Auctions and Bidding: A Primer

Paul Milgrom

The Journal of Economic Perspectives is currently published by American Economic Association.

Auctions and Bidding: A Primer

Paul Milgrom

A painting contractor once complained to me that the jobs put up for competitive bids are unlike other painting jobs.

> I do most of my work for a few builders that I've known for years. My estimates of what it will cost to do a job for one of them come out about right. Sometimes a little high, sometimes low, but about right overall. Occasionally, when business is slow, I bid on a big job for another builder, but those jobs are different: They always run more than I expect.

Maybe the contractor was right to think bid jobs are different, but it is more likely that he suffered from too simple a view of what is involved in preparing a competitive bid. Our analysis will show that even an experienced estimator working in familiar terrain can lose money if he doesn't understand the subtleties of competitive bidding.

The phenomenon experienced by the painting contractor, known as the "Winner's Curse," is just one of the surprising and puzzling conclusions that have been turned up by modern research into auctions. Another is the theoretical proposition (supported also by some experimental evidence) that, for example, a sealed-bid Treasury bill auction in which each buyer pays a price equal to the highest rejected bid would yield more revenue to the Treasury than the current procedure in which the winning bidder pays the seemingly higher amount equal to his own bid. There are also subtle results that demonstrate the equivalence of such apparently different institutions as the standard sealed-bid auction, in which the auctioneer/seller sells the goods to the highest bidder for a price equal to his bid, and the Dutch auction, in which the auctioneer/seller begins by asking a high price and gradually lowers the price until

■ *Paul Milgrom is Professor of Economics, Stanford University, Stanford, California.*

some bidder shouts "Mine" to claim the item. Other results explain the use of standard auctions as the selling schemes that maximize the welfare of the bid–taker, or as schemes that lead to efficient allocations, minimize transaction costs, guard against corruption by the bid-taker's agents or mitigate the effects of collusion among the bidders. Finally, for some environments, the theory makes sharp, testable predictions about the bids and profits of various classes of bidders. This paper relies mainly on theory to study these issues, but it will also review some experimental evidence and recent empirical studies testing the predictions of the theory.

Pitfalls for Bidders

One of the earliest operations research studies of competitive bidding was made by Lawrence Friedman (1956). He argued that a bidder should study the past behavior of its competitors to discover the patterns that governed their bidding. This information could be used, in any particular competition, to estimate the probability distribution of any particular competitor i's bid b_i, as represented by the function describing how likely the bid b_i is to be less than any particular amount b. This is the cumulative distribution function: $F_i(b) = \text{Prob}(b_i \leq b)$.

Suppose that the bidders are vying for a road construction contract and that the usual rules of sealed bid auctions apply, so that it is the lowest bid that wins. Examine the problem of just one of the contractors. If it bids b and wins and if its cost of completing the contract is c, then its profits will be $b - c$; if it loses the auction, its profits will be zero. The contractor's bid of b will win precisely when *all* the other contractors make higher bids. The probability that the ith competing contractor makes a higher bid is $1 - F_i(b)$. Thus, if there are N other bidders, the probability that a bid of b beats them all is $P(b) = (1 - F_1(b)) \ldots (1 - F_N(b))$ and the contractor's expected profit is equal to that probability times the profit margin in the bid: $P(b)(b - c)$. Friedman recommends that the bid b be chosen to maximize that expression.

This expression for expected profits depends on two important assumptions. One is that the bids made by competitors are statistically independent (the *independence assumption*) and that their distributions can be somehow estimated from history. The other is that the bidder actually knows the amount c that it will cost him to complete the contract (the *private values assumption*). The independence assumption means that there is no unobserved common factor affecting all of the competitor's bids while the private values assumption allows the contractor to ignore the competitors' information in forming its cost estimate. These assumptions make collecting data, constructing the model, and solving the optimization problem easy, but they may often fail to portray the auction environment accurately.

An alternative model that violates both of these assumptions was used by ARCO in preparing bids for offshore oil tracts sold by the U.S. government. The detailed logic of their model has been described by Capen, Clapp, and Campbell (1971). As applied to the problem of bidding for construction contracts, ARCO's model would

say that a contractor doesn't really know what the job will cost, and regards the actual cost C of the job as a random variable. Uncertainties about cost arise from uncertainties about factors that will affect all bidders like the number of tons of concrete that will be required, blasting difficulties, cold weather construction delays, changing factor prices, and so on in addition to factors that are idiosyncratic and vary from bidder to bidder. Each contractor's cost estimate is just an estimate, subject to error. No contractor knows what its cost will be and each realizes that the other bidders may possess information or analyses that the contractor would find useful for its own cost estimation, so the private values assumption fails.

