Economics 703
Advanced Microeconomics **Prof. Peter Cramton**

Lecture Note 6: Auctions, Reputations, and Bargaining

Outline

1. Auctions and Competitive Bidding
   a. McAfee and McMillan (*Journal of Economic Literature*, 1987)
   b. Milgrom and Weber (*Econometrica*, 1982)
2. Reputations and Cooperation
   a. Entry Deterrence
   b. Evolution of Cooperation, Axelrod and Hamilton (*Science*, 1981)
   c. Rational Cooperation in the Prisoner's Dilemma
      Kreps, Milgrom, Roberts, and Wilson (*Journal of Economic Theory*, 1982)
3. Reputations and Bargaining
   a. Game Theoretic Models of Trading Processes
      Wilson (*Advances in Economic Theory*, 1987)
   b. Bargaining, Dynamic Monopoly, and the Coase Conjecture
      Gul, Sonnenschein, and Wilson (*Journal of Economic Theory*, 1986)
   c. Reputations in Bargaining
      Ausubel and Deneckere (*Econometrica*, 1989)

1. Auctions and Competitive Bidding

a. McAfee and McMillan (Journal of Economic Literature, 1987)

Why study auctions? Three reasons come to mind. First, auctions are an important means of exchange. It has been estimated that roughly 450% of the world GNP is traded each year by auction (this surprisingly large number is a result of all the financial securities that are traded by auction each year). Second, understanding auctions should help us understand the formation of markets by modeling the competition on one side of the market. Third, auctions represent an excellent application of game theory, since in an auction the rules of the game are made explicit.

<u>Simple Auctions</u>

Auctions typically take one of four simple forms:

<u>Oral</u>		<u>Sealed Bid</u>
English (↑ price)		2nd Price
Dutch (↓ price)	≡	1st Price

English: price increases until only one bidder is left; the remaining bidder gets the good and pays the highest bid.

Dutch: prices decreases until a bidder accepts the price; this bidder gets the good and pays the price at acceptance.

Second Price: each bidder submits a bid in a sealed envelope; the highest bidder gets the good and pays the second highest bid.

First Price: each bidder submits a bid in a sealed envelope; the highest bidder gets the good and pays the amount of his bid.

The Dutch auction is strategically equivalent to the first price auction. [Can you see why?]

Much of the theory on auctions is interested in determining equilibrium behavior in these simple auctions and comparing how much revenue the seller generates under each of the simple auction forms. Common variations to the simple auctions are the use of either a reserve price or an entry fee by the seller.

A central feature of the auction environment is the private information the bidders have about the value of the item being sold. This private information typically comes in one of three forms:

<u>Models of Private Information</u>

(1) Independent Private Value: $v_i \sim F_i$ independently of v_j

(2) Common Value: $e_i = v + \varepsilon_i$, $\varepsilon_i \sim F_i$ w/ mean 0.

(3) Affiliated Value: $v_i(x,s)$, my value depends on private information $x = (x_1,...,x_n)$ and state of world s.

In the independent private value model, each bidder has a privately known value v_i for the item, which is drawn from a distribution F_i independently of v_j for $j \neq i$. The independent private value model makes sense if differences in value arise from heterogeneous preferences over the attributes of the item. At the other extreme is the common value model, where it is assumed that each bidder has a private (unbiased) estimate of the value of the good, but the true value is the same for all bidders. This model makes sense if the bidders have homogeneous preferences, so they value the item the same ex post, but have different estimates of this true value. The common value model applies for goods bought for resale, such as oil drilling rights. The affiliated value model is the general model, allowing both private value uncertainty and common value uncertainty. In this model, each bidder has private information that is positively correlated with the bidder's value of the good.

<u>Benchmark Model: Independent Private Values, Symmetric, Risk Neutral Bidders</u>

The easiest case to analyze and hence a natural benchmark is the independent private value auction with risk neutral, symmetric bidders, no bidding costs, and where the payment depends only on the bids.

buyer values $v_1,...,v_n \sim F$ on $[0,\infty)$

seller values v_0 (common knowledge)

order statistics $v_{(1)} \geq v_{(2)} \geq ... \geq v_{(n)}$.

Both the English and second price auction have a unique equilibrium in dominant strategies.

English: bid up to your value or until others stop.

2nd Price: bid your value.

The outcome is the same in either case: the bidder with the highest value wins and pays the second highest value.

The winner gets $v_{(1)} - v_{(2)}$ ex post and expects in the interim state to get:

$$E_{v_{(1)}}[v_{(1)} - v_{(2)}] = E\left(\frac{1 - F(v_{(1)})}{f(v_{(1)})}\right).$$

This represents the information rent received by the winner of the auction. The seller's expected revenue is the expectation of the highest value less the informational rent paid to the winner:

$$E[J(v_{(1)})] = E\left(v_{(1)} - \frac{1 - F(v_{(1)})}{f(v_{(1)})}\right).$$

It is common to assume that J is increasing, which implies that bidders with higher values pay more for the good.

1st Price ($\equiv$ Dutch)

The symmetric equilibrium bidding strategy in the first price or Dutch auction can be derived as follows. A bidder's expected profit as a function of his value v and his bid b(v) is given by

$\pi(v,b(v)) = (v - b(v))\Pr(Win|b(v))$.

