

13 Auctions

13.1 Values Private and Common, Continuous and Discrete

Bargaining and auctions are two extremes in the many ways to sell goods, as Bulow & Klemperer (1996) explain. In typical bargaining, one buyer faces one seller and they make offers and counteroffers free from the formal rules we impose in theoretical modelling (though as in Bulow and Klemperer, it is worth considering what happens if the bargaining does have rules to which the players commit). In typical auctions, many bidders face one seller and make offers according to formal rules as rigid as those of the theorist. Bargaining is slow but flexible; auctions are fast but rigid.

Bargaining models generate different results with different assumptions, but since it is usually hard to match the assumptions with particular real situations, the practical implications come from the simplest models— ideas such as the importance of avoiding misunderstanding, determining and then concealing one's own reservation price, bluffing, and manipulating the timing of offers.

Auction models, on the other hand, may also generate different results with different assumptions but it is easier to match the assumptions with particular real situations, or even to create a real situation to match the model. That is because auctions vary not only in the underlying preferences of the players, as in bargaining, but in the specific rules used to play the game, and those rules are chosen by one of the players, usually with legal commitment to them. Thus, auction theory lends itself to what Alvin Roth (2002) calls “the economist as engineer”: the use of technical economic theory to design institutions for specific situations. As with other engineering, the result may not be of general interest, but tailoring the model to the situation is both tricky and valuable, requiring the same kind of talent and care as developing the general theory.

Because auctions are stylized markets with well-defined rules, modelling them with game theory is particularly appropriate. Moreover, several of the motivations behind auctions are similar to the motivations behind the asymmetric information contracts of Part II of this book. Besides the mundane reasons such as speed of sale that make auctions important, auctions are useful for a variety of informational purposes. Often bidders know more than the seller about the value of what is being sold, and the seller, not wanting to suggest a price first, uses an auction as a way to extract information. Art auctions are a good example, because the value of a painting depends on the bidder's tastes, which are known only to himself. Efficient allocation of resources is a goal different from profit maximization, but auctions are useful for that too. A good example is told in Boyes and Happel's 2001 article, “Auctions as an Allocation Mechanism in Academia: The Case of Faculty Offices.” At their business school, the economics department used an auction to allocate new offices, whereas the management department used seniority, the statistics department used dice,

and the finance department posted a first-come first-serve sign-up sheet without warning. The auction has the best chance of coming up with an immediate efficient allocation. (Why do I say “immediate”? Why would the long-run probably stay the same?)

Auctions are also useful for agency reasons, because they hinder dishonest dealing. If the mayor were free to offer a price for building the new city hall and accept the first contractor who showed up, the lucky contractor would probably be the one who made the biggest political contribution. If the contract is put up for auction, cheating the public is more costly, and the difficulty of rigging the bids may outweigh the political gain.

We will spend most of this chapter on the effectiveness of different kinds of auction rules in extracting surplus from bidders, which will require finding the strategies with which bidders respond to the rules. Section 13.1 classifies auctions based on the relationships between different bidders’ estimates of the value of what is being auctioned. Section 13.2, a necessarily very long section, explains the possible auction rules and the bidding strategies optimal for each rule. Section 13.3 compares the outcomes under the various rules, proving the Revenue Equivalence Theorem and showing how it becomes invalid if bidders are risk averse. Section 13.4 shows how to choose an optimal reserve price using the similarity between optimal auctions and monopoly pricing. Section 13.5 analyzes common-value auctions, which can lead bidders into “the winner’s curse” if they are not careful. Section 13.6 discusses asymmetric equilibria (The Wallet Game) and information affiliation.

Private-Value and Common-Value Auctions

Auctions differ enough for an intricate classification to be useful. One way to classify auctions is based on differences in the values bidders put on what is being auctioned. We will 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 .

In a **private-value auction**, a bidder can learn nothing about his value from knowing the values of the other bidders. An example is the sale of antique chairs to people who will not resell them. Usually a bidder’s value equals his value estimate in private-value auction models. If an auction is to be private-value, it cannot be followed by costless resale of the object. If there were resale, a bidder’s value would depend on the price at which he could resell, which would depend on the other bidders’ values. What is special about a private-value auction is that a bidder cannot extract any information about his own value from the value estimates of the other bidders. Knowing all the other values in advance would not change his estimate. It might well change his bidding strategy, however, so we distinguish between the **independent private-value auction**, in which knowing his own value tells him nothing about other bidders’ values, and other situations such as the **affiliated private-value auction** (affiliation being a concept which will be explained later) in which he might be able to use knowledge of his own value to deduce something about other players’ values.

In a **pure common-value auction**, the bidders have identical values, but each bidder forms his own estimate on the basis of his own private information. An example is bidding for U. S. Treasury bills. A bidder’s estimate would change if he could sneak a look at the other bidders’ estimates, because they are all trying to estimate the same true value.

The values in most real-world auctions are a combination of private value and common value, because the value estimates of the different bidders are positively correlated but not identical. As always in modelling, we trade off descriptive accuracy against simplicity. It is common for economists to speak of mixed auctions as “common-value” auctions, since their properties are closer to those of common-value auctions. Krishna (2002) has used the term **interdependent value** for the mixed case.

The private-value/common-value dichotomy is about what a bidder knows about his own value, but a separate dimension of auctions is what a bidder knows about other bidders’ values. One possibility is that the values are common knowledge. This makes optimal bidding simple. In a private-value auction, the highest-valuing bidder can bid just above the second-highest value. In a common-value auction, the bidders will take into account each others’ information and their value estimates will instantly converge to a single common estimate, the best one given all available information. In either case, bidders don’t have to worry about cleverly deducing each others’ information by observing how the bidding proceeds.

It would be odd, however, to observe an auction in the real world in which the seller knew the value estimates. In that case, the seller shouldn’t be using an auction. He should just charge a price equal to the highest value estimate, a price he knows will be accepted by one of the bidders.

More commonly, only the bidder himself knows his value, not other bidders or the seller. Most simply, the estimated values are statistically independent, the independent private-value case that we will be analyzing for the first half of this chapter. If the estimates are independent, then a bidder’s only source of information about his value is his private information and his only source of information about other bidders’ value estimates is what he observes of their bidding. If the value estimates are not independent, a bidder can use his private information to help estimate other bidders’ values (and thus how they will bid), and if he does not know his own value perfectly he can use observations of their value estimates to improve his estimate of his own value— which takes us to the common value case.

Lack of statistical independence can be present even in private-value auctions, however. Smith might know his own value perfectly, and so would not be able to learn anything about it by learning Jones’s value. That is what makes the auction a private-value auction. But if the values are not independent— if, say they are positively correlated— then if Smith knows that his own value is unusually high he could predict that Jones’s value was also high. This would affect Smith’s bidding strategy, even though it would not affect the maximum he would be willing to pay for the object being sold. I will not be analyzing private-value auctions with correlated values here, but as explained in Riley (1989) the difference can be important. Although bidding in the ascending and second-price auction rules that we will soon discuss is unchanged when values are correlated, bidding becomes lower in the descending and first-price auctions. In contrast, we will see that bidding in a common-value auction is generally more cautious than in a private-value auction even in ascending and second-price auctions.

To look at these various auction rules, we will use the following two games, which I

have set up using the language of mechanism design since choosing auction rules fits well into that paradigm. The Ten-Sixteen Auction will be our running example for when values are discrete ($v = 10$ or $v = 16$), and the Continuous-Value Auction will be our example for when they are on a continuum ($v \sim f(v)$ on $[\underline{v}, \bar{v}]$).

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.)
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.).
2. Each bidder simultaneously chooses to participate in the auction or to stay out.
3. The bidders and the seller choose \tilde{v} according to the mechanism procedure.
4. The object is allocated and transfers are paid according to the mechanism.

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

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

To get around the open-set problem, we will assume for any auction rule that ties are broken in favor of whoever has the highest value, or randomly if the values are equal. Otherwise, if, for example, we said that ties simply split the probability of winning, then if $v_1 = 10$ and $v_2 = 16$ and this were known to both bidders, it would not be even a weak equilibrium to have them bid $p_1 = 10$ and $p_2 = 10$, because Bidder 2 would deviate to a slightly higher bid— but the smallest bid strictly greater than 10 does not exist if bid increments can be infinitesimal.

The second game will be our example for when values lie on a continuum.

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 $[v, \bar{v}]$.

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.).
2. Each bidder simultaneously chooses to participate in the auction or to stay out.
3. The bidders and the seller choose \tilde{v} according to the mechanism procedure.
4. The object is allocated and transfers are paid according to the mechanism, if it was accepted by all bidders.

Payoffs:

The seller's payoff is

$$\pi_s = \sum_{i=1}^n t(\tilde{v}_i, \tilde{v}_{-i}) \quad (3)$$

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

$$\pi_i(v_i) = G(\tilde{v}_i, \tilde{v}_{-i})v_i - t(\tilde{v}_i, \tilde{v}_{-i}) \quad (4)$$

Many possible auction procedures fit the mechanism paradigm, even ones that are never used in practice. The mechanism could allocate the good with 70% probability to the high bidder and with 30% probability to the lowest bidder, for example; or each bidder could be made to pay the amount he bids, even if he loses; or t could include an entry fee; or there could be a "reserve price," a minimum bid for which the seller will surrender the good. In this analysis, the seller will choose a direct mechanism that satisfies a participation constraint for each bidder type v_i (Bidder i will join the auction, so, for example, the entry fee is not too large), and an incentive compatibility constraint (the bidder will truthfully reveal his type; $\tilde{v}_i = v_i$).

13.2 Optimal Strategies under Different Rules in Private-Value Auctions

Auctions have the same bewildering variety of rules as poker does. We will look at five different auction rules, using the private-value setting since it is simplest. In teaching this material, I ask each student to pick a value between 80 and 100, after which we conduct the various kinds of auctions. I advise the reader to try this. Pick two values and try out sample strategy profiles for the different auctions as they are described. Even though the values are private, it will immediately become clear that the best-response bids still depend on the strategies the bidder thinks other bidders have adopted.

The five auction rules we will consider (with common synonyms for them) are :

- 1 Ascending (English, open-cry, open-exit);
- 2 First-Price (first-price sealed-bid);
- 3 Second-Price (second-price sealed-bid, Vickrey);
- 4 Descending (Dutch)
- 5 All-Pay

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.

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. This is optimal because he always wants to buy the object if the price is less than its value to him, but he wants to pay the lowest price possible. All bidding ends when the price reaches the second-highest value of any bidder present at the auction. The optimal strategy is independent of risk neutrality if bidders know their own values with certainty rather than having to estimate them, although risk-averse bidders who must estimate their values should be more conservative in bidding.

The optimal private-value strategy is simple enough that details of the ascending auction usually do not make much difference, but there are a number of possibilities.

- (1) The bidders offer new prices using pre-specified increments such as dollars or thousands of dollars.
- (2) The **open-exit** auction, in which the price rises continuously and bidders show their willingness to pay the price by not dropping out, where a bidder's dropping out is publicly announced to the other bidders.
- (3) The **silent-exit** auction (my neologism), in which the price rises continuously and bidders show their willingness to pay the price by not dropping out, but a bidder's dropping out is not known to the other bidders.
- (4) The **Ebay auction**, in which a bidder submits his "bid ceiling," the maximum price he is willing to pay. During the course of the auction the seller uses the bid ceilings to raise the current winning bid only as high as necessary, and the winner is the player whose bid is highest at a prespecified ending time.
- (5) The **Amazon auction**, in which a bidder submits his bid ceiling. During the course of the auction the seller uses the bid ceilings to raise the current winning bid only as high as necessary, and the winner is the player whose bid is highest at a prespecified ending time or ten minutes after the last increase in the current winning bid, whichever is later.

The precise method can be quite important in common-value auctions, where knowing what other players are doing alters a bidder's own value estimate. If the auction is open-exit, for example, a bidder who observed that most of the other bidders dropped out at a low price would probably revise his own value estimate downwards, something he would not know to do in a silent-exit 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). 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 $t(\tilde{v}_{(i)}, \tilde{v}_{-i}) = 0$. 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 $t(\tilde{v}_{(1)}, \tilde{v}_{-1}) = \text{Max}\{\tilde{v}_{(2)}, p^*\}$ 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. A bidder's optimal strategy in the ascending auction would be to set his bid or bid ceiling to $p(v = 10) = 10$ and $p(v = 16) = 16$. His expected payoff would be

$$\begin{aligned}\pi(v = 10) &= 0 \\ \pi(v = 16) &= 0.5(16 - 10) + 0.5(16 - 16) = 3\end{aligned}\tag{5}$$

The expected price, the payoff to the seller, is

$$\pi_s = 0.5^2(10) + 0.5^2(16) + 2(0.5)^2(10) = 2.5 + 4 + 5 = 11.5\tag{6}$$

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.

