Eric Rasmusen, erasmuse@indiana.edu, Nov 9, 2006

13 Auctions

This chapter is a big one. It really would take 3 75-minute sessions– and don't try to use all these slides even for 3 sessions. Pick and choose which ones you want to cover.

I plan to cover, in one 75-minute session, parts of 13.1, 13.2, 13.5. I will pick just a few derivations to do, e.g., the first-price private value auction optimal strategy.

I don't think I'll use overheads–I will write on the board. I will use selected overheads as notes for myself.

Sell a soft drink using the 5 auction rules. Say I will collect the money from the lowest winning bid.

13.1 Values Private and Common, Continuous and Discrete

Private-Value and Common-Value Auctions

Call the dollar value of the utility that bidder i receives from an object its **value** to him, v_i, and we will denote his estimate of the value by $\hat{v}_i$.

Private-value auction: a bidder can learn nothing about his value from knowing the values of the other bidders. (antique chairs– not for resale)

Independent private-value auction: knowing his own value tells him nothing about OTHER bidders' values.

Affiliated private-value auction: he can use knowledge of his own value to deduce something about other players' values.

Pure common-value auction: the bidders have identical values, but each bidder forms his own estimate on the basis of his own private information.

The Ten-Sixteen Auction

Players: One seller and two bidders.

Order of Play:
0. Nature chooses Bidder i's value for the object to be either $v_i = 10$ or $v_i = 16$, with equal probability. (The seller's value is zero.)

The Continuous-Value Auction

Players: One seller and two bidders.

Order of Play:
0. Nature chooses Bidder i's value for the object, v_i, using the strictly positive, atomless density $f(v)$ on the interval $[\underline{v}, \overline{v}]$.

A Mechanism Interpretation

1. The seller chooses a mechanism $[G(\tilde{v}_i, \tilde{v}_{-i})v_i - t(\tilde{v}_i, \tilde{v}_{-i})]$ that takes payments t and gives the object with probability G to player i (including the seller) if he announces that his value is $\tilde{v}_i$ and the other players announce $\tilde{v}_{-i}$. He also chooses the procedure in which bidders select $\tilde{v}_i$ (sequentially, simultaneously, etc.).

Payoffs: The seller's payoff is

$$\pi_s = \sum_{i=1}^{n} t(\tilde{v}_i, \tilde{v}_{-i}) \tag{1}$$

Bidder i's payoff is zero if he does not participate, and otherwise is

$$\pi_i(v_i) = G(\tilde{v}_i, \tilde{v}_{-i})v_i - t(\tilde{v}_i, \tilde{v}_{-i}) \tag{2}$$

The mechanism could allocate the good with 70% probability to the high bidder and with 30% probability to the lowest bidder.

Each bidder could be made to pay the amount he bids, even if he loses.

The payment t could include an entry fee.

There could be a "reserve price," a minimum bid for which the seller will surrender the good.

13.2 Optimal Strategies under Different Rules in Private-Value Auctions

Ascending (English, open-cry, open-exit)

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

Payoffs
The winner's payoff is his value minus his highest bid ($t = p$ for him and $t = 0$ for everyone else). The losers' payoffs are zero.

Types of Ascending Auctions

(1) The bidders offer new prices using pre-specified increments such as thousands of dollars.

(2) The **open-exit** auction.

(3) The **silent-exit** auction.

(4) The **Ebay auction**.

(5) The **Amazon auction**.

The ascending auction can be seen as a mechanism in which each bidder announces his value (which becomes his bid), the object is awarded to whoever announces the highest value (that is, bids highest), and he pays the second-highest announced value (the second-highest bid).

Discussion
A bidder's dominant strategy in a private-value ascending auction is to stay in the bidding until bidding higher would require him to exceed his value and then to stop.

First-Price (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 bidder's strategy is his bid as a function of his value.

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

Strategies in the First-Price Auction

In the first-price auction what the winning bidder wants to do is to have submitted a sealed bid just enough higher than the second-highest bid to win.

If all the bidders' values are common knowledge and he can predict the second- highest bid perfectly, this is a simple problem.

If the values are private information, then he has to guess at the second-highest bid, however, and take a gamble.

His tradeoff is between bidding high–thus winning more often–and bidding low–thus benefiting more if the bid wins.

His optimal strategy depends on his degree of risk aversion and beliefs about the other bidders, so the equilibrium is less robust to mistakes in the assumptions of the model than the equilibria of ascending and second-price auctions.

The First-Price Auction with a Continuous Distribution of Values

Suppose Nature independently assigns values to n risk-neutral bidders using the continuous density $f(v) > 0$ (with cumulative probability $F(v)$) on the support $[0, \bar{v}]$.

A bidder's payoff as a function of his value v and his bid function $p(v)$ is, letting $G(p(v))$ denote the probability of winning with a particular $p(v)$:

$$\pi(v, p(v)) = G(p(v))[v - p(v)]. \tag{3}$$

Thus,

$$p(v) = v - \frac{\pi(v, p(v))}{G(p(v))}. \tag{4}$$

Lemma 1: If a player's equilibrium bid function is differentiable, it is strictly increasing in his value: $p'(v) > 0$.

Lemma 1 implies that the bidder with the greatest v will bid highest and win.

Using the Envelope Theorem

The probability $G(p(v))$ that a bidder with price p_i will win is the probability that v_i is the highest value of all n bidders.

The probability that a bidder's value v is the highest is $F(v)^{n-1}$, the probability that each of the other $(n-1)$ bidders has a value less than v. Thus,

$$G(p(v)) = F(v)^{n-1}. \tag{5}$$

The Envelope Theorem says that if $\pi(v, p(v))$ is the value of a function maximized by choice of $p(v)$ then its total derivative with respect to v equals its partial derivative, because $\frac{\partial \pi}{\partial p} = 0$:

$$\frac{d\pi(v,p(v))}{dv} = \frac{\partial \pi(v,p(v))}{\partial p}\frac{\partial p}{\partial v} + \frac{\partial \pi(v,p(v))}{\partial v} = \frac{\partial \pi(v,p(v))}{\partial v}. \tag{6}$$

Then

$$\frac{d\pi(v,p(v))}{dv} = G(p(v)). \tag{7}$$

Substituting from equation (5) gives us π's derivative, if not π, as a function of v:

$$\frac{d\pi(v,p(v))}{dv} = F(v)^{n-1}. \tag{8}$$

Integrate over all possible values from zero to v and include the base value of $\pi(0)$ (=0) as the constant of integration:

$$\pi(v,p(v)) = \pi(0) + \int_0^v F(x)^{n-1}dx = \int_0^v F(x)^{n-1}dx. \tag{9}$$