By the envelope theorem, the (total) derivative of π with respect to v is equal to the partial derivative of π with respect to v, ($\partial\pi/\partial b = 0$ since b is chosen to maximize π):

$$\frac{d\pi}{dv} = \frac{\partial \pi}{\partial b}\frac{\partial b}{\partial v} + \frac{\partial \pi}{\partial v} = \frac{\partial \pi}{\partial v}$$

But then $d\pi/dv = \Pr(Win|b(v)) = \Pr(\text{highest bid})$

$= \Pr(\text{highest value}) = F(v)^{n-1}$ a.e.

So by the Fundamental Theorem of Calculus,

$$\pi(v) = \pi(0) + \int_0^v F(u)^{n-1}du = \int_0^v F(u)^{n-1}du,$$

where the second equality follows since $\pi(0) = 0$ by individual rationality. But then substituting into $\pi(v,b(v)) = (v - b(v))\Pr(Win|b(v))$ yields

$$b(v) = v - \frac{p(v)}{\Pr(\text{Win})} = v - F(v)^{-(n-1)} \int_0^v F(u)^{n-1} du.$$

As an example if v is uniformly distributed on [0,1], then F(v) = v, so b(v) = v - v/n = v(n-1)/n. Notice that the optimal bid converges to the value as $n \to \infty$, so in the limit the seller is able to extract the full surplus. In equilibrium, the bidder bids the expected value of the second highest value given that the bidder has the highest value.

Revenue Equivalence

For our benchmark auction, the seller's revenue is the same for each of the four simple auctions: English = 2nd = 1st = Dutch. This follows since all four have the same probability of winning (the highest value bidder wins the auction), $d\pi/dv$ = Pr(Win), and $\pi(0) = 0$, so the bidder gets the same profit in each, and hence so must the seller. [This follows from the analysis of the direct revelation game.]

Optimal Auction

In the optimal auction, the seller sets a reserve price r s.t. J(r) = v_0 where J(v) = v - [1-F(v)]/f(v), so r = $J^{-1}(v_0)$. The purpose of the reserve is to increase competition for the highest value bidder. To understand how this works, consider the case of a single bidder. By setting a reserve r, the seller gets $\pi(r)$ = r[1 - F(r)] + v_0F(r). The first-order-condition for maximization is

1 - F(r) - rf(r) + v_0f(r) = 0, so v_0 = r - [1-F(r)]/f(r). It is interesting that this optimal reserve does not change with the number of bidders. [But the importance of the reserve decreases as the number of bidders increases.]

Asymmetric Bidders

If the bidders are asymmetric (i.e., there values are drawn from different distributions F_i), then Myerson (1982) showed that the seller should employ an asymmetric auction by awarding the good to the bidder i with the highest value of

$$J_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}.$$

For the special case where the distributions are the same except the means differ, then the optimal auction favors bidders with the lower expected values ex ante. That is, the seller should discriminate in favor of the disadvantaged buyer, so as to increase competition among bidders.

Royalty

Should the seller use a royalty? Suppose $\tilde{v}$ is an ex post observation about the true value. Then under a linear royalty scheme the winner's payment is a function of the bid b, the royalty rate r, and the ex post value $\tilde{v}$:

$$p = b + r\tilde{v}.$$

It makes sense to use a royalty to the extent that the ex post value v is not subject to moral hazard. If $\tilde{v}$ is exogenous (no moral hazard), then r should be 100%. But if $\tilde{v}$ is influenced by the costly effort of the winner, then a lower royalty rate must be used to preserve the correct incentives for the winning bidder to develop the good (this is the case in the extraction of oil). With moral hazard, the optimal contract is linear in $\tilde{v}$. If the bidders are risk averse, a royalty can also serve to shift risk from the bidders to the seller.

Risk Aversion

If the bidders are risk averse, then the revenue equivalence theorem no longer applies. In particular, the seller's revenue from the 1st Price > English. In the English auction, it still is a dominant strategy to bid your value (why?), so the outcome is the same as in the risk neutral case. But in the first price auction, a bidder has an incentive to increase his bid from the risk neutral bid, since by increasing the bid, his risk is reduced: he gets a higher probability of winning a smaller prize. This lottery with reduced risk is preferable for the risk averse bidder. Hence, competition is greater with the first price auction, as bidders attempt to limit risk by bidding higher.

The optimal auction with risk averse bidders is derived from similar intuition. The seller increases bidder competition by subsidizing high bidders that lose and penalizing low bidders.

With risk averse bidders should the seller try to keep the number of bidders a secret? The answer is yes. Keeping the number of bidders secret increases competition ex ante by reducing the bid dispersion and the informational rents.

Collusion

What if the seller is faced with a group of colluding bidders. This is called a bidding ring: n bidders colluding as 1 (and essentially submitting a single bid). In this case, the seller should set a higher reserve r using the distribution $F(v)^n$. The seller's profit as a function of the reserve is

$\pi(r) = r[1 - F(r)^n] + v_0F(r)^n$,

so the first-order-condition for maximization is satisfied if r is such that

$$v_0 = r - \frac{1 - F(r)^n}{nF(r)^{n-1} f(r)}.$$

Many Items

What if the seller is not selling a single item, but many items? Consider the extreme case where the seller can produce an unlimited amount of the good at a marginal cost of c. What is the optimal selling mechanism? The seller should not use an auction at all, but instead should post a fixed price r such that $c = r - [1 - F(r)]/f(r)$.

Costs of Entering Bidding

In many real auctions it is costly to bid. How can this affect the bidding behavior?. Fishman (1988) considers a model of preemptive bidding in takeovers, where it is costly for bidder i to acquire his private value information v_i. If the bidders enter sequentially, then the initial bidder may make a preemptive bid if his value is high enough, to signal that he has a high value, and thereby, discourage other bidders from entering.