Discussion

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. As we will see later, however, there are good reasons why sellers so often choose to use first-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 $p(v)$:

$$\pi(v, p(v)) = G(p(v))[v - p(v)]. \quad (7)$$

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 (7) is

$$\frac{d\pi(v)}{dp} = G'(v - p) - G = 0. \quad (8)$$

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

$$\frac{d\pi(v_1)}{dp_1} = G'(p)(v_1 - p) - G(p) = 0 = \frac{d\pi(v_2)}{dp_2} = G'(p)(v_2 - p) - G(p), \quad (9)$$

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 (7), it is

$$p(v) = v - \frac{\pi(v, p(v))}{G(p(v))}. \quad (10)$$

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,

$$G(p(v)) = F(v)^{n-1}. \quad (11)$$

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 $\frac{\partial \pi}{\partial p} = 0$:

$$\frac{d\pi(v, p(v))}{dv} = \frac{\partial \pi(v, p(v))}{\partial p} \frac{\partial p}{\partial v} + \frac{\partial \pi(v, p(v))}{\partial v} = \frac{\partial \pi(v, p(v))}{\partial v}. \quad (12)$$

We can apply the Envelope Theorem to equation (7) 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,

$$\frac{d\pi(v, p(v))}{dv} = G(p(v)). \quad (13)$$

Substituting from equation (11) gives us π 's derivative, if not π , as a function of v :

$$\frac{d\pi(v, p(v))}{dv} = F(v)^{n-1}. \quad (14)$$

To get $\pi(v, p(v))$ from its derivative, (14), integrate over all possible values from zero to v and include the a base value of $\pi(0)$ as the constant of integration:

$$\pi(v, p(v)) = \pi(0) + \int_0^v F(x)^{n-1} dx = \int_0^v F(x)^{n-1} dx. \quad (15)$$

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 (10) and substitute for $G(p(v))$ and $\pi(v, p(v))$ from equations (11) (15):

$$p(v) = v - \frac{\int_0^v F(x)^{n-1} dx}{F(v)^{n-1}}. \quad (16)$$

Suppose $F(v) = v/\bar{v}$, the uniform distribution. Then (16) becomes

$$\begin{aligned} p(v) &= v - \frac{\int_0^v \left(\frac{x}{\bar{v}}\right)^{n-1} dx}{\left(\frac{v}{\bar{v}}\right)^{n-1}} \\ &= v - \frac{\int_{x=0}^v \left(\frac{1}{\bar{v}}\right)^{n-1} \left(\frac{1}{n}\right) x^n}{\left(\frac{v}{\bar{v}}\right)^{n-1}} \\ &= v - \frac{\left(\frac{1}{\bar{v}}\right)^{n-1} \left(\frac{1}{n}\right) v^n - 0}{\left(\frac{v}{\bar{v}}\right)^{n-1}} \\ &= v - \frac{v}{n} = \left(\frac{n-1}{n}\right) v. \end{aligned} \quad (17)$$

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 $v < 1$.

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

The result in equations (??) and (17) depended crucially on the value distribution having a continuous support. When this is not true, 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:

$$\begin{aligned} \pi(v = 10) &= 0 \\ \pi(v = 16) &= 3 \\ \pi_s &= 11.5. \end{aligned} \quad (18)$$

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

This will serve as an illustration of how to find an equilibrium mixed strategy when bidders mix over a continuum of pure strategies rather than just between two. The first step is to see why the equilibrium cannot be in pure strategies (though some portions of the equilibrium strategy can be pure, e.g. if $v = 10$, a bidder will bid $p = 10$ with probability one).

First, there is no equilibrium in pure strategies, either symmetric or asymmetric. In any equilibrium, $p(v = 10) = 10$, because if either bidder used the bid $p < 10$, it would cause the other player to deviate to $(p + \epsilon)$, and a bid above 10 exceeds the object's value.

If $v = 16$, however, a player will randomize his bid, as I will now show. Suppose the two bidders are using the pure strategies $p_1(v_1 = 16) = z_1$ and $p_2(v_2 = 16) = z_2$. The values of z_1 and z_2 would lie in $(10, 16]$ because a bid of exactly 10 would lose to the positive-probability bid of $p(v = 10) = 10$ given our tie-breaking assumption and a bid over 16 would exceed the object's value. yield a negative payoff. Either $z_1 = z_2$, or $z_1 \neq z_2$. If $z_1 = z_2$, then each bidder has incentive to deviate to $(z_1 - \epsilon)$ and win with probability one instead of tying. If $z_1 < z_2$, then Bidder 2 will deviate to bid $(z_1 + \epsilon)$. If he does that, however, Bidder 1 would deviate to bid $(z_1 + 2\epsilon)$, so he could win with probability one at trivially higher cost. The same holds true if $z_2 < z_1$. Thus, there is no equilibrium in pure strategies.

The second step is to figure out what pure strategies will be mixed between by a bidder with $v = 16$. It turns out that they form the interval $[10, 13]$. As just explained, the bid $p(v = 16)$ will be no less than 10 (so the bidder can win if his rival's value is 10) and no greater than 16 (which would always win, but unprofitably). The pure strategy of $(p = 10)|(v = 16)$ will win with probability of at least 0.50 (when the other bidder happens to have $v = 10$, given our tie-breaking rule), yielding a payoff of $0.50(16 - 10) = 3$. This rules out bids in $(13, 16]$, since even if they always win, their payoff is less than 3. Thus, the upper bound \bar{p} must be no greater than 13.

The lower bound \underline{p} must be exactly 10. If it were at $(10 + \epsilon)$ then a bid of $(10 - 2\epsilon)$ would have an equal certainty of winning the auction, but would have ϵ higher payoff. Thus, $\underline{p} = 10$.

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 $p = 10$. In a mixed-strategy equilibrium, though, the payoff from any of the strategies mixed between must be equal. Thus, \bar{p} cannot be less than 13.

We are not quite done looking at the strategies mixed between. 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 $[10, 13]$? No, it turns out.

(a) Bidder 2's mixing density does not have an atom at any point a in $[10, 13]$ — no point a has positive probability, as opposed to positive density. An example of such an atom would be if the mixing distribution were the density $m(p) = 1/6$ over the interval $[10, 13]$ plus an atom of probability $1/2$ at $p = 13$, so the cumulative probability would be $M(p) = p/6$ over $[10, 13)$ and $M(13) = 1$. Using $M(p)$, a point such as 11 would have zero probability even though the interval, say, of $[10.5, 12.5]$ would have probability $2/6$.

If there were an atom at a , Bidder 1 would respond by putting positive probability on $(a + \epsilon)$ and zero probability on a . But then Bidder 2 would respond by putting zero probability on a and shifting that probability to $(a + 2\epsilon)$.

(b) Bidder 2's mixing density does not have a gap $[g, h]$ anywhere with $g > 10$ and $h < 13$. If it did, then Bidder 1's payoff from bidding g and h would be

$$\pi_1(g) = \text{Prob}(p_2 < g)v_1 - g \quad (19)$$

and

$$\pi_1(h) = \text{Prob}(p_2 < h)v_1 - h = \text{Prob}(p_2 < g)v_1 - h, \quad (20)$$

where the second equality in $\pi_1(h)$ is true because there is zero probability that p_2 is between g and h . Bidder 1 will put zero probability on $p_1 = h$, since its payoff is lower than the payoff from $p_1 = g$ and will put zero probability on slightly larger values of p_1 too, since by continuity their payoffs will also be less than the payoff from $p_1 = g$. This creates a gap $[h, h^*]$ in which $p_1 = 0$. But then Bidder 2 will want to put zero probability on $p_2 = h^*$ and slightly higher values, by the same reasoning, which means that our original hypothesis of only a gap $[g, h]$ is false.

Thus, we can conclude that the mixing density $m(p)$ is positive over the entire interval $[10, 13]$, with no atoms. What will it look like? Let us confine ourselves to looking for a symmetric equilibrium, in which both bidders use the same function $m(p)$. We know the expected payoff from any bid p in the support must equal the payoff from $p = 10$ or $p = 13$, which is 3. Therefore, 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

$$0.5(16 - p) + 0.5M(p)(16 - p) = 3. \quad (21)$$

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

$$M(p) = \frac{6}{16 - p} - 1, \quad (22)$$

which has the density

$$m(p) = \frac{6}{(16 - p)^2} \quad (23)$$

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

Since each bidder type has the same expected payoff in this first-price auction as in the ascending auction, and the object is sold with probability one, it must be that the seller's payoff is the same, too, equal to 11.5, as we found in equation (6).

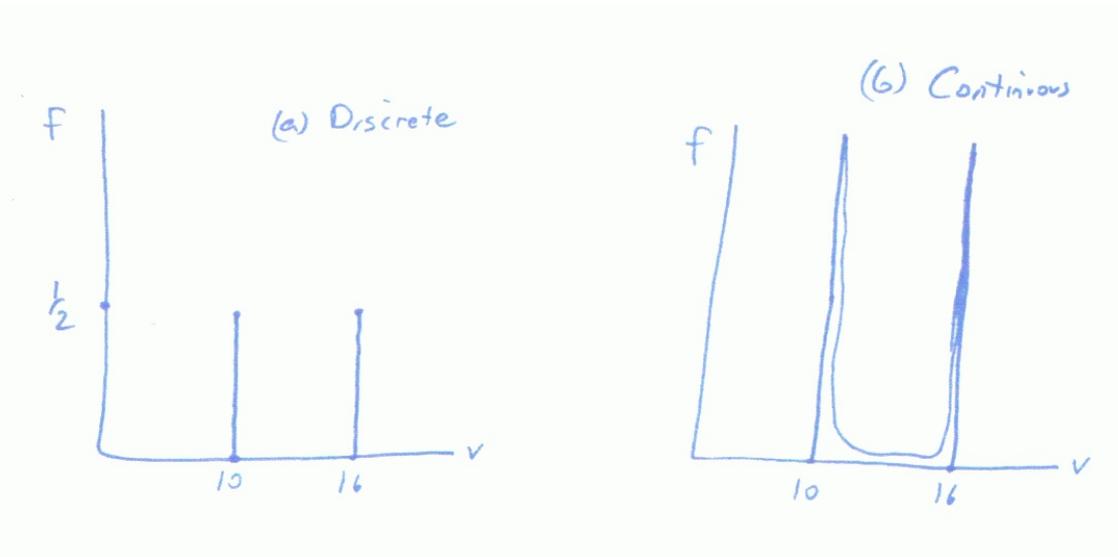


Figure 1: Smoothing a Discrete Distribution

You may find it odd that the general continuous-value auction has a pure-strategy equilibrium but our particular discrete-value auction does not. Usually if a game lacks a pure-strategy equilibrium in discrete type space, it also lacks one if we “smooth” the probability distribution by making it continuous but still putting almost all the weight on the old discrete types, as in Figure 1.

This is related to a remarkable feature of private-value auctions with discrete values: the mixed-strategy equilibria do not necessarily block efficiency (and the revenue equivalence we study later). When players randomize, it would seem that sometimes by chance the highest-valuing player would be unlucky and lose the auction, which would be inefficient. Not so here. As explained in Riley (1989) and Wolfstetter (1999, p. 204), if the values of each player are distributed discretely over some set $\{0, v_a, v_b, \dots, v_w\}$ then in the symmetric equilibrium mixed strategy, the supports of the mixing distributions are $v_a : [0, p_1], v_b : [p_1, p_2], v_w : [p_{w-1}, p_w]$, where $p_1 < p_2 < \dots < p_w$. The supports do not overlap. Each type of bidder acts as if he was in competition just with his own type (since he will surely win over the lower types and will surely lose to the higher types) and the object is allocated to a bidder who values it most. The mixing only determines who wins when two players happen to have the same type.

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.

Discussion

Second-price auctions are similar to ascending auctions. They are rarely used in reality, but are useful for modelling. Bidding one's value is a weakly dominant strategy: a bidder who bids less is no more likely to win the auction (and probably less likely, depending on $f(v)$), but he pays the same price—the second-highest-valuing player's bid—if he does win. The structure of the payoffs is reminiscent of the Groves Mechanism of Section 10.4, because in both games a bidder's strategy affects some major event (who wins the auction, or whether the project is undertaken), but his strategy affects his own payoff only via that event. In the auction's symmetric equilibrium, each bidder bids his value and the winner ends up paying the second-highest value. If bidders know their own values, the outcome does not depend on risk neutrality.

