KELLOGG SCHOOL OF MANAGEMENT

Strategic Decision-Making

DECS-452 Professor Bob Weber
Week #4

1. In class today, you will be assigned a partner/opponent. You should arrange to meet with that other individual (in person, or by telephone) before our next class, and play through the "Repeated Prisoners' Dilemma" (attached). (Bring your result sheet with you.)

2. Read the "missionary" problem (problem (2) on the "Common Knowledge" handout), and be prepared to discuss your thoughts in our next meeting. (Believe it or not, there is a direct connection with the "Prisoners' Dilemma" experiment.)

3. Before the subsequent class (two classes from now), write up your strategy for problem (3). I will collect the strategies, shuffle them, and then hand them out again, so be sure that yours is written up clearly enough for a classmate to be able to carry it out.

4. Next week, we will carry out our first out-of-class negotiating exercise. If you will not be in class, please arrange for a classmate to get you paired (on Tuesday/Thursday) with someone.

5. Also next week, you'll receive the midterm project (which can be done in groups of 1-3). It will involve the analysis of a few small cases, will count for 20% of the course grade, and will be due a week later.

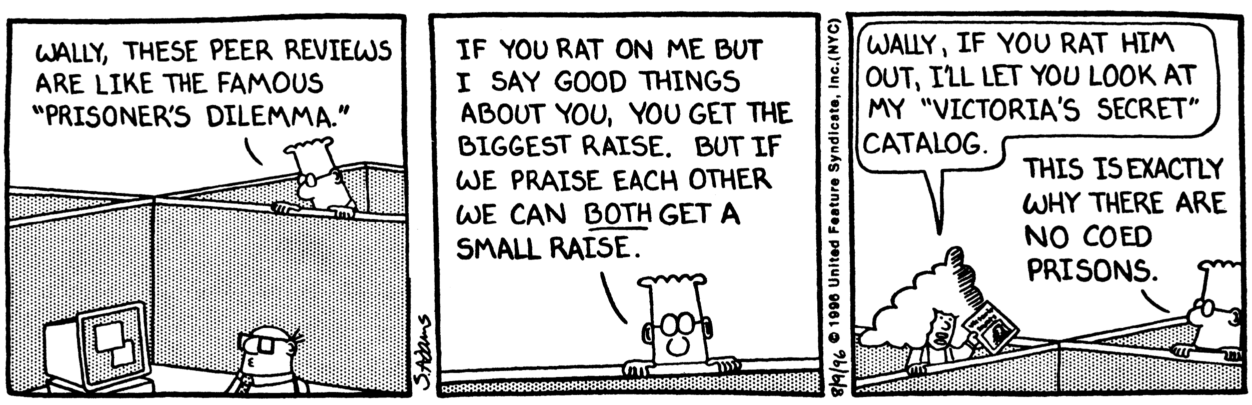

Payoffs in the Dragon/Quantum Case

Dragon	cut → cut disk → cut	cut → cut disk → new	cut → new disk → cut	cut → new disk → new
W → cut, S → cut	0.1300	0.1300	0.0050	0.0050
W → cut, S → disk	0.1235	0.1190	0.0285	0.0240

P11: +STRONG*(PD+PE)+WEAK*PE
P12: +STRONG*(PD+PE)+WEAK*PE
P13: +STRONG*PE
P14: +STRONG*PE

P21: +STRONG*DISK*(PB+PC+PD+PE)+WEAK*PE
+STRONG*DISK*(PC+PD+PE)+WEAK*PE
+STRONG*DISK*(PB+PC+PD+PE)
+STRONG*DISK*(PC+PD+PE)

Quantum	cut → cut disk → cut	cut → cut disk → new	cut → new disk → cut	cut → new disk → new
W → cut, S → cut	($10.80)	($10.80)	($12.30)	($12.30)
W → cut, S → disk	($10.35)	($10.35)	($13.20)	($13.20)