Correlated Values

The independent private value model is a handy benchmark, but in many environments the bidders values are correlated. This can be modeled as each bidder i having private information x_i with higher x_i implying a higher value v. In this case, $E(v|x_i) \geq E(v|x_i, x_i > x_j \ \forall\ j \neq i)$, so that winning is bad news about value (this problem of adverse selection in exchange is known as the winner's curse). We will treat this case in the next section, but one result is worth mentioning.

If the bidders have private values, but these values are not statistically independent, then Crémer & McLean (1985) show that the seller can extract all rents in the optimal auction for the case of a finite number of values. This is accomplished by offering each bidder a choice between a lottery and an second price auction. This works whenever i's value can be recovered from the values of the other bidders.

b. Milgrom and Weber (*Econometrica*, 1982)

Our analysis of the independent private value model is now complete. Recall that the seller's revenue is the same from any of the simple auctions with independent private values and risk neutral bidders:

English = 2nd Price = 1st Price = Dutch.

But if the bidders are risk averse then:

Dutch = 1st Price ≥ English = 2nd Price.

How then can we explain the frequent use of the English auction? One explanation is that values are positively correlated, as would be the case if there was some chance that the winning bidder would resell the object at a later date.

Common Value

The extreme case is where the bidders have the same value for the item, but this value is uncertain. Each bidder has a private estimate $x_i = v + \varepsilon_i$ of the value v, where ε_i represents the noise in the estimate.

You can imagine the item is a glass jar filled with coins. Each bidder in the room estimates the value of the coins in the jar, but the estimate is imperfect: some overestimate the value of the coins, others underestimate the value. In a first price auction, you would expect each bidder to make a bid that is an increasing function of the estimate. If everyone adopts the same bidding strategy, then the winner of the auction is going to be the bidder that overestimated the value of the coins the most. A bidder that does not condition his bid on the assumption that his estimate is the most optimistic among all the bidders will lose money quickly as a result of the winner's curse.

Winner's Curse: I won. Therefore, I overestimated the most. My bid only matters when I win, so I should condition my bid on winning (i.e., that I overestimated the most).

Winning is bad news about my estimate of value. This is a form of adverse selection that arises in any exchange setting: if you want to trade with me, it must be that no one else offered more, because they did not think that the item is worth what I am willing to pay. Understanding the winner's curse is essential for success in auctions and other trading games.

A second result that appears in common value auctions concerns the value of private information to the bidders. Here rents to the bidders come solely from the privacy of the information and not the quality of the information. For example, in an auction with three bidders, if two have the same information and a third has poorer but independent information, then the two with the same information will get a payoff of zero in equilibrium, whereas the one with poorer information gets a positive payoff. Rents come from the extra information you bring to auction. One implication of this result is that the seller should reveal all his private information.

A third result of common value auctions is that price tends to aggregate information: as $n \to \infty$, the price converges to the true value if the monotone-likelihood-ratio-property is satisfied.

General Symmetric Model with Affiliated Values

Milgrom and Weber's substantial accomplishment is the analysis of auctions with affiliated information. Their model has n bidders, each with private info x_i, which can be thought of as i's estimate of the value. Let $x = \{x_1,...,x_n\}$ be the vector of estimates. Bidder i's value of the good depends on the state of world s and the private information x. Let f(s,x) be the joint density, which is symmetric in x. Bidder i's value $v_i = u(s,x_i,x_{-i})$ is assumed to be symmetric in x, increasing, and continuous. Assume values are affiliated: if you have high value, it is more likely that I have high value.

Definition. The random variables $Z = \{S,X_1,...,X_n\}$ are *affiliated* if the joint density f(z) is such that $\forall$ z,z'

$$f(z \vee z')f(z \wedge z') \geq f(z)f(z') \text{ where } z \vee z' = \max\{z,z'\} \text{ and } z \wedge z' = \min\{z,z'\}.$$

Affiliation is graphically presented below:

Pr(all high)Pr(all low) $\geq$ Pr(high,low)Pr(low,high)

Pr(HH)Pr(LL) $\geq$ Pr(HL)Pr(LH)

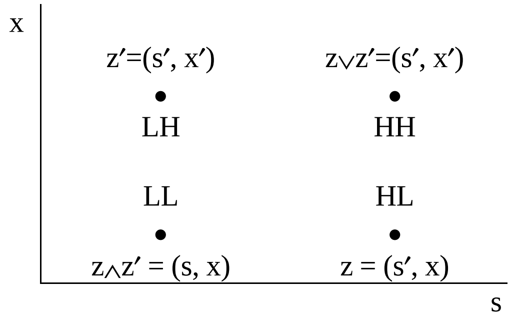

Affiliation can be stated in terms of the conditional density g(x|s). In this form it is called the monotone likelihood ratio property (MLRP).

Definition. The conditional density g(x|s) satisfies the MLRP if $\forall$ s < s' and x < x'

$$\frac{g(x|s)}{g(x|s')} \geq \frac{g(x'|s)}{g(x'|s')}$$ (i.e., the likelihood ratio is decreasing in x).

Properties of Affiliated RVs

$z = \{z_1,...,z_n\} \in$ A (affiliated)

1. E[g(z)h(z)|s] $\geq$ E[g(z)|s]E[h(z)|s]

2. $f \in A \Leftrightarrow \partial^2 \ln(f)/\partial z_i \partial z_j \geq 0$.

3. $f = g \cdot h$, $g,h \geq 0$, $g,h \in A \Rightarrow f \in A$.