The simplest way to illustrate the consequence of this sort of estimation error is to replace the private values assumption by the *common values assumption* that the contractors are all equally capable and that each could, if called upon, do the job at the same cost C. Although this assumption is special, it makes it possible to illustrate some general phenomena rather simply. For additional simplicity, suppose that the bidders make unbiased estimates $X_i = C + \tilde{\varepsilon}_i$, where the estimation errors are independent.

Despite the independence of the estimation *errors*, the bidders' *estimates* are not independent, because the estimates in this Bayesian model are the sums of the common random term C and the independent errors.[1] As we shall later argue, this failure of the independence assumption has important consequences for the comparative performance of alternative auctioning rules.

Now comes a crucial observation. Even though each contractor's individual estimate is unbiased (that is, equal on average to the expected cost), the *lowest* estimate is biased downward. Indeed, because the expected value of the individual estimation errors is zero, the expected value of the minimum estimation error must be less than zero, and that implies the claimed estimation bias.

Now suppose all bidders determine their bids by adding the same fixed and/or percentage markup to their estimated costs, or using any other markup rule for which higher costs lead to higher bids. Then the winning bidder will be the one with the lowest estimate of project completion costs, and the winner's cost estimate will be too low on average.

The phenomenon just described is known as the winner's curse. It forms the basis of the explanation and advice that I might have offered to the painting contractor quoted in the introduction: "There may be nothing unusual about the painting jobs on which you have bid and nothing terribly wrong with your cost estimates. The problem is that in competitive bidding, your bid usually loses when you overestimate your actual costs. Often, when you win a job, it will be because your cost estimates were too low. To make money in competitive bidding, you will need to mark up your bids twice: once to correct for the underestimation of costs on the projects you win and a second time to include a margin for profit. Don't let the presence of several competing bidders push you into making too aggressive a bid. The markup to adjust for underestimation will have to be larger the larger is the number of your competitors

---

[1] Indeed, the covariance of any two different estimates X_i and X_j is equal to the variance of C.

and the more you respect the accuracy of their cost estimation; you may, however, want to make the profit markup smaller when there are more competitors. Also, the payoff to careful cost estimation in competitive bidding is great, because it allows you to bid aggressively without great risk. If you can also develop a reputation among your competitors for being an unusually savvy estimator, that's even better for you, because it will compel sensible competitors to bid more cautiously against you and allow you either to increase your profit markup or to win more bids."[2]

Since the contractor (who was my father) was retired by the time I understood these lessons, I have not tried my explanation on him.

Students are quite rightly reluctant to accept these results as proof that it is always best to bid cautiously. "You can't make any money if you never win a bid, and you can't win if you are too cautious" is a common response. The most important lessons to be learned from both the theory and the experiments are that the returns in bidding come from cost and information advantages, that naive bidding strategies can squander these advantages, and that bidders without some advantage have little hope of earning much profit, but could with a little bit of carelessness suffer large losses.

Equivalences Among Auction Institutions

To fix the terminology for this section, let us assume that the auctioneer is selling some goods and the bidders are the potential buyers. The first general question to ask about auction markets is to what extent the details of the institution matter. Should sealed bids be used, with the contract being awarded to the highest bidder for a price equal to its bid? Would the price be higher or the outcome more efficient if an open outcry auction of the kind used by the English auction houses were adopted, where bidders call out increasing bids until only the highest bidder remains? How do these alternatives compare with the Dutch auction, in which the auctioneer initially calls a high price and then lowers it continuously until some bidder claims the goods?

In one of the earliest and most remarkable economic analyses of auctions, William Vickrey (1961) studied those questions (and others). Let us review first Vickrey's analysis of the Dutch and sealed bid auctions.

In a sealed bid auction, each bidder independently and privately picks a price and offers to buy the goods at that price. The one who bids the highest price wins. The Dutch auction is seemingly quite different. The auctioneer calls the prices, beginning with a high price and proceeding to successively lower ones. The bidder listens to the prices called, notices whether any other bidder has accepted a price, and finally accepts some price if no other bidder has done so first.