-CUT-P11*SUCCESS
-CUT-P12*SUCCESS
-NEW-P13*SUCCESS
-NEW-P14*SUCCESS

-CUT*(1-STRONG*(1-DISK))-P21*SUCCESS
-STRONG*DISK*(NEW+(PC+PD+PE)*SUCCESS)-WEAK*(CUT+PE*SUCCESS)
-STRONG*DISK*(CUT+(PB+PC+PD+PE)*SUCCESS)-WEAK*NEW
-STRONG*DISK*(NEW+(PC+PD+PE)*SUCCESS)-WEAK*NEW

0.05 = Pr(**strong**)
0.95 = Pr(**weak**)

0.60 = Pr(new **disk** technology feasible)

Quantum's Costs (net present value, $millions):

$60.000 = cost of Dragon's **success**
$3.000 = cost of price **cut**
$12.000 = cost of price cut + **new** product

	Change in CEI	Prob
PA	0%	0.05
PB	(0%,1%)	0.15
PC	(1%,2%)	0.10
PD	(2%,5%)	0.60
PE	5%	0.10

Dragon vs. Quantum

There are three key words to keep in mind when facing a strategic problem where a party has private information: ***Represent. Analyze. Implement.***

Represent: Dragon holds private information, about which Quantum has beliefs. We can represent this by imagining Nature dealing Dragon a single card from a 20-card deck, where 19 of the cards say "Weak," and one says "Strong." Dragon sees the card, which determines Dragon's financial strength. Quantum sees only the deck. Through this artifice, we establish a situation where Dragon might be in either the "weak" or the "strong" state.

A strategy for Dragon must indicate how Dragon would act in either state. – This, after all, is what Quantum is speculating upon. A strategy can be thought of as combining two half-strategies: One specifies how Dragon will act if weak, and the other specifies how they will act if strong.

Dragon's expected payoffs are computed across all possible states, using Quantum's beliefs.

Analyze: Carry out whatever analysis seems appropriate with respect to the strategic representation of the "game." Ultimately, a strategy for Dragon is selected.

Implement: Now – for the very first time – Dragon looks at its "card." This tells Dragon which part of the strategy – which "half"-strategy – to carry out.

Why does this work? If the selected strategy for Dragon is truly "best," it must consist of two half-strategies, each of which is best on its own. (If, for example, there were a better half-strategy when strong, that half-strategy, combined with the weak half-strategy from the originally-selected full strategy, would yield a better full strategy.)

From the perspective of an outside strategic consultant, the representation and analysis can be done without the consultant ever needing to determine whether Dragon actually *is* weak or strong. The consultant rides off into the sunset, and only then does Dragon look at its actual situation, and implement the appropriate half-strategy.

J.L. KELLOGG GRADUATE SCHOOL OF MANAGEMENT

Strategic Decision-Making

DECS-452 Professor Bob Weber

Common Knowledge

1. There are two individuals, A and B, facing one another in a room. Two consecutive positive integers are selected (by a referee; the choice probabilities are not relevant to this problem, except that *any* two such numbers might have been chosen). One of the numbers is written (in charcoal) on A's forehead, and the other on B's: Each sees the other's number, but not his own. Alternately, starting with A, each is asked the question, "Do you know what your number is?" If he answers, "No," the question is then asked of the other. (That is, A is asked first, then B, then A again, and so on.) For example, if A sees a 1, he answers, "Yes," when first asked (and the game ends). If B sees a 2, and hears A say, "No," then he can say, "Yes," when first asked. (He knows he has a 3.)

   (a) Assume the numbers are 3 (on A) and 4 (on B). What, if anything, is common knowledge before A answers the first question?

   (b) Both know (on the basis of what they see in the case mentioned in (a)) that the first two answers will be "No." Given this, what information is conveyed by A's first answer? By B's?

   (c) In what cases will the game eventually end in a "Yes"? You might first consider the case in (a).