4. $z \in A$, $g_1,...,g_k \uparrow \Rightarrow g_1(z_1),...,g_k(z_k) \in A$.

5. $z_1,...,z_k \in A \Rightarrow z_1,...,z_{k-1} \in A$.

6. $z \in A$, $H \uparrow \Rightarrow E[H(z)|a \leq z \leq b] \uparrow$ as a,b $\uparrow$.

7. $E[V_i \mid X_i = x, Y_1=y_1,...,Y_{n-1}=y_{n-1}] \uparrow$ in x, where $Y_1 \geq ... \geq Y_{n-1}$ are the order statistics of the n-1 other bidders.

Throughout the analysis, assume

1. no collusion.

2. the choice of auction doesn't reveal info.

3. the choice of auction doesn't affect who plays.

For each auction, Milgrom and Weber do the following:

1. Find the symmetric equilibrium bidding function.

2. Determine how the seller should use any private information.

3. Find the order of the simple auctions with respect to the seller's revenue.

The main result is that in terms of seller revenue:

English $\geq$ 2nd Price $\geq$ 1st Price = Dutch

The intuition for this result is as follows. The equilibrium bid function depends on everyone's information. The more (affiliated) information you condition on, the higher the bid.

How does the price depend on the bids in the simple auctions?

1. 1st Price: only 1st bid
2. 2nd Price: 1st and 2nd bids
3. English: all bids

Hence, the English auction does best because it involves conditioning on the most information. Here it is assumed that the bidders in an English auction observe the point at which each bidder drops out of the auction. This is an extreme assumption, but the result still holds so long as the players can infer some information as the bidding occurs.

<u>2nd Price</u>

Let $v(x,y) = E[V_i \mid X_i = x, Y_1 = y]$, which is increasing in x,y.

<u>Claim</u>. In a second price auction, the optimal strategy is to bid $b(x) = v(x,x)$; that is, bid your expected value given your signal is the same as the second highest.

<u>Proof Sketch</u>. Maximize the probability of winning whenever it is profitable, since your bid does not affect your payment. Hence, b(x) is chosen to maximize $E[(V_i - P)1_{\{P<b\}}|x]$, where $P = \max_{j\neq i} b(x_j)$. The solution to this maximization can be found by applying the revelation principle. Suppose the bidder reports x'. Then $P < b$ iff $b(y) < b(x')$ iff $y < x'$. Hence,

$$\text{select } x' \text{ to } \max \int_{-\infty}^{x'} [v(x,y) - b(y)]f(y|x)dy.$$

The FOC evaluated at $x' = x$ is $[v(x,x) - b(x)]f(x|x) = 0$.

Hence, $b(x) = v(x,x)$. QED

Other results:

1. If n+1st bidder's information is a garbling of X_1,Y_1, then n+1st bidder's profit = 0.

2. If $(X_0,S,X_1,...,X_n) \in A$ and the seller announces x_0, then $E[V(Y_1,Y_1) \mid X_i>Y_1] \leq E[V(Y_1,Y_1,X_0) \mid X_i>Y_1]$. This says that the expected sale price is greater if the seller announces x_0. The proof (like most of the proofs in this paper), uses a property of affiliation: the expected value is increasing when one conditions on something affiliated. The intuition is that by revealing affiliated information the seller is able to increase competition by reducing the winner's curse effect, which is driving bids down. It is assumed here that the information x_0 is verifiable, so there is no adverse selection problem for the seller.

English (assuming each bidder observes the other bidders dropping out)

Here the optimal strategy after k bidders have dropped out at prices $p_1,...,p_k$ is

$b_k(x|p_1,...,p_k) = E[V_i \mid x, (y_1,...,y_{n-k-1})=x,\{b_{j-1}(Y_{n-j}|p_1,...,p_{j-1})=p_j\}_{j=1,k}]$

Bid to the point where you are indifferent between winning and losing given your information. The seller's revenue from English ≥ 2nd Price, since the English auction reveals affiliated information (the n - 2 low bidders reveal their signals by the bid at which they drop out).

1st Price (Dutch)

To determine the equilibrium bidding strategy in the first price or Dutch auction, we again employ the revelation principle and analyze the direct revelation game. Suppose the bidder reports x'. Then

x' chosen to $\max \int_{\underline{x}}^{x'} [v(x,y) - b(x')]f(y|x)dy.$

The FOC evaluated at x' = x (the equilibrium condition) is

$$0 = [v(x,x) - b(x)]f(x,x) - \int_{\underline{x}}^{x} b'(x)]f(y|x)dy.$$

Hence,

$$b'(x) = [v(x,x) - b(x)]\frac{f(x|x)}{F(x|x)},$$

which is a first-order linear differential equation with boundary condition $b(\underline{x}) = v(\underline{x},\underline{x})$.

Claim. The seller's revenue from 2nd Price ≥ 1st Price.

Proof Sketch. Let $R(x',x) = E[V_i 1_{\{Y1<x'\}}|X_i = x]$. In either auction, the bidder will report x' to max R(x',x) - P(x',x)F(x'|x), where P(x',x) is the expected price and F(x'|x) is the probability of winning given x' is reported and i has information x. The first order condition evaluated at x' = x is

$0 = R_1 - P_1F - Pf$. Let P^1 = 1st price and P^2 = 2nd price, and

$P_1 = \partial P/\partial x'$; $P_2 = \partial P/\partial x$.

Suppose $P^1 > P^2$ (higher rev in 1st price).

1st price: $P_2^1 = 0$ (price only depends on x'). 2nd price: $P_2^2 \geq 0$ by affiliation (the higher x, the more likely second highest (y) is high).

