Econ 805 – Advanced Micro Theory I
Dan Quint
Fall 2009
Lecture 9

We're back in the Milgrom-Weber affiliated interdependent values model. Today, we're going to compare first-price to ascending auctions.

Milgrom and Weber spend a bunch of time and notation demonstrating the existence of a symmetric equilibrium in the first-price auction. For our purposes, we're going to just accept that it exists, since we don't need a closed-form expression for bidding functions to prove the result.

First, a useful lemma:

Lemma 1. *Suppose g, h are continuous differentiable functions, and $g(0) \geq h(0)$. If for all $x \geq 0$, $g(x) \leq h(x)$ implies $g'(x) \geq h'(x)$, then $g(x) \geq h(x)$ for all x.*

Proof. Suppose not: that is, suppose $g(y) < h(y)$ for some $y > 0$. Let $z = \sup\{z' < y \mid g(z') \geq h(z')\}$; by continuity of g and h, $g(z) = h(z)$. So $g(z) = h(z)$, and $g(z') < h(z')$ for $z' \in (z, y]$. By assumption, $g'(z') \geq h'(z')$ for z' within this range; integrating gives $g(y) - g(z) \geq h(y) - h(z)$, which contradicts the joint claim that $g(z) = h(z)$ and $g(y) < h(y)$. □

Our main result is this:

Theorem 1. *For each type of bidder, the expected price paid in the second-price auction, conditional on winning, is at least as high as the bid in the first-price auction.*

Recall that

$$\hat{v}(x, y) = E(v_1 | t_1 = x, t^{(1)} = y)$$

and define $F(t^{(1)}|t_1)$ as the conditional cumulative distribution function of $t^{(1)}$, the highest opponent type, conditional on a realization t_1 of my own type, and $f(t^{(1)}|t_1)$ as the corresponding density.

Let β^F denote the equilibrium bidding function in the first-price auction. Recall that the equilibrium bid in the second-price auction is $\hat{v}(t, t)$, so our theorem is that

$$E\left(\hat{v}(t^{(1)}, t^{(1)}) | t^{(1)} < t_1 = s\right) \geq \beta^F(s)$$

Define

$$\beta^S(s, t) = E\left(\hat{v}(t^{(1)}, t^{(1)}) | t^1 = s, t^{(1)} \leq t\right)$$

as the expected price that bidder 1 expects to pay in the second-price auction, conditional on having a true type of s but bidding as if he were type t. In equilibrium, then, a bidder with type s will bid $\beta^S(s, s)$, so we want to show that $\beta^S(s, s) \geq \beta^F(s)$.

Let

$$\tilde{v}(s, t) = E\left(\hat{v}(s, t^{(1)}) | t^1 = s, t^{(1)} < t\right)$$

be the expected value of the object to bidder 1, conditional on bidder 1 having true type s, bidding as if he were type t, and winning. Then his expected payoff is

$$V^F(s) = \max_t \left(\tilde{v}(s,t) - \beta^F(t)\right) F(t|s)$$

where F is once again the cumulative distribution function of his highest opponent's type. In equilibrium, this max is achieved by setting $t = s$, so by the envelope theorem, V' is the s-derivative of the right-hand side, evaluated at $t = s$, which is (using numbers to denote partial derivatives)

$$(V^F)'(s) = \tilde{v}_1(s,s)F(s|s) + \left(\tilde{v}(s,s) - \beta^F(s)\right) F_2(s|s)$$

In the second-price auction,

$$V^S(s) = \max_t \left(\tilde{v}(s,t) - \beta^S(s,t)\right) F(t|s)$$

and so

$$(V^S)'(s) = \tilde{v}_1(s,s)F(s|s) - \beta_1^S(s,s)F(s|s) + \left(\tilde{v}(s,s) - \beta^S(s,s)\right) F_2(s|s)$$

Now, we want to apply the ranking lemma to $g(x) = V^F(x)$ and $h(x) = V^S(x)$. We can write these as (plugging $t = s$ into the objective functions above) $V^F(s) = \left(\tilde{v}(s,s) - \beta^F(s)\right) F(s|s)$ and $V^S(s) = \left(\tilde{v}(s,s) - \beta^S(s,s)\right) F(s|s)$.