The Bid Function

We can now return to the bid function in equation (4) and substitute for $G(p(v))$ and $\pi(v, p(v))$ from equations (5) (9):

$$p(v) = v - \frac{\int_0^v F(x)^{n-1} dx}{F(v)^{n-1}}. \tag{10}$$

Suppose $F(v) = v/\bar{v}$, the uniform distribution. Then (10) becomes

$$\begin{aligned} p(v) &= v - \frac{\int_0^v \left(\frac{x}{\bar{v}}\right)^{n-1} dx}{\left(\frac{v}{\bar{v}}\right)^{n-1}} \\ &= v - \frac{\left. \left(\frac{1}{\bar{v}}\right)^{n-1} \left(\frac{1}{n}\right) x^n \right|_{x=0}^{v}}{\left(\frac{v}{\bar{v}}\right)^{n-1}} \\ &= v - \frac{\left(\frac{1}{\bar{v}}\right)^{n-1} \left(\frac{1}{n}\right) v^n - 0}{\left(\frac{v}{\bar{v}}\right)^{n-1}} \\ &= v - \frac{v}{n} = \left(\frac{n-1}{n}\right) v. \end{aligned} \tag{11}$$

The First-Price Auction: A Mixed-Strategy Equilibrium in the Ten-Sixteen Auction

When the value distribution does not have a continuous support, the equilibrium in a first-price auction may not even be in pure strategies.

Now let each of two bidders' private value v be either 10 or 16 with equal probability and known only to himself.

In a first-price auction, a bidder's optimal strategy is to bid $p(v = 10) = 10$, and if $v = 16$ to use a mixed strategy, mixing over the support $[\underline{p}, \bar{p}]$, where it will turn out that $\underline{p} = 10$ and $\bar{p} = 13$, and the expected payoffs will be:

$$\begin{aligned} \pi(v = 10) &= 0 \\ \pi(v = 16) &= 3 \\ \pi_s &= 11.5. \end{aligned} \tag{12}$$

These are the same payoffs as in the ascending auction, an equivalence we will come back to in a later section.

The Equilibrium

$p(v = 10) = 10$. If either bidder used the bid $p < 10$, the other player would deviate to $(p + \epsilon)$, and a bid above 10 exceeds the object's value.

The bid $p(v = 16)$ will be between 10 (so the bidder can win if his rival's value is 10) and 16 (which would always win, but unprofitably).

The pure strategy of $(p = 10)|(v = 16)$ will win with probability of at least 0.50, yielding payoff $0.50(16 - 10) = 3$. This rules out bids in $(13, 16]$, because their payoff is less than 3.

The upper bound $\bar{p}$ must be exactly 13. If it were any less, then the other player would respond by using the pure strategy of $(\bar{p} + \epsilon)$, which would win with probability one and yield a payoff of greater than the payoff of 3 $(= 0.5(16 - 10))$ from $p = 10$.

When a player mixes over a continuum, the modeller must be careful to check for
(a) atoms (some particular point which has positive probability, not just positive density), and
(b) gaps (intervals within the mixing range with zero probability of bids). Are there any atoms or gaps within the interval [10,13]?

The Mixing Density

The mixing density $m(p)$ is positive over the entire interval $[10, 13]$, with no atoms. Since if our player has value $v = 16$ there is probability 0.5 of winning because the other player has $v = 10$ and probability $0.5M(p)$ of winning because the other player has $v = 16$ too but bid less than p, the payoff is

$$0.5(16 - p) + 0.5M(p)(16 - p) = 3. \tag{13}$$

This implies that $(16 - p) + M(p)(16 - p) = 6$, so

$$M(p) = \frac{6}{16 - p} - 1, \tag{14}$$

which has the density

$$m(p) = \frac{6}{(16 - p)^2} \tag{15}$$

on the support $[10, 13]$, rising from $m(10) = \frac{1}{6}$ to $m(13) = \frac{4}{6}$.

Second-Price Auctions (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 bidder's strategy is his bid as a function of his value.

Payoffs: The winning bidder's payoff is his value minus the second-highest bid. The losing bidders' payoffs are zero. The seller's payoff is the second-highest-bid.

Asymmetric Equilibria

Consider a variant of the Ten-Sixteen Auction, in which each of two bidders' values can be 10 or 16, but where the realized values are common knowledge.

Bidding one's value is a **symmetric equilibrium**, meaning that the bid function $p(v)$ is the same for both bidders: $\{p(v = 10) = 10, p(v = 16) = 16\}$.

But consider the following equilibrium:

$$\begin{array}{ll} p_1(v = 10) = 10 & p_1(v = 16) = 16 \\ p_2(v = 10) = 1 & p_2(v = 16) = 10 \end{array} \tag{16}$$

Descending Auctions (Dutch)

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

Strategies
A bidder's strategy is when to stop the bidding as a function of his value.

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

The descending auction is **strategically equivalent** to the first-price auction.

All-Pay Auctions

Rules: Each bidder places a bid simultaneously. The bidder with the highest bid wins, and each bidder pays the amount he bid.

Strategies: A bidder's strategy is his bid as a function of his value.

Payoffs: The winner's payoff is his value minus his bid. The losers' payoffs are the negative of their bids.

Discussion: The winning bid will be lower in the all-pay auction than under the other rules, because bidders need a bigger payoff when they do win to make up for their negative payoffs when they lose. At the same time, since even the losing bidders pay something to the seller it is not obvious that the seller does badly (and in fact, it turns out to be just as good an auction rule as the others, in this simple risk-neutral context).

The Equal-Value All-Pay Auction

Suppose each of the n bidders has the same value, v.

Under the all-pay auction rule, this game is quite interesting.

The equilibrium is in mixed strategies.

Either the maximum bid is less than v, in which case someone could deviate to $p = v$ and increase his payoff;

or one bidder bids v and the rest bid at most $p' < v$, in which case the high bidder will deviate to bid just above p'.

A Symmetric Equilibrium

Suppose we have a symmetric equilibrium, so all bidders use the same mixing cumulative distribution $M(p)$. Let us conjecture that $\pi(p) = 0$, which we will later verify. The payoff function for each bidder is the probability of winning times the value of the prize minus the bid, which is paid with probability one, and if we equate that to zero we get

