September 6, 1999.Overheads October 6, 1999.

Eric Rasmusen, Erasmuse@indiana.edu

12 Auctions

12.1 Auction Classification and Private-Value Strategies

Private, Common, and Correlated Values

We will call the dollar value of the utility that player i receives from an object its **value** to him, V_i, and we will call his *estimate* of its value his **valuation**, $\hat{V}_i$.

In a **private-value** auction, each player knows his value with certainty, although he may still have to estimate the values of the other players.

In a **common-value** auction, the players have identical values, but each player forms his own valuation by estimating on the basis of his private information.

Auction Rules and Private-Value Strategies

The types of auctions to be described are:

(1) English (first-price open-cry).
(2) First-price sealed-bid.
(3) Second-price sealed-bid (Vickrey).
(4) Dutch (descending).

(1) English (first-price open-cry)

Rules. Each bidder is free to revise his bid upwards. When no bidder wishes to revise his bid further, the highest bidder wins the object and pays his bid.

Strategies. A player's strategy is his series of bids as a function of (1) his value, (2) his prior estimate of other players' valuations, and (3) the past bids of all the players. His bid can therefore be updated as his information set changes.

Payoffs. The winner's payoff is his value minus his highest bid.The losers' payoffs are zero.

A player's dominant strategy in a private-value English auction is to keep bidding some small amount ϵ more than the previous high bid until he reaches his valuation, and then to stop.

(3) Second-price sealed-bid (Vickrey)

Rules. Each bidder submits one bid, in ignorance of the other bids. The bids are opened, and the highest bidder pays the amount of the second-highest bid and wins the object.

Strategies. A player's strategy is his bid as a function of his value and his prior belief about other players' valuations.

Payoffs. The winner's payoff is his value minus the second-highest bid that was made. The losers' payoffs are zero.

(4) Dutch (descending)

Rules. The seller announces a bid, which he continuously lowers until some buyer stops him and takes the object at that price.

Strategies. A player's strategy is when to stop the bidding as a function of his valuation and his prior beliefs as to other players' valuations.

Payoffs. The winner's payoff is his value minus his bid. The losers' payoffs are zero.

(2) First-price sealed-bid

Rules. Each bidder submits one bid, in ignorance of the other bids. The highest bidder pays his bid and wins the object.

Strategies. A player's strategy is his bid as a function of his value and his prior beliefs about other players' valuations.

Payoffs. The winner's payoff is his value minus his bid.The losers' payoffs are zero.

Suppose Smith's value is 100. If he bid 100 and won when the second bid was 80, he would wish that he had bid only less. If it is common knowledge that the second-highest value is 80, Smith's bid should be $80+\epsilon$. If he is not sure about the second-highest value, the problem is difficult and no general solution has been discovered. The tradeoff is between bidding high— thus winning more often— and bidding low— thus benefiting more if the bid wins.

Suppose that there are N risk-neutral bidders, and that Nature assigns them values independently using a uniform density from 0 to some amount $\bar{v}$. Denote player i's value by v_i, and let us consider the strategy for player 1. If some other player has a higher value, then in a symmetric equilibrium, player 1 is going to lose the auction anyway, so we can ignore that possibility in finding his optimal bid. Player 1's equilibrium strategy is to bid */epsilon* above his expectation of the second-highest value, conditional on his bid being the highest (i.e., assuming that no other bidder has a value over v_1).

