16 Feb. 2006.
Eric Rasmusen, Erasmuse@indiana.edu, http://www.rasmusen.org. This replaces pieces of the 22 Nov. 2005 draft.

Anything incorporated is so marked. Everything is incorproated right now.

XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX

INCORP.

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). In the Continuous-Value Auction, denote the highest announced value by $\tilde{v}_{(1)}$, the second-highest by $\tilde{v}_{(2)}$, and so forth. the highest bidder gets the object with probability $G(\tilde{v}_{(1)}, \tilde{v}_{-1}) = 1$ at price $t(\tilde{v}_{(1)}, \tilde{v}_{-1}) = \tilde{v}_{(2)}$, and for $i \neq 1$, $G(\tilde{v}_{(i)}, \tilde{v}_{-i}) = 0$ and $t(\tilde{v}_{(1)}, \tilde{v}_{-1}) = 0$. This is incentive compatible, since a player's value announcement only matters if his value is highest, and he then wants to win if and only if the price is less than or equal to his value. It satisfies the participation constraint because his lowest possible payoff following that strategy is zero, and his payoff is higher if he wins and $\tilde{v}_{(1)} > \tilde{v}_{(2)}$.

Since each bidder's expected payoff is strictly positive, the *optimal* mechanism for the seller would be more complicated. As we will discuss later, it would include a **reserve price** p^* below which the object would remain unsold, changing the first part of the mechanism to $G(\tilde{v}_{(1)}, \tilde{v}_{-1}) = 1$ and $t(\tilde{v}_{(1)}, \tilde{v}_{-1}) = Max\{\tilde{v}_{(2)}, p^*\}$ if $\tilde{v}_{(1)} \geq p^*$ but $G(\tilde{v}_{(1)}, \tilde{v}_{-1}) = 0$ if $\tilde{v}_{(1)} < p^*$. We have seen in Chapter 10 that optimal mechanisms are not always efficient, and this is an example: the object will go unsold if $\tilde{v}_{(1)} < p^*$.

In the Ten-Sixteen Auction, the seller's value is $v_s = 0$, and each of two bidders' private values v_1 and v_2 is either 10 or 16 with equal probability, known only to the bidder himself.

incorporated

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 $p(v)$:

$$\pi(v, p(v)) = G(p(v))[v - p(v)]. \tag{1}$$

Let us first prove a lemma.

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

Proof: The first-order condition from payoff (1) is

$$\frac{d\pi(v)}{dp} = G'(v - p) - G = 0. \tag{2}$$

The optimum is an interior solution because at $p_i = 0$ the payoff is increasing and if p_i becomes large enough, π is negative. Thus, $\frac{d^2\pi(v_i)}{dp_i^2} \leq 0$ at the optimum. Using the implicit function theorem and the fact that $\frac{d^2\pi(v_i)}{dp_i dv_i} = G' \geq 0$ because a higher bid does not yield a lower probability of winning, we can conclude that $\frac{dp_i}{dv_i} \geq 0$, at least if the bid function is differentiable. But it cannot be that $\frac{dp_i}{dv_i} = 0$, because then there would be values v_1 and v_2 such that $p_1 = p_2 = p$ and then

$$\frac{d\pi(v_1)}{dp_1} = G'(p)(v_1 - p) - G(p) = 0 = \frac{d\pi(v_2)}{dp_2} = G'(p)(v_2 - p) - G(p), \tag{3}$$

which cannot be true. So the bidder bids more if his value is higher. Q.E.D.

Now let us try to find an equilibrium bid function. From equation (1), it is

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

That is not very useful in itself, since it has $p(v)$ on both sides. We need to find ways to rewrite π and G in terms of just v.