2. Forty couples live on a remote island. In their society, three rules are rigidly adhered to:

   1. Whenever a woman is unfaithful to her husband, the other thirty-nine men meet and commonly share this information.

   2. No one ever tells a man that his wife has been unfaithful to him.

   3. If a man ever goes to bed with his wife in the evening, *knowing* that she has been unfaithful to hum at some time in the past, he kills her before morning (and his action becomes public knowledge in the morning).

   The actual state of "affairs" is that all forty wives have at some time in the past been unfaithful (although each has — to herself — foresworn further affairs). Each man therefore knows that there are at least thirty-nine unfaithful women on the island (the wives of the other men), but is unsure about his own wife's fidelity. For years, everyone has been living happily (if a bit uneasily) in this setting.

2. The missionary problem (continued)

One day, a missionary comes to the island for a visit. He stays for several weeks, during which he visits privately with each resident. On the day of his departure, as the eighty residents gather to watch him row away, he calls out, “I enjoyed my visit, but was disappointed to learn that there is at least one unfaithful wife on the island.” He then passes beyond earshot, before any questions can be asked of him. Of course, none of the husbands are surprised to hear his statement. And indeed, the next morning no wives are found to have been killed. Yet on the fortieth night after his departure, just when the wives are beginning to relax, all forty men kill their spouses.

(a) Explain what happened. (You might start by considering a simpler situation, in which only two or three couples reside on the island.)

(b) What did the missionary tell the men, that they didn’t know before?

Note: There’s no sexual bias intended in this problem, Indeed, last year (and next), the sexes are switched.

---

3. Assume that you are playing the game at right against an unspecified opponent. You are Player I, the “row” player; your payoffs are listed first.

	A	B
A	3,3	0,12
B	12,0	3,3

In each round, the two players move simultaneously, but both will learn the outcome of that round before playing the next one. Each player’s final payoff will be the sum of his or her payoffs from the twenty rounds. Write down a precise strategy (involving randomization, if you wish), to describe how you would play this game. Your strategy should be complete and clear enough that someone else in class could act as your agent and implement your strategy based on your written description.

3. **The Nash bargaining model**

The first formal game-theoretic analysis of bargaining was presented by John Nash in the early 1950's. He considered situations in which two individuals must choose how to coordinate their actions to mutual advantage, when each is fully aware of the set of potential agreements, and of the preferences of the other over those agreements.

Each party has available a list of actions; any chosen pair of actions yields an outcome. Furthermore, the preferences of each over the possible outcomes, as well as over probabilistic mixtures of outcomes, are commonly-known, and satisfy the standard von Neumann-Morgenstern axioms, i.e., both individuals are expected utility maximizers.

Nash began his analysis by assuming that the bargaining problem under investigation had some pre-specified "conflict" outcome, which would occur in the absence of agreement. For example, in the Battle of the Sexes the natural conflict outcome is for each to go to his or her more-favored destination.

Example 4. Alfred, who is near-broke, possesses a $100 bill. The serial number of the bill, mmddyyyy, happens to be the birthdate of Burton, a wealthy eccentric. Burton wishes to acquire the bill as a keepsake, and would be willing to pay as much as $500 for it. Alfred's utility for money is proportional to the square-root of the amount he holds (i.e., he is risk-averse); Burton's utility for money is linear (i.e., he is risk-neutral). How much should Burton pay Alfred for the bill?

In this example, again, the conflict outcome appears obvious: Alfred spends the $100 bill.

Nash next noted that many different agreements might be "stable," in the sense that, were an agreement reached and further discussion impossible, both parties would voluntarily carry out their roles in the agreement. In the Battle of the Sexes, if the agreement is to use a specific weighted coin to select randomly a joint destination, then, even *after* the coin flip, neither party can gain by unilaterally deviating from the agreement and going elsewhere. In the Alfred-Burton example, an agreement on any price between $100 and $500 is an agreement from which neither gains by walking away.