$$M(p)^{n-1}v = p, \tag{17}$$

so

$$M(p) = \sqrt[n-1]{\frac{p}{v}}. \tag{18}$$

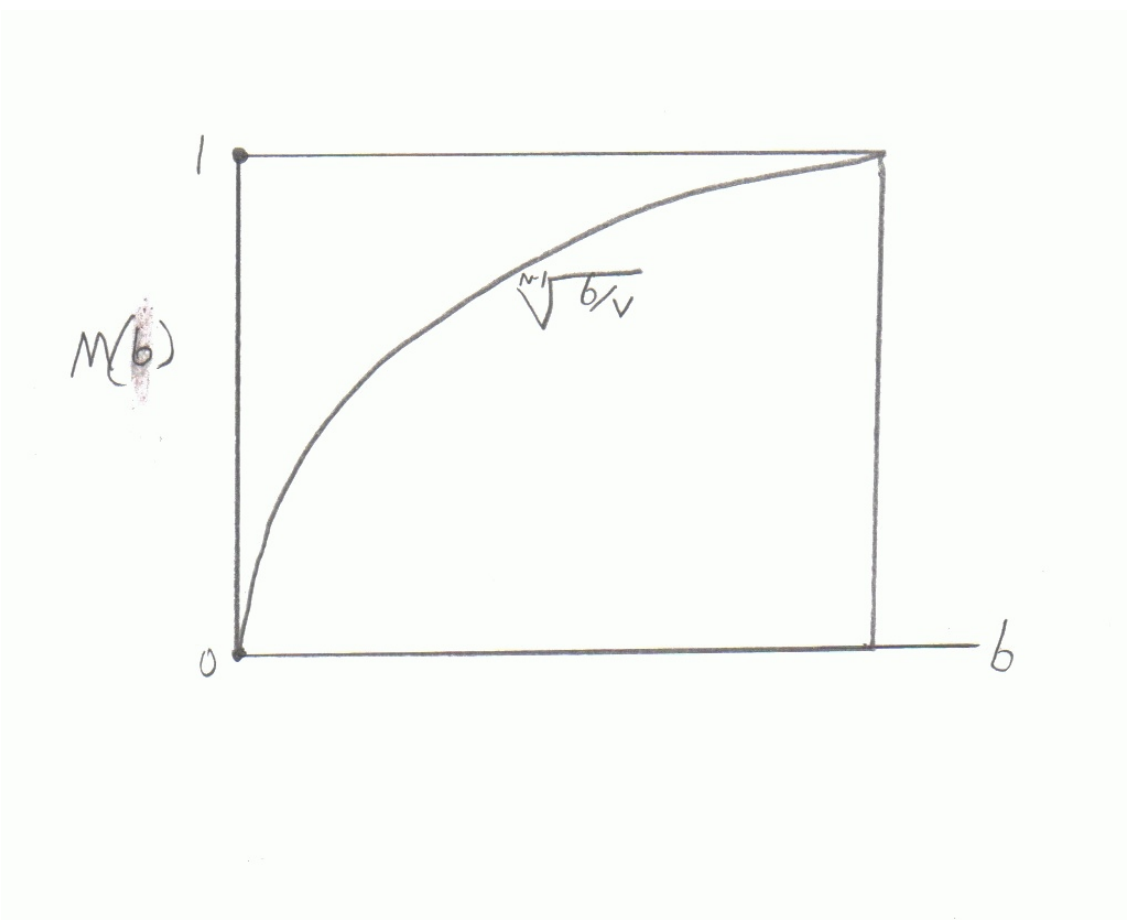

The Continuous-Value All-Pay Auction

Suppose each of the n bidders picks his value v from the same density $f(v)$. Conjecture that the equilibrium is symmetric, in pure strategies, and that the bid function, $p(v)$, is strictly increasing. The equilibrium payoff function for a bidder with value v who pretends he has value z is

$$\pi(v, z) = F(z)^{n-1}v - p(z), \tag{19}$$

since if our bidder bids $p(z)$, that is the highest bid only if all $(n-1)$ other bidders have $v < z$, a probability of $F(z)$ for each of them.

Finding the Equilibrium

The function $\pi(v, z)$ is not necessarily concave in z, so satisfaction of the first-order condition will not be a sufficient condition for payoff maximization, but it is a necessary condition since the optimal z is not 0 (unless $v = 0$) or infinity and from (19) $\pi(v, z)$ is differentiable in z in our conjectured equilibrium. Thus, we need to find z such that

$$\frac{\partial \pi(v, z)}{\partial z} = (n-1)F(z)^{n-2}f(z)v - p'(z) = 0 \qquad (20)$$

In the equilibrium, our bidder does follow the strategy $p(v)$, so $z = v$ and we can write

$$p'(v) = (n-1)F(v)^{n-2}f(v)v \qquad (21)$$

Integrating up, we get

$$p(v) = p(0) + \int_0^v (n-1)F(x)^{n-2}f(x)x dx \qquad (22)$$

This is deterministic, symmetric, and strictly increasing in v, so we have verified our conjectures.

The Outcome

Suppose values are uniformly distributed over [0,1], so $F(v) = v$. Then equation (22) becomes

$$\begin{aligned} p(v) &= p(0) + \int_0^v (n-1)x^{n-2}(1)x dx \\ &= p(0) + \Big|_{x=0}^{v} (n-1)\frac{x^n}{n} \qquad (23) \\ &= 0 + \left(\frac{n-1}{n}\right) v^n, \end{aligned}$$

where we can tell that $p(0) = 0$ because if $p(0) > 0$ a bidder with $v = 0$ would have a negative expected payoff.

If there were $n = 2$ bidders, a bidder with value v would bid $v^2/2$, win with probability v, and have expected payoff $\pi = v(v) - v^2/2 = v^2/2$. If there were $n = 10$ bidders, a bidder with value v would bid $(9/10)v^{10}$, win with probability v^9, and have expected payoff $\pi = v(v^9) - (9/10)v^{10} = v^{10}/(10)$.

The Dollar Auction

Consider an ascending auction to sell a dollar bill in which the players offer higher and higher bids, and the highest bidder wins– but both the first- and second-highest bidders pay their bids.

What happens?

THE REVENUE EQUIVALENCE THEOREM.

Let all players be risk-neutral with private values drawn independently from the same atomless, strictly increasing distribution $F(v)$ on $[\underline{v}, \bar{v}]$. If under either Auction Rule A_1 or Auction Rule A_2 it is true that:

(a) The winner of the object is the player with the highest value; and

(b) The lowest bidder type, $v = \underline{v}$, has an expected payment of zero;

then the symmetric equilibria of the two auction rules have the same expected payoffs for each type of bidder and for the seller.

Proving the Theorem

THE REVENUE EQUIVALENCE THEOREM. Let all players be risk-neutral with private values drawn independently from the same atomless, strictly increasing distribution $F(v)$ on $[\underline{v}, \bar{v}]$. If under either Auction Rule A_1 or Auction Rule A_2 it is true that:

(a) The winner of the object is the player with the highest value; and

(b) The lowest bidder type, $v = \underline{v}$, has an expected payment of zero;

then the symmetric equilibria of the two auction rules have the same expected payoffs for each type of bidder and for the seller.

Proof. Let us represent the auction as the truthful equilibrium of a direct mechanism in which each bidder sends a message z of his type v and then pays an expected amount $p(z)$. (The Revelation Principle says that we can do this.) By assumption (a), the probability that a player wins the object given that he chooses message z equals $F(z)^{n-1}$, the probability that all $(n-1)$ other players have values $v < z$. Let us denote this winning probability by $G(z)$, with density $g(z)$. Note that $g(z)$ is well defined because we assumed that $F(v)$ is atomless and everywhere increasing.

