Econ 805 – Advanced Micro Theory I
Dan Quint
Fall 2007
Lecture 12 – Oct 16 2007

Last week, we relaxed both private values and independence of types, using the Milgrom-Weber setting of affiliated signals. We found that "maximal-information" button auctions revenue-dominated "minimal-information" button auctions, which dominated sealed-bid first-price auctions; and that revealing his own information was, in expectation, also beneficial to the seller.

Recalling that private values are a special case of the Milgrom-Weber setup, we've now found that

- Risk-aversion favors the first-price over the ascending auction...
- ...but affiliation of signals favors the ascending over the first price

To try to capture the benefits of both – and to avoid a couple other problems we'll look at shortly – Paul Klemperer proposed a hybrid of the two auction types, called an Anglo-Dutch Auction. An Anglo-Dutch auction works like this. Suppose there are N bidders. Hold an ascending auction until all but two of the bidders have dropped out. Then hold a first-price auction among between the two remaining bidders, with the reserve price equal to the price at which the third-to-last bidder dropped out.

The first stage can be thought of as "information extraction," since the actual types of the $N-2$ lowest bidders are revealed by their equilibrium play. (This information revelation was what led the max-information button auction to dominate the others under risk neutrality.) But then holding a sealed-bid auction among the last two bidders gives some of the benefits of a first-price auction under risk-aversion. (Using a sealed-bid auction for the final stage also solves two other problems we'll come to later: the existence of low-revenue alternative equilibria in the ascending auction, and the relative ease of collusion.)

I sent out a link last week to a short note by Dan Levin and Lixin Ye, "Hybrid Auctions Revisited," which is forthcoming in Economics Letters. They show an example of a setting with affiliated private values and risk-averse bidders. They also generalize Klemperer's Anglo-Dutch auction by allowing the ascending portion to end with more than two bidders remaining (holding a sealed-bid first-price auction among anywhere between 2 and $N-1$ survivors).

The Levin and Ye example uses a common technique for generating positively-correlated signals: conditional independence. They assume an unobserved but stochastic "state of the world" V, drawn from the uniform distribution on the interval $[\underline{v}, \overline{v}]$. Conditional on a realized state $V = v$, bidder valuations x_i are then uniform draws on the interval $[v-\epsilon, v+\epsilon]$.

They implement risk-aversion using a particular utility function, the Constant Relative Risk Aversion utility

$$u(y) = y^{1-\rho}$$

for $\rho \in [0, 1)$. It is so named because the coefficient of relative risk aversion,

$$-y \frac{u''(y)}{u'(y)}$$

turns out to be equal to ρ everywhere, therefore constant.

Levin and Ye consider hybrid (Anglo-Dutch) auctions where an ascending auction is held until k bidders have dropped out, and then a first-price auction is held among the remaining $N - k$ bidders.

The results they find are as follows:

- For any realization of v, the expected revenue in any hybrid auction is at least as high as it would be in an English (pure ascending) auction.
- In their example, the optimal hybrid auction is when $k = 1$, that is, the ascending portion should only continue until one bidder drops out. (This doesn't seem to be very general – seems related to the assumption of uniform distributions.)
- For ρ sufficiently large large (bidders sufficiently risk-averse), the first-price auction dominates the optimal hybrid auction for "most" realizations of v, and therefore in expectation
- On the other hand, for ρ sufficiently small, the optimal hybrid auction dominates the first-price auction
- (This last point: we know the optimal hybrid auction revenue-dominates the ascending auction with max info; and we know from Milgrom-Weber that when $\rho = 0$, the ascending dominates the first-price. So this last result is by continuity.)

So they find that the optimal hybrid auction dominates the pure ascending auction, but that whether the hybrid or the first-price is better depends on the degree of risk aversion.

Common Value Auctions

Another special case of the Milgrom and Weber setup is the case of common values. That is, settings where at any realized type profile, all bidders have the same valuation for the object

$$v_i(t) = v_j(t)$$

(These are sometimes called "pure common values," to differentiate them from settings with common and private components.)

Common values are commonly used to model

- Natural resource auctions. The value of leases to drill for oil and natural gas depend on the amount underground, which is not known precisely until it is extracted. The value of rights to log in federally-owned forests depends on the number of type of trees, which can similarly only be estimated.

- Corporate takeovers. The value of a company is the discounted value of future profits, and may not depend much on who owns the company.

- Spectrum auctions. The value of rights to parts of the wireless spectrum may depend more on the overall profitability of the cell phone and other industries, more than the particulars of which company controls what.

