal achievement can't be less than 4, and Cathy's can't be less than 3. So, they exclude movie A (for this gives Rob 3: point (3,8)), and they exclude movie D (for this gives Cathy 2: point (8.2)). The bargaining region becomes reduced to the new triangle area:

(4.3), (4,9), (10,3). Notice that the points for the outcome utilities for A and for D are outside this new bargaining space.

(2) By principle 2, a very important principle for the kind of bargaining problem Rob and Cathy have, the solution must be a point on the bargaining line; that is, it must be Pareto-optimal. There is only one possible point on this line: (7,6). The point (6,6) is in the interior of the triangle, and is sub-optimal. To choose movie B, then, would violate principle 2, one of the principles of rationality.

3) By principle 3, the initial bargaining inequality between Rob and Cathy must be proportionally respected in selecting a solution point. Rob specific bargaining power – his need to do well on the paper assignment for the film course – is represented in the (4,3) origin point of the smaller triangle. Cathy has nothing in her bargaining power to offset Rob's ability to require, as a matter of fairness, that she concedes more to him than he is required, as a matter of fairness, to concede to her.

4) By principle 4, Rob and Cathy must treat each other as equal, once they are operating within the bargaining space bounded by the triangle satisfying principles 1 - 3. The point (7,6) lies midway between the bargaining line (4,9), (10,3). Alternatively, subtracting the inequality between Rob and Cathy leaves them equal: 7 - 4 = 3, and 6 - 3 = 3.

To verify that this solution is correct, Rob and Cathy apply Nash's formula: (7 - 4) x (6 - 3) yields the largest product of any other point on the bargaining line, and it's the largest product if the formula is applied to the outcomes of the other options.

In this bargaining problem, the goal can't be literally divided. The rational choice is that they both see movie C for a (7,6) payoff, but this does not mean that Rob should see 7/13 of the movie and Cathy should see 6/13 of it (as if the movie were like the lake in example 2, or like income as in example 1). It means that seeing the whole movie C gives Rob more goal achievement than seeing it gives Cathy. This solution says that Rob "deserves" (in this case, justifiably "needs") to get slightly more out of the movie they decide to see than Cathy "deserves"; and seeing C will give Rob greater rewards than it provides Cathy, in the proportion 7-to-6.

Pause:

The (7,6) solution, movie C, may have seemed obvious to you all along, in which case you might be thinking: why all this trouble? Why all this processing? Why all this analysis? Isn't this over-analyzing the obvious to the extreme? This is a fair reaction, and provides the opportunity to recall what's going on in the big picture. The purpose of practical reasoning, as presented in this text, is not just to arrive at a decision. For this you don't need reasoning; you can simply go with what seems obvious, or just follow what others do in similar situations. What we are trying to do is learn how to use a system of principles and methods that will *justify* our decisions as good ones. If you think about your decision problem in a systematic way (put it into the right rational choice model), the wheels of analysis and evaluation (practical reasoning) will do their work and produce a choice that can be accounted for step by step as rational. What is philosophically important is the reasoning behind a decision, that is: the principles and methods that justify it as a good decision, not the fact that a decision is arrived at which all along looked obvious.

EXERCISE:

1) Analyze (frame) and arbitrate these bargaining problems. Feel free to add appropriate details to these stories that you would naturally expect to find in such decision situations.

a) 2 agents want to go out to a restaurant together, but their food likes and favorite restaurants are different. They have 5 options: A - F. If each goes alone, Row picks B just because it's near (U(3)), and Col picks E (U(3) because it has take out. Here are the outcome utilities (scale: 1 - 15) for dining together:

	A	B	C	D	E	F
Row:	1	4	7	10	12	14
Col:	15	12	9	6	4	2

Each is trying to get the other to go to the one's more favorite restaurant: Row favors they both eat at F, Col desires they both eat at A. What is the option closest to the "ideal" restaurant for these agents? Justify your arbitration by making the practical reasoning steps, and the principles you apply, clear.

b) 2 agents, 1 has a car with no gas (U(1)), the other has a lot gas but no car (U(4)). If they pool their resources, they have a working car (U(20)). The free ride (taking off with the car alone) = 20, and the sucker (being stuck with no car and no gas) = 0. How do they negotiate their potential for cooperation into a bargaining problem? What kinds of assurance against exploitation could they use? What is the rational choice for sharing the working car fairly?