If we assume that v_1 is the highest value, the probability that player 2's value, which is uniformly distributed between 0 and v_1, equals v is $1/v_1$, and the probability that v_2 is less than or equal to v is v/v_1. The probability that v_2 equals v and is the second-highest value is

$$Prob(v_2 = v) \cdot Prob(v_3 \leq v) \cdot Prob(v_4 \leq v) \cdots Prob(v_N \leq v), \tag{1}$$

which equals

$$\left(\frac{1}{v_1}\right)\left(\frac{v}{v_1}\right)^{N-2}. \tag{2}$$

Since there are $N-1$ players besides player 1, the probability that one of them has the value v, and v is the second-highest is $N-1$ times expression (12.2). The expectation of v is the integral of v over the range 0 to v_1,

$$\begin{aligned} E(v) &= \int_0^{v_1} v(N-1)(1/v_1)[v/v_1]^{N-2} dv \\ &= (N-1)\tfrac{1}{v_1^{N-1}} \int_0^{v_1} v^{N-1} dv \\ &= \tfrac{(N-1)v_1}{N}. \end{aligned} \tag{3}$$

Thus we find that player 1 ought to bid a fraction $\frac{N-1}{N}$ of his own value, plus ϵ.

The previous example is an elegant result, but it is not a general rule. Suppose Smith knows that Brown's value is 0 or 100 with equal probability, and Smith's value of 400 is known by both players. Brown bids either 0 or 100 in equilibrium, and Smith always bids $(100+\epsilon)$, because his value is so high that winning is more important than paying a low price.

If Smith's value were 102 instead of 400, the equilibrium would be much different. Smith would use a mixed strategy, and while Brown would still offer 0 if his value were 0, if his value were 100 he would use a mixed strategy too. No pure strategy can be part of a Nash equilibrium, because if Smith always bid a value $x < 100$, Brown would always bid $x + \varepsilon$, in which case Smith would deviate to $x + 2\varepsilon$, and if Smith bid $x \geq 100$ he would be paying 100 more than necessary half the time.

Equivalence Theorems

In all four kinds of private independent-value auctions discussed, the seller's expected price is the same. This fact is the biggest result in auction theory: the **revenue equivalence theorem** (Vickrey [1961]).

Hindering Buyer Collusion

As I mentioned at the start of this chapter, one motivation for auctions is to discourage collusion between players. Some auctions are more vulnerable to this than others. Robinson (1985) has pointed out that whether the auction is private-value or common-value, the first-price sealed-bid auction is superior to the second-price sealed-bid or English auctions for deterring collusion among bidders.

Consider a buyer's cartel in which buyer Smith has a private value of 20, the other buyers' values are each 18, and they agree that everybody will bid 5 except Smith, who will bid 6. (We will not consider the rationality of this choice of bids, which might be based on avoiding legal penalties.) In an English auction this is self-enforcing, because if somebody cheats and bids 7, Smith is willing to go all the way up to 20 and the cheater will end up with no gain from his deviation. Enforcement is also easy in a second-price sealed-bid auction, because the cartel agreement can be that Smith bids 20 and everyone else bids 6.

In a first-price sealed-bid auction, however, it is hard to prevent buyers from cheating on their agreement in a one-shot game. Smith does not want to bid 20, because he would have to pay 20, but if he bids anything less than the other players' value of 18 he risks them overbidding him. The buyer will end up paying a price of 18, rather than the 6 he would receive in an English auction with collusion. The seller therefore will use the first-price sealed-bid auction if he fears collusion.

12.3 Risk and Uncertainty over Values

In a private value auction, does it matter what the seller does, given the Revenue Equivalence Theorem? Yes, because of risk aversion, which invalidates the Theorem.

If the seller can reduce bidder uncertainty over the value of the object being auctioned, should he do so?

Suppose there are N bidders, each with a private value, in an ascending open cry auction. Each measures his private value v with an independent error. This error is with equal probability $-x, +x$ or 0. Let us denote the measured value by $\hat{v}$, which is an unbiased estimate of v. What should bidder i bid up to?