In all these cases, uncertainty about the object's actual value may dwarf differences in ex-post value to the different bidders, so common values might be a reasonable approximation of the truth. (Thursday, we'll come back to the question of whether "close" is close enough.)

Drainage Tract Auctions

For today, we'll focus on an early common-values model, the Drainage Tract model. The model was created in the late 1960s, to model first-price auctions for offshore oil drilling rights. A **drainage tract** is a tract next to one that is already being drilled by an oil company. Thus, in drainage tract auctions, one company is expected to have better information about the true value of the land than its competitors. (The company with the neighboring tract, and therefore the better information, is referred to as the neighbor.)

The model is credited to Bob Wilson, but with a different bidding procedure. Equilibrium bidding in a conventional first-price auction is analyzed in a paper by Richard Englebrecht-Wiggans, Paul Milgrom, and Robert Weber, and further in another paper by Milgrom and Weber. Drainage tract auctions were also examined empirically by Ken Hendricks, Robert Porter, and Charles Wilson, who found that nearly all the theoretical predictions of the model held empirically, with one exception; they introduced a change to the model that fixed that theoretical prediction to match the data.

The model assumes there is just one neighbor, that is, just one firm with an informational advantage. They assume that the non-neighbors have no private information, only access to the public information that all of them (and the neighbor) can see; and that the neighbor has access to some additional relevant private information.

For the analysis of equilibrium play, it is without loss of generality to assume that the neighbor knows the value exactly, since all that matters is its expected value conditioned on all available information, all of which the neighbor is assumed to have.

Non-neighbors, who know nothing, will turn out to play mixed strategies in equilibrium. However, it helps to assume that they get an uninformative private signal, and play an equilibrium strategy that is increasing in that signal.

Assume the neighbor is bidder 1. We assume that all bidders get independent signals t_i, which are normalized to be drawn from the uniform distribution on $[0, 1]$, and that the value of the object to whoever wins it is $v(t_1)$, where v is increasing. (Independence of signals is not a meaningful assumption here, since all signals but t_1 are simply randomizing devices for bidders playing mixed strategies.)

Equilibrium is easiest to describe with just one non-neighbor:

Theorem 1. *Suppose v is continuous, differentiable, nondecreasing, and has a positive right-hand derivative at 0. If $N = 2$, the first-price auction in this setting has a unique equilibrium. The neighbor bids*

$$\beta(s) = E(v(t_1)|t_1 < s) = \frac{1}{s}\int_0^s v(r)dr$$

The non-neighbor bids uses **the same bid function** *as a function of his own signal, and gets expected profit of* 0 *regardless of his bid.*

That is, the unique equilibrium of the drainage tract has the informed bidder bidding an increasing function of his information, and the uninformed bidder playing a mixed strategy with exactly the same distribution of bids.

First, we prove this is an equilibrium; then that it is unique. Consider first the uninformed bidder. At any bid $x = \beta(y)$, the uninformed bidder wins whenever $\beta(t_1) < \beta(y)$, for $t_1 < y$. Thus, when he wins, he pays $\beta(y)$ for an object worth, in expectation,

$$E\left(v(t_1)|t_1 < y\right) = \beta(y)$$

so his expected payoff is 0. Bidding less than $\beta(0)$ wins with probability 0, and bidding above $\beta(1)$ is dominated by bidding $\beta(1)$, so bidding $\beta(t_2)$ is a best-response.

Next, the informed bidder. Assume his true type is t_1, his opponent bids according to β, and he chooses to bid $\beta(t)$ for some $t \in [0, 1]$. Then his expected payoff is

$$t(v(t_1) - \beta(t))$$

since he wins whenever $t_2 < t$. Differentiating gives his first-order condition

$$v(t_1) - \beta(t) - t\beta'(t)$$

This should equal 0 at $t = t_1$, so

$$v(t_1) = \beta(t_1) + t_1\beta'(t_1) = \frac{d}{dt}(t_1\beta(t_1))$$

Integrating from 0 with respect to t_1 gives

$$\int_0^x v(t_1)dt_1 = x\beta(x)$$

giving the same bid function β as the solution to the FOC. β is increasing, and since it satisfies the first-order condition, it satisfies the envelope condition; so it is a best-response among possible bids $\beta(y)$. Bids below $\beta(0)$ or above $\beta(1)$ are dominated, and so it's an equilibrium strategy.