First, tackle $G(p(v))$. Monotonicity of the bid function (from Lemma 1) implies that the bidder with the greatest v will bid highest and win. Thus, 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,

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

Next think about $\pi(v, p(v))$. 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 $\frac{\partial \pi}{\partial p} = 0$:

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

We can apply the Envelope Theorem to equation (1) to see how π changes with v assuming $p(v)$ is chosen optimally, which is appropriate because we are characterizing not just any bid function, but the optimal bid function. Thus,

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

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

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

To get $\pi(v,p(v))$ from its derivative, (8), integrate over all possible values from zero to v and include the a base value of $\pi(0)$ 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 last step is true because a bidder with $v = 0$ will never bid a positive amount and so will have a payoff of $\pi(0,p(0)) = 0$.

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

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

What a happy ending to a complicated derivation! If there are two bidders and values are uniform on $[0,1]$, a bidder should bid $p = v/2$, which since he has probability v of winning yields an expected payoff of $v^2/2$. If $n = 10$ he should bid $\frac{9}{10}v$, which since he has probability v^9 of winning yields him an expected payoff of $v^{10}/10$, quite close to zero if $v < 1$.

INCORP

THE REVENUE EQUIVALENCE THEOREM. *Let all players be risk-neutral with private values drawn 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.

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

$$\pi(z, v) = G(z)v - p(z). \tag{12}$$

The first-order condition with respect to the player's choice of type message z is

$$\frac{d\pi(z; v)}{dz} = g(z)v - \frac{dp(z)}{dz} = 0, \tag{13}$$

so

$$\frac{dp(z)}{dz} = g(z)v. \tag{14}$$

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

$$\frac{dp(v)}{dv} = g(v)v. \tag{15}$$

Next, we integrate (15) over all values from zero to v, adding $p(\underline{v})$ as the constant of integration:

$$p(v) = p(\underline{v}) + \int_{\underline{v}}^{v} g(x)x dx. \tag{16}$$

We can use (16 to substitute for $p(v)$ in the payoff equation (12), which becomes, after replacing z with v and setting $p(\underline{v}) = 0$ because of assumption (b),

$$\pi(v, v) = G(v)v - \int_{\underline{v}}^{v} g(x)x dx. \tag{17}$$

Equation (17) 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.

There are many versions of the revenue equivalence theorem, and the name of the theorem comes from a version that just says that the seller's revenue is the same across auction rules rather than including bidders too. The version proved above is adapted from Proposition 3.1 of Krishna (2002, p. 30). Other versions, which use different proof approaches, can be found in Klemperer (1996, p. 40), and Milgrom (2004, p. 74). Two assumptions that are standard across versions are that the bidders are risk neutral and that their values are drawn from the same distribution.

It is only when we apply the Revenue Equivalence Theorem to the diverse auction rules we laid out earlier that its remarkable nature can be appreciated. The symmetric equilibria of the ascending, first-price, second-price, descending, and all-pay auctions with continuous values all satisfy the two conditions of the Theorem: (a) the winner is the bidder with the highest value, and (b) the lowest type makes an expected payment of zero. Thus, the following corollary is true.

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.*

p. 448. INCROPROATED

The revenue ranking is thus that the ascending open-exit auction has the highest expected revenue for the seller, the second-price auction is in the middle, and the first-price auction is lowest. The revenue depends on how intensely bidders compete up the price under each auction rule, which in turn depends on how much of an informational advantage the highest-signal-bidder has. In the ascending auction, all the bidders come to know $s_{(n)}$, and the winning price and who wins depends on $s_{(2)}$ and $s_{(1)}$, so the bidder with the highest signal has a relatively small advantage. In the second-price auction, the winning price and who wins depends on $s_{(2)}$ and $s_{(1)}$, and that information comes to be known only to those two bidders. In the first-price auction, the winning price and who wins depend only on $s_{(1)}$, so the bidder with the highest signal has the only relevant information. Thus, his informational rent is greatest under that auction rule. Paradoxically, the bidders prefer an auction rule which makes it harder for them to pool their information and accurately estimate v. Or perhaps this is not paradoxical. What the seller would like best would be for every bidder to truthfully announce his signal publicly, because then every bidder would have the same estimate, that would be the amount each would bid, and the informational rent would fall to zero.