Despite the seeming complexity of the Dutch auction bidder's task, a bidder who plans his actions in advance will find that his problem is identical to that facing a bidder in a sealed bid auction. For regardless of how the bidder approaches the calculations, the only genuine choice open to him is to select the highest price at which

---

[2]Formal propositions to this effect are provided by Milgrom and Weber (1982a).

he will be willing to claim the goods. An example will help to illustrate this point. Suppose the bidder decides to proceed as follows. First, he will wait until the price has fallen to some level p_0 and then infer what he can from the failure of the others to bid at that point. Then, depending on the outcome of some calculation, he may choose to claim the item or he may choose to wait. In the latter case, he will perform some other calculation given the then available information to select another (lower) price p_1 at which to reevaluate, and the process then repeats itself. The outcome of this algorithm is entirely predictable. The upshot is that there is a single number p that represents the first price at which the bidder will claim the item. Let us call p the "bid." Using this language, we find that under the rules of the Dutch auction, the goods will be awarded to the "highest bidder" at a price equal to his bid. But those are precisely the rules of the standard sealed bid auction! The apparent complexity of the possible bidding strategies in the Dutch auction is a chimera; the only real choice a bidder has is to select his "bid" p.

Putting this conclusion somewhat differently, what this argument shows is that when the Dutch and sealed bid auctions are each modeled as "strategic form games," the games are identical.[3] That is, the sets of strategies are identical and the outcome rules that transform strategies into allocations are identical. Since solution concepts like the Nash equilibrium work on strategic forms, these concepts predict powerfully that the identity of the winner and the price the winner pays will always be the same for these two kinds of auctions.

In small stakes laboratory experiments, however, this prediction appears not to hold: winning bidders in these experiments tend to pay a lower price in the Dutch auction than in the sealed bid auction. Why might this happen? One hypothesis, which I favor, is that the Dutch auction format discourages planning by the subjects: Dutch auctions are not "played" in the normal form. This hypothesis could be tested by manipulating the experimental conditions to encourage articulation of a plan before bidding in the Dutch auction to see if that reduces or eliminates the price disparity. For example, the experiment could have pairs of subjects who must place a joint bid; this is likely to encourage discussion of how to bid and may provide direct evidence on whether the subjects express their strategies as single numbers. An alternative hypothesis is that the subjects do not maximize expected utility, but instead use some other decision criterion. Finally, given the small stakes in most experimental conditions, the excitement of playing the game may lead the subjects in the Dutch auction game to prolong the game by holding out longer, which would account for the lower price. In this case, raising the stakes would tend to diminish the differences in the outcomes between the two alternative auction forms.[4]

Next, consider the common form of auction used by English auction houses, often called the English open outcry auction. Here, the auctioneer begins with the lowest acceptable price—the *reserve price*—and proceeds to solicit successively higher bids

---

[3] In the language of game theory, the Dutch and sealed-bid auctions have the same "reduced normal form."
[4] See Cox, Smith, and Walker (1983) for a report of an experiment along these lines.

from the customers until no one will increase the bid. Then the item is "knocked down" (sold) to the highest bidder. Suppose that the bidders know the value to themselves of the item being auctioned. This is the private values assumption: it rules out the possibilities that the value of the goods to one bidder depends on its resale value to other bidders, or on the availability of substitutes regarding which the other bidders may have private information, or on how much others may admire the item, and so on. With the private value assumption, the bidder's dominant strategy in an English auction is to bid until the price exceeds his willingness to pay. Evidently, at equilibrium, the item will be awarded to the bidder who values it most highly for a price equal to the second highest valuation. This outcome is efficient.

Vickrey observed that the outcome of the English auction could also be achieved by means of a sealed-bid auction with the following rules: Each bidder submits a bid. The item is awarded to the highest bidder at a price equal to the second highest bid, or to the reserve price if that is higher. The bidders in this Vickrey *second-bid* auction are price takers—the price that the winning bidder pays is determined by the competitors' bids alone and does not depend on any action the bidder undertakes. From a single bidder's point of view, a bid in the Vickrey auction determines which "price offers" the bidder would be willing to accept; he will accept any "offers" up to the price he has bid. It is thus a dominant strategy for a bidder in this auction to submit a bid equal to his true reservation price, for he then accepts all offers which are below his reservation price and none that are above. When each bidder adopts his dominant strategy, the outcome will be that the item is awarded to the bidder with the highest valuation for a price equal to the second highest valuation. The existence of a *dominant strategy* in this auction means that the bidder can choose his sealed bid without regard for how others bid. Thus, the second-bid auction duplicates the principal characteristics of the English open outcry auction. For that reason, it is customary to model the English auction as a second-bid auction, a custom we shall respect in this essay.