One difference between ascending and second-price auctions is that second-price auctions have peculiar asymmetric equilibria because the actions in them are simultaneous. 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 there are asymmetric equilibria.

Consider the following equilibrium.

$$\begin{aligned} p_1(v = 10) &= 10 & p_1(v = 16) &= 16 \\ p_2(v = 10) &= 1 & p_2(v = 16) &= 10 \end{aligned} \tag{24}$$

Since Bidder 1 never bids less than 10, Bidder 2 knows that if $v_2 = 10$ he can never get a positive payoff, so he is willing to choose $p_2(v = 10) = 1$. Doing so results in a sale price of 1, for any $p_1 > 1$, which is better for Bidder 1 and worse for the seller than a price of 10, but Bidder 2 doesn't care about their payoffs. In the same way, if $v_2 = 16$, Bidder 2 knows that if he bids 10 he will win if $v_1 = 10$, but if $v_2 = 16$ he would have to pay 16 to win and would earn a payoff of zero. He might as well bid 10 and earn his zero by losing.¹

Perhaps the seller's fear of asymmetric equilibria like this is why Second-price auctions are so rare. They have actually been used, though, in a computer operating system. An operating system must assign a computer's resources to different tasks, and researchers at Xerox Corporation designed the Spawn system, under which users allocate "money" in a second-price auction for computer resources. See "Improving a Computer Network's Efficiency," *The New York Times*, p. 35 (29 March 1989).

Descending Auctions (Dutch)

¹Trembling-hand perfectness, however, would rule out this kind of equilibrium. If Bidder 1 might tremble and bid, for example, 4 by accident, Bidder 2 would not want to ever bid less than 4. Bidding less than one's value is weakly dominated by bidding exactly one's value—but Nash equilibrium strategies can be weakly dominated, as we saw with the Bertrand Game in Chapter 3.

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.

Discussion

The typical descending auction is **strategically equivalent** to the first-price auction, which means there is a one-to-one mapping between the strategy sets and the equilibria of the two games. The reason for the strategic equivalence is that no relevant information is disclosed in the course of the auction, only at the end, when it is too late to change anybody's behavior. In the first-price auction a bidder's bid is irrelevant unless it is the highest, and in the descending auction a bidder's stopping price is irrelevant unless it is the highest. The equilibrium price is calculated the same way for both auctions.

A descending auction does not have to be like a first-price auction as a matter of logic, though. Vickrey (1961) notes that a descending auction could be set up as a second-price auction. When the first bidder presses his button, he primes an auction-ending buzzer that does not go off until a second bidder presses his button. In that case, the descending auction would be strategically equivalent to a second-price auction. Economists almost always mean "first-price descending auctions" when they use the term, however.

Descending—"Dutch"—auctions have been used in the Netherlands to sell flowers—see the Aalsmeer auction website at <http://www.vba.nl> for information and photos. They have also been used in Ontario to sell tobacco using a clock four feet in diameter marked with quarter-cent gradations. Each of six or so bidders has a stop button. The clock hand drops a quarter-cent at a time, and the stop buttons are registered so that ties cannot occur (tobacco bidders need reflexes like race-car drivers). The farmer sellers watch from an adjoining room and can later reject the bids if they feel they are too low (a form of reserve price) The clock is fast enough to sell 2,500,000 lb. per day (Cassady [1967, p. 200]).

Descending auctions are common in less obvious forms. Filene's is one of the biggest stores in Boston, and Filene's Basement is its most famous department. In the basement are a variety of marked-down items formerly in the regular store, each with a price and date attached. The price customers pay at the register is the price on the tag minus a discount which depends on how long ago the item was dated. As time passes and the item remains unsold, the discount rises from 10 to 50 to 70 percent. The idea of predictable time discounting has also been used by bookstores too ("Waldenbooks to Cut Some Book Prices in Stages in Test of New Selling Tactic," *The Wall Street Journal*, March 29, 1988, p. 34).

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).

I do not know of the all-pay rule ever being used in a real auction, but it is a useful modelling tool because it models rent-seeking very well. When a number of companies lobby a politician for a privilege, they are in an all-pay auction because even the losers have paid by incurring the cost of lobbying. When a number of companies pursue a patent, it is an all-pay auction because even the losers have incurred the cost of doing research.

The Equal-Value All-Pay Auction

Suppose each of the n bidders has the same value, v . That is not a very interesting game for most of the auction rules, though it is true that for the second-price auction there exists the strange asymmetric equilibrium $\{v, 0, 0, \dots, 0\}$. Under the all-pay auction rule, however, this game is quite interesting. The equilibrium is in mixed strategies. This is easy to see, because in any pure-strategy profile, 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' .

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.² 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

$$M(p)^{n-1}v = p, \tag{25}$$

so

$$M(p) = \sqrt[n-1]{\frac{p}{v}}, \tag{26}$$

as shown in Figure 2. At the extreme bids that a bidder with value v might offer, $M(0) = \sqrt[n-1]{\frac{0}{v}} = 0$ and $M(v) = \sqrt[n-1]{\frac{v}{v}} = 1$, so we have found a valid distribution function $M(p)$. Moreover, since the payoff from one of the strategies between which it mixes, $p = 0$, equals zero, we have verified our conjecture that $\pi(p) = 0$ in the equilibrium.

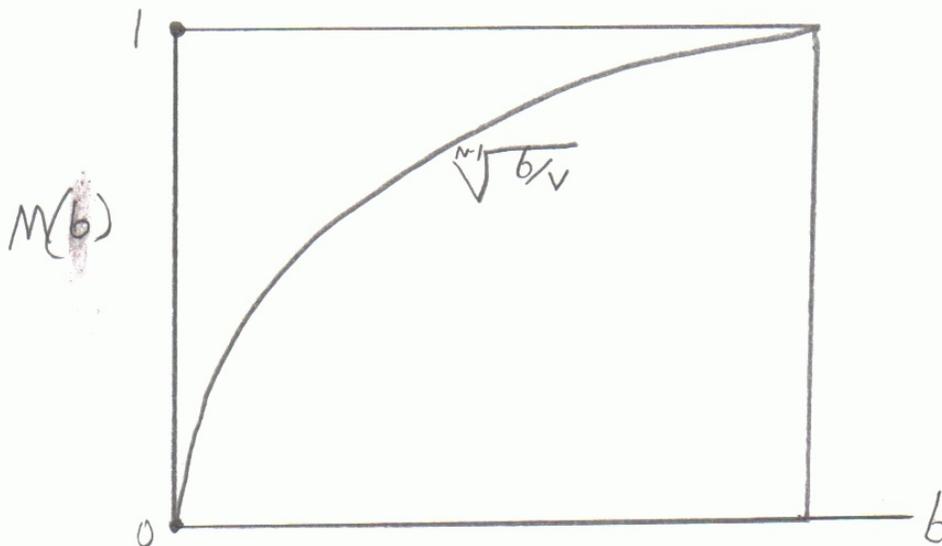


Figure 2: The Bid Function in an All-Pay Auction with Identical Bidders

Consider now what happens with the all-pay rule in the Continuous-Value Auction Game, using an explanation adapted from Krishna (2002, ch. 3.2).

²It turns out there is a continuum of asymmetric equilibria in this game if $n > 2$, but a unique equilibrium if $n = 2$. See Kovenock, Baye & de Vries (1996) for a full characterization of all-pay auctions with complete information.

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

$$\pi(v, z) = F(z)^{n-1}v - p(z), \quad (27)$$

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.

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 (27) $\pi(v, z)$ is differentiable in z in our conjectured equilibrium. Thus, we need to find z such that

$$\frac{\partial \pi(v, z)}{\partial z} = (n - 1)F(z)^{n-2}f(z)v - p'(z) = 0 \quad (28)$$

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

$$p'(v) = (n - 1)F(v)^{n-2}f(v)v \quad (29)$$

Integrating up, we get

$$p(v) = p(0) + \int_0^v (n - 1)F(x)^{n-2}f(x)xdx \quad (30)$$

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

We can verify that truthelling is a symmetric equilibrium strategy by substituting for $p(z)$ from (30) into payoff equation (27).

$$\begin{aligned} \pi(v, z) &= F(z)^{n-1}v - p(z) \\ &= F(z)^{n-1}v - p(0) - \int_0^z (n - 1)F(x)^{n-2}f(x)xdx \\ &= F(z)^{n-1}v - p(0) - F(z)^{n-1}z + \int_0^z F(x)^{n-1}dx, \end{aligned} \quad (31)$$

where the last step uses integration by parts ($\int gh' = gh - \int hg'$, where $g = x$ and $h' = (n - 1)F(x)^{n-2}f(x)$). Maximizing (31) with respect to z yields

$$\frac{\partial \pi(v, z)}{\partial z} = (n - 1)F(z)^{n-2}f(z)(v - z), \quad (32)$$

which is maximized by setting $z = v$. Thus, if $(n - 1)$ of the bidders are using this $p(v)$ function, so will the remaining bidder, and we have a Nash equilibrium.

Let's see what happens with a particular value distribution. Suppose values are uniformly distributed over $[0,1]$, so $F(v) = v$. Then equation (30) becomes

$$\begin{aligned}
 p(v) &= p(0) + \int_0^v (n-1)x^{n-2}(1)x dx \\
 &= p(0) + \left[(n-1) \frac{x^n}{n} \right]_{x=0}^v \\
 &= 0 + \left(\frac{n-1}{n} \right) v^n,
 \end{aligned} \tag{33}$$

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 $\pi = v(v^9) - (9/10)v^{10} = v^{10}/(10)$. As we will see when we discuss the Revenue Equivalence Theorem, it is no accident that this is the same payoff as for the first-price auction when values were uniformly distributed on $[0,1]$, in equation (??).

The Dollar Auction

A famous example of an auction in which not just the winner pays is the dollar auction of Shubik (1971). This is 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. If the players begin with infinite wealth, the game illustrates why equilibrium might not exist if strategy sets are unbounded. Once one bidder has started bidding against another, both of them do best by continuing to bid, so as to win the dollar as well as pay the bid. (If there are three or more players, all but the top two will be happy to stop bidding early in the game.) This auction may seem absurd, but as a variant on the all-pay auction it has considerable similarity to patent races and arms races. See Baye & Hoppe (2003) for more on the equivalence between innovation games and auctions.

The all-pay auction and the dollar auction are just two examples of auctions in which a player must pay something even though he loses. Chapter 3's War of Attrition is another example, which is something like a second-price all-pay auction. Even odder is the **loser-pays auction**, a two-player auction in which only the loser pays. As we have seen with the all-pay auction and its revenue equivalence to other auction forms, however, the fact that an auction's rules are strange does not mean it is necessarily worse for bidders.

All-pay auctions are a standard way to model rentseeking: imagine that n players each exert e in effort simultaneously to get a prize worth V , the winner being whoever's effort is highest. Another common way to model rentseeking is as an auction in which the highest bidder has the best chance to win but lower bidders might win instead. Tullock (1980) started a literature on this in an article which was valuable despite a mistaken claim that the expected amount paid by the bidders might exceed the value of the prize. (See Baye, Kovenock & de Vries (1999) for a more recent analysis of this **rent dissipation**). There is no obvious way to model contests, and the functional form does matter to behavior, as Jack Hirshleifer (1989) tells us. In the most popular functional form, P_1 and P_2 are the probabilities of winning of the two players, e_1 and e_2 are their efforts, and R and θ are parameters which can be used to increase the probability that the high bidder wins or to give one player an advantage over the other. The victory function is then assumed to be

$$P_1 = \frac{\theta e_1^R}{\theta e_1^R + e_2^R} \quad \text{and} \quad P_2 = \frac{e_2^R}{\theta e_1^R + e_2^R} \quad (34)$$

If $\theta = 1$ and R becomes large, this becomes close to the simple all-pay auction, because neither player has an advantage and the highest bidder wins with probability near one.

Once we depart from true auctions, however, the modeller must be careful about some seemingly obvious assumptions. Two bidders can simply refuse to enter the dollar auction, for example, but two countries have a harder time refusing to enter a situation in which an arms race is tempting, though they can try to collude in keeping the bids small. It is also often plausible that the size of the prize rises or falls with the bids— for example, when the contest is a mechanism used by a team to motivate its members to produce more (see Chung [1996]) or when the prize shrinks with effort because rent-seeking hurts the economy (see Alexeev & Leitzel [1996]).