3.1 **Mediation**

In other cases, there are mutually beneficial, stable agreements which require help from the outside.

Example 5. Consider the following two-person game, in which each party must choose between two actions; the utility payoffs from the various pairs of selections are indicated in the diagram. (A's payoff is the first in each pair.)

		B: left	B: right
A	top	6,6	2,7
	bottom	7,2	0,0

There are three equilibrium points in this game: (top, right), (bottom, left), and a mixed-strategy equilibrium in which each independently and randomly chooses an action, with his first action being twice as likely to be chosen as his second. The corresponding expected payoffs are (2,7), (7,2), and (4 2/3, 4 2/3). Through the use of a joint randomizing device which, with some triple of agreed-upon probabilities, chooses one of the three equilibria, they can achieve any expected-payoff pair in the small central triangular region in Figure 1 as the outcome of a stable agreement. (That is, once the device is agreed upon, neither gains by unilaterally failing to carry out his role in the equilibrium point chosen by the device.)

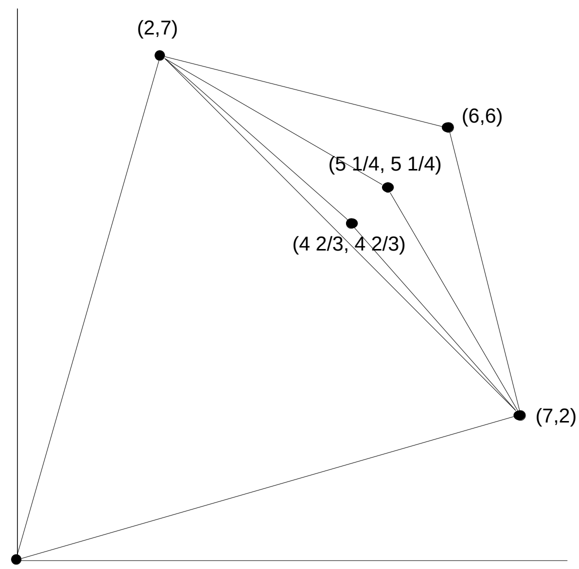

Example 5

A mediator could help them create other, mutually-preferred stable agreements. For example, they could agree that the mediator would leave the room and flip a fair coin. If it comes up “heads,” he will return and *privately* whisper “top” to A, and “left” to B. If it comes up tails, he will re-flip it: On “heads,” he will return and whisper “top” to A and “right” to B ; on “tails,” he will whisper “bottom” to A , “left” to B. If each party knows nothing about the mediator’s out-of-the-room actions except for his own whispered message, and if he expects the other to obey the mediator’s message, he can do no better than to obey his own. This mediated procedure yields them expected payoffs of (5 1/4, 5 1/4). Other stable agreements obtainable through mediation yield expected payoffs in the middle triangular region of the figure.

Generally, our view is that a **mediator** is an intervenor who can (publicly *or* privately) receive information from the parties, and transmit information back to them, according to rules upon which the parties themselves have agreed. Under this definition, a mediator can also be delegated the responsibility of carrying out public randomizations, as in the case of a randomized selection of a joint destination in the Battle of the Sexes.

3.2 Regulation

In Example 5, even better payoffs (such as (6,6)) are available to the parties, but no stable agreement allows them to obtain these payoffs. Similarly, in the Prisoners' Dilemma, both would prefer the ("don't confess", "don't confess") outcome to the equilibrium outcome, but neither can expect the other to adhere to an unenforced verbal agreement.

However, if an outside agency can be brought into the situation, and empowered to exact sufficiently high penalties from any deviator, either of these agreements *becomes* stable. Penalties of 1 or more are sufficient to "stabilize" the agreement (top, left) in Example 5; "Honor among thieves", when backed up by physical retribution, makes a D.A.'s task much more difficult than the Prisoners' Dilemma would suggest. Formally, we view a **regulator** as an intervenor to whom the parties may voluntarily grant the ability to force certain actions upon them (through the setting of appropriately-large penalties for noncompliance). In many instances, the civil courts play a regulatory role in enforcing contractual provisions.