From the FOC, $P_1 = [R_1 - Pf]/F$. Hence,

$$\frac{dP^2}{dx} = P_1^2 + P_2^2 \geq P_1^2 = \frac{R_1 - P^2 f}{F} \geq \frac{R_1 - P^1 f}{F} = P_1^1 = \frac{dP^1}{dx}$$

So $P^1 > P^2$ and $dP^2/dx \geq dP^1/dx$, but this is a contradiction, since P^1 and P^2 start at the same point (they have the same boundary condition). Hence, $P^1 \leq P^2$. QED

Linkage Principle

The main results about seller revenue can be understood in terms of the linkage principle: a higher price is obtained if the price is linked to more affiliated information.

Auction	Condition on
1st Price	winner's estimate $X_i \geq Y_1$
2nd Price	1st & 2nd estimate $X_i \geq Y_1$ & $Y_1 = x_1$
English	all estimates $X_i \geq Y_1$, $Y_1 = x_1$, $b(Y_j) = p_j$

Hence, English ≥ 2nd Price ≥ 1st Price and the seller should always reveal information. The more information you condition on the higher is the price.

2. Reputation and Cooperation

A reputation arises, whenever the other's belief about me is based on my past actions, and this belief affects what happens in future. Reputational effects can be important in dynamic games and repeated games. The two most common types of reputations are

(a) "I'm a tough guy; look out."

(b) "I'll be nice to you if your nice to me."

We will now illustrate how each of these reputations works.

(a) One vs. Sequence: "I'm tough, look out."

Our example is entry deterrence. There is an incumbent firm (I) in a market facing a sequence of potential entrants E_n. The stage game played between I and each E_n is as follows:

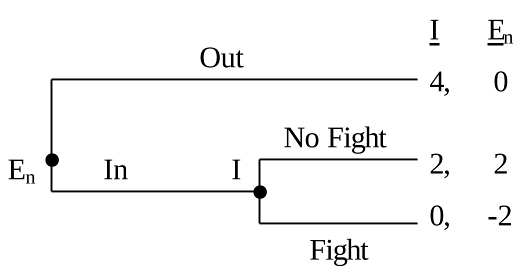

Notice that there is a unique subgame-perfect equilibrium in the one-shot game: the potential entrant enters and the incumbent does not fight, resulting in (2,2). To facilitate reputations it is assumed that entrant E_n knows what happened in the previous contests. But even then, in the finite-horizon we get (2,2) as the unique outcome by backward induction. The incumbent is not able to establish a reputation for fighting to deter other potential entrants.

But suppose that there is a small probability that I's payoffs are not as above, but instead it is a dominant strategy for I to fight. This is modelled as a game with incomplete information in which there are two types of incumbents, a rational incumbent that maximizes profit and a biological competitor that fights off any entry. Then we can get (4,0) in the early periods. In equilibrium, the rational incumbent imitates the incumbent type that always fights, creating a reputation for fighting.

This model has been studied by Kreps-Wilson and Milgrom-Roberts to resolve Selton's chain store paradox. They are able to show that (4,0) is the unique Nash equilibrium if there exists a small probability that fighting is a dominant strategy for I.

In the ∞-horizon game, one can get the reputation story to work without introducing incomplete information. In particular, the following strategies constitute a subgame-perfect equilibrium in which the entrants always stay out, resulting in the (4,0) payoff:

I: F when challenged, unless didn't fight in past, then never fight.

E_n: Stay out unless I failed to fight in past, then always enter.

One difficulty with this equilibrium is that the reputation is never tested. The reputation story is made more convincing if the reputation is created by I fighting off entrants. This is possible if some E_n's like to enter. The reputation still works so long as the value of the reputation is worth the cost of maintaining it: I sacrifices a short-run gain to preserve a reputation, resulting in future gains, which exceed the short-run loss. Let α = Pr(always enter) and δ = Pr(continue another period).

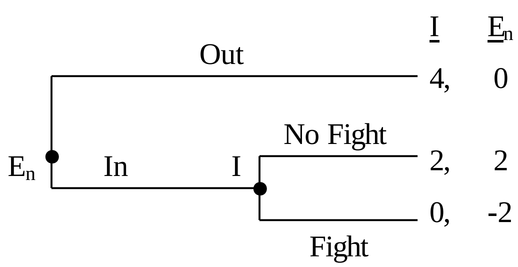

Should I fight? By fighting, I gets δV vs. $2/(1-\delta)$ if I does not fight, where V is the continuation value of maintaining the reputation as a fighter. Then

$V = 4(1-\alpha) + \delta V$, so $V = 4(1-\alpha)/(1-\delta)$. Hence, fighting is better than not fighting so long as $\alpha \leq 1 - 1/(2\delta)$; that is, if there are not too many entrants that always enter and the future is sufficiently important.

(b) Cooperation "I'll be nice to you if..."

We now turn to the second type of reputation, a reputation for being cooperative. Consider the following repeated Prisoner's Dilemma:

		Col C	Col D	
Row	C	1, 1	b, a	a > 1, b < 0, a + b < 2
	D	a, b	0, 0	

Here a > 1 and b < 0, so that D is a dominant strategy, and a + b < 2, so that (C,C) is the best outcome.

Can the parties establish a reputation for being cooperative (playing C), so as to get the (1,1) outcome? It is essential that the party that puts other at risk endure. Since both are at risk here, we cannot get the reputation story to work if only one party endures. (Unless there are sequential moves with the enduring party moving second.) But if both endure, then we can get cooperation at least in the infinite horizon game.