We cannot specify "optimal" bidding strategies for the Dutch auction and the standard sealed-bid ("first-bid") auction in the same way that we did for the English and second-bid auctions, because the profit-maximizing bid in these auctions depends on what bids the competitors make. To analyze this problem, let us shift attention from "bids" to "bidding strategies." A (pure) *strategy* for a bidder specifies what bid to make as a function of the information the bidder has. As illustrations, the construction contractor in our earlier example makes its bid as a function of its estimate of the cost of performing the contract and of any information it has about its competitors, and the oil company makes its bid for drilling rights as a function of its geologists' estimates of hydrocarbon potential. If each bidder correctly anticipates the bidding strategies his competitors will use and selects his own optimal strategy accordingly, then the collection of strategies form a *Nash equilibrium* of the bidding game. To say that a bidder correctly anticipates his competitors' *strategies* is not to say that he correctly anticipates their *bids*, but only how they would bid if they had some particular information. We will use the Nash equilibrium solution concept to analyze the auction games here.

The first thing to note about the first-bid auction, whether run using sealed bids or using a Dutch descending auction, is that there is no assurance that the equilibrium outcome will be efficient. In fact, in any environment where the bidders have observably different characteristics, the equilibrium outcome of sealed bidding is inefficient with some positive probability. To illustrate, suppose there are two bidders, one who is known to have a personal reservation price, or *valuation*, of $101 and a second whose valuation is either $50 with probability 4/5 or $75 with probability of 1/5. The first bidder is assumed not to know the valuation of the second, but he knows its distribution. If the first bidder bids $51, he will win $50 (his valuation minus his bid) at least 4/5 of the time, yielding an expected profit of at least $40. If he bids $62 or more, he can win no more than $39 (= $101 − $62), so he will never make that choice. Since the first bidder never bids as much as $62, an optimizing second bidder must win sometimes when his valuation is $75, and the allocation then is inefficient. This stands in contrast to the English auction, which would always lead to an efficient allocation in this environment.

Much of auction theory has analyzed "symmetric" environments where the bidders cannot discern differences among their competitors. For example, one popular formulation assumes that the bidders' private valuations are independent and identically distributed random variables. In this case, as Vickrey first showed, there is an equilibrium in which each bidder adopts the same *strategy*; that is, each bids the same increasing function of his personal valuation. As a consequence, the bidder with the highest personal valuation will make the highest actual bid and the equilibrium allocation will be efficient.

Given that the English and second-price auctions and the Dutch and sealed bid (first-bid) auctions are both efficient in this environment, what can be said about how the total surplus will be divided between profits for the bidders and revenue for the seller? If the auctioneer/seller has the power to set the rule, the answer to this distributional question will predict which type of auction will be seen in practice. Will the first-bid auction, in which a winning bidder pays the amount of his own bid, lead to higher payments on average than the second-bid auction, in which the price is set equal to the second highest bid? The matter is not an obvious one, because the bidder in the first-bid auction will optimally shade his bid down to allow a margin for profit, while (as argued earlier) the bidder in a second-bid auction will find it optimal to bid the full amount of his valuation. Is the profit margin deducted by the bidder in the first-bid auction greater or less than his expected profit when he wins in the second-bid auction?

To illustrate the surprising answer, let us adopt an indirect approach. In any auction, the bids made by the bidders can be viewed as labels that are processed through the rules of the auction to determine the outcome. The actual items of interest to a bidder here are the probability P that his bid will win and the expected payments that he will be required to make if he wins or loses. Let us assume for simplicity that only winners pay, and denote the winning bidder's expected payment, given his bid, by E. Then, a bid is really just an indirect way to choose among the real items of interest—the pairs (P, E) where the set of possible pairs is determined jointly by the

rules of the auction and the strategies adopted by one's competitors. At an equilibrium of the bidding game, the bidder correctly perceives how his bids map into (P, E) pairs.