13.3 Revenue Equivalence, Risk Aversion, and Uncertainty

We can now collect together the outcomes of these various auction rules and compare them. We have seen that the first-price and descending auctions are strategically equivalent, so the payoffs to the bidders and seller will be the same under each rule regardless of whether values are private or common and whether the players are risk neutral or risk averse.

When values are private and independent, the second-price and ascending auctions are the same in the sense that the bidder who values the object most highly wins and pays the second-highest of the values of all the bidders present, but the strategies are different in the two auctions. In all five kinds of auctions, however, the seller's expected revenue is the same. This is the biggest result in auction theory: the **Revenue Equivalence Theorem** of Vickrey (1961). This has been variously generalized from Vickrey's original statement (e.g., Klemperer (2004, p. 40), so there is no single Revenue Equivalence Theorem, but they all share the same idea of payoffs being the same under various auction rules. We will look at two versions, a general one for auctions with particular properties and a more specific one— a corollary, really— for the five auction rules just analyzed.

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.

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

$$\pi(z, v) = G(z)v - p(z). \quad (35)$$

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

$$\frac{d\pi(z; v)}{dz} = g(z)v - \frac{dp(z)}{dz} = 0, \quad (36)$$

so

$$\frac{dp(z)}{dz} = g(z)v. \quad (37)$$

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

$$\frac{dp(v)}{dv} = g(v)v. \quad (38)$$

Next, we integrate (38) over all values from zero to v , adding $p(\underline{v})$ as the constant of integration:

$$p(v) = p(\underline{v}) + \int_{\underline{v}}^v g(x)xdx. \quad (39)$$

We can use (39) to substitute for $p(v)$ in the payoff equation (35), which becomes, after replacing z with v and setting $p(\underline{v}) = 0$ because of assumption (b),

$$\pi(v, v) = G(v)v - \int_{\underline{v}}^v g(x)xdx. \quad (40)$$

Equation (40) 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.*

Although the different auctions have the same expected payoff for the seller, they do not have the same realized payoff. In the first-price auction, for example, the winning bidder's payment depends entirely on his own value. In the second-price auction the winning bidder's payment depends entirely on the second-highest value, which is sometimes close to his own value and sometimes much less. Thus, we will see that first-price auctions are better if players are risk averse.

Remember, too, that the revenue equivalence theorem requires not just that the bidders have private values, not common, but that the private values be independent. To see why, consider what happens if there are two bidders, both with values drawn uniformly from $[0,10]$, but interdependently, with $v_2 = 10 - v_1$. If we put aside equilibria with weakly dominated strategies (e.g., for a player to bid 0 if his value is less than 5), the second-price auction yields revenue equal to $p = v_{(2)}$, the second-highest value. The seller can extract more revenue, however, by using the auction rule that the winner is the highest bidder, and he pays 10 minus the second-highest bid. One equilibrium under that rule has both players bidding their values, and $p = 10 - v_{(2)} = v_{(1)} > v_{(2)}$.

Risk Aversion in Private-Value Auctions³

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.

³This explanation is adapted from Chapter 8 of the English draft of the Chinese version of Wolfstetter (1999).

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

$$0.5U(6) = U(16 - \bar{p}). \quad (41)$$

Since $0.5U(6) < U(16 - 13)$ by concavity of U , it must be that $\bar{p} > 13$. We found the mixing distribution function $M(p)$ by equating $\pi(p)$ to the payoff from bidding 10, which is $0.5U(6)$, so

$$\pi(p) = 0.5U(16 - p) + 0.5M(p)U(16 - p) = 0.5U(6), \quad (42)$$

which can be solved to yield

$$M(p) = \frac{U(6)}{U(16 - p)} - 1, \quad (43)$$

which has the density

$$m(p) = \frac{U(6)'}{U} (16 - p)U(16 - p)^2, \quad (44)$$

compared with the risk-neutral density $m(p) = \frac{6}{(16-p)^2}$ from equation (23). 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.

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, let us 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

$$\begin{aligned}\pi(v, z) &= G(z)U[v - p(z)] \\ &= F(z)^{n-1}U[v - p(z)],\end{aligned}\tag{45}$$

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

$$\frac{\partial \pi(v, z)}{\partial z} = (n-1)F(z)^{n-2}f(z)U[v - p(z)] + F(z)^{n-1}U'[v - p(z)][-p'(z)] = 0,\tag{46}$$

In equilibrium, $z = v$. Using that fact, for all $v > \underline{v}$ (since $F(\underline{v}) = 0$) we can solve equation (46) for $p(z)$ to get

$$p(v) = \left(\frac{(n-1)f(v)}{F(v)} \right) \left(\frac{U[v - p(v)]}{U'[v - p(v)]} \right)\tag{47}$$

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

$$\frac{U[v - p(v)]}{U'[v - p(v)]} = v - p(v),\tag{48}$$

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

$$\frac{U[v - p(v)]}{U'[v - p(v)]} > v - p(v).\tag{49}$$

Thus, for a given v , the bid function in (47) 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.

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

We have seen that when bidders are risk averse, revenue equivalence fails because the second-price auction is riskier than the first-price auction. By the same reasoning, the ascending auction is riskier, and by strategic equivalence the descending auction is the same as the first-price. Risk aversion matters for different reasons if the bidders do not know their values precisely. As we will see later, in the context of common values, uncertainty over one's own value will generate conservative bidding behavior for the strategic reason of the "Winner's Curse" — a reason which applies whether bidders are risk averse or not — and because of the "Linkage Principle" of Milgrom & Weber (1982). Value uncertainty also has a simpler effect that is driven by risk aversion and applies even in independent private-value auctions: buying a good of uncertain value is intrinsically risky, whether buying it by auction or by a posted price.

Consider the following question:

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

Let us assume that the seller can precommit to reveal both favorable information and unfavorable information, since of course he would like best to reveal only information that raises bidder estimates of the object's value (though the unravelling effect of Chapter 10 might undo such a strategy). It is often plausible that the seller can set up an auction system which reduces uncertainty — say, by a regular policy of allowing bidders to examine the goods before the auction begins. Let us build a model to show the effect of such a policy.

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 $U(v - p)$?

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

$$\pi(p) = \frac{U([\hat{v} - x] - p)}{3} + \frac{U(\hat{v} - p)}{3} + \frac{U([\hat{v} + x] - p)}{3} \quad (50)$$

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. This is true because we can write $U(v - p) = v - p$ and

$$\pi(\text{risk neutral}, p = \hat{v}) = \frac{([\hat{v} - x] - \hat{v})}{3} + \frac{(\hat{v} - \hat{v})}{3} + \frac{(\hat{v} + x - \hat{v})}{3} = 0. \quad (51)$$

Thus, under risk neutrality uncertainty over one's own value does not affect the optimal strategy, except that the bidder's bid ceiling is his expected value for the object rather than his value. Adverse selection aside, there is no reason for the seller to try to improve bidder information.

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

$$\frac{U([\hat{v} - x] - p)}{3} + \frac{U([\hat{v} + x] - p)}{3} < \left(\frac{2}{3}\right) U(\hat{v} - p), \quad (52)$$

so his expected payoff in equation (50) is less than $U(\hat{v} - p)$, and if $p = \hat{v}$ his payoff is less than $U(0)$. A risk-averse bidder will have a negative expected payoff from paying his bid ceiling unless it is strictly less than his estimated value.

Notice that the seller does not have control over all the elements of the model. The seller can often choose the auction rules unilaterally. This includes not just how bids are made, but such things as whether the bidders get to know how many potential bidders are in the auction, whether the seller himself is allowed to bid, and so forth. Also, the seller can decide how much information to release about the goods. The seller cannot, however, decide whether the bidders are risk averse or not, or whether they have common or private values, no more than he can choose what their values are for the good he is selling. All of those assumptions concern the utility functions of the bidders. At best, the seller can do things such as choose to produce goods to sell at the auction which have common values instead of private values.

An error I have often seen is to think that the presence of uncertainty over one's value always causes the Winner's Curse that we will shortly examine. It does not, unless the auction is in common values; uncertainty over one's value is a necessary but not sufficient condition for the Winner's Curse. It is true that risk-averse bidders should not bid as high as their value estimates if they are uncertain about them, even if the auction is in private values. That sounds a lot like a Winner's Curse, but the reason for the discounted bids is completely different, depending as it does on risk aversion. If bidders are uncertain about value estimates but they are risk neutral, their dominant strategy is still to bid up to their value estimates. If the Winner's Curse is present, even if a bidder is risk-neutral he discounts his bid because if he wins, on average his estimate will be greater than the value.

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, an approach developed in Bulow & Roberts (1989) for risk-neutral private values, and Bulow & Klemperer (1996) for common values and risk aversion. This approach compares the seller in an auction to an ordinary monopolist who sells using a posted price. We start with an auction to just one bidder, then extend the idea to an auction with multiple bidders, and finally return to the surprising similarity between an auction with one bidder and a monopoly selling to a continuum of bidders using a posted price.

1. One Bidder. If there is just one bidder, the seller will do badly in any of the auction rules we have discussed so far. The single bidder would bid $p_1 = 0$ and win.

The situation is really better suited to bargaining or simple monopoly than to an auction. The seller could use an auction, but a standard auction yields him zero revenue. so posting a price offer to the bidder makes more sense. If the auction has a reserve price, however, it can be equivalent to posting a price, just as in bargaining the making of a single take-it-or-leave-it offer of p^* is equivalent to posting a price.

What should the 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

$$\begin{aligned}\pi(p^*) &= Pr(p^* < v)(p^* - v_s) + Pr(p^* > v)(0) \\ &= [1 - F(p^*)](p^* - v_s).\end{aligned}\tag{53}$$

This has first-order-condition

$$\frac{d\pi(p^*)}{dp^*} = [1 - F(p^*)] - f(p^*)[p^* - v_s] = 0.\tag{54}$$

On solving (54) for p^* we get

$$p^* = v_s + \left(\frac{1 - F(p^*)}{f(p^*)} \right).\tag{55}$$

The optimal take-it-or-leave-it offer, the “reserve price” p^* satisfies equation (55). 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. One reason to use a reserve price is so the seller does not sell an object for a price worth less than its value to him, but that is not all that is going on.⁴

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

$$\begin{aligned}\pi(p^*) &= Pr(p^* > v_1)(0) + Pr(v_2 < p^* < v_1)(p^* - v_s) + Pr(p^* < v_2 < v_1)(v_2 - v_s) \\ &= \int_{v_1=\underline{v}}^{p^*} f(v_1)(0)dv_1 + \int_{v_1=p^*}^{\bar{v}} \left(\int_{v_2=\underline{v}}^{p^*} (p^* - v_s)f(v_2)dv_2 + \int_{v_2=p^*}^{v_1} (v_2 - v_s)f(v_2)dv_2 \right) f(v_1)dv_1\end{aligned}\tag{56}$$

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.

⁴The second-order condition for the problem is $\frac{d^2\pi(p^*)}{dp^{*2}} = -2f(p^*) + f'(p^*)[p^* - v_s] \leq 0$. This might well be false, and in any case several values of p^* might satisfy equation (55) so it is only a necessary condition, not a sufficient one. Another way to see that $p^* > v_s$ is to observe that $\frac{d\pi(p^*=v_s)}{dp^*} = [1 - F(v_s) - f(v_s)[v_s - v_s]] > 0$, so p^* should be increased beyond v_s .

Now differentiate equation (56) to find the optimal reserve price p^* using Leibniz's integral rule (given in the Mathematical Appendix):

$$\begin{aligned}
\frac{d\pi(p^*)}{dp^*} &= 0 + -f(p^*) \left(\int_{v_2=\underline{v}}^{p^*} (p^* - v_s) f(v_2) dv_2 + \int_{v_2=p^*}^{p^*} (v_2 - v_s) f(v_2) dv_2 \right) \\
&+ \int_{v_1=p^*}^{\bar{v}} \left((p^* - v_s) f(p^*) - (p^* - v_s) f(p^*) + \int_{v_2=\underline{v}}^{p^*} f(v_2) dv_2 \right) f(v_1) dv_1 \\
&= -f(p^*) \left(\int_{v_2=\underline{v}}^{p^*} (p^* - v_s) f(v_2) dv_2 + 0 \right) + \int_{v_1=p^*}^{\bar{v}} \left(\int_{v_2=\underline{v}}^{p^*} f(v_2) dv_2 \right) f(v_1) dv_1 \\
&= -f(p^*) F(p^*) (p^* - v_s) + (1 - F(p^*)) F(p^*) = 0
\end{aligned} \tag{57}$$

Dividing by F , the last line of expression (57) implies that

$$p^* = v_s + \frac{1 - F(p^*)}{f(p^*)}, \tag{58}$$

just what we found in equation (55) for the one-bidder case. Remarkably, the optimal reserve price is unchanged! Moreover, equation (58) applies to any number of bidders, not just $n = 2$. Only Bidders 1 and 2 show up in the equations we used in the derivation, but that is because they are the only ones to affect the result in a second-price auction.