The figure for Example 5 illustrates the variety of stable agreements which can be maintained at different levels of communication or intervention. If the parties can only communicate by telephone, they can achieve only the three outcomes corresponding to equilibrium payoffs. If they can meet (to observe the result of joint randomization), they can achieve an expected outcome anywhere within the triangle determined by these three. Through mediation, they can add the lower four-sided region to the range of stable agreements. And through regulation, the upper four-sided region can also be added.

3.3 Arbitration

Assume that the parties have agreed to use a regulator, in order to expand the set of possible (i.e., stable) agreements to its fullest. There still remains the problem of choosing from among these possible agreements. The parties can do this through open debate, always facing the possibility that they will fail to reach a settlement. Alternatively, they can invite yet another intervenor to enter the picture, and ask him to suggest a particular agreement, on the grounds, for instance, of his perception of "equity." Indeed, if they simultaneously agree to empower a regulator to enforce that suggestion, a settlement is guaranteed. (The regulator-arbitrator combination is what is frequently described as "binding arbitration.")

In order to distinguish between the roles of third-party intervenors in bargaining, we choose to define an **arbitrator** as an intervenor who is invited to suggest an agreement. Much of the rest of this chapter will focus on the procedures by which an arbitrator might choose his suggestion.

3.4 The Nash solution

Given the multitude of potential agreements, Nash suggested a set of rules (formally, "axioms") which determine a unique suggested agreement for every problem. The rules are stated as conditions an arbitration *procedure* should satisfy, where by "procedure" we mean a consistent philosophy to be applied across all bargaining problems. First, these rules require that the suggested agreement be feasible, Pareto-efficient (i.e., no alternative feasible agreement should be better for both parties), and individually rational (i.e., the suggested agreement should offer to each party at least as much as he would obtain at the conflict outcome). Second, the suggested settlement should depend on the parties' underlying preferences, and not on the utility functions chosen to represent those preferences. Third, in symmetric situations (that is, situations where the range of feasible agreements is symmetric, and the parties receive equal utility payoffs at the conflict outcome), the suggested agreement should offer equal payoffs to the two parties. Finally (and most controversially), if after an agreement is suggested, it is found that some alternative, unsuggested agreement was in fact not feasible, the original suggestion should still stand (i.e., the procedure should be "independent of irrelevant alternatives").

Nash showed that there is only one agreement-selection procedure which has all of these properties; thus, an arbitrator who accepts these rules as compelling *must* follow this unique procedure. The procedure selects, in every problem, the agreement which maximizes the product of the parties' respective utility gains from agreement, measured relative to their conflict payoffs.

In the Battle of the Sexes, the selected outcome under the Nash procedure is for the parties to jointly randomize their choice of destination, assigning probability $1/2 + 1/2\ (t_B/b_B - t_A/d_A)$ (if this is greater than 1, go to the mountains for certain; if it is less than 0, go to the beach). Notice that the mountains (A's more-favored destination) are selected more frequently when it is B who favors togetherness over destination relatively more than A.

In the Alfred-Burton example, the selected outcome is for Burton to pay Alfred the amount \$x which maximizes $(\sqrt{x} - 10)(500 - x)$, and hence Burton should pay \$277.78. This is less than the split-the-difference payment of \$300; Alfred's aversion to risk works against him in the arbitrated solution.

3.5 Optimal threats

Having dealt with the question of how to select a final agreement, Nash turned back to the question of how the conflict outcome (which forms, in a sense, the starting point for the arbitrator's considerations) should be identified in situations where the result of disagreement is not obvious (for example, when the parties have available a variety of retributive strategies). He proposed that the parties, *knowing how their dispute will be arbitrated once the conflict point is determined*, simultaneously write down the actions they will take if agreement is not reached, and empower a regulator to force them to carry out these actions in the absence of agreement. Nash then showed that in every case both parties will have *optimal* threatened actions, i.e., threats which leave them optimally positioned for the arbitration stage.