Evolution of Cooperation

Axelrod and Hamilton (1982) are interested in how cooperation can evolve. They explored this question by asking game theorists to determine what is a successful strategy? They then collected these strategies and let them compete in a tournament. The best strategy turned out to be the simplest one: tit-for-tat (play C and then do whatever the other did in the previous period).

Their explanation for why tit-for-tat did so well hinges on a number of the properties of the tit-for-tat strategy. Tit-for-tat is nice, provokable, forgiving, and easily recognizable. Although tit-for-tat never beats anything, it never loses by much (always nearly ties).

They then go on to provide an evolutionary argument for the success of tit-for-tat. tft is evolutionary stable: it cannot be invaded by a mutant. Defective mutants will die; cooperative mutants will live but won't harm tft.

But all D is evolutionary stable too. Mutant cooperators get clobbered and die; tft survives.

How then does tft get started? tft is initially viable: it survives in any population. Members of the same kinship play each other, then tft grows and prospers. All D dies.

Rational Cooperation in Finitely Repeated PD

Kreps-Milgrom-Roberts-Wilson (JET, 1982)

Turning back to the repeated game:

		Col C	Col D	
Row	C	1, 1	b, a	a > 1, b < 0, a + b < 2
	D	a, b	0, 0	

In the ∞-horizon game, any individually rational payoff (better than 0) can be supported as a SPE as $\delta \to 1$ by the folk theorem. But if the game is finitely repeated, only (0,0) occurs in any Nash equilibrium.

How can we explain cooperation in the finitely repeated prisoner's dilemma? In experiments the vast majority of play involves playing C now so get (1,1) in the future. What we need is some uncertainty about the preferences of one of the two parties. In particular, suppose that there is some small probability that Row always plays tft. Can the rational Row player sustain cooperation by pretending to be the tft player? The answer is yes, so long as there are sufficiently many periods.

N = total # of stages

n = # of stages to go

q_n = Pr(Row is tft), $\delta = q_N > 0$.

C_n = {C,D} Col's strategy in n; payoff π_n^C

R_n = {C,D} Row's strategy in n; payoff π_n^R

<u>Theorem</u>: In any SE, the number of D's is bounded above, independent of N. Hence, we get C in all but the last few rounds.

This is accomplished by Row imitating the irrational tft player. The proof is in seven steps.

<u>Proof</u>: 1. $q_n=0 \Rightarrow C_n=D, R_n=D$ (by induction)

2. $C_{n+1}=D \Rightarrow R_n=D$ (o.w. $q_{n-1}=0$)

3. $C_{n+1}= C \Rightarrow \pi_n^C \geq qn + b$ (Col uses C till D strategy)

($q=q_n$) $\qquad qn + (1-q)b \geq qn + b$

4. $C_{n+1} = D \Rightarrow \pi_n^C \geq q(n-1) + 2b$ (Worst C payoff.)

$\qquad R_n=D$ by (2); $C_n=C$ then (3).

5. $\pi_n^R \geq q(n-1) + 3b - a$ (Worst R payoff)

rat Row plays tft; gets within b - a of Col's pay.

6. R_n=tft if $q(n-2) > 2a - 4b$

(a) $C_{n+1} = D \Rightarrow R_n = D$ by (1).

(b) $C_{n+1} = C$: What is R_n?

$R_n = D \Rightarrow q_{n-1}=0 \Rightarrow \pi_n^R \leq a$.

$R_n = C \Rightarrow \pi_n^R \geq b + [q(n-2) + 3b - a]$

$\qquad$ if $C_n=D$ from (5) in pd n-1

so $R_n=C$ if $q(n-2) > 2a - 4b$, or $n > n^*=(2a-4b+2q)/q$

7. Col can't benefit from D before n^*, since Row plays tft.

No D before $n^* + 1$. QED

The proof constructs a very weak upper bound on the number of periods of D. In fact, it is often possible to get cooperation in all but the last period if δ large.

rat Row: play D in last; play D in next to last.

tft Row: C if C by Col in previous period; D if D by Col in previous period.

Col: D in last.

Col next to last. By playing D, Col gets: $\delta a + (1-\delta)(0 + 0) = \delta a$

By playing C, Col gets: $\delta(1+a)+(1-\delta)(b+0) = \delta a+\delta+b(1-\delta)$

Hence, $C \geq D$ if $\delta \geq -b(1-\delta)$ or $\delta \geq -b/(1-b)$. So long as δ is sufficiently large and b is sufficiently near 0, Col will play C in all but the last round, so long as Row plays C.

3. Reputations and Bargaining

a. Game Theoretic Models of Trading Processes (Wilson, 1987)

Three questions are addressed in this survey:

1. Are standard trading rules efficient?

   standard trading rules: auctions, double auctions, bid/ask market

   interim efficiency: not common knowledge dominated.

2. Dynamic Trading Models

   Bargaining: signal strength through delay

   Monopoly: Coase conjecture

3. Theories of Markets from Micro Models

As an example of how we can introduce competitive pressure in a bargaining model, consider the following model of a seller trading a single good with several buyers. The seller announces a high price and then gradually lowers it until one of the buyers accepts, as in a Dutch auction. There are two sources of impatience driving the buyers to accept before the price falls too much: (1) the discounting of future gains, and (2) the fear that one of the other buyers will accept first. Let p(t) be the seller's price path, which is a decreasing function of time t. Suppose each buyer's valuation v is drawn independently from the distribution F. A buyer's payoff is $[v - p(t)]e^{-rt}$ if he trades at time t at the price p(t).