In fact, only the highest-valuing bidder matters to the optimal reserve price. The reserve price only affects the winning price if the second-highest-valuing bidder happens to have a rather low value. Conditioning on that, Bidder 1 would bid low if there were no reserve price— it is almost as if he faced no competition. But the seller, conditioning on that value being low, would want to set a reserve price so that Bidder 1 must pay more. Since the reserve price only matters if all but one of the bidders have low values, it doesn't matter whether “all but one” is 99 or 0. Conditioning on all but one bidder having low values, the seller is conditioning on the game having only one bidder who matters.

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

$$q(p) = 1 - F(p) \tag{59}$$

and

$$\text{Revenue} \equiv pq = p(1 - F(p)) \tag{60}$$

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

$$\begin{aligned} \text{Marginal Revenue} \equiv \frac{dR}{dq} &= p + \left(\frac{dp}{dq}\right) q \\ &= p + \left(\frac{1}{\frac{dq}{dp}}\right) q \\ &= p + \left(\frac{1}{-f(p)}\right) (1 - F(p)) \\ &= p - \frac{1 - F(p)}{f(p)} \end{aligned} \tag{61}$$

Setting marginal revenue to marginal cost, the profit-maximizing monopoly price is the one at which the marginal revenue in (61) equals c .

Does equation (61) look familiar?⁵ Equating (61) to marginal cost, thinking of marginal cost as v_s , the seller's opportunity cost, and moving p to the left-hand-side yields the optimal price equation we found for one bidder in equation (55). That is because the mathematics of the problem is identical whether the seller is facing a continuum of bidders on distribution $F(v)$ or one one bidder drawn randomly from the continuum $F(v)$. The problem is just like that in a take-it-or-leave-it-offer bargaining model where the bidder's type is unknown to the seller. In all three situations— the continuum of bidders, the auction to one seller with a reserve price, and the single offer to a single bidder, the seller is in effect using the basic monopoly pricing rule of setting quantity so that marginal revenue equals marginal cost. The difference is in interpretation. In the auction and bargaining contexts, the marginal change in the number of units of quantity becomes the marginal change in the probability of selling one unit. The seller is still picking quantity, but he is picking it in the interval $[0,1]$ instead of $[0, \infty]$ when he is selling just one unit to one bidder. To increase that probability, and thus the expected number of units sold, the seller must reduce his price, just as an ordinary monopolist must reduce his price to increase the number of units he sells.

⁵As with monopoly in general, it might happen here that marginal revenue equals marginal cost at more than one quantity. The $MR = MC$ rule is only a necessary condition, not a sufficient one, for profit maximization.

Figure 3a shows this. In the auction context, c could represent the seller's production cost, or it could be any other kind of opportunity cost that creates the minimum price at which the seller would part with the good, v_{seller} . The auction seller should act like a monopolist with constant marginal cost of v_{seller} (constant because he is producing just one unit), facing the demand curve based on $f(v)$, which means he should set a reserve price for the quantity where marginal revenue equals marginal cost.

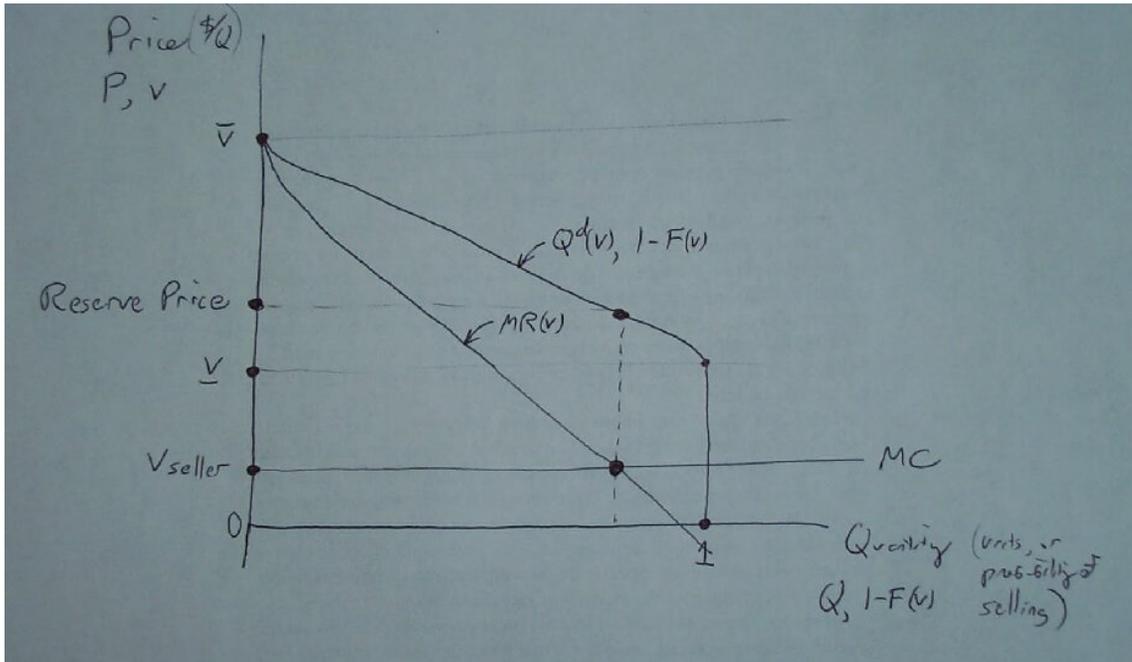


Figure 3a: Auctions and Marginal Revenue: Reserve Price Needed

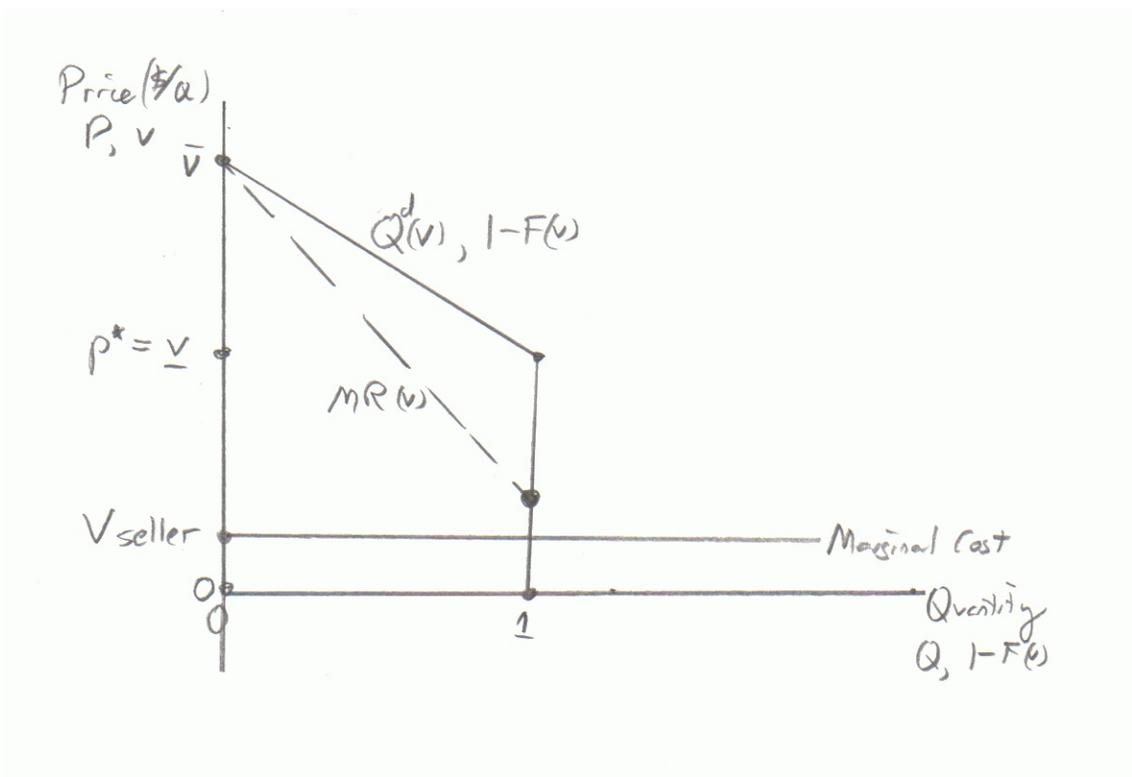


Figure 3b: Auctions and Marginal Revenue: No Reserve Price Needed

The optimal reserve price will always be positive. In Figure 3a, even if v_s were zero instead of positive, the curves are such that the reserve price should be positive, even though the probability of making a sale would then equal one. This corresponds to the idea that a monopolist will always raise the price above marginal cost, even though in certain situations (such as the curves in Figure 3b) he will not reduce output below the competitive level and a reserve price is redundant.

If output is reduced below the competitive level, the outcome is inefficient, something true both in conventional monopoly and here. Here, output is inefficiently low if no sale takes place of the one unit even though $v > v_s$ for some bidder. In that case, what the seller has done is to inefficiently reduce the expected output to below the one unit he has available, resulting in an expected welfare loss equal to the area of a triangle, just as in conventional monopoly.

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 $p^* = v_s$. We have also shown, however, that without a reserve price greater than v_s , none of the five auction rules is optimal for the seller. With the addition of an optimal reserve price, though, it can be shown (though we will not do so here) that in simple settings the seller need use no more complicated auction rules than one of the five we have studied.

In more complicated settings, of course, things do get more complicated, and there is the possibility of inefficiency not just because the object is not sold at all, but because it might be sold to the “wrong” bidder (that is, not to the bidder who values it most). I have already mentioned that this can happen in asymmetric auctions, where the bidders have values drawn from different distributions instead of just one $F(v)$. It could happen that for two bidders, $\left(\frac{1-F_1(v)}{f_1(v)}\right) > \left(\frac{1-F_2(v)}{f_2(v)}\right)$, in which case our rule for setting p^* becomes ambiguous— which bidder’s F function should we use? In such cases, the seller does best if he biases the auction rules in favor of Bidder 2, who will then sometimes win even if $v_2 < v_1$. A similar problem can arise if F is an unusual distribution, for which higher v does not imply that $\left(v - \frac{1-F(v)}{f(v)}\right)$ is higher. For more on the intricacies of such situations, see Myerson (1981) or Bulow & Roberts (1989).

Hindering Bidder Collusion

The choice of auction rule can matter for another reason which comes to mind now that we have been discussing the seller as monopolist: the bidders may act as a monopsonist by colluding. If the bidders can cooperate as if they were one bidder, we are in the situation just described of an auction to one bidder, where use of a reserve price is critical and where the information-eliciting benefit of an auction to the seller evaporates. If the bidders try to collude but remain selfish, however, the choice of an auction rule can make the difference between successful collusion and failure for them. Some auction rules are more vulnerable to collusion than others.

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. (See, too, Graham & Marshall [1987]). 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. (We will not consider the rationality of this choice of bids, which might be based on avoiding legal penalties.) 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. Enforcement is also easy in a second-price auction, because the cartel agreement can be that Smith bids 20 and everyone else bids 6, and if anyone cheats and bids higher than 6 he still loses unless he bids 20 or more.

Unlike in an ascending or second-price auction, however, 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.

Thus, the seller should use a first-price auction if he fears collusion. Even then, cheating will be a problem if the game is repeated, and where collusive "bidding rings" are most often found is in markets where the same bidders meet over and over in auctions for similar objects. Thus, antique auctions are notorious for collusion among bidders who are professional antique dealers (and thus are "regulars" at the auctions and meet repeatedly). A similar problem arises when the same highway contractors repeatedly compete with each other in bidding for who will carry out a government job at the lowest cost. A similar example, described in the book by Sultan (1974), is the Electric Conspiracy of the 1950's, in which antitrust authorities found and prosecuted collusion among executives in a few large companies in bidding for electrical generating equipment contracts.