A Second Step in the Proof

The expected payoff of any player of type v is the same, since we are restricting ourselves to symmetric equilibria. It equals

$$\pi(z,v) = G(z)v - p(z). \tag{24}$$

The first-order condition with respect to the player's choice of type message z (which we can use because neither $z = 0$ nor $z = \bar{v}$ is the optimum if condition (a) is to be true) is

$$\frac{d\pi(z;v)}{dz} = g(z)v - \frac{dp(z)}{dz} = 0, \tag{25}$$

so

$$\frac{dp(z)}{dz} = g(z)v. \tag{26}$$

We are looking at a truthful equilibrium, so we can replace z with v:

$$\frac{dp(v)}{dv} = g(v)v. \tag{27}$$

Finishing the Proof

Next, we integrate (27) over all values from zero to v, adding $p(\underline{v})$ as the constant of integration:

$$p(v) = p(\underline{v}) + \int_{\underline{v}}^{v} g(x)x dx. \tag{28}$$

We can use (28 to substitute for $p(v)$ in the payoff equation (24), which becomes, after replacing z with v and setting $p(\underline{v}) = 0$ because of assumption (b),

$$\pi(v, v) = G(v)v - \int_{\underline{v}}^{v} g(x)x dx. \tag{29}$$

Equation (29) says the expected payoff of a bidder of type v depends only on the $G(v)$ distribution, which in turn depends only on the $F(v)$ distribution, and not on the $p(z)$ function or other details of the particular auction rule. But if the bidders' payoffs do not depend on the auction rule, neither does the seller's. Q.E.D.

A REVENUE EQUIVALENCE COROLLARY

Let all players be risk-neutral with private values drawn from the same strictly increasing, atomless distribution $F(v)$. The symmetric equilibria of the ascending, first-price, second-price, descending, and all-pay auctions all have the same expected payoffs for each type of bidder and for the seller.

Risk Aversion in Private-Value Auctions

When bidders are risk averse, the Revenue Equivalence Theorem fails.

Consider Bidder 1 in The Ten-Sixteen Auction when he knows his own value is $v_1 = 16$ but does not know v_2. In the second-price auction, he has an equal chance of a payoff of either 0 (if $v_2 = 16$) or 6 (if $v_2 = 10$), regardless of whether the bidders are risk averse or not, because bidding one's value is a weakly dominant strategy.

Compare that with his payoff in the first-price auction, in which the equilibrium is in mixed strategies. If the bidders are risk neutral, then as we found earlier, if the bidder has value 16 he wins using a bid in the mixing support [10,13] and achieves a payoff in [3,6] with probability 0.75, and he loses and earns payoff of zero with probability 0.25. The (0,6) gamble of the second-price auction is riskier than the (0, 3 to 6) gamble of the first-price auction. The (0,6) gamble is simpler, but it has more dispersion.

Risk Aversion in the First Price Auction

If the bidders are risk averse, then the optimal strategies in the first-price auction change. It remains true that the bidders mix on an interval $[10, \bar{p}]$. We derived $\bar{p}$ and the optimal mixing distribution by equating expected payoffs, however, and a certain win at a price of 10 will now be worth more to a bidder than a 50% chance of winning at a price of 13. Let us denote the concave utility function of each bidder by $U(v - p)$ and normalize by defining $U(0) \equiv 0$. The expected payoff from $p = 10$, which wins with probability 0.5, must equal the expected payoff from the upper bound $\bar{p}$ of the mixing support, so

$$0.5U(6) = U(16 - \bar{p}). \tag{30}$$

Since $0.5U(6) < U(16 - 13)$ by concavity of U, it must be that $\bar{p} > 13$.

The Mixing Density

We found the mixing distribution function $M(p)$ by equating $\pi(p)$ to the payoff from bidding 10, which is $0.5U(6)$, so

$$\pi(p) = 0.5U(16-p) + 0.5M(p)U(16-p) = 0.5U(6), \quad (31)$$

which can be solved to yield

$$M(p) = \frac{U(6)}{U(16-p)} - 1, \quad (32)$$

which has the density

$$m(p) = \frac{U(6)'}{U}(16-p)U(16-p)^2, \quad (33)$$

compared with the risk-neutral density $m(p) = \frac{6}{(16-p)^2}$ from equation (15). Thus, risk aversion of the bidders actually spreads out their equilibrium bids (the support is broader than [10,13]), but it remains true that the first-price auction is less risky than the second-price auction.

Risk Aversion in the Continuous-Value Auction

What happens in the Continuous-Value Auction? In the second-price auction ,the optimal strategies are unchanged, so seller revenue does not change if bidders are risk averse.

To solve for the equilibrium of the first-price auction, look at a given bidder's incentive to report his true type v as z in an auction in which the payment is $p(z)$ and the probability of winning the object is $G(z)$.

The bidder maximizes by choice of z

$$\begin{aligned} \pi(v,z) &= G(z)U[v - p(z)] \\ &= F(z)^{n-1}U[v - p(z)], \end{aligned} \tag{34}$$

where $\pi(v,0) = 0$ because $F(0) = 0$.

At the optimum,

$$\begin{aligned} \frac{\partial \pi(v,z)}{\partial z} &= (n-1)F(z)^{n-2}f(z)U[v - p(z)] \\ &+F(z)^{n-1}U'[v - p(z)][-p'(z)] = 0, \end{aligned} \tag{35}$$

The Ultimate Effect of Risk Aversion

In equilibrium, $z = v$. Using that fact, for all $v > \underline{v}$ (since $F(\underline{v}) = 0$) we can solve equation (35) for $p(z)$ to get

$$p(v) = \left(\frac{(n-1)f(v)}{F(v)} \right) \left(\frac{U[v - p(v)]}{U'[v - p(v)]} \right) \tag{36}$$

Now let's look at the effect of risk aversion on $p(v)$. If U is linear, then

$$\frac{U[v - p(v)]}{U'[v - p(v)]} = v - p(v), \tag{37}$$

but if the bidder is risk averse, so U is strictly concave,

$$\frac{U[v - p(v)]}{U'[v - p(v)]} > v - p(v). \tag{38}$$

Thus, for a given v, the bid function in (36) makes the bid higher if the bidder is risk averse than if he is not. The bid for every value of v except $v = \underline{v}$ increases ($p(\underline{v}) = \underline{v}$, regardless of risk aversion).

By increasing his bid from the level optimal for a risk-neutral bidder, the risk- averse bidder insures himself. If he wins, his surplus is slightly less because of the higher price, but he is more likely to win and avoid a surplus of zero.

Seller Revenue and Risk Aversion

As a result the seller's revenue is greater in the first-price than in the second-price auction if bidders are risk averse.

But since under risk neutrality the first-price and second-price auctions yield the same revenue, under risk aversion the first-price auction must yield greater revenue, both in expectation and conditional on the highest v present in the auction.