We know that $V^F(0) = V^S(0) = 0$. If $V^F(s) \leq V^S(s)$, then $\beta^F(s) \geq \beta^S(s,s)$. But we know (we proved last time) that $F(t|s)$ is decreasing in s, so $F_2(s|s)$ is negative. And by affiliation, β_1^S is positive. Comparing the envelope theorem equations gives $V^F(s) \leq V^S(s) \ \rightarrow \ (V^F)'(s) \geq (V^S)'(s)$, so by the ranking lemma, $V^F(s) \geq V^S(s)$ for all s; so $\beta^F(s) \leq \beta^S(s,s)$.

We also have the same result as before:

Theorem 2. *For each type of bidder, the expected payment by that bidder, conditional on winning, when the seller reveals t_0 is at least as high as the expected payment when the seller does not reveal t_0.*

The Linkage Principle

We've now shown that from an expected revenue point of view,

max info ascending t_0 revealed	$\geq$	min-info ascending t_0 revealed	$\geq$	first-price sealed-bid t_0 revealed
$\geq$		$\geq$		$\geq$
max info ascending t_0 hidden	$\geq$	min-info ascending t_0 hidden	$\geq$	first-price sealed-bid t_0 hidden

(Start with bottom row, then "similarly show" top row.)

It's also clear that, as you move up and to the left on this chart, more information is revealed, and in particular, the winning bidder's payment (in equilibrium) is a function of more information.

This is an illustration of the linkage principle:

"The more closely the winning bidder's payment is linked to his actual type (as opposed to his bid), the greater the expected revenue will be"

which, in this case, can be rephrased as,

"The more things the winning bidder's payment depends on that are positively correlated with his type, the greater the expected revenue will be"

(One way to think about this is that bidder's expected payoff can be thought of as his "information rents," that is, the extra surplus he is able to get by having private information. But in auction formats where information is revealed which is correlated with his private information, his private information becomes "less private" in a sense, so he gets a smaller surplus, and therefore more goes to the seller. In the logical limit – where information revealed over the course of the auction fully reveals the highest type – the seller could simply make a take-it-or-leave-it offer to extract full surplus.)

Other Stuff

Milgrom and Weber also give some results on auctions with reserve prices, and auctions with reserve prices and entry fees; as well as some results on risk-averse bidders (but only with constant absolute risk aversion).

Risk Aversion and Correlation

We found today that with risk-neutral bidders, when types are correlated (affiliated) or values are interdependent, the ascending auction outperforms the first-price auction. On the other hand, you'll show on the second homework that with independent private values, when bidders are risk-averse, the first-price auction outperforms the ascending auction. The obvious question is, what to do when there is both risk-aversion and correlation.

To try to capture the benefits of both the first-price and ascending auctions, as well as to avoid other problems we'll talk about later, Paul Klemperer proposed a hybrid of the two auction types, which he called the Anglo-Dutch Auction. An Anglo-Dutch auction works as follows. We begin with an ascending auction, which goes on until all but two bidders drop out; then the remaining two bidders participate in a sealed-bid auction with a reserve price equal to the price at which the third-to-last bidder dropped out. This way, the ascending portion of the auction is used for information extraction – which is valuable if signals are affiliated – while the sealed-bid portion is used to increase revenue if bidders are risk-averse.

There's a nice note by Dan Levin and Lixin Ye which works through a numerical example, and shows that in their environment (with private values, affiliated types, and risk-averse bidders), the hybrid auction gives strictly higher revenue than a pure ascending auction. They use conditional independence to generate correlated private values – there is some unobserved "state of the world"

v, and once that is realized, bidders have independent private values drawn from the uniform distribution on the interval $[v - \epsilon, v + \epsilon]$. They assume bidders have Constant Relative Risk Aversion utility $u(y) = y^{1-\rho}$ for $\rho \in [0, 1)$.