3.6 Summary

In light of the above discussion, we can interpret Nash's approach as separating negotiations into two stages: a threat-making stage, which is strictly competitive (in the sense that each party is attempting to stake as strong a claim as possible prior to the second stage) and determines the conflict outcome, followed by an arbitration stage, in which the gains from agreement are allocated between the parties.

Of course, if the parties both accept the principles presented above, they can determine for themselves the agreement which an arbitrator would suggest, and thus avoid formal arbitration. However, the final agreement might still require a regulator, at least in the form of a judicial system, in order to guarantee that both parties carry through with their responsibilities under the agreement.

It is important to note that the threats made by both parties in the first stage need never be carried out: The procedure always leads to agreement. This will not necessarily be the case when, in the next chapters, we consider bargaining problems in which the parties are not perfectly informed about the situation they face.

[Schelling and others have noted a tactic available in bargaining, even under conditions of complete information, which is not considered in Nash's analysis. One of the parties can attempt to make a preemptive precommitment which changes the set of feasible agreements. For example, in the Battle of the Sexes, one of the parties could make a nonrefundable prepayment on a weekend for two at his or her more-favored destination, and present this action to the other as a *fait accompli*. If both should do so, an inferior outcome must result.]

4. Bargaining under uncertainty

The difficulties which can arise when parties hold private information are dramatically illustrated in the following well-known example.

Example 6 (the Akerlof "lemon" problem). An owner of a used car is negotiating with a prospective buyer. The quality of the car is known only to the seller; expressed in terms of the car's value to the seller, the buyer believes it equally likely to be worth any amount between $0 and $500. The buyer, who would utilize the car to a greater extent, would derive 50% more value from its ownership. At what price might a sale take place?

Only if the car is worth less than $x to the seller would he agree to a sale at $x. But then, from the buyer's perspective, given that the seller agrees to accept a price of $x, the expected value of the car to the seller is no more than $x/2, and therefore, its expected value to the buyer is at most $3x/4. Hence, the buyer should refuse to buy the car at *any* price the seller is willing to accept! (Classroom experiments consistently bear out the empirical validity of this analysis — Subjects argue interminably, but trade never occurs.) Even though both parties know that a mutually advantageous trade exists, trade cannot take place unless someone acts irrationally.

4.1 Auditing

How might the seller and buyer work around this impasse? One possibility is to have a mechanic inspect the car, and provide an appraisal to them. This would convert the problem to one of complete information, amenable to the type of analysis described in Chapter 3.

Another possibility is to write a warranty into the sales contract, providing for payment adjustments after the buyer learns, through use, the quality of the car. Such a contract is actually a spectrum of contingent contracts, each written under the assumption of complete information, one for every possible quality level of the car. (Clearly, a regulator is required to implement a warranteed sale.)

In the first case, the mechanic acts as an auditor; in the second, post-sale observation plays an auditing role. We generally view an **auditor** as an individual (or procedure) through which information held by one party can be made public. Unless specific mention to the contrary is made, we shall assume throughout the remainder of this paper that auditing is not available, and will consider instead how, through their actions, parties provide information to one another, or to intervenors.

4.2 Games with incomplete information

Beginning in 1965 with research sponsored by the U.S. Arms Control and Disarmament Agency, game theorists and economists have focused substantial effort on attempts to understand bargaining under uncertainty. Most of this research falls into one of two categories: studies of what can conceivably be accomplished by the appropriate choice of a format for negotiations, and studies of what can be expected to occur in the context of some specific format.

Consider a general view of two-party bargaining. The parties make statements, true or false; they bluff, threaten, bluster, and otherwise interact in attempts to convince one another of their respective preferences and constraints. Finally, *something* happens — either an agreement is reached, or conflict ensues.