The seller, whether risk neutral or risk averse, will prefer the first-price auction when bidders are risk averse.

Uncertainty over One's Own Value

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 auction. Each measures his private value v with an independent error $\epsilon > 0$. This error is with equal probability $-x, +x$ or 0.

The bidders have diffuse priors, so they take all values of v to be equally likely, ex ante.

Let us denote a bidder's measured value by $\hat{v} = v + \epsilon$, which is an unbiased estimate of v.

In the ascending auctions we have been studying so far, where $\epsilon = 0$, the optimal bid ceiling was v.

Now, when $\epsilon > 0$, what bid ceiling should be used by a bidder with utility function $U(v - p)$?

Optimal Strategies for Uncertain Bidders

If the bidder wins the auction and pays p for the object, his expected utility at that point is

$$\pi(p) = \frac{U([\hat{v}-x]-p)}{3} + \frac{U(\hat{v}-p)}{3} + \frac{U([\hat{v}+x]-p)}{3} \quad (39)$$

If he is risk neutral, this yields him a payoff of zero if $p = \hat{v}$, and winning at any lower price would yield a positive payoff. Under risk neutrality uncertainty over one's own value does not affect the optimal strategy.

If the bidder is risk averse, however, then the utility function U is concave and

$$\frac{U([\hat{v}-x]-p)}{3} + \frac{U([\hat{v}+x]-p)}{3} < \left(\frac{2}{3}\right) U(\hat{v}-p), \quad (40)$$

so his expected payoff in equation (39) is less than $U(\hat{v}-p)$, and if $p = \hat{v}$ his payoff is less than $U(0)$.

13.4 Reserve Prices and the Marginal Revenue Approach

A **reserve price** p^* is a bid put in by the seller, secretly or openly, before the auction begins, which commits him not to sell the object if nobody bids more than p^*.

The seller will often find that a reserve price can increase his payoff. If he does, it turns out that he will choose a reserve price strictly greater than his own value: $p^* > v_s$.

To see this, we will use the **marginal revenue approach** to auctions [Bulow & Roberts (1989)]

This approach compares the seller in an auction to an ordinary monopolist who sells using a posted price.

An Auction with One Bidder and a Reserve Price

The seller will do badly in any of the auction rules we have discussed so far.

What should the seller's offer p^* be?

Let the bidder have value distribution $F(v)$ on $[\underline{v}, \bar{v}]$ which is differentiable and strictly increasing, so the density $f(v)$ is always positive. Let the seller value the object at $v_s \geq \underline{v}$. The seller's payoff is

$$\begin{aligned} \pi(p^*) &= Pr(p^* < v)(p^* - v_s) + Pr(p^* > v)(0) \\ &= [1 - F(p^*)](p^* - v_s). \end{aligned} \tag{41}$$

This has first-order-condition

$$\frac{d\pi(p^*)}{dp^*} = [1 - F(p^*)] - f(p^*)[p^* - v_s] = 0. \tag{42}$$

On solving (42) for for p^* we get

$$p^* = v_s + \left(\frac{1 - F(p^*)}{f(p^*)} \right). \tag{43}$$

The reserve price is strictly greater than the seller's value for the object ($p^* > v_s$) unless the solution is such that $F(p^*) = 1$ because the optimal reserve price is the greatest possible bidder value, in which case the object has probability zero of being sold.

2. Multiple Bidders.

Now let there be n bidders, all with values distributed independently by $F(v)$. Denote the bidders with the highest and second-highest values as Bidders 1 and 2.

The seller's payoff in a second-price auction is

$$\begin{aligned} \pi(p^*) &= Pr(p^* > v_1)(0) + Pr(v_2 < p^* < v_1)(p^* - v_s) \\ &+ Pr(p^* < Ev_2 < v_1)(v_2 - v_s) \\ \\ &= \int_{v_1=\underline{v}}^{p^*} f(v_1)(0)dv_1 + \int_{v_1=p^*}^{\bar{v}} \left(\int_{v_2=\underline{v}}^{p^*} (p^* - v_s) f(v_2) dv_2 \right. \\ &+ \left. \int_{v_2=p^*}^{v_1} (v_2 - v_s) f(v_2) dv_2 \right) f(v_1) dv_1 \end{aligned} \tag{44}$$

This expression integrates over two random variables. First, it matters whether v_1 is greater than or less than p^*, the outer integrals. Second, it matters whether v_2 is less than p^* or not, the inner integrals.

It turns out that

$$p^* = v_s + \frac{1 - F(p^*)}{f(p^*)}, \tag{45}$$

just what we found in equation (43) for the one-bidder case. Remarkably, the optimal reserve price is unchanged!

3. A Continuum of Bidders: The Marginal Revenue Interpretation

Now think of a firm with a constant marginal cost of c facing a continuum of bidders along the same distribution $F(v)$ that we have been using. The quantity of bidders with values above p will be $(1 - F(p))$, so the demand equation is

$$q(p) = 1 - F(p), \quad Revenue \equiv pq = p(1 - F(p)) \tag{46}$$

The marginal revenue is then (keeping in mind that $\frac{dq}{dp} = -f(p)$)

$$\begin{aligned} Marginal\ Revenue \equiv \frac{dR}{dq} &= p + \left(\frac{dp}{dq}\right) q \\ &= p + \left(\frac{1}{\frac{dq}{dp}}\right) q \\ &= p - \frac{1 - F(p)}{f(p)} \end{aligned} \tag{47}$$

Profit-maximizing monopoly price is the one at which the marginal revenue equals c.