Suppose a bidder's valuation (reservation price) for the goods being offered is X. If he wins the auction, his expected profits are just the difference between this valuation and his expected payment: $X - E$. So if he selects the point (P, E), his expected profits corresponding will be:

$$U[P, E; X] = P \cdot (X - E) \tag{1}$$

The bidder's optimal choice of (P, E) will clearly depend on his reservation level X. Let $(P^*(X), E^*(X))$ denote the optimal choice for the bidder and let the corresponding maximal expected profits by $U^*(X) = U[P^*(X), E^*(X); X]$. By studying the function U^*, we learn how the expected gains from trade are divided among the seller and the various possible types of bidders in each auction format. The main conclusion is this:

Revenue Equivalence Theorem: The English and sealed bid auctions yield exactly the same expected profit for every bidder valuation and the same expected revenue for the seller. Indeed, every auction that allocates the goods efficiently and offers no profit to a zero valuation bidder has the same expected profits for every bidder valuation and the same expected revenue for the seller.

Proof is given below. The proofs in this paper highlight the general techniques which are used repeatedly in auction theory. The proofs may be omitted by readers without any loss of continuity in the development.

Proof.   Applying the Envelope Theorem to (1), we have:

$$U^{*\prime}(X) = U_X[P^*(X), E^*(X); X] = P^*(X) \tag{2}$$

Given the hypothesis of the Theorem that $U^*(0) = 0$, we may integrate (2) to obtain:

$$U^*(X) = \int_0^X P^*(s) \, ds. \tag{3}$$

For any auction where the allocation is always efficient, $P^*(X)$ is just the probability that the other bidders' valuations are less than X. Thus, by (3), all such auctions yield identical expected profits for the bidders. Since the total surplus generated by trade is the same in all such auctions and the bidders' profits are the same, the seller's expected revenues must be the same, too.   QED

The conclusion in the foregoing model that standard auctions like the first-bid auction, the Dutch auction, and the English auction all lead to the same expected revenues for the seller and expected profits for the bidders has spawned a number of analyses that tweak the model in some way to explain why one auction rule or another might be expected to perform better in practice. Two variations are reviewed in the next sections which seem especially interesting because they help to explain in plausible ways why open outcry auctions like the English auction are by far the most

prevalent auctions in the world, yet industrial procurement auctions are almost always first-bid auctions.

Auctions with Endogenous Quantities

This section presents and analyzes the auction model introduced by Hansen (1988) to explain the use of first-bid auctions for industrial procurement. For this application, we flip back to the contracting perspective: The bidders are again sellers, the bid-taker is a buyer, and the lowest bidder will be declared the winner.

The model retains all of the assumptions of the preceding model but one: Rather than purchasing a single unit, the buyer who solicits bids can purchase as many units as he wishes at the price fixed by the auction. The buyer's quantity decision is modelled by a demand function $q(p)$; if the price determined by the auction is p then the bidder will actually purchase $q(p)$ units. Each bidder's private valuation in this model is its unit cost parameter c; it is assumed that the bidder can produce as many units as are demanded at cost c per unit.

The rules of the variable-quantity first-bid auction are as follows. Privately and simultaneously, each bidder submits a price bid; the lowest bidder is declared the winner; the price p is set equal to the lowest bid; and the buyer purchases the quantity $q(p)$. The rules of the second-bid auction are the same, except that the price p is set equal to the second lowest bid. One could recover the model analyzed in the preceding section as a special case of this model by setting $q(p) = 1$. However, the objective here will be to analyze the extra effects caused by elastic demand, so we shall assume that the demand curve slopes downward: $q' < 0$ and that there is a "choke price" $\bar{p}$ such that $q(\bar{p}) = 0$. As in the simpler model, there again exists an equilibrium for each auction game in which all the bidders adopt the same strategy and the bids are an increasing function of the production cost parameter. So, parallel to our earlier finding that the goods are sold to the highest evaluator, we find in this model that the contract is awarded to the low cost producer. However, because the equilibrium price will generally exceed marginal cost, the final allocation will *not* be efficient. Our first task is therefore to analyze and compare the expected prices, quantities, and surplus in the two auction formats. Then, we will look at how the gains to the more efficient format are distributed between the bidders and the bid-taker.

In the second-bid auction, even with variable quantities, bidders are price takers; the price a winning bidder will receive depends not on his own bid but on the bid of his nearest competitor. From the perspective of an individual bidder, the price he names merely specifies the lowest price at which he will be willing to undertake production. So it is a dominant strategy for each bidder to name a price equal to his marginal cost c; that way he accepts all offers to produce output at a price exceeding his cost per unit, and no other offers. Thus far, the analysis is unchanged from the case in which the quantity demanded is always one unit.