To show uniqueness, note that any equilibrium bid functions β_1 and β_2 for the two bidders must have the same range, and that range must be convex – otherwise some type could lower his bid without lowering his chance of winning. Bidder 2 must play a mixed strategy with no point masses, or else bidder 1 will have no best-response at certain types (he'll want to outbid the mass by the minimum amount). Since bidder 2 mixes, he must get the same expected payoff from every bid; but his lowest bid must win with probability 0, so he must get expected payoff of 0 whatever he bids. This means that the expected value of $v(t_1)$, conditional on $\beta_1(t_1) < x$, must be equal to x; or

$$\int_0^{\beta_1^{-1}(x)} (v(t_1) - x)dt_1 = 0$$

Plug in $x = \beta_1(r)$, since x must be within the range of β_1, and we find

$$\int_0^r (v(t_1) - \beta_1(r))dt_1 = 0$$

or $\beta_1(r) = \frac{1}{r}\int_0^r v(t_1)dt_1$. So this is the only possible equilibrium strategy for the informed bidder.

The first-order condition of bidder 1's optimization problem, along with the boundary condition $\beta_1(1) = \beta_2(1)$, establish that this is also the only strategy for the uninformed bidder that makes β_1 the neighbor's best-response.

When $N > 2$, equilibrium is very similar:

Theorem 2. *When $N > 3$, a set of bidding functions $(\beta_1, \beta_2, \ldots, \beta_N)$ are an equilibrium if and only if*

$$\beta_1(t_1) = \frac{1}{t_1}\int_0^{t_1} v(s)ds$$

and the probability distribution of

$$\max\{\beta_2(t_2), \ldots, \beta_N(t_N)\}$$

is the same as the probability distribution of $\beta_1(t_1)$.

The proofs are basically the same. The non-neighbors still get expected profits of 0 conditional on winning, so they can mix in any way that collectively gives the neighbor the right incentives.

This equilibrium characterization establishes some very strong theoretical predictions about drainage tract auctions:

- The distributions of the neighbor's bids should be the same as the distribution of the highest non-neighbor bid, and should not depend on the number of non-neighbors participating in the auction
- Non-neighbors should average 0 profits when they win; neighbors should earn positive profits when they win

Hendricks, Porter and Wilson examined a large number of actual drainage tract auctions held between 1959 and 1979, and found the following:

- When there were multiple neighbors, generally only one of them bid on a particular tract, so the model with just one neighbor fit fairly well (likely due to collusion)
- Winning neighbors got average ex-post profits of 180% of their bid; winning non-neighbors got average ex-post profits of approximately their bid (therefore 0 net profits)
- The distribution of the neighbor's bid and the highest non-neighbor's bid match up almost perfectly at the high end, but the non-neighbor's distribution is "too high" at the low end

To "fix" the low end of the distribution, Hendricks Porter and Wilson modified the model to allow for a random reserve price, which could be correlated with the true value; and found that the data fit the model fairly well. Given the somewhat surprising predictions, and the fact that they were borne out by the data, the drainage tract model is generally considered to be a smashing success.

The Value of Information

Milgrom and Weber, in "The Value of Information in a Sealed-Bid Auction," analyze the sensitivity of expected payoffs in drainage tract auctions to changes in the information setup. Many of the results in one of my working papers (which we'll look at in a couple weeks) are extensions of these results to more general common value auctions. Everything here is done assuming only one non-neighbor, but many of the results extend.

First, we can calculate the expected payoffs of the neighbor. If the non-neighbor bids using β, the probability of winning given a bid of b is $\Pr(\beta(t_2) < b) = \Pr(t_2 < \beta^{-1}(b)) = \beta^{-1}(b)$, so

$$V_1(t_1) = \max_b (v(t_1) - b)\beta^{-1}(b)$$

and therefore

$$V_1'(t_1) = v'(t_1)\beta^{-1}(b)$$

evaluated at $b = \beta(t_1)$, or $V_1' = t_1 v'(t_1)$. By symmetry, a neighbor with type 0 wins with probability 0 in equilibrium, so $V_1(0) = 0$, so

$$V(t_1) = \int_0^{t_1} s v'(s) ds$$

We can do a change of variables with $w = v(s)$, and then

$$V(t_1) = \int_0^{v(t_1)} F(w) dw$$

where F is the distribution of $v(t_1)$. We can also take the expected value of this (by integrating over $v(t_1)$ with its corresponding density) and find that

$$E_{t_1} V(t_1) = \int_0^{v(t_1)} (1 - F(w))F(w) dw$$

Now suppose that there are two signals available, X and Y, about the value the object.