Figure 3a: Auctions and Marginal Revenue: Reserve Price Needed or Not Needed

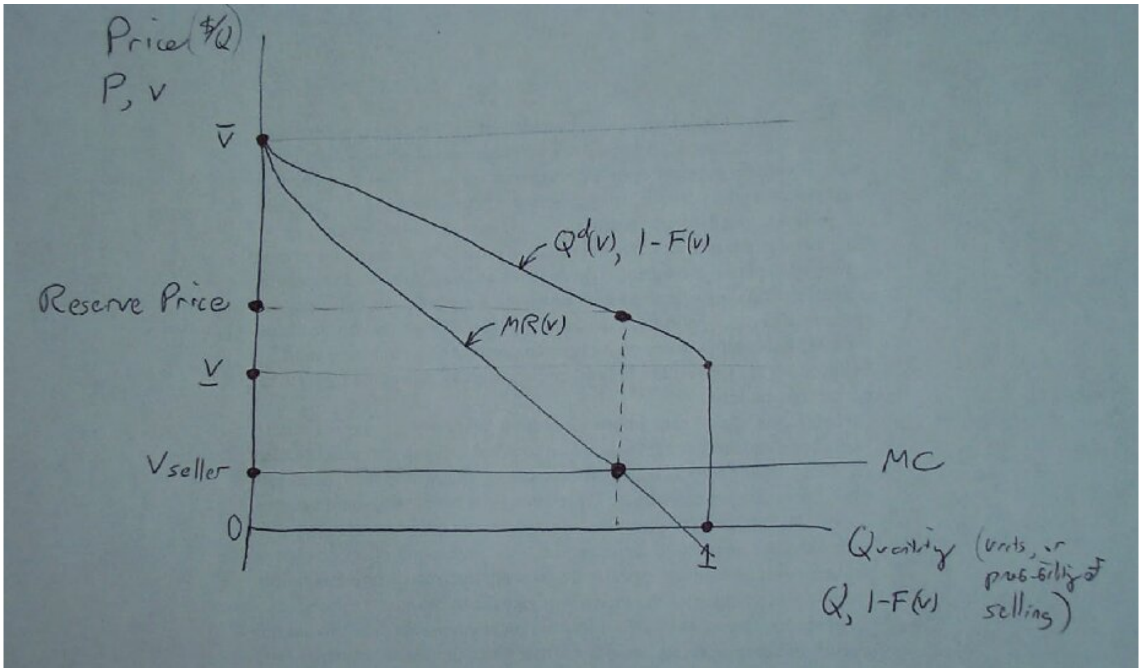

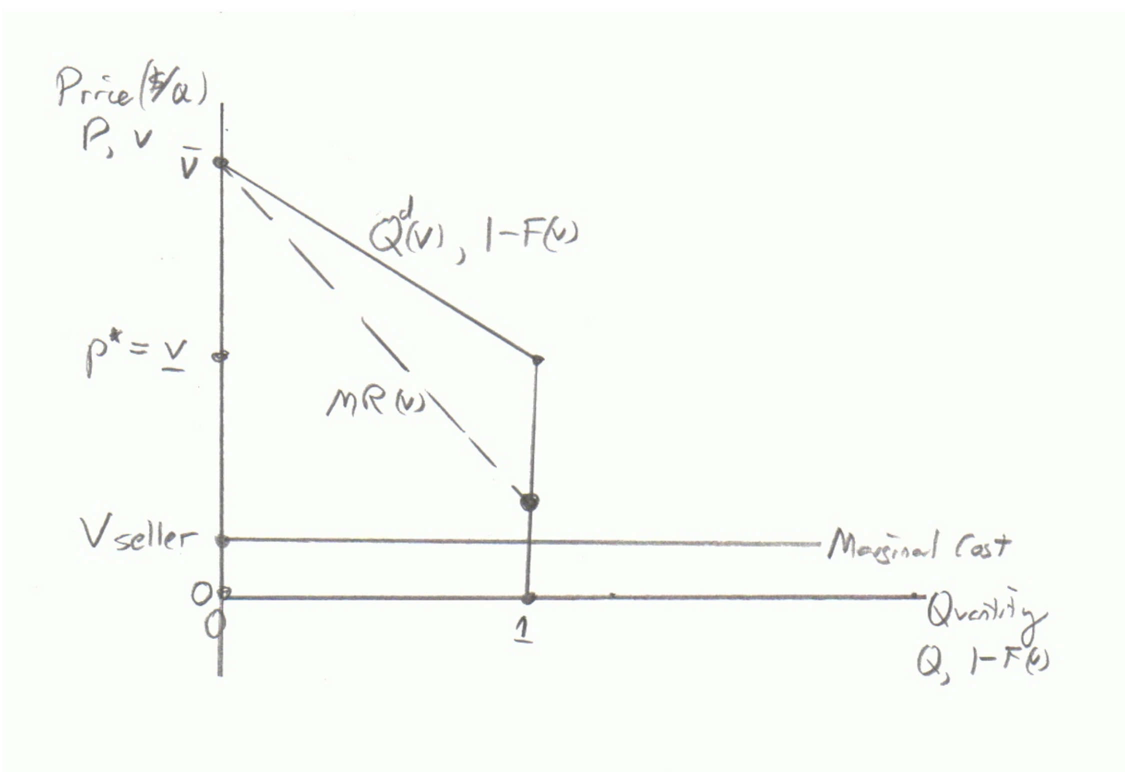

Monopoly and Auctions

If output is reduced below the competitive level, the outcome is inefficient, as in conventional monopoly. This happens if no sale takes place of the one unit even though $v > v_s$ for some bidder.

Unlike a conventional monopoly, there is a possibility of inefficient "overproduction" in an auction.

That happens if the sale takes place even though no bidder values the good as much as the seller: $v < v_s$ for the winning bidder.

A positive reserve price, therefore, can help efficiency rather than hurt it.

All the five auction forms — first-price, second-price, descending, ascending, and all-pay— can be efficient in a private-value setting, but only if the reserve price is set not at the profit-maximizing level but at $p^* = v_s$.

Hindering Bidder Collusion

Robinson (1985) has pointed out that whether the auction is private-value or common-value, the first-price auction is superior to the second-price or ascending auctions for deterring collusion among bidders.

Consider a bidder's cartel in which bidder Smith has a private value of 20, the other bidders' values are each 18, and they agree that everybody will bid 5 except Smith, who will bid 6.

In an ascending 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.

In a first-price auction the bidders have a strong temptation to cheat. The bid p' that the colluders would choose for Smith would be lower than $p' = 20$, since he would have to pay his bid, but if p' is anything less than the other bidders' value of 18 any one of them could gain by deviating to bid more than p' and win.

Common Value Auctions and the Winner's Curse

In a pure common-value auction, all players have the same value, but they estimate it with different errors.

What happens if everybody bids his best estimate of the value?

Thus: shade your bid till you expect zero profit even if you overestimated the most.

If Smith is more risk averse than Brown, then Smith should be more cautious for two reasons.

The gamble is worth less to Smith– the reason analyzed above in the private-value setting.

Also, when Smith wins against a rival like Brown who regularly bids more, Smith probably overestimated the value.

If there is a private value component, and it is bigger for Brown than for Smith, Smith should also be extra-cautious.

Table 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	

Strategies in Common-Value Auctions

Milgrom & Weber (1982) found that when there is a common-value element in an auction and signals are "affiliated" then revenue equivalence fails.

The first-price and descending auctions are still identical, but they raise less revenue than the ascending or second-price auctions.

If there are more than two bidders, the ascending auction raises more revenue than the second-price auction.

If signals are affiliated then even in a private value auction, in which each bidder knows his own value with certainty, the first-price and descending auctions will do worse.

Signals from a Uniform Distribution

Suppose n signals are independently drawn from the uniform distribution on $[\underline{s}, \overline{s}]$.

Denote the j^{th} highest signal by $s_{(j)}$.

The expectation of the kth highest value is

$$Es_{(k)} = \underline{s} + \left(\frac{n+1-k}{n+1} \right) (\overline{s} - \underline{s}) \tag{48}$$

Let n risk-neutral bidders, $i = 1, 2, ...n$ each receive a signal s_i independently drawn from the uniform distribution on $[v - m, v + m]$, where v is the true value of the object to each of them.