What is new is that, in the first-bid auction, the winner will have an additional incentive to shade his bid. For consider the problem facing an individual bidder.

Suppose that $P(p)$ is the probability that a price bid of p will be lowest and suppose that the bidder's production cost per unit is c. Then he will choose his bid p to maximize expected profits, calculated as the probability the bid will win times the quantity sold in that case times the profit per unit sold. Compared with the case where $q' = 0$, the bidder in a first-bid auction with $q' < 0$ finds it more profitable to reduce his bid price p, because there is the additional effect that reducing the bid causes the quantity sold to rise when he wins.[5] The upshot is that the equilibrium bids in the first-bid auction are reduced by the introduction of elastic demand. When demand is inelastic, we already know that the first-bid and second-bid auctions lead to the same average price. *In this model, for every level of cost of the winning bidder, the first-bid auction leads to a lower average price than the second-bid auction.*[6]

Recalling that bids are really just ways of parameterizing the choice among the objects of real interest, we can perform the same analysis on the bidders' choice of expected quantities sold. The upshot is that *for every level of cost of the winning bidder, the quantity $q_F(c)$ sold in the first-bid auction is greater than the expected quantity $q_S(c)$ sold in the second-bid auction.*

Finally, to compare the efficiency of the two auction formats, observe that whenever demand is downward sloping, society is risk averse about the quantity to be produced; total surplus is higher when a fixed quantity q is produced than when a random quantity with mean q is produced. We have already argued that, for every level of c, the expected quantity is higher in the first-bid auction and, unlike the second-bid auction, the quantity given c in the first-bid auction is not random. So, *for every level c of the cost of the winning bidder, the expected total surplus is greater in the first-bid auction than in the second-bid auction.*

Who captures the gains generated by switching to a first-bid auction? One might guess that since the first-bid auction generates lower average prices, the bidders are made worse off. In fact, the reverse is true.

Proposition. Although the average price is lower and the expected quantity is greater in the first-bid auction than in the second-bid auction, the bidder/sellers' expected profit is greater as well.

---

[5] The constrained optimization problem looks this way:

$$\underset{p}{\mathrm{Max}}\, P(p)q(p)(p - c).$$

The derivative of this objective function consists of two terms that do not depend on $q'(p)$ and one term that does:

$$P'(p)q(p)(p - c) + P(p)q(p) + q'(p)P(p)(p - c)$$

Moving from the case $q'(p) = 0$ to $q'(p) < 0$ reduces the slope of the objective everywhere and therefore reduces the maximizing bid p.

[6] There are more steps to the formal analysis than indicated here. If other competitors lower their bid prices, then one must account for the change in the first bidder's choice problem. With price competition, one might expect that more aggressive bidding by competitors reinforces the incentive to cut one's own bid price, and the full equilibrium analysis given by Hansen confirms that intuition.

Proof. The proof is quite similar to that of the Revenue Equivalence Theorem. Let Q denote the expected quantity sold by a bidder who makes any particular bid, defined as the probability of making sale times the expected quantity given that a sale is made, and let E be the expected price per unit. Then if a bidder with marginal production cost c selects a bid corresponding to (Q, E), his expected profits are:

$$U[Q, E; c] = Q(E - c) \tag{4}$$

This expression is analogous to the one given earlier (in (1)) for the case where quantity was fixed at a single unit. Let $U^*(c) = U[Q^*(c), E^*(c); c]$.

Applying the Envelope Theorem to (4), we have:

$$U^{*\prime}(c) = U_c[Q^*(c), E^*(c); c] = Q^*(c) \tag{5}$$

Observing that $U^*(\bar{p}) = 0$, we may integrate (5) to obtain:

$$U^*(c) = \int_c^{\bar{p}} Q^*(s) \, ds. \tag{6}$$

For the first-bid auction, $Q^*(s) = q_F(s)$; for the second-bid auction, $Q^*(s) = q_s(s)$. Since we have argued above that $q_F(s) > q_s(s)$, it is evident from (6) that $U_F^*(c) > U_S^*(c)$, and the Proposition is proved.   QED

One can show, as well, that under any of a variety of mild restrictions on the demand function, the buyer enjoys greater expected consumer surplus with the first-bid auction. When these conditions hold, both sides of the market prefer this auction mechanism. This relatively simple model therefore provides a straightforward explanation of why first-bid auctions are so widely used in procurement.