13.5 Common-Value Auctions and the Winner's Curse

In Section 13.1 we distinguished private-value auctions from common-value auctions, in which the bidders all have the same value for the object but their value estimates may differ. All five sets of rules discussed there can be used for common-value auctions, but the optimal strategies change. In common-value auctions, each bidder can extract useful information about the object's value to himself from the bids of the other bidders. Surprisingly enough, a bidder can use the information from other bidders' bids even in a sealed-bid auction, as will be explained below.

A common-value auction in which all the bidders knew the value would not be very interesting—or very different from a private-value auction—but more commonly the bidders must estimate the common value. The obvious strategy, especially following our discussion of private-value auctions, is for a risk-neutral bidder to bid up to his unbiased estimate of the value. But this strategy makes the winner's payoff negative, because the winner is the bidder who has made the largest positive error in his estimate. The bidders who underestimated the number of pennies lose the auction, but their payoff equals zero, which they would receive even if the true value were common knowledge. Only the winner suffers from his estimation error: he has stumbled into the **Winner's Curse**, a phenomenon first described in Rothkopf (1969) and Wilson (1969).

When some bidders are better informed than others, the Winner's Curse becomes even more severe. Naturally if a bidder knows that he is the worst informed, he should be cautious. Anyone, for example, who outbids 50 experts on the value of an object to win an auction should worry about why all the experts bid less. But the presence of the poorly informed bidder also increases the danger for the experts. Any expert who wins not just against 49 equally well-informed experts but also against a naive bidder who might well have made a large overestimation error and bid too high should worry too.

Once bidders recognize the possibility of the Winner's Curse and adjust their bidding strategies, the winner will no longer have to regret his victory. Having adjusted by scaling down their bids to be lower than their unbiased estimates, the winner may still be the bidder with the biggest overestimation error, but the winning bid can still be less than the true value. Thus, the problem is to decide how much less than one's value estimate to bid.

The mental process is a little like deciding how much to bid in a private-value, first-price auction. There, a bidder wants to bid less than his value, but he also wants to win if he can do so cheaply enough. He therefore tries to estimate the value of the second-highest bid conditional upon himself having the highest value and winning. In a common-value auction a bidder's first step is to use similar mathematics, but to estimate his own value conditional upon winning the auction, not the second-highest value.

One way to think about a bidder's conditional estimate is to think about it as a conditional bid. The bidder 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. Instead, he would like to submit a bid of [X if I lose, but $(X - Y)$ if I win], where X is his value estimate conditional on losing and $(X - Y)$ is his estimate conditional on winning—a lower estimate, since winning implies his overestimation error was the biggest of anybody's. 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 when he loses, the size of his bid is irrelevant.

Another way to look at the Winner’s Curse is based on the Milgrom definition of “bad news” (Milgrom [1981b], appendix b). Suppose the government is auctioning off the mineral rights to a plot of land with common value v and that bidder i has value estimate \hat{v}_i . Suppose also that the bidders are identical in everything but their value estimates, which are based on the various information sets Nature has assigned them, and that the equilibrium is symmetric, so the equilibrium bid function $p(\hat{v}_i)$ is the same for each bidder. If Bidder 1 wins with a bid $p(\hat{v}_1)$ that is based on his prior value estimate \hat{v}_1 , his posterior value estimate \tilde{v}_1 is

$$\tilde{v}_1 = E(V|\hat{v}_1, p(\hat{v}_2) < p(\hat{v}_1), \dots, p(\hat{v}_n) < p(\hat{v}_1)). \quad (62)$$

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

$$\tilde{v}_1 < E(V|\hat{v}_1) = \hat{v}_1, \quad (63)$$

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

Deciding how much to scale down the bid is a hard problem because the amount depends on how much all the other bidders scale down. In a second-price auction a bidder calculates the value of \tilde{v}_1 using equation (62), but that equation hides considerable complexity under the disguise of the term $p(\hat{v}_2)$, which is itself calculated as a function of $p(\hat{v}_1)$ using an equation like (62).

Oil Tracts and the Winner’s Curse

The best known example of the Winner’s Curse is from bidding for offshore oil tracts. Offshore drilling can be unprofitable even if oil is discovered, because something must be paid to the government for the mineral rights. Physicists Capen, Clapp & Campbell suggested in a 1971 paper in the *Journal of Petroleum Engineering* that bidders’ ignorance of what they termed “the Winner’s Curse” caused overbidding in US government auctions of the 1960s. If the oil companies had bid close to what their engineers estimated the tracts were worth, rather than scaling down their bids, the winning companies would have lost on their investments. The hundredfold difference in the sizes of the bids in the sealed-bid auctions shown in Table 1 lends some plausibility to the view that this is what happened.

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	

Later studies such as Mead, Moseidjord & Sorason (1984) that actually looked at profitability concluded that the rates of return from offshore drilling were not abnormally low, so perhaps the oil companies did scale down their bids rationally. The spread in bids is surprisingly wide, but that does not mean that the bidders did not properly scale down their estimates. Although expected profits are zero under optimal bidding, realized profits could be either positive or negative. With some probability, one bidder makes a large overestimate which results in too high a bid even after rationally adjusting for the Winner's Curse. The knowledge of how to bid optimally does not eliminate bad luck; it only mitigates its effects.

Another consideration is the rationality of the other bidders. If bidder Apex has figured out the Winner's Curse, but bidders Brydcox and Central have not, what should Apex do? Its rivals will overbid, which affects Apex's best response. Apex should scale down its bid even further than usual, because the Winner's Curse is intensified against overoptimistic rivals. If Apex wins against a rival who usually overbids, Apex has very likely overestimated the value.

Risk aversion affects bidding in a surprisingly similar way. If all the bidders are equally risk averse, the bids would be lower, because the asset is a gamble, whose value is lower for the risk averse. If Smith is more risk averse than Brown, then Smith should be more cautious for two reasons. The direct reason is that the gamble is worth less to Smith—the reason analyzed above in the private-value setting. The indirect reason is that when Smith wins against a rival like Brown who regularly bids more, Smith probably overestimated the value. Parallel reasoning holds if the bidders are risk neutral, but the private value of the object differs among them.

In fact, Capen, Clapp, and Campbell all worked for the oil company Arco, and developed a bidding strategy for oil leases that took into account the Winner's Curse. Company executives realized that they were bidding too high, but were not clear about whether the solution was to reduce the geological estimates, raise the discount rate, or something else, before the three worked out their strategy. There was debate within the company as to whether to make the idea of the Winner's Curse public. The idea was a private advantage for Arco— but other oil company's ignorance of it meant that they were bidding too high, which may have hurt them most, but hurt Arco too. After some years, the company decided to let them reveal the secret (See "The Tale of the 'Winner's Curse' " at <http://www.aapg.org/explorer/2004/12dec/capen.cfm>)

The Winner's Curse crops up in situations seemingly far removed from auctions. An employer should beware of hiring a worker passed over by other employers. Someone renting an apartment should worry that he is the first potential renter who arrived when the neighboring trumpeter was asleep. A firm considering a new project should be concerned that the project has been considered and rejected by competitors. The Winner's Curse can even be applied to political theory; certain proposals for innovations keep reappearing in the political arena over time. Will the first electorate to adopt them fall prey to the Winner's Curse?

On the other hand, if information is revealed in the course of an auction, the fact that values are common can lead to higher bidding, not lower. "Getting carried away" may be a rational feature of a common-value auction. Suppose that the setting is not pure common value, but a mix of a common value and private values. If a bidder has a high private value and then learns in the course of the bidding that the common value is larger than he thought, he may well end up bidding more than he had planned and winning, but he would not regret it afterwards. Especially in these mixed situations, when bidders have different private values and when some but not all of each bidder's information is revealed by his bidding, common-value auctions can become very complicated indeed.

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. (In fact, 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.)

We will not prove the ranking of revenues by auction generally, but we will go through an example with estimation errors that are uniformly distributed. This example will be tractable because of some special properties of optimal estimates when errors are uniformly distributed, so let's start with discussion of those properties.

Suppose n signals are independently drawn from the uniform distribution on $[\underline{s}, \bar{s}]$. Denote the j^{th} highest signal by $s_{(j)}$. The expectation of the k th highest value happens to be

$$Es_{(k)} = \underline{s} + \left(\frac{n+1-k}{n+1} \right) (\bar{s} - \underline{s}) \quad (64)$$

This means the expectation of the very highest value is

$$Es_{(1)} = \underline{s} + \left(\frac{n}{n+1} \right) (\bar{s} - \underline{s}) \quad (65)$$

The expectation of the second-highest value is

$$Es_{(2)} = \underline{s} + \left(\frac{n-1}{n+1} \right) (\bar{s} - \underline{s}) \quad (66)$$

The expectation of the lowest value, the n 'th highest, is

$$Es_{(n)} = \underline{s} + \left(\frac{1}{n+1} \right) (\bar{s} - \underline{s}). \quad (67)$$

Let n risk-neutral bidders, $i = 1, 2, \dots, n$ 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. The best estimate of the value given the set of n signals is

$$Ev|(s_1, s_2, \dots, s_n) = \frac{s_{(n)} + s_{(1)}}{2}. \quad (68)$$

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. A density that had a peak, like the normal density, would yield a different result, but here all we can tell from the data is that all values of v between $(6 + m)$ and $(24 - m)$ are equally probable.

Figure 4 illustrates why this is true. Someone who saw just signals $s_{(n)}$ and $s_{(1)}$ could deduce that v could not be less than $(s_{(1)} - m)$ or greater than $(s_{(n)} + m)$. Learning the signals in between— would be unhelpful, because the only information that, for example, $s_{(2)}$ conveys is that $v \leq (s_{(2)} + m)$ and $v \geq (s_{(2)} - m)$, facts which our observer had already figured out from $s_{(n)}$ and $s_{(1)}$.

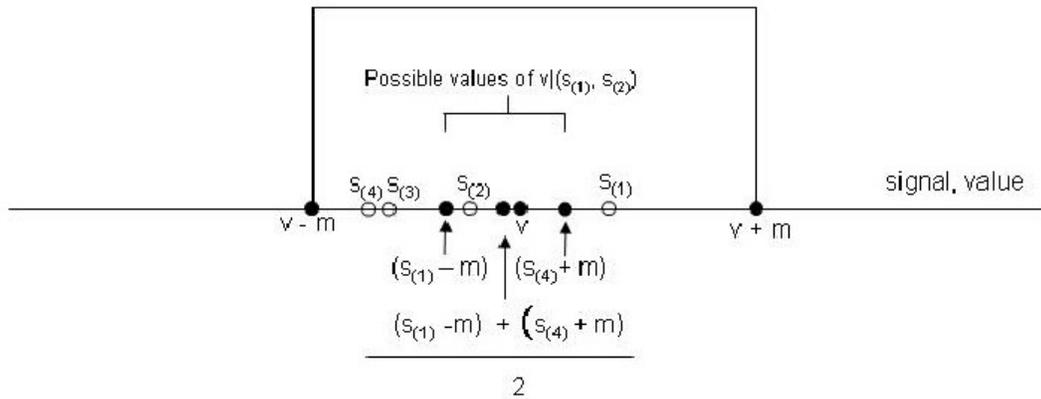


Figure 4: Extracting Information From Uniformly Distributed Signals

Now let us go to the game itself, which is based on the introductory chapter of Paul Klemperer's 2000 book.

The Uniform-Signal Common-Value Auction

Players:

One seller and n bidders.

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.

Payoffs:

Payoff depends on the particular rules, but if the object is sold, the payoff to the seller is the sum of all the payments and the value to a bidder is the value of the object, if he wins, minus his payments.

What are the strategies in symmetric equilibria for the different auction rules? (we will ignore possible asymmetric equilibria.)

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.

The price paid by the winner will be the price at which the second-highest bidder drops out, which is $\frac{s_{(n)} + s_{(2)}}{2}$. The expected values are, from equations (66) and (67),

$$\begin{aligned} E s_{(n)} &= (v - m) + \left(\frac{n+1-n}{n+1}\right) ((v + m) - (v - m)) \\ &= v + \left(\frac{1-n}{n+1}\right) m \end{aligned} \tag{69}$$

and

$$\begin{aligned} E s_{(2)} &= (v - m) + \left(\frac{n+1-2}{n+1}\right) ((v + m) - (v - m)) \\ &= v + \left(\frac{n-3}{n+1}\right) m. \end{aligned} \tag{70}$$

Averaging them yields the expected winning price,

$$\begin{aligned} E p_{(2)} &= \frac{[v + \left(\frac{1-n}{n+1}\right)m] + [v + \left(\frac{n-3}{n+1}\right)m]}{2} \\ &= v - \left(\frac{1}{2}\right) \left(\frac{1}{n+1}\right) 2m. \end{aligned} \tag{71}$$

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

$$E p_{(2)} = v - \left(\frac{1}{10}\right) (100) = v - 10. \tag{72}$$

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). This will be true for all three auction rules we examine here.