(In their example, the optimal hybrid auction turns out to have a shorter ascending portion than in Klemperer's Anglo-Dutch auction – rather than waiting for all but two bidders to drop out, they find it's better to wait for just one bidder to drop out, and then switch to a sealed-bid auction immediately. But I think that's just a feature of their distributional assumptions, not a generalizable result.) Levin and Ye They show that when there is affiliation and a moderate amount of risk-aversion, this hybrid auction can dominate both a pure ascending auction and a pure first-price sealed-bid auction. However, when bidders are sufficiently risk-averse, a pure first-price auction still generates higher revenue.

Peter Eso, in a paper I just added to the syllabus recently, analyzes another example, analyzes another example of a setting with correlation and risk-aversion. He uses discrete types – each bidder's private value has just two possible realizations – and constant absolute risk aversion, $u(w) = -e^{-rw}$, $r > 0$. Recall the example from Myerson – with risk-neutrality and correlated types, the seller can extract all the bidders' surplus, leaving them with nothing. Eso finds that full surplus extraction is still possible in his setting when the correlation between bidder types is sufficiently strong.

Common Value Auctions

One 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. (We'll come back later to the question of whether "close" is close enough.)

Drainage Tract Auctions

We'll talk about the drainage tract model next lecture.

Ex Post Equilibria

Equilibrium Multiplicity

In the drainage tract model, we found that the first-price auction had a unique equilibrium. (Well, actually unique when $N = 2$, functionally unique when $N \geq 3$.) In the Milgrom Weber paper, when equilibrium strategies were defined for the two types of ascending auctions, however, it was claimed only that these strategies were an equilibrium, not that there were no others. Which brings us to the following problem: with common values, there are often lots of equilibria.

Let's look at a simple example of this: two bidders, with types t_1 and t_2 which are independently $U[0, 1]$, and

$$v_1(t_1, t_2) = v_2(t_1, t_2) = t_1 + t_2$$

It's not hard to see that each bidder bidding $b_i(t_i) = 2t_i$ is one ex post equilibrium of the second price auction. If my opponent bids $2t_j$, this is the price I have to pay if I want the object, which

I value at $t_i + t_j$. So if $t_i + t_j > 2t_j$, or $t_i > t_j$, I want to outbid him, and when $t_i < t_j$, I don't; bidding $2t_i$ accomplishes this.

In the symmetric equilibrium, the bidder with the higher type wins. (In common-value auctions, there is no question of efficiency, since the bidders value the object the same.) A winning bidder pays on average $E(2t_j | 2t_j < 2t_i) = t_i$ for an object worth, on average, $t_i + E(t_j | t_j < t_i) = \frac{3}{2} t_i$. Expected revenue is $E \min\{2t_1, 2t_2\} = \frac{2}{3}$.

Now consider the following alternative strategies:

$$b_1(t_1) = \frac{3}{2} t_1, \quad b_2(t_2) = 3t_2$$

If I'm player 1, I now have to pay $3t_2$ if I win, so I want to win when $t_1 + t_2 > 3t_2$, or $t_1 > 2t_2$, or $\frac{3}{2} t_1 > 3t_2$. So even if I know his bid and type, if player 2 bids $3t_2$, I can't do better than by bidding $\frac{3}{2} t_1$. Similarly, if 1 bids $\frac{3}{2} t_1$, bidder 2 wants to outbid him whenever $t_1 + t_2 > \frac{3}{2} t_1$, or $t_2 > \frac{1}{2} t_1$, or $3t_2 > \frac{3}{2} t_1$, which he accomplishes by bidding $3t_2$.

So these strategies form another ex post equilibrium. But in this one, bidder 1 wins less often than before, and pays more, on average, when he does win. On the other hand, bidder 2 wins more often, and pays less on average.