Basically, each party, knowing his own preferences, adopts a "private strategy," which specifies how he will act (or respond) at any stage of the negotiations, given what has transpired prior to that stage. (One can view this private strategy as a complete, explicit set of instructions given to an agent who will represent the party in the negotiations.)

An important (and frequently overlooked) consideration in choosing our own private strategy is that the opposing party does *not* know our own preferences and constraints (i.e., he does not know our "type"), and therefore continually updates his perception of us on the basis of our observed behavior. He does this by assessing the likelihood that we would act the way we have, for each of the possible types of opponent we might be. Therefore, he bases his responses (i.e., portions of his own private strategy) on his assumptions of how each of our possible types would act (i.e, on the private strategies he assumes our various potential types would adopt). It follows that, in order to decide upon our own appropriate actions, we must anticipate the conclusions he will draw: We must ask ourselves how we *would* have acted, had we been any type *other* than the type we actually are. (The Scottish poet Robert Burns anticipated our need to take this view when he wrote, in his *To a Louse*, "O wad some Pow'r the giftie gie us, to see oursels as others see us!")

Example 7 (the Walkenhorst Chemical case). Jack Walkenhorst, a young inventor, is preparing for a court hearing. Lakeland Chemical, a large conglomerate, has filed a patent application on a production process similar to one he has previously patented. If the court validates Lakeland's application and Lakeland begins to compete with Jack, he will suffer substantial short-term losses. However, as a result of his recent research he has an important piece of private information: Another process, significantly different from and much cheaper than either of the two contested processes, is commercially feasible. If Lakeland wins the suit, and engages Jack in competition, they will ultimately lose money, and Jack will eventually recoup his losses.

If Lakeland knew the true situation, they would freely choose to withdraw their patent application. But Jack cannot reveal any details of the new process without jeopardizing the new patent, for which he will not be prepared to apply for another six months. At this point, the outcome of the court case appears to be a toss-up. What can Jack do to improve his situation?

In this example, Jack would like to say to Lakeland, "Believe me — If you pursue the suit, win, and engage me in competition, you will eventually regret it." However, Lakeland cannot know whether Jack truly has something up his sleeve, or is merely bluffing in order to protect his position should he lose the case; that is, they don't know Jack's "type." If the making of this statement would convince them to stay out of competition, then his nonexistent, but potential, "bluffing" type would certainly make the statement. Therefore, Lakeland's perception of the situation will not be changed by Jack's statement: *Either* of Jack's types (his true type, or his bluffing type) would make it. Consequently, if Lakeland originally considers it unlikely that Jack has the ability to hurt them, his statement will not deter their entry.

The moral of this story is that, when preparing for negotiations, we must not merely focus on the private strategy our actual type will follow: We must also consider which private strategies we would follow, were we any type other than our actual one. One can view the preparation for negotiations as a roundtable discussion among a party and his various *alter egos*, in which the participants must decide upon the coordinated face they will present to the outside world. Some types might wish to "bluff," i.e., to mimic the private strategy of some other type in hope of persuading the outside world that they *are* that type. Other types might wish, in turn, to "signal," i.e., to take actions which clearly reveal their actual situation. (Jack Walkenhorst might choose to drop his current suit against Lakeland as a token of faith. If this would convince Lakeland to delay competition, his actual type would gain; if the delay would be of less value to his "bluffing" type than the current 50% chance of winning the suit, then that type would not make the same offer — Dropping the suit is a signal of his true type which Lakeland can believe. Indeed, if Jack is not clever enough to think of this signal, Lakeland (or an intervenor) can suggest it to him. A formal agreement, in which Jack drops the suit in exchange for a six-month delay in Lakeland's entry, works to the advantage of both parties and should be acceptable to both.)

As we have seen, a party's types might find themselves with conflicting desires; some resolution of this internal conflict must be reached before the one "true" type can decide upon his private strategy. An important note, to which we will return shortly, is that the final resolution of the inter-type conflict cannot involve binding agreements across types. Only one type actually exists; the others can't penalize him for breaking any agreement.