If bidder i is risk neutral, he should bid up to $\hat{v}$. His expected utility is, if he pays$\hat{v}$,

$$Ew = 1/3(\hat{v} + x - \hat{v}) + 1/3(\hat{v} - \hat{v}) + 1/3(\hat{v} - x - \hat{v}) = 0. \quad (4)$$

If bidder i is risk averse, however, and wins with bid v_{bid}, his expected utility is

$$EU(w) = 1/3U(\hat{v}+x-v_{bid})+1/3U(\hat{v}-v_{bid})+1/3U(\hat{v}-x-v_{bid}) \quad (5)$$

Note that if the utility function U is concave,

$$1/3U(\hat{v} + x - v_{bid}) + 1/3U(\hat{v} - x - v_{bid}) < 2/3U(\hat{v} - v_{bid}), \quad (6)$$

The implication is that a fair gamble of x has less utility than no gamble. This means that the middle term in equation (12.5) must be positive if it is to be true that $EU(w) = U(0)$, which means that $\hat{v} - v_{bid} > 0$. In other words, bidder i will have a negative expected payoff unless his maximum bid is strictly less than his valuation.

12.4 Common-Value Auctions and the Winner's Curse

To avoid the winner's curse, players should scale down their estimates in forming their bids.

The mental process is a little like deciding how much to bid in a private-value, first-price sealed-bid auction, in which bidder Smith estimates the second-highest value conditional upon himself having the highest value and winning.

In the common-value auction, Smith estimates his own value, not the second-highest, conditional upon himself winning the auction.

He knows that if he wins using his unbiased estimate, he probably bid too high, so after winning with such a bid he would like to retract it. Ideally, he would submit a bid of [X *if I lose, but* $(X-Y)$ *if I win*], where X is his valuation conditional upon losing and $(X-Y)$ is his lower valuation conditional upon winning. If he still won with a bid of $(X-Y)$ he would be happy; if he lost, he would be relieved. But Smith can achieve the same effect by simply submitting the bid $(X-Y)$ in the first place, since the size of losing bids is irrelevant.

Another explanation of the winner's curse can be devised from the Milgrom definition of "bad news" (Milgrom [1981b], Appendix B). Suppose that the government is auctioning off the mineral rights to a plot of land with common value V and that bidder i has valuation $\hat{V}_i$. Suppose also that the bidders are identical in everything but their valuations, which are based on the various information sets Nature has assigned them, and that the equilibrium is symmetric, so the equilibrium bid function $b(\hat{V}_i)$ is the same for each player. If Bidder 1 wins with a bid $b(\hat{V}_1)$ that is based on his prior valuation $\hat{V}_1$, his posterior valuation $\tilde{V}_1$ is

$$\tilde{V}_1 = E(V|\hat{V}_1, b(\hat{V}_2) < b(\hat{V}_1), \ldots, b(\hat{V}_n) < b(\hat{V}_1)). \tag{7}$$

The news that $b(\hat{V}_2) < \infty$ would be neither good nor bad, since it conveys no information, but the information that $b(\hat{V}_2) < b(\hat{V}_1)$ is bad news, since it rules out values of b more likely to be produced by large values of $\hat{V}_2$. In fact, the lower the value of $b(\hat{V}_1)$, the worse is the news of having won. Hence,

$$\tilde{V}_1 < E(V|\hat{V}_1) = \hat{V}_1, \tag{8}$$

and if Bidder 1 had bid $b(\hat{V}_1) = \hat{V}_1$ he would immediately regret having won. If his winning bid were enough below $\hat{V}_1$, however, he would be pleased to win.

Oil Tracts and the Winner's Curse

The best known example of the winner's curse is from bidding for offshore oil tracts. The hundredfold difference in the sizes of the bids in the sealed-bid auctions shown in Table 12.1 lends some plausibility to the view that this is what happened.

Table 12.1 Bids by Serious Competitors in Oil Auctions

Offshore Louisiana 1967 Tract SS 207	**Santa Barbara Channel 1968** Tract 375	**Offshore Texas 1968** Tract 506	**Alaska North Slope 1969** Tract 253
32.5	43.5	43.5	10.5
17.7	32.1	15.5	5.2
11.1	18.1	11.6	2.1
7.1	10.2	8.5	1.4
5.6	6.3	8.1	0.5
4.1		5.6	0.4
3.3		4.7	
		2.8	
		2.6	
		0.7	
		0.7	
		0.4	