One more:

$$b_1(t_1) = \frac{101}{100} t_1, \quad b_2(t_2) = 101 t_2$$

The logic is the same; this is indeed an ex post equilibrium. Now bidder 1 wins whenever $t_1 > 100 t_2$, which is hardly ever, and earns hardly any surplus when he does win. On the other hand, bidder 2 almost always wins, and pays on average just barely more than $\frac{1}{2}$. And the expected revenue is very close to $\frac{1}{2}$, lower than before.

These are all examples of a more general result:

Theorem 3. *Suppose $N = 2$ and $v_1 = v_2 = v(t_1, t_2)$. Then for any continuous and strictly increasing function $f : \Re \rightarrow \Re$,*

$$b_1(t_1) = v(t_1, f^{-1}(t_1)), \quad b_2(t_2) = v(f(t_2), t_2)$$

is an ex post equilibrium.

The proof looks just like the logic we've already done. If bidder 2 bids $v(f(t_2), t_2)$, 1 wants to win whenever

$$v(t_1, t_2) > v(f(t_2), t_2)$$

which is whenever $t_1 > f(t_2)$, which is whenever $f^{-1}(t_1) > t_2$, which is exactly the times when

$$v(t_1, f^{-1}(t_1)) > v(f(t_2), t_2)$$

so there you go.

So what does this mean? That there's a massive number of equilibria in ascending (or second-price) common-value auctions.

In the example we did before, $v = t_1 + t_2$, and bidders' bids are bounded below by t_i, so while revenue is lower than in the symmetric auction, it isn't crazy low. This isn't always the case, though. Suppose types are within $[0,1]$, and consider the equilibrium where $f(s) = s^\alpha$, where α is large. Suppose instead of being additive, though, types are multiplicative, that is, $v(t_1, t_2) = t_1 t_2$. Then

$$b_2(t_2) = v(f(t_2), t_2) = t_2^\alpha t_2 = t_2^{1+\alpha}$$

As α gets large, this goes to 0 for all $t_2 < 1$. So for α large, this equilibrium says that bidder 1 always wins, and always pays 0.

Equilibria like this are sometimes referred to as "collusive" equilibria – bidder 2 appears to be letting bidder 1 win. But the strategies still form an equilibrium (and, in fact, and ex post equilibrium).

Almost Common Values

So far, we've shown that with pure common values, there are symmetric equilibria that look "normal," as well as asymmetric equilibria where one bidder may get nearly all of the surplus. Paul Klemperer's paper, "Auctions with Almost Common Values" (1998, European Economic Review 42) makes the case that when one player has even a slight private-value advantage over the other in an otherwise common-values setting, this type of outcome is the only one we should expect in an ascending auction.

He models the wallet game – an independent additive common-values ascending auction like we were just looking at. His story: two people look in their wallets, count the money there, and then hold an ascending bid for a prize equal to the combined money in both their wallets. But player 1, if he wins, also gets an extra dollar.

Klemperer's point is that this small advantage has an explosive effect on the outcome. "Well-behaved" equilibrium play in an ascending auction reveals the loser's type. Continuing to bid in an ascending, common value auction is predicated on the following logic: "given my own information, and given your equilibrium bid strategy, I'm comfortable paying this price for the object, even if you drop out right now." Suppose I'm the advantaged player, and I drop out at some point during the auction; equilibrium play implies that knowing that, my opponent doesn't mind paying the current price for the object. But if that's true, then since the object is worth even more to me, I shouldn't have dropped out.

This means that, if I have even a very small private-value advantage in a common-value ascending auction, I should never lose! Here's one way to put this formally:

Theorem 4. *Suppose that $v_1 = v(t_1, t_2) + \epsilon$ and $v_2 = v(t_1, t_2)$. In any equilibrium of an ascending auction with continuous, undominated strategies, bidder 1 wins with probability 1.*

Let types be between 0 and 1. Using undominated strategies, bidder 1 given type 1 must bid at least $v(1, 0) + \epsilon$, while bidder 2 given type 0 must bid at most $v(1, 0)$, so there is no such equilibrium where 2 always wins.