In view of the previous considerations, game theorists have chosen to define a **strategy** for a party in a bargaining environment as a joint specification of private strategies, one for each of his possible types. The private strategy of the true type is implemented; an opponent updates his beliefs about the party on the basis of observed actions, together with that opponent's guess as to the full strategy which was selected. (The standard rule of probability theory used for this updating is known as "Bayes' Rule.")

In a rational world, in which each party considers his opponent's strategic problems as well as his own, it is reasonable to expect that each party will believe his opponent's strategy to be optimal for each of the opponent's types, given the opponent's belief about the party's own choice of strategy. (This is because no type can be compelled by the other types to adopt a non-optimal strategy.) A pairing of such strategies, in which each believes correctly, is formally known as a **(Bayesian) equilibrium point** of the bargaining "game." The analysis of any specific dispute begins (for a game theorist) with a game-model of the communication and commitment abilities of the parties, and proceeds with a study of the Bayesian equilibria of the game.

DECS-452 Course Outline

F. The Dragon/Quantum case

1. Strategic form, with incomplete information
   a) Strategies, as functions of private information: Specification of what you will do, given your actual type, as well as what you *would* do, were your type different.
   b) Incorporating uncertainty concerning outcome of disk research

2. Dominant strategy analysis

3. Interpretation of equilibrium: Quantum will respond differently to price cut, or to introduction of disk player

G. Walkenhorst Chemical

1. **Signalling**: Taking an action which your actual type can afford, but which would be unthinkable for the type your competitor believes you may be.

2. The importance (and sufficiency) of at least one party being "clever"

The Repeated Prisoners' Dilemma

You have already seen the Prisoners' Dilemma as a one-shot game (both in class, and in the first section of the "Negotiation and Arbitration" notes). The description is repeated below:

The Prisoners' Dilemma. Two men have been arrested for a minor offense. However, the district attorney is certain (although he has no hard evidence) that they are also responsible for a much more serious crime. He separates the criminals, and offers each the same deal: If neither confesses to the more serious crime (such a confession would implicate both), he will ask for two-year sentences on the lesser offense. If both confess, he will request five-year sentences. But, if only one confesses, that one will go free, and the maximum penalty of the law (an eight-year sentence) will be requested for the other.

In this example, each has a dominant strategy: to confess. Indeed, the strategy pair ("confess", "confess") is the unique equilibrium point of the game. At this equilibrium point, both are worse off than at the outcome of the strategy pair ("don't confess", "don't confess"). Yet, in the absence of any external, enforceable agreement, it must be expected that each will confess.

In order to be consistent with tradition, let me relabel the strategies: "don't confess" = "cooperate" (with your accomplice), "confess" = "defect". Also let me shift the payoff scale somewhat: Please note that D is still a dominant strategy for each of the two parties.

	cooperate (C)	defect (D)
cooperate (C)	5,5	-5,10
defect (D)	10,-5	-2,-2

The analysis given above deals with the one-shot game. For this experiment, I'd like you to play a twenty-times repeated version of the Prisoners' Dilemma.

At each stage of the game, each player should write either "C" or "D" on his/her copy of the attached sheet. When you have both made your decisions, *simultaneously* show each other your sheets. **You should not communicate in any way, other than to show your choices at the end of each stage.** (If you must do this by telephone, I suggest that one player read his/her choice first for the odd-numbered stages, and the other read first for the even-numbered stages.)

Repeated Prisoners' Dilemma: Decision Sheet

Remember: Your goal is to do as well as you can **for yourself**.

Stage Number	Your Choice (C or D)	Counterpart's Choice	Your Payoff
1			
2			
3			
4			
5			
6			
7			
8			
9			
10			
11			
12			
13			
14			
15			
16			
17			
18			
19			
20			
Your name:		Your Total:	