A buyer's strategy is to accept at time t(v) (decreasing in v). We now determine the best response of a buyer to the price path p(t) given that the other buyers are waiting until t(v). Suppose our buyer waits until s. The buyer chooses s to max $[v - p(s)]e^{-rs}F(t^{-1}(s))^{n-1}$. The first-order-condition, after substituting the equilibrium condition s = t(v), is

$$|p'(t(v))| = [v-p(t(v))]\{r + (n-1)|t'(v)|f(v)/F(v)\}$$

↑ interest rate ↑ hazard rate another 1st

We are left with a first-order differential equation that determines the buyer's equilibrium response. The point I wish to make is that the two sources of impatience enter the equation as the two terms in curly brackets: (1) the discount rate r, and (2) the hazard rate that another buyer will accept first.

b. Bargaining, Dynamic Monopoly, and the Coase Conjecture

Gul, Sonnenschein, and Wilson (*Journal of Economic Theory*, 1986)

Consider a monopolist selling a durable good to a set of buyers. The static monopoly theory, has the monopolist set a price so that marginal revenue is equal to marginal cost. But this leads the monopolist to charging a price much greater than the competitive price, and the monopolist reaps monopoly profits. What if the monopolist can change his price over time. Would the monopolist want to price discriminate by selling at a high price to buyers with high values first and then lowering price for buyers with low values? The answer is no. The monopolist is best off if he can commit to selling only at the monopoly price.

But what if the monopolist cannot commit to not cutting price after all purchases have been made at the monopoly price? Clearly, the monopolist has an incentive to cut the price so as to sell to the remaining buyers. Coase (1972) conjectured that if the monopolist cannot commit to not dropping price, then the will quickly fall to his marginal cost. The purpose of this paper is to determine under what conditions the Coase conjecture is true.

The model is sketched below.

Seller's cost = 0 (without loss of generality).

Buyers $q \in [0,1]$; values (demand) f(q)

B's payoff if trade at t at p_t: $\delta^t[f(q) - p_t]$

The equilibrium conditions are:

(a) consumers correctly anticipate prices p(t)

(b) at every pt, monop can't increase profits by deviating from p(t).

There are actually two interpretations of this model:

(1) a durable goods monopolist selling to a market of buyers as described above, and (2) Bargaining with One-Sided Uncertainty: a seller with known valuation selling to a buyer, whose valuation v = f(q) is known privately.

The buyer's acceptance strategy is said to be <u>stationary</u> if it depends only on the current price and not the past history of prices.

Two cases

Gap Case (easy) $f(1) = \underline{v} > 0$.

∃! SPE by backward induction. The seller successively skims the high-value buyers from the market by offering lower and lower prices. The price falls to $\underline{v}$ in finite time, since eventually the remaining set of buyers is so small that the seller is offers $\underline{v}$ and all remaining buyers accept. The seller must randomize off the equilibrium path. A version of the Coase conjecture holds, even without requiring that the buyer's strategy is stationary: the initial price is less than $\underline{v} + \varepsilon$, as $\delta \to 1$.

No Gap Case $f(1) = 0$

In the no gap case, there is a marginal buyer. Notice that we cannot use backward induction in this case, since the game does not end in finite time. As a result, we get many equilibria. Indeed, there is a continuum of equilibria even with the buyer using a stationary strategy. The Coase conjecture, however, is true if buyer's strategy is stationary:

Coase Conjecture

$\forall\ \varepsilon>0, \exists\ \delta'<1$ s.t. $\forall\ \delta>\delta'$ and $\forall$ stationary equilibria σ, initial price is less than ε.

Intuition for Coase Conjecture

I now provide some intuition for why the Coase conjecture is true. First consider the problem in continuous time. Consider any price path p(t) with the buyer using a stationary strategy. If the price path isn't flat, then the seller can make more money by running the clock twice as fast.

In discrete time, the idea of the proof is to show if one prices much above 0, then ∃ lower price path that is better. Buyers use a stationary strategy. Hence, the seller can accelerate process: offer tomorrow's price today.

today: offer p accepted by q.

tomorrow: offer $p' < p$ accepted by $q' > q$.

cost: foregone profit on those that accept today

$$(p - p')(q' - q)$$

benefit: continuation value a day earlier

$$(1-\delta)v(q') \leq (p - p')(q' - q).$$

$$\frac{p-p'}{1-d}\frac{q'-q}{1-d} \geq \frac{v(q')}{1-d} \quad \text{and } t' - t = \Delta \approx 1-d$$

Let $\Delta \to 0$.

$$\left(\frac{p-p'}{\Delta}\right) \quad \left(\frac{q'-q}{\Delta}\right) \quad \geq \quad \left(\frac{v(q')}{\Delta}\right)$$

$$\left(\begin{matrix}\text{rate of}\\ \text{price}\\ \text{drop}\end{matrix}\right) \quad \left(\begin{matrix}\text{rate of}\\ \text{purchase}\end{matrix}\right) \quad \geq \quad \left(\frac{\text{continuation value}}{\text{length of day}}\right)$$

$$\begin{matrix}\uparrow & \uparrow & & \uparrow\\ \text{bounded or no one buys} & \text{assume bounded} & & \text{so } v(q') \to 0\end{matrix}$$

The rate of price drop must be bounded or no one buys. Hence if we assume that the rate of purchase is bounded as well, then the continuation value must go to zero as $\Delta \to 0$. Therefore, the price is near 0 in the future, so it must be near 0 today too (or else the buyers would wait for the price near 0).