Assume that they have "diffuse priors" on v, which means they think any value from $v = -\infty$ to $v = \infty$ is equally likely and we do not need to make use of Bayes's rule.

Figure 4: Extracting Information From Uniformly Distributed Signals

The best estimate of the value given the set of n signals is

$$Ev|(s_1, s_2, ..., s_n) = \frac{s_{(n)} + s_{(1)}}{2}. \tag{49}$$

The estimate depends only on two out of the n signals— a remarkable property of the uniform distribution.

If there were five signals $\{6, 7, 7, 16, 24\}$, the expected value of the object would be 15 (=[6+24]/2), well above the mean of 12 and the median of 7, because only the extremes of 6 and 24 are useful information.

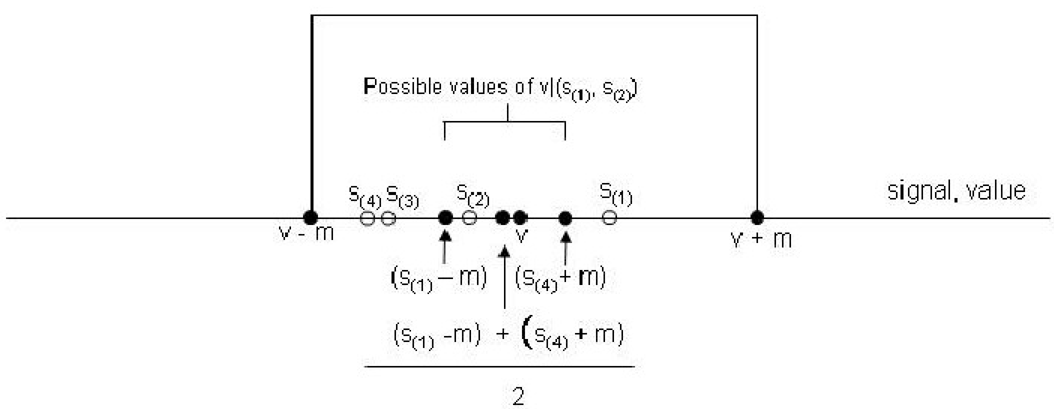

The Uniform-Signal Common-Value Auction

Order of Play: 0. Nature chooses the common value for the object v using the uniform density on $[-\infty, \infty]$ (the limit of $[-x, x]$ as x goes to infinity), and sends signal s_i to Bidder i using the uniform distribution on $[v - m, v + m]$.
1. The seller chooses a mechanism that allocated the object and payments based on each player's choice of p. He also chooses the procedure in which bidders select p (sequentially, simultaneously, etc.).
2. Each bidder simultaneously chooses to participate in the auction or to stay out.
3. The bidders and the seller choose value of p according to the mechanism procedure.
4. The object is allocated and transfers are paid according to the mechanism.

The Ascending Auction (open-exit)

Equilibrium: If no bidder has quit yet, Bidder i should drop out when the price rises to s_i. Otherwise, he should drop out when the price rises to $p_i = \frac{p_{(n)}+s_i}{2}$, where $p_{(n)}$ is the price at which the first dropout occurred.

Explanation: If no other bidder has quit yet, Bidder i is safe in agreeing to pay his signal, s_i. Either (a) he has the lowest signal, or (b) everybody else has the same signal value s_i too, and they will all drop out at the same time. In case (a), having the lowest signal, he will lose anyway. In case (b), the best estimate of the value is s_i, and that is where he should drop out.

Once one bidder has dropped out at $p_{(n)}$, the other bidders can deduce that he had the lowest signal, so they know that signal $s_{(n)}$ must equal $p_{(n)}$. Suppose Bidder i has signal $s_i > s_{(n)}$. Either (a) someone else has a higher signal and Bidder i will lose the auction anyway and dropping out too early does not matter, or (b) everybody else who has not yet dropped out has signal s_i too, and they will all drop out at the same time, or (c) he would be the last to drop out, so he will win. In cases (b) and (c), his estimate of the value is $p_{(i)} = \frac{p_{(n)}+s_i}{2}$, since $p_{(n)}$ and s_i are the extreme signal values and the signals are uniformly distributed, and that is where he should drop out.

Seller Revenue

The price paid by the winner will be the price at which the second-highest bidder drops out, which is $\frac{s_{(n)}+s_{(2)}}{2}$.

$$
\begin{aligned}
Ep_{(2)} &= \frac{\left[v+\left(\frac{1-n}{n+1}\right)m\right]+\left[v+\left(\frac{n-3}{n+1}\right)m\right]}{2} \\
&= v-\left(\frac{1}{2}\right)\left(\frac{1}{n+1}\right)2m.
\end{aligned} \tag{50}
$$

If $m = 50$ and $n = 4$, then

$$
Ep_{(2)} = v - \left(\frac{1}{10}\right)(100) = v - 10. \tag{51}
$$

Expected seller revenue increases in n, the number of bidders (and thus of independent signals) and falls in the uncertainty m (the inaccuracy of the signals).

That this is an open-exit auction is crucial. Other bidders need to learn the lowest signal.

Common Values: The Second-Price Auction

Equilibrium: Bid $p_i = s_i - \left(\frac{n-2}{n}\right) m$.

Explanation: Bidder i should think of himself as being tied for winner with one other bidder, and so having to pay exactly his bid. Thus, he imagines himself as the highest of $(n-1)$ bidders drawn from $[v-m, v+m]$ and tied with one other. Then,

$$p_i = s_i - \left(\frac{n-2}{n}\right)(m)$$

On average, the second-highest bidder actually has the signal $Es_{(2)} = v + \left(\frac{n-3}{n+1}\right) m$, so

$$\begin{aligned} Ep_{(2)} &= [v + \left(\tfrac{n-3}{n+1}\right) m] - \left(\tfrac{n-2}{n}\right)(m) \\ &= v - \left(\tfrac{n-1}{n}\right)\left(\tfrac{1}{n+1}\right) 2m. \end{aligned} \tag{52}$$

If $m = 50$ and $n = 4$, then

$$Ep_{(2)} = v - \left(\frac{3}{4}\right)\left(\frac{1}{5}\right)(100) = v - 15. \tag{53}$$

If there are at least three bidders, expected revenue is lower in the second- price auction. (We found revenue of $(v-10)$ with $n = 4$ in the ascending auction.)

If $n = 2$, however, the expected price is the same. $v_{(n)} = v_{(2)}$, so the winning price is based on the same information in both auctions.

Common Values: The First-Price Auction

Equilibrium: Bid $(s_i - m)$.

Explanation: Bidder i bids $(s_i - z)$ for some amount z that does not depend on his signal, because given the assumption of diffuse priors, he does not know whether his signal is a high one or a low one.