It is not always true that the bidders can deduce the lowest signal in an ascending auction and use that to form their bid. Their ability to discover $s_{(n)}$ depended crucially on the open-exit feature of the auction— that the player with the lowest signal had to openly drop out, rather than lurk quietly in the background. A secret-exit ascending auction would behave like a second-price auction instead.

The Second-Price Auction

Equilibrium: Bid $p_i = s_i - \binom{n-2}{n} m$.

Explanation: In forming his strategy, 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. On average, if this happens,

$$\begin{aligned} s_i &= (v-m) + \binom{((n-1)+1-1)}{[n-1]+1} ([v+m] - [v-m]) \\ &= (v-m) + \binom{n-1}{n} (2m) \\ &= v + \binom{n-2}{n} (m). \end{aligned} \tag{73}$$

He will bid the value v which solves equation (73), yielding the optimal strategy, $p_i = s_i - \binom{n-2}{n} (m)$.

On average, the second-highest bidder actually has the signal $Es_{(2)} = v + \binom{n-3}{n+1} m$, from equation (70). So the expected price, and hence the expected revenue from the auction, is

$$\begin{aligned} Ep_{(2)} &= [v + \binom{n-3}{n+1} m] - \binom{n-2}{n} (m) \\ &= v + \binom{n(n-3)-(n+1)(n-2)}{(n+1)n} m, \\ &= v - \binom{n-1}{n} \binom{1}{n+1} 2m. \end{aligned} \tag{74}$$

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

$$Ep_{(2)} = v - \binom{3}{4} \binom{1}{5} (100) = v - 15. \tag{75}$$

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 in the second-price and ascending auctions. Then, $v_{(n)} = v_{(2)}$, so the winning price is based on the same information in both auctions.

The First-Price Auction

Equilibrium: Bid $(s_i - m)$.

Explanation: This is the simplest strategy of the three, but the hardest to derive. 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

$$T_i \equiv s_i - (v - m) \quad (76)$$

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 $(z + m - T_i)$.

If, instead, 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

$$G(T_i + \epsilon) - G(T_i) \approx (n - 1)T_i^{n-2} \left(\frac{1}{2m}\right)^{n-1} \epsilon. \quad (77)$$

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

$$\int_{T_i=0}^{2m} \left[\left((n - 1)T_i^{n-2} \left(\frac{1}{2m}\right)^{n-1} \epsilon \right) (z + m - T_i) - \epsilon \left(\frac{T_i}{2m}\right)^{n-1} \right] dT_i = 0. \quad (78)$$

This implies that

$$\epsilon \left(\frac{1}{2m}\right)^{n-1} \int_{T_i=0}^{2m} [(n - 1)T_i^{n-2} (z + m) - (n - 1)T_i^{n-1} - T_i^{n-1}] dT_i = 0. \quad (79)$$

which in turn implies that

$$\epsilon \left(\frac{1}{2m}\right)^{n-1} \Big|_{T_i=0}^{2m} (T_i^{n-1} (z + m) - T_i^n) = 0, \quad (80)$$

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 $p_i = s_i - m$.

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

$$\begin{aligned}
 E s_{(1)} &= \underline{s} + \left(\frac{n+1-1}{n+1}\right) (\bar{s} - \underline{s}) \\
 &= v - m + \left(\frac{n}{n+1}\right) ((v + m) - (v - m)) \\
 &= v - m + \left(\frac{n}{n+1}\right) (2m)
 \end{aligned} \tag{81}$$

The expected revenue is therefore

$$E p_{(1)} = v - (1) \left(\frac{1}{n+1}\right) 2m. \tag{82}$$

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

$$E p_{(1)} = v - \left(\frac{1}{5}\right) (100) = v - 20. \tag{83}$$

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

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.

13.6 Asymmetric Equilibria, Affiliation, and Linkage: The Wallet Game

Asymmetric Equilibria in Common-Value Auctions

Besides the symmetric equilibria I have been discussing so far, asymmetric equilibria are typical, robust, and plausible in common-value auctions. That is because the severity of the winner's curse facing Bidder i depends on the bidding behavior of the other bidders. If other bidders bid aggressively, then if i wins anyway, he must have a big overestimate of the value of the object. So the more aggressive are the other bidders, the more conservative ought Bidder i to be— which in turn will make the other bidders more aggressive. The Wallet Game of Klemperer (1998) illustrates this.

The Wallet Game

Players

Smith and Jones.

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

$$s_1 + s_2 - \text{Min}(p_1, p_2) \tag{84}$$

A symmetric 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 $p > s_i + s_j$.

This equilibrium clearly illustrates how bidders should base their strategy on the strategy they expect the other bidders to use. Note that Bidder's i 's strategy is unrelated to his prior beliefs about Bidder j 's value. It might be, for example, that using $f_2(s_2)$, the expected value of s_2 is 100, but if Bidder 1's wallet contains $s_1 = 8$, he should just bid 16. So doing, he will probably lose, but if he bids 108 and wins, it will only be because Bidder 2's wallet contains 54 or less. That is fine if it contains 8 or less, but if it contains, for example, $s_2 = 9$, then Bidder 2 will bid up to 18 and stop, Bidder 1 will win with his bid ceiling of 108, and his payoff will be $8 + 9 - 18 = -1$.

The Wallet Game has both independent values and pure common values, a curious combination. The value of the two wallets is the same for both bidders, but the signal each receives is independent. Knowing s_1 is useless in predicting s_2 , though it is useful in predicting the common value.

The independence of the signals makes a bidder's optimal strategy particularly sensitive to what he thinks the other bidder's strategy is, since his own signal tells him nothing about the other player's signal and he must rely on the other player's bidding for any information about it. Thus, there are many asymmetric equilibria. One of them is $(p_1 = 10s_1, p_2 = \frac{10}{9}s_2)$. If the two players tie, having chosen $p = p_1 = p_2$, then $10s_1 = \frac{10}{9}s_2$, which implies that $s_1 = \frac{1}{9}s_2$, so $s_1 + s_2 = 10s_1 = p$, and $v = p$. This is a bad equilibrium for Bidder 2 because he hardly ever wins and when he does win it's because s_1 was very low—so there is hardly any money in Bidder 1's wallet. Being the aggressive bidder in an equilibrium is valuable. If there is a sequence of auctions, this means establishing a reputation for aggressiveness can be worthwhile, as shown in Bikhchandani (1988).

Asymmetric equilibria can even arise when the players are identical. Second-price, two-person, common-value auctions usually have many asymmetric equilibria besides the symmetric equilibrium we have been discussing (see Milgrom [1981c] and Bikhchandani [1988]). Suppose that Smith and Brown have identical payoff functions, but Smith thinks Brown is going to bid aggressively. The winner's curse is intensified for Smith, who would probably have overestimated if he won against an aggressive bidder like Brown, so Smith bids more cautiously. But if Smith bids cautiously, Brown is safe in bidding aggressively, and there is an asymmetric equilibrium. For this reason, acquiring a reputation for aggressiveness is valuable.

Oddly enough, if there are three or more bidders the second-price, common-value auction has a unique equilibrium, which is also symmetric. The open-exit ascending auction is different: it has asymmetric equilibria, because after one bidder drops out, the two remaining bidders know that they are alone together in a subgame which is a two-bidder auction. Regardless of the number of bidders, first-price auctions do not have this kind of asymmetric equilibrium. Threats in a first-price auction are costly because the high bidder pays his bid even if his rival decides to bid less in response. Thus, a bidder's aggressiveness is not made safer by intimidation of another bidder.

Affiliation, the Monotone Likelihood-Ratio Property, and the Linkage Principle⁶

Milgrom & Weber (1982) introduced the idea of **affiliation**: a formal definition of two variables tending to move upwards together that is useful in the auction context. Suppose bidder 1 has a value v_1 which is an increasing function of a private signal x_1 that he receives, and which might also depend on the private signal of bidder 2, x_2 . What we need to know to analyze the auction is what happens to bidder 1's estimate of his value as he observes or deduces more about bidder 2's signal. A simple and plausible situation is that whenever he learns that x_2 takes a large value, his estimate of his value v_1 rises. Thus, if he observes bidder 2 bid more and he deduces that x_2 is large, he should increase his estimate of v_1 . Affiliation is defined so this will happen if the signals x_1 and x_2 are strongly affiliated, and if they are weakly affiliated, bidder 1's estimate of v_1 at least will not fall as x_2 rises. For simplicity, I will define affiliation for the case of two signals.

⁶Much of this discussion is based on Cramton (undated), Klemperer (2004, p. 51) and Wolfstetter (1999).

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 ,

$$f(x_1 = Small, x_2 = Low)f(x_1 = Big, x_2 = High) \geq f(x_1 = Small, x_2 = High)f(x_1 = Big, x_2 = Low) \quad (85)$$

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. Notice that this allows a joint probability distribution such as the following, in which even if x_2 is *Low*, x_1 is probably *Big*.

$$\begin{aligned} Prob(x_1 = Big, x_2 = High) &= 0.5 \\ Prob(x_1 = Small, x_2 = Low) &= 0.1 \\ Prob(x_1 = Big, x_2 = Low) &= 0.2 \\ Prob(x_1 = Small, x_2 = High) &= 0.2 \end{aligned}$$

Imagine that Bidder 3, who is ignorant of both x_1 and x_2 , can deduce about x_1 if he learns that $x_2 = Low$. Bidder 3's prior is

$$\begin{aligned} Prob(x_1 = Big) &= Prob(x_1 = Big, x_2 = High) + Prob(x_1 = Big, x_2 = Low) \\ &= 0.5 + 0.2 = 0.7. \end{aligned} \quad (86)$$

Suppose Bidder 3 finds out that x_2 is *Low*. We will use the same kind of manipulation as in bayes's rule, though not bayes's rule itself. Since

$$Prob(x_1 = Big, x_2 = Low) = Prob(x_1 = Big|x_2 = Low)Prob(x_2 = Low), \quad (87)$$

we can write Bidder 3's posterior as

$$\begin{aligned} Prob(x_1 = Big|x_2 = Low) &= \frac{Prob(x_1=Big,x_2=Low)}{Prob(x_2=Low)} \\ &= \frac{0.2}{Prob(x_1=Big,x_2=Low)+Prob(x_1=Small,x_2=Low)} \\ &= \frac{0.2}{0.2+0.1} \end{aligned} \quad (88)$$

Thus, observing $x_2 = Low$ leaves Bidder 3 still thinking that probably $x_1 = Big$, but the probability has fallen from 0.7 to 0.66. $x_2 = Low$ is bad news about the value of x_1 .

The big 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.

This is very much like the Monotone Likelihood-Ratio Property, which 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 ,

$$\frac{g(Big|Low)}{g(Big|High)} \leq \frac{g(Small|Low)}{g(Small|High)}. \quad (89)$$

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 .

This no doubt leaves the reader's head spinning quite as much as it does the author's, despite my attempt at simplification. A final, equivalent definition of affiliation, applicable when the signals are distributed according to a joint density $f(x_1, x_2)$ that is continuous and twice differentiable is that x_1 and x_2 are affiliated if

$$\frac{\partial \log(f)^2}{\partial x_1 \partial x_2} \geq 0. \quad (90)$$

One of the rewards of establishing that bidders' signals are affiliated is **the linkage principle**, which says that when the amount of affiliated information available to bidders increases, the equilibrium sales price becomes greater. Thus, the seller should have a policy of disclosing any affiliated information he possesses. Also, 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. We saw this in the Uniform-Signal Common-Value Auction examples. The result is not restricted to common-value auctions, however; seller revenue rises when affiliated information is released even in an affiliated private value auction, in which the bidders each know their own private values but not those of other bidders. The intuition behind the linkage principle is hard to grasp, because it applies even when there is no winner's curse (as in the affiliated private value auction), and it does not always apply in common value auctions (when there are just two bidders, the ascending auction is not superior to the second-price auction, as we saw above). The best intuition I have seen is that on page 128 of Klemperer (2004): that bidder profits arise from their private information, and release of affiliated information reveals something about the private information of each bidder, including, especially, the one who will win, and thus heightens competition.