If the rate of purchase unbounded, then price must be flat (so a positive fraction is willing to accept), but price is tied down at end (must be 0), so price is near 0 today.

Determination of Stationary Equilibrium

A stationary equilibrium is calculated by solving the following dynamic programming problem:

buyer type $q \in [0,1]$; value f(q)

stationary strategy P(q): q accepts $p \leq P(q)$

continuation value V(q): S's value when B on [q,1].

S strategy y(q): offer price P(y(q)) accepted by y(q)

The best response conditions are then:

(1) buyer q is indifferent between P(q) today and P(y(q)) tomorrow.

B best response: $f(q) - P(q) = \delta[f(q) - P(y(q))]$.

(2) S selects y to maximize profit.

S best response: $v(q) = \max_{y \geq q} P(y)(y-q) + \delta V(y)$

Bargaining Interpretation

Fudenberg, Levine, and Tirole (1985)

How does the problem change if under the bargaining interpretation?

one seller, one buyer w/ uncertain valuation v

$v \sim F(v)$

seller makes all offers.

Same problem, but must find a SE not SPE, since this is a game with incomplete information. One difficulty with this model is that in bargaining one typically thinks that both trades can make offers.

What if the bargainers alternate making offers?

Two papers address this problem:

(1) Grossman and Perry (JET, 1986)

There is a unique PSE for δ not too large (range of B's pool).

As $\delta \to 1$, B's that make serious offer shrink to 0, and there does not exist a PSE.

(2) Gul and Sonnenschein (Econometrica, 1987):

Prove Coase conjecture for bargaining with alternating offers and stationary buyer strategy.

<u>What if drop stationarity of buyer's strategy?</u>

Ausubel and Deneckere (1989)

B on $q \in [0,1]$ w/ value $v = f(q)$

preference: $\delta^t[f(q) - p_t]$; accept p if $v \geq \beta(p)$

S w/ cost 0 offers price $\sigma_t(h^t)$ w/ history h^t.

S max PV of profits: $\pi(v)$ = continuation value.

q_t = residual demand after t periods.

stationarity: cutoff value β only depends on p.

so B can't punish S for accelerating decline.

<u>weak Markov (payoff relevant + last period)</u>

S: $p_{t+1} = P(q_t,p_t)$ B: $\beta(p)$

<u>Strong Markov (payoff relevant)</u>

S: $p_{t+1} = P(q_t)$ B: $\beta(p)$

<u>Work with p or q? Doesn't matter</u>

<u>w/ p</u> $\phi(p,\beta(p))$ offered tomorrow.

B: $\beta(p) - p = \delta[\beta(p) - \phi(p,\beta(p))]$

S: $\pi(v) = \max_p p[f_{-1}(\beta(p)) - f^{-1}(v)] + \delta\pi(\beta(p))$

<u>w/ q</u> P(y(q)) offered tomorrow.

B: $f(q) - P(q) = \delta[f(q) - P(y(q))]$

S: $V(q) = \max_y P(y)(y - q) + \delta V(y)$

When can we solve this dynamic program?

Want distribution to stay same after skimming.

closed form solution if $f(q) = 1 - q^{-m}$ for $m > 0$.

<u>Linear Case</u> $f(q) = 1 - q$.

$\phi(v) = \alpha v$ $\pi(v) = (\alpha/2)v^2$

$\beta(p) = p/(1-\delta)^{1/2}$ $\alpha = 1 - 1/\delta + (1-\delta)^{1/2}/\delta$

Note: as $\delta \to 1$, $\alpha \to 0$, and $p_0 \to 0$.

Suppose with commitment (fixed price p^*) get profit π^*.

<u>Folk Theorem</u>: $\forall\ \varepsilon>0\ \exists\ \delta'<1$ s.t. $\forall\ \delta > \delta'\ \exists$ SPE w/ payoff $\pi \in [\varepsilon, \pi^*-\varepsilon]$.

<u>Proof Sketch</u>. Use reputational strategy: (B punishes S)

If ever deviate from price path that yields π, play Coase equilibrium. Price path must decrease to 0, but only so fast as to generate Coase payoff which decreases to 0 as $\delta \to 1$. QED

Best S can do as function of δ is U shaped.

get π^* w/ $\delta = 0$:

weak punishment, but committed to price.

get π^* w/ $\delta = 1$:

strong punishment

get $\pi < \pi^*$ w/ $\delta \in (0,1)$:

future matters and limited punishment.

<u>What if more than one firm?</u>

Gul (Rand,87), Ausubel and Deneckere (Rand,87):

Get folk theorem even with stationarity: Can get other to punish if deviate from price path (so don't need B to punish). Both price at MC is SPE.

Competition can yield monop profit but not monopoly!

<u>What if monopoly but potential entrant?</u>

Again get folk theorem. Equilibrium with three phases:

P1. I charges close to π^*

P2. Else entry w/ SPE supporting P1.

P3. Else MC pricing if P2 not followed.

One is almost enough for monopoly!

<u>Does stationarity make sense?</u>

In favor of stationary is that $\exists!$ SPE in gap case, <u>regardless</u> of how small the gap is, and this SPE satisfies the Coase conjecture. There are, however, a number of arguments against stationarity:

1. <u>History does matter</u>: S doesn't accelerate price cuts, because B will wait for lower prices. B uses current concession rate to infer future rate.

2. <u>Stationarity may lead to crazy outcomes</u>: e.g. ∞ repeated prisoner's dilemma: stationarity $\Rightarrow$ (D,D)

3. <u>Strong Markov equilibria may not exist</u>:

S strategy not stationary; only B.