Correlated Bidder Information

In many auction settings, the quantity to be supplied is specified in advance, so that the effects explained in the last section are absent. In these settings, the English (open, ascending) auctions are much more common than any other form. Here we relax the independence assumption of the standard model in order to explain the prevalence of English auctions and of certain other auction practices.

What are we to use in place of the independence assumption? When the bidders' costs or valuations depend on some common random factors, so that all the bidders are estimating the same variables, their estimates will be positively correlated even if their estimation *errors* are independent. Positive correlation has been especially prominent in models of auctions for oil and gas drilling rights, where the rights being acquired are, to a first approximation, of equal value to each of the bidders, and the main uncertainties concern such common factors as the quantities of recoverable hydrocarbons, the costs of recovery, the costs of transporting the product to market (perhaps through as yet undeveloped pipelines over the Arctic Slope), future world

energy prices, and so on. The common uncertainties found in these auctions also play a large role in the sale of items like wine or art which are purchased at least partly for their savings or investment value, as the parties estimate what it would cost to purchase the same vintage in the future or what the eventual resale price for the painting will be. So there is good reason to believe that positive correlations among value estimates will often be present.

The actual equilibrium analysis of auctions relies on a stronger notion than positive correlation. The appropriate concept, known as *affiliation*, was introduced by Milgrom and Weber (1982). Affiliation of the bidders' value estimates in the auction model captures the idea that as a bidder's value estimate rises, he expects others' estimates to rise as well, in the sense that higher values for other bidders' estimates become relatively more likely.[7]

Turning now to the model formulation, since English auctions are mostly used by sellers of goods, it is most natural to switch back to the model formulation where there is a single item or bundle of goods being offered for sale. Thus, bidders are once again buyers, and the auctioneer is the seller or his representative. The bidders each have an estimate of the value of the goods to themselves; these estimates are assumed to be symmetrically distributed and affiliated. The value of the goods may be a common value, identical to all the bidders, or it may be a private value, or some mixture of these. Each bidder chooses his bid as a function of his value estimate to maximize the expected excess of value received minus payments made.

Parts of the analysis are unchanged by the new assumption of affiliation. First, the equilibrium bids are still increasing functions of the bidders' valuations, so the winner is the bidder with the highest valuation in both auction formats. Consequently, the total surplus does not differ between the two auction forms. Second, the profits of a marginal bidder who is just willing to participate in the auction is zero in both formats. Let $U^*(X)$ be the expected profits of a bidder with valuation X in the auction. Then if the valuation of a marginal bidder is X, the second conclusion can be stated as $U_F^*(\underline{X}) = U_S^*(\underline{X}) = 0$. The point of the present analysis is to compare how affiliation affects the auctioneer's ability to extract the incremental profits associated with larger value estimates, that is, how it affects the slopes $U_F^{*\prime}$ and $U_S^{*\prime}$ of the profit functions.

Heuristically, the comparison is easy. In a first-bid auction, when a bidder with valuation X submits a bid of b and wins, the price he pays is b; it does not depend on X. In an English or second-bid auction, the expected payment by a winning bidder whose private valuation of the good is X depends on his bid b, *but it also depends on his private valuation X to the extent that the distribution of the remaining bids depends on X.* This extra effect raises the slope $U_S^{*\prime}$ but not $U_F^{*\prime}$, tending to result in higher average prices in the English auction. One still needs to verify, though, that these effects are not eliminated by equilibrium adjustments.

---

[7]A collection of random variables $(X_1, \ldots, X_N)$ with symmetric joint density f is affiliated if for any $y \geq y'$ in $\mathscr{R}^{n-1}$, the ratio of densities $f(y|x)/f(y'|x)$ rises in x; this is the basis of the statement in the text that higher values y become relatively more likely as x increases.

Proposition. The expected total surplus is the same in the English and sealed-bid auctions in the model, but for every value of X the bidder's expected profit $U^*(X)$ is smaller in the English auction and the seller's expected revenue is correspondingly higher.