Define T_i to be how far the signal s_i is above its minimum possible value, $(v - m)$, so

$$T_i \equiv s_i - (v - m) \tag{54}$$

and $s_i \equiv v - m + T_i$. Bidder i has the highest signal and wins the auction if T_i is big enough, which has probability $\left(\frac{T_i}{2m}\right)^{n-1}$, which we will define as $G(T_i)$, because it is the probability that the $(n-1)$ other signals are all less than $s_i = v - m + T_i$. He earns v minus his bid of $(s_i - z)$ if he wins, which equals $(z + m - T_i)$.

The Epsilon Argument

If Bidder i deviated and bid a small amount ϵ higher, he would win with a higher probability, $G(T_i + \epsilon)$, but he would lose ϵ whenever he would have won with the lower bid. Using a Taylor expansion, $G(T_i + \epsilon) \approx G(T_i) + G'(T_i)\epsilon$, so

$$G(T_i + \epsilon) - G(T_i) \approx (n-1)T_i^{n-2} \left(\frac{1}{2m} \right)^{n-1} \epsilon. \tag{55}$$

The benefit from bidding higher is the higher probability, $[G(T_i + \epsilon) - G(T_i)]$ times the winning surplus $(z + m - T_i)$. The loss from bidding higher is that the bidder would pay an additional ϵ in the $\left(\frac{T_i}{2m}\right)^{n-1}$ cases in which he would have won anyway.

A Bidder's Optimal Strategy

In equilibrium, he is indifferent about this infinitesimal deviation, taking the expectation across all possible values of his "signal height" T_i, so

$$\int_{T_i=0}^{2m} \left[\left((n-1) T_i^{n-2} \left(\frac{1}{2m} \right)^{n-1} \epsilon \right) (z + m - T_i) - \epsilon \left(\frac{T_i}{2m} \right)^{n-1} \right] d \tag{56}$$

This implies that

$$\begin{array}{l} \epsilon \left(\frac{1}{2m}\right)^{n-1} \int_{T_i=0}^{2m} \left[\left((n-1) T_i^{n-2} \right) (z+m) \right. \\ \left. -(n-1) T_i^{n-1} - T_i^{n-1} \right] dT_i = 0. \end{array} \tag{57}$$

which in turn implies that

$$\begin{array}{l} \epsilon \left(\frac{1}{2m} \right)^{n-1} \Bigg|_{T_i=0}^{2m} \left(T_i^{n-1} (z+m) \right. \\ \left. -T_i^n \right) = 0, \end{array} \tag{58}$$

so $(2m)^{n-1}(z+m) - (2m)^n - 0 + 0 = 0$ and $z = m$. Bidder i's optimal strategy in the symmetric equilibrium is to bid $p_i = s_i - m$.

Winning Bid and Expected Revenue

The winning bid is set by the bidder with the highest signal, and that highest signal's expected value is

$$\begin{aligned} Es_{(1)} &= \underline{s} + \left(\tfrac{n+1-1}{n+1}\right)(\overline{s} - \underline{s}) \\ &= v - m + \left(\tfrac{n}{n+1}\right)(2m) \end{aligned} \tag{59}$$

The expected revenue is therefore

$$Ep_{(1)} = v - (1)\left(\frac{1}{n+1}\right)2m. \tag{60}$$

If $m = 50$ and $n = 4$, then

$$Ep_{(1)} = v - \left(\frac{1}{5}\right)(100) = v - 20. \tag{61}$$

Here, the revenue is even lower than in the second-price auction, where it was $(v - 15)$ (and the revenue is lower even if $n = 2$).

Revenue ranking: Ascending highest, then second-price, then first-price.

The Wallet Game

Order of Play
(0) Nature chooses the amounts s_1 and s_2 of the money in each player's wallet using density functions $f_1(s_1)$ and $f_2(s_2)$. Each player observes only his own wallet's contents.
(1) Each player chooses a bid ceiling p_1 or p_2. An auctioneer auctions off the two wallets by gradually raising the price until either p_1 or p_2 is reached.

Payoffs:
The player who bids less has a payoff of zero. The winning player pays the bid ceiling of the loser and hence has a payoff of

$$s_1 + s_2 - Min(p_1, p_2) \tag{62}$$

A symmmetric equilibrium is for Bidder i to choose bid ceiling $p_i = 2s_i$.

This is an equilibrium because if he wins at exactly that price, Bidder j's signal must be $s_j = s_i$ and the value of the wallets is $2s_i$.

If Bidder i bids any lower, he might pass up a chance to buy the wallet for less than its value.

If he bids any higher, he would only win if $p > 2s_j$ too, which implies that $p > s_i + s_j$.

Affiliation

Definition. The signals x_1 and x_2 are affiliated if for all possible realizations $Small < Big$ of x_1 and $Low < High$ of x_2, the joint probability $f(x_1, x_2)$ is such that z_1 and z_2,

$$\begin{aligned} & f(x_1 = Small, x_2 = Low) f(x_1 = Big, x_2 = High) \\ & \geq f(x_1 = Small, x_2 = High) f(x_1 = Big, x_2 = Low). \end{aligned} \tag{63}$$

Thus, affiliation says that the probability the values of x_1 and x_2 move in the same direction is greater than the probability they move oppositely.

The implication of two signals being affiliated is that the expected value of the winning bid conditional on the signals is increasing in all the signals.

When one signal rises, that has the positive direct effect of increasing the bid of the player who sees it, and non-negative indirect effects once the other players see his bid increase and deduce that he had a high signal.

The Monotone Likelihood-Ratio Property

The Monotone Likelihood-Ratio Property is the same thing expressed in terms of the conditional densities, the posteriors.

Definition. The conditional probability $g(x_1|x_2)$ satisfies the Monotone Likelihood Ratio Property if the likelihood ratio is weakly decreasing in x_1, that is, for all possible realizations $Small < Big$ of x_1 and $Low < High$ of x_2,

$$\frac{g(Big|Low)}{g(Big|High)} \leq \frac{g(Small|Low)}{g(Small|High)}. \tag{64}$$

The Monotone Likelihood Ratio Property says that as x_2 goes from Low to High, the Big value of x_1 becomes relatively more likely.

It can be shown that this implies that for any value z, the conditional cumulative distribution of x_1 up to $x_1 = z$ given x_2 weakly increases with x_2, which is to say that the distribution $G(x_1|x_2)$ conditional on a larger value of x_2 stochastically dominates the distribution conditional on a smaller value of x_2.

The Linkage Principle:

The linkage principle: when the amount of affiliated information available to bidders increases, the equilibrium sales price becomes greater.

The seller should have a policy of disclosing any affiliated information he possesses.

Auction rules which reveal affiliated information in the course of the auction (e.g., open-exit auctions) or use it in determining the winner's payment (e.g., the second-price auction) will result in higher prices.