Since by assumption bid functions are continuous, if both bidders win with positive probability, then there must be some interval of bids within the range of both bidders' bid functions. Pick b some bid in the interior of this region, and define t_1 and t_2 such that

$$b = \beta_1(t_1) = \beta_2(t_2)$$

If $v(t_1, t_2) > b$, then bidder 2 can gain by increasing his bid slightly, since the additional times he wins are profitable. Similarly, if $v(t_1, t_2) < b$, then bidder 2 can gain by decreasing his bid slightly, since he is now winning unprofitably. So $b = v(t_1, t_2)$. But by the same logic, $b = v(t_1, t_2) + \epsilon$. Since this is impossible, there can't be any range where bid distributions overlap, so bidder 1 always wins.

(The caveat of undominated strategies is important, though. Suppose $v = t_1 + t_2$, and t_1 and t_2 are between 0 and 1. Bidder 1 always bidding 0 and bidder 2 always bidding 50 is an equilibrium as well; but both bidders are playing weakly-dominated strategies at every stage.)

Klemperer gives several examples of auctions where common values are a reasonable approximation, but one participant is expected to have a slight edge over the others in terms not of information but of the ex-post value of the object:

- The Los Angeles PCS license in one of the spectrum auctions, where one bidder, Pacific Telephone, was already established as a local brand and had a database of potential local customers

- Leveraged Buyout firms competing to acquire a company, but when one of them already has a small ownership stake, or "toehold", making the company slightly cheaper for it to acquire

- Glaxo bidding against other pharmaceutical companies to acquire Wellcome, where most people felt Glaxo had slightly greater synergies with Wellcome than its competitors

In fact, Pacific Telephone did win the LA PCS license, and at a price many felt was very low. In the Glaxo-Wellcome case, bidders faced bidding costs which were substantial (but dwarfed by the potential gains); Glaxo submitted the first bid, and although two other companies gave serious consideration to bidding higher, neither did, and Glaxo acquired Wellcome at their initial price.

The Milgrom book gives our theorem in a slightly more general version – he requires only that $v_1(t_1, t_2) > v_2(t_1, t_2)$ everywhere, not that they differ by a constant amount ϵ – and he offers the following as a representative ex post equilibrium: $b_1(t_1) = v_1(t_1, 1)$, and $b_2(t_2) = v_2(0, t_2)$. That is, bidder 1 bids his value if bidder 2 is his best possible type, and bidder 2 bids his value if bidder 1 is his worst possible type. Bidder 1 always wins; and the average revenue is $E_{t_2} v_2(0, t_2)$, which could be pretty low.

Klemperer's point, then, is that when one bidder has a private-value advantage over another, sealed-bid (first-price) auctions are preferable. In a first-price auction, the advantage does not get magnified in the same way; and so even the disadvantaged bidders earn substantial profit, and are therefore undeterred by small bidding or entry costs.

(Klemperer is very concerned with getting bidders who are likely to lose to participate anyway, since without enough competitors, auctions can fail spectacularly. We mentioned earlier that in an asymmetric private-values setting, even when value distributions overlap, “strong” bidders prefer ascending auctions, while “weak” bidders prefer first-price auctions. This is sort of an analogous result. This time, values are mostly common, but with a small private component; in an ascending auction, the “strong” bidder always wins, often at a low price, and the “weak” bidder gets 0 expected profits. This is another way to interpret the failed Swiss 3G spectrum auction: once an ascending format was announced, the marginal competitors chose not to participate, leading to an insufficient number of bidders to generate true competition.)

Klemperer points out, however, that in spectrum auctions, there is a gain to holding ascending auctions, since with multiple (possibly complementary) prizes up for auction simultaneously, firms can figure out which ones they may be be able to win before committing too much money to buying complementary ones; he proposes the Anglo-Dutch auction we discussed earlier as a way to capture much of this benefit without falling prey to low revenues due to private-value differences among firms. He also suggests a solution when selling a company.