Theorem 3. *Let*

- *A be expected payoffs to the neighbor if it's common knowledge that he knows only X*
- *B be expected payoffs to neighbor if he knows X and Y but the non-neighbor bids as if he only knows X*
- *C be expected payoffs to neighbor if it's common knowledge he knows X and Y*

Then $C \geq B \geq A$.

$B \geq A$ is obvious, since the neighbor could pretend he didn't know Y if it doesn't help him.

To see that $C \geq B$, we'll actually show that expected payoffs are weakly higher at each realized value of

$$v = E(V|X, Y)$$

Let F_X be the distribution of $E(V|X)$, and F_{XY} the distribution of $E(V|X, Y)$.

Expected payoffs in case B, at a given realization v, are

$$\int_0^v F_X(w)dw$$

This is because once I calculate $E(V|X, Y)$, my expected payoff is just based on that value and the bid distribution I'm up against, which is the same as in auction A. On the other hand, given v, my expected payoff is

$$\int_0^v F_{XY}(w)dw$$

in case C. Iterated expectations show that

$$E(V|X = x) = E_{y|X=x}E(V|X = x, Y = y)$$

or to put it another way, if we define

$$Z = E(V|X = x) - E(V|X = x, Y = y)$$

then $E(Z|X) = 0$ for all X. So F_{XY} is a mean-preserving spread of F_X; and therefore by second-order stochastic dominance,

$$\int_0^v F_{XY}(s)ds \geq \int_0^v F_X(s)ds$$

so $C \geq B$ at any realized value v. Taking the expectation over v finishes the proof.

So more information is better for the neighbor; and it's even better when the neighbor knows you have better information. Intuitively, this is because when the neighbor is better informed, the non-neighbor faces a more severe winner's curse, and therefore has to bid more cautiously, which is good for the neighbor.

Now look at it from the uninformed bidder's side:

Theorem 4. *Suppose it's common knowledge that the neighbor knows X and Y. The non-neighbor does better learning Y and keeping it secret than learning Y publicly.*

If the non-neighbor learns Y publicly, we're back in the situation where the neighbor has some private information (X) and the non-neighbor has only public information, so the non-neighbor still gets expected payoff 0.

On the other hand, if the non-neighbor learns Y secretly, then in some cases, he can get positive expected payoffs. For instance, if

$$E(V|Y) > E(V)$$

that is, if Y increases the expected value of V, the non-neighbor can guarantee an expected payoff of the difference $E(V|Y) - E(V)$ by bidding at the top of the neighbor's range and always winning.

Finally, suppose the neighbor knows X and the seller knows Y. Here, the result is more ambiguous. Milgrom and Weber decompose the seller's expected gain from revealing Y to both bidders into two effects, a "publicity effect" and a "weighting effect". The publicity effect is that, if X and Y are correlated, then revealing Y gives away some information about X as well, making it less private and therefore less valuable. So the publicity effect always favors revealing the information.

The weighting effect is that knowing Y can make $E(V)$ either more or less sensitive to X, and can therefore increase or decrease the value of knowing X. For instance, suppose $V = X + \epsilon$. Knowing ϵ may make knowledge of X more valuable, which would make X and ϵ "informational complements". On the other hand, X and Y can be informational substitutes if knowing Y makes V less sensitive to X. (Formal definition in Milgrom book, or in Milgrom Weber.)

The publicity effect always favors revealing Y, just like it did in the ascending auctions we studied last week. The weighting effect can go in either direction. But:

Theorem 5. *If (X, Y, V) affiliated, then publicly announcing Y increases the seller's expected revenue.*

(If bidder 1 already knows both X and Y, then revealing Y always increases the seller'e expected revenue.)

(Like last week's results with revealing t_0, all these are ex ante results – that is, committing ahead of time to always revealing Y dominates (or in some cases, doesn't dominate) planning to never reveal Y. If the seller could selectively reveal Y, or better yet, lie about Y, things change.)

Today, we saw a particular common values setting with a unique equilibrium in the first-price auction. Thursday, we'll see that second-price auctions with common values tend to admit multiple equilibria, and that a small private-value advantage for one bidder can have an explosive effect on equilibrium outcomes (Klemperer's "Almost Common Values" paper).

In addition, the drainage tract auction with one non-neighbor featured a surprising type of symmetry: even though the bidders were asymmetrically informed, they ended up having the same equilibrium bid distributions. This turns out to hold in some other types of common-value auctions as well, and we'll look at a couple of those.