c) Two police departments are experiencing an increase crime. Each separately gathers data on suspected criminals. On the basis of the data it gathers, one department is able to make on average 10 criminal arrests per year (U(1)). On the basis of the data gathered by the other department, it makes on average 25 criminal arrests per year (U(2.5). If these two departments cooperate by pooling their intelligence gathering technology, however, they can together arrest on average 40 criminals per year, a mutual cooperation gain of 5 additional criminal arrests, worth U(15) in public approval and new equipment. The free ride (one department grabs all the public credit and rewards for the increase in criminal arrests) = U(15), and the sucker's outcome (public loss of face for lack of improvement during increased crime activity) = U(-10). How would they negotiate guarantees against the sucker's outcome, and how do they fairly share the cooperative outcome?

d) It's the start of the semester and Nan needs several expensive nursing texts: the bookstore's price is $450. Meanwhile, Jim will be graduating and has these very texts. The bookstore will buy them for \$50. Jim posts a notice on the campus web-site listing the books and a request: make an offer. He doesn't mention a price. Here is the wrinkle: Nan needs these texts a lot more than Jim needs to sell them; this clearly gives Jim more bargaining power than Nan.

e) (True story!) Louis Tiffany (b. 1848, d. 1933), the famous stained glass artist, faced a difficult decision. He had established his own stained glassmaking company. At the time, glass blowers, colorers, and designers belonged to a powerful union and were all men. However, Tiffany had hired a few women who were outstanding glass-artists to design and make some of his colored glass. Tiffany and his company clearly benefited from the work of these women; their works were big sellers and were making Tiffany stained-glass works famous. But the women were a threat to the glass-worker's union (letting women work in the art-glass industry would mean fewer jobs for men). The union threatened to strike unless these women were fired, and a strike would severely hurt the company; but so would firing the women artists. Without filling in the complete story, let's suppose that Tiffany had 2 options: R1 = he can keep the women or R2 = fire them; and suppose that the union had 2 options: C1 = continue working or C2 = strike. If Tiffany keeps his women artists, the union will go on strike, and if the union goes on strike Tiffany will eventually be forced to fire them. The utility values, let's suppose, are these on an interval scale (-10…0…20): (R1 = 20, C1 = -5), (R1 = -6, C2 = -3), (R2 = -5, C1 = 8), (R2 = -3, C2 = 5). However, if this potentially cooperative game can be made into a bargaining problem, let's say that the joint gain for Tiffany and the union is U(20), but without cooperation the union strikes U(5) and Tiffany fires the women U(-3). How might negotiations transform this conflict into a bargaining problem, and how should these two agents share the mutual gain?

Sources and suggested readings:

This Chapter draws primarily from Binmore (2005) Chapter 2, Davis (1983) Chapter 5, Gauthier (1986) Chapter V, Mullen and Roth (2002) Chapter 9, Resnik (1987) Chapter 5-5, and Straffin (1993) Chapter 16. Bargaining and negotiation theory is a part of rational choice theory that tends to be technically demanding, but Fisher and Ury's (1991) bestseller is highly recommended for practical insights and guidelines on negotiations. Straffin is especially clear on the Nash arbitration scheme, and his Chapter 17 is recommended to anyone interested in how a complex labor-management dispute might be arbitrated using Nash's scheme. Kahneman and Tversky explore the connection between concession and loss aversion and present interesting empirical results in Tversky (2004) Chapter 29 "*Conflict Resolution*". For the importance of bargaining and negotiations in the social sciences generally see Elster (1089a) Ch. XIV, and in more depth (1989b) Chapter 2. Binmore, and Rapoport (1966) Chapter 8, offer good discussions of alternatives to the Nash bargaining scheme, the latter in more depth. Bargaining and justice (and more broadly: ethics) are deeply interrelated. For the rational choice (rationalist) approach to their connection, see Braithwaite's (1963) inaugural address (only 55 pages!) for a sustained analysis and solution of a single bargaining problem (involving musicians) from a moral philosophy perspective, and Gauthier's Chapter V discussion of the connection between bargaining and morality. For an evolutionary (naturalist) approach, Skyrms (1996) and (2004) are recommended. However, if you are able to read just one thing on bargaining and negotiation, don't miss Chapter 2 of Schelling (1980), it's a classic.