The linkage principle provides a reason why sellers may wish to extend auctions over time in multiple rounds, why they should encourage active bidding throughout rather than let bidders "lurk" and suddenly bid near the end, why they should lay out their own information as clearly and early as possible, and why they should let bidders know each others' identities. The idea is a slippery one, however, and Perry & Reny (1999) show that the linkage principle actually can fail if more than one unit of the good is being auctioned, so players submit bids for the possible purchase of multiple units. Multiple-unit auctions and "package auctions," in which not just multiple units but multiple objects are sold in a single auction, are active areas of research.

Notes

N13.1 Values Private and Common, Continuous and Discrete

- Milgrom & Weber (1982) is a classic article that covers many aspects of auctions. McAfee & McMillan (1987) is an excellent older survey of auction theory which takes some pains to relate it to models of asymmetric information. More recent is Maskin (2004). Klemperer (2000) collects many of the most important articles in an edited volume. Vijay Krishna's 2002 *Auction Theory*, Paul Klemperer's 2004 *Auctions: Theory and Practice* (which is relatively nontechnical), and Paul Milgrom's 2004 *Putting Auction Theory to Work* (which is very good on the mathematical assumptions) are good textbook treatments. I particularly like the auction chapters in Elmar Wolfstetter's 1999 *Topics in Microeconomics: Industrial Organization, Auctions, and Incentives*. Paul Milgrom's consulting firm, Agora Market Design, has a website with many good working papers that can be found via <http://www.market-design.com>.
- Cassady (1967) is an excellent source of pre-Web institutional detail. The appendix to his book includes advertisements and sets of auction rules, and he cites numerous newspaper articles. The rise of the Web, an ideal setting for auctions, fortuitously occurred at the same time as the rise of auction theory. See Bajari & Hortacsu's 2004 survey, "Economic Insights from Internet Auctions."

N13.2 Optimal Strategies under Different Rules in Private-Value Auctions

- Many (all?) leading auction theorists were involved in the seven-billion dollar spectrum auction by the United States government in 1994, either helping the government choose an auction rule to sell spectrum or helping bidders decide how to buy it. Paul Milgrom's 1999 book, *Auction Theory for Privatization*, tells the story. See also McAfee & McMillan (1996). Interesting institutional details have come in the spectrum auctions and stimulated new theoretical research. Ayres & Cramton (1996), for example, explore the possibility that affirmative action provisions designed to help certain groups of bidders may have actually increased the revenue raised by the seller by increasing the amount of competition in the auction.
- One might think that an ascending second-price, open-cry auction would come to the same results as an ascending first-price, open-cry auction, because if the price advances by ϵ at each bid, the first and second bids are practically the same. But the second-price auction can be manipulated. If somebody initially bids \$10 for something worth \$80, another bidder could safely bid \$1,000. No one else would bid more, and he would pay only the second price: \$10.
- Auctions are especially suitable for empirical study because they are so stylized and generate masses of data. Hendricks & Porter (1988) is a classic comparison of auction theory with data. See Bajari, Hong & Ryan (2004) or the Athey & Haile (2005) and Hendricks & Porter (forthcoming) surveys of empirical work on auctions.
- After the last bid of an open-cry art auction in France, the representative of the Louvre has the right to raise his hand and shout "pre-emption de l'etat," after which he takes the painting at the highest price bid (*The Economist*, May 23, 1987, p. 98). How does that affect the equilibrium strategies? What would happen if the Louvre could resell?

- **Share Auctions.** In a share auction each bidder submits a bid for both a quantity and a price. The bidder with the highest price receives the quantity for which he bid at that price. If any of the product being auctioned remains, the bidder with the second-highest price takes the quantity he bid for, and so forth. The rules of a share auction can allow each bidder to submit several bids, often called a **schedule** of bids. The details of share auctions vary, and they can be either first-price or second-price. Modelling is complicated; see Wilson (1979).

N13.3 Revenue Equivalence, Risk Aversion, and Uncertainty

- Che & Gale (1998) point out that if bidders differ in their willingness to pay in a private value auction because of budget constraints rather than tastes then the revenue equivalence theorem can fail. The following example from page 2 of their paper shows this. Suppose two budget-constrained bidders are bidding for one object. In Auction 1, each bidder has a budget of 2 and knows only his own value, which is drawn uniformly from $[0,1]$. The budget constraints are never binding, and it turns out that the expected price is $1/3$ under either a first-price or a second-price auction. In Auction 2, however, each bidder knows only his own budget, which is drawn uniformly from $[0,1]$, and both have values for the object of 2. The budget constraint is always binding, and the equilibrium strategy is to bid one's entire budget under either set of auction rules. The expected price is still $1/3$ in the second-price auction, but now it is $2/3$ in the first-price auction. The seller therefore prefers to use a first-price auction.
- **A Mechanism To Extract All the Surplus.** Myerson (1981) shows that if the bidders' private information is correlated, the seller can construct something akin to a Maskin matching scheme mechanism of the kind discussed in Chapter 10 that extracts all the information and all the surplus. In the Uniform-Signal Common-Value Auction, where signals are uniform in $[v - m, v + m]$ ask Bidder i to bid s_i , allocate the good to the high bidder at the price $\frac{s_{(1)} + s_{(n)}}{2}$, which is an unbiased estimate of v , and ensure truth-telling by the boiling-in-oil punishment of a large negative payment if the reports are such that $s_{(n)} < s_{(1)} - m$, which cannot possibly occur if all bidders tell the truth.

N13.4 Common-Value Auctions and the Winner's Curse

Rothkopf (1969) and Wilson (1969) seem to be the first published accounts of the Winner's Curse. An article on Edward Capen, "The Tale of the 'Winner's Curse'" at <http://www.aapg.org/explore> says that the term was first published in Capen, Clapp & Campbell (1971). I recommend the article for its tale of the use of theory in business practice.

- The Winner's Curse and the idea of common values versus private values have broad application. The Winner's Curse is related to the idea of "regression to the mean" discussed in Section 2.4 of this book. Kaplow & Shavell (1996) use the idea to discuss property versus liability rules, one of the standard rule choices in law-and-economics. If someone violates a property rule, the aggrieved party can undo the violation, as when a thief is required to surrender stolen property. If someone violates a liability rule, the aggrieved party can only get monetary compensation, as when someone who breaches a contract is required to pay damages to the aggrieved party. Kaplow and Shavell argue that if a good has independent values, a liability rule is best because it gives efficient incentives for rule violation; but if it has common value and courts make errors in measuring the common value, a property rule may be better. See especially page 761 of their article.

N13.6 Asymmetric Equilibria, Affiliation, and Linkage: The Wallet Game

- Even if value estimates are correlated, the optimal bidding strategies can still be the same as in private-value auctions if the values are independent. If everyone overestimates their values by 10 percent, a bidder can still extract no information about his value by seeing other bidders' value estimates.

Problems

13.1. Rent Seeking (medium)

Two risk-neutral neighbors in sixteenth century England, Smith and Jones, have gone to court and are considering bribing a judge. Each of them makes a gift, and the one whose gift is the largest is awarded property worth £2,000. If both bribe the same amount, the chances are 50 percent for each of them to win the lawsuit. Gifts must be either £0, £900, or £2,000.

- What is the unique pure-strategy equilibrium for this game?
- Suppose that it is also possible to give a £1500 gift. Why does there no longer exist a pure-strategy equilibrium?
- What is the symmetric mixed-strategy equilibrium for the expanded game? What is the judge's expected payoff?
- In the expanded game, if the losing litigant gets back his gift, what are the two equilibria? Would the judge prefer this rule?

13.2. The Founding of Hong Kong (medium)

The Tai-Pan and Mr. Brock are bidding in an ascending auction for a parcel of land on a knoll in Hong Kong. They must bid integer values, and the Tai-Pan bids first. Tying bids cannot be made, and bids cannot be withdrawn once they are made. The direct value of the land is 1 to Brock and 2 to the Tai-Pan, but the Tai-Pan has said publicly that he wants it, so if Brock gets it, he receives 5 in "face" and the Tai-Pan loses 10. Moreover, Brock hates the Tai-Pan and receives 1 in utility for each 1 that the Tai-Pan pays out to get the land.

- First suppose there were no "face" or "hate" considerations, just the direct values. What are the equilibria if the Tai-pan bids first?
- Continue supposing there were no "face" or "hate" considerations, just the direct values. What are the three possible equilibria if Mr. Brock bids first? (Hint: in one of them, Brock wins; in the other two, the Tai-pan wins.)
- Fill in the entries in Table 2, including the "face" and "hate" considerations.

Table 2: The Tai-Pan Game

Winning bid:	1	2	3	4	5	6	7	8	9	10	11	12
If Brock wins:												
π_{Brock}												
$\pi_{Tai-Pan}$												
If Brock loses:												
π_{Brock}												
$\pi_{Tai-Pan}$												

- (d) In equilibrium, who wins, and at what bid?
- (e) What happens if the Tai-Pan can precommit to a strategy?
- (f) What happens if the Tai-Pan cannot precommit, but he also hates Brock, and gets 1 in utility for each 1 that Brock pays out to get the land?

13.3. Government and Monopoly (medium)

Incumbent Apex and potential entrant Brydox are bidding for government favors in the widget market. Apex wants to defeat a bill that would require it to share its widget patent rights with Brydox. Brydox wants the bill to pass. Whoever offers the chairman of the House Telecommunications Committee more campaign contributions wins, and the loser pays nothing. The market demand curve for widgets is $P = 25 - Q$, and marginal cost is constant at 1.

- (a) Who will bid higher if duopolists follow Bertrand behavior? How much will the winner bid?
- (b) Who will bid higher if duopolists follow Cournot behavior? How much will the winner bid?
- (c) What happens under Cournot behavior if Apex can commit to giving away its patent freely to everyone in the world if the entry bill passes? How much will Apex bid?

13.4. An Auction with Stupid Bidders (hard)

Smith's value for an object has a private component equal to 1 and another component Z that is common with Jones and Brown. Jones's and Brown's private components both equal zero. Each bidder estimates Z independently. Bidder i 's estimate is either x_i above the true value or x_i below, with equal probability. Jones and Brown are naive and always bid their value estimates. The auction is ascending. Smith knows all three values of x_i , but not whether his estimate is too high or too low.

- (a) If $x_{Smith} = 0$, what is Smith's dominant strategy if his estimate of Z is 20?
- (b) If $x_i = 8$ for all bidders and Smith estimates that $Z = 20$, what are the probabilities that he puts on different possible values of Z ?
- (c) If $x_i = 8$ for Jones and Brown but $x_{Smith} = 0$, and Smith knows that $Z = 12$ with certainty, what are the probabilities he puts on the different combinations of bids by Jones and Brown?
- (d) Why is 9 a better upper limit on bids for Smith than 21, if his estimate of Z is 20, and $x_i = 8$ for all three bidders?
- (e) Suppose Smith could pay amount 0.001 to explain optimal bidding strategy to his rival bidders, Jones and Brown. Would he do so?

13.5. A Teapot Auction with Incomplete Information (easy)

Smith believes that Brown's value v_b for a teapot being sold at auction is 0 or 100 with equal probability. Smith's value of $v_s = 400$ is known by both bidders.

- (a) What are the bidders' equilibrium strategies in an open cry auction? You may assume that in case of ties, Smith wins the auction.
- (b) What are the bidders' equilibrium strategies in a first-price sealed-bid auction? You may assume that in case of ties, Smith wins the auction.
- (c) Now let $v_s = 102$ instead of 400. Will Smith use a pure strategy? Will Brown? You need not find the exact strategies used.

Auctions: A Classroom Game for Chapter 13

The instructor will bring to class a glass jar of pennies to be auctioned off. There are between 0 and 100 pennies in the jar, the number being drawn from a uniform distribution. Twenty percent of the students in the class must close their eyes while the jar is being displayed. The instructor will pass the jar around to the other students in the class to let them try to figure out how many pennies are inside. He will then auction off the jar five times, using five different sets of rules. (The instructor will decide for himself whether to play for real money or not.)

The first auction will be a first-price auction. Each student submits a bid, and also record his estimate of the number of pennies.

The second auction will be a second-price auction. Each student submits a bid, and also record his estimate of the number of pennies.

The third auction will be an all-pay auction. Each student submits a bid, and also record his estimate of the number of pennies.

The fourth auction will be a descending auction. After the auction, each student submits his estimate of the number of pennies.

The fifth auction will be an ascending auction. After the auction, each student submits his estimate of the number of pennies.

At this point, the instructor will announce the number of pennies in the jar, and the results of each auction.