Proof. To emphasize the parallel with the preceding proofs, we add the private values assumption to the list of hypotheses. We continue, however, to assume that the bidders' estimates are affiliated. Then for a bidder whose value estimate is X and who bids just enought to win when the highest estimate of an opposing bidder is Z, the expected payoff is:

$$U[Z, E; X] = P(Z|X)(X - E) \qquad (7)$$

where E is the bidder's expected payment when he wins and $P(Z|X)$ is the probability that all the other bidders have value estimates less than Z, given that the present bidder has value estimate X. Given the strategies of the other bidders, the expected payment E is some function $E(Z, X)$ of the bid (parameterized by Z) and the information X of the bidder. So,

$$U[Z, E(Z, X); X] = P(Z|X)(X - E(Z, X)) \qquad (8)$$

Let $U^*(X)$ be the maximized value of (8). At equilibrium, we know that $Z^* = X$ is the optimal choice of Z. So, the Envelope Theorem implies that:

$$U^{*\prime}(X) = U_X = P(Z^*|X)(1 - E_X(Z^*, X)) + P_X(Z^*|X)(X - E(Z^*, X))$$

$$= P(X|X)(1 - E_X(X, X)) + P_X(X|X)(X - E(X, X))$$

$$= P \cdot (1 - E_X) + (P_X/P)U^* \qquad (9)$$

where the arguments X are suppressed in the last line to improve readability. Affiliation implies that P_X is negative. It is then a standard exercise in differential equations to show that, given the fixed boundary condition $U^*(\underline{X}) = 0$, all the values of $U^*(X)$ fall if the coefficient function E_X is increased.

Notice, however, that for the first-bid auction $E_X = 0$, that is, the bidder's expected payment depends only on his bid and not on his value estimate. For the second-bid auction, $E_X \geq 0$, because (holding the bid fixed) a higher value estimate makes higher bids by the other bidders relatively more likely. Hence, raising E_X by switching to the second-bid (English) auction causes $U^*(X)$ to fall for all X, and the Proposition is proved.   QED

Our formal proof directly mirrors the intuitive argument given earlier. The difference between the expected payoffs in the two kinds of auctions is directly traceable to the fact that the bidder's expected payment in the second-bid auction rises

with his value estimate, because the payment depends on the value estimate of a competitor which is affiliated with his own estimate. No such effect is present in the first-bid auction, which is therefore less effective in extracting surplus from the bidders. If one supposes that the auctioneer (representing the seller) sets the rules, then the Proposition provides a possible explanation of why English auctions are so much more prevalent than sealed-bid auctions.

The sort of analysis we have just employed can be applied in a variety of ways. The main general insight of the analysis is the *Linkage Principle*. According to this informal principle, the bidders are made worse off and the seller better off if the price paid by the buyer can be more effectively linked to exogenous variables that are affiliated with the bidder's private information. For example, the Principle implies that the use of royalties in the selling of mineral rights or publication rights—a practice that links the price paid to the actual value—will increase the seller's average receipts. It also implies that if the auctioneer/seller has private information about the item being sold that is affiliated with the bidders' estimates of value, then a policy of always revealing that information increases average receipts compared to a policy of never revealing information. For if a policy of revealing information is adopted, the price becomes linked to the seller's information, which extracts more of the winning bidder's surplus.[8]

Another application of the Linkage Principle arises in a simple model of the weekly Treasury bill auction. In the model, there are N bidders and $M < N$ items for sale, and each buyer wants to buy only one unit. The brokers who bid in the T-bill auction all estimate the future market price at which they will be able to resell the bills to their customers. So, it may be sensible to assume that their estimates are affiliated. That suggests that, in place of the present "discriminatory" auction rule where each bidder pays the amount of his own bid, the Treasury might do better to adopt an auction in which the price paid by each bidder is linked to the bids made by others. Among the possibilities are the uniform price auctions in which the price charged to all bidders is, for example, the lowest accepted bid, or the highest rejected bid, or the average accepted bid. Each is in fact predicted by the theory to generate a higher average price than the auction at which each bidder pays the amount of his own bid.

The analysis in this section is subject to a number of important qualifications. First, we have developed the arguments only for symmetric bidding models. There was a good reason for that: The theory makes no sharp predictions about the outcome in asymmetric environments. Second, we assumed that the bidders were risk neutral. In "private values" models, risk aversion raises receipts in the first-bid auction but not in the second-bid auction; that could reverse some of the auction rankings. However, in general value environments, the effects of risk aversion are ambiguous. Although most auctions do not require the bidder to commit a significant fraction of his wealth, some (like construction bids) do often carry the risk of insolvency, and risk aversion may be important in these cases. Third, we assumed that equilibrium strategies would

---

[8]For experimental evidence that tends to confirm this prediction, see Kagel and Levin (1986).

be adopted by the bidders. The computation of equilibrium strategies for first-bid auctions in the environments we have studied is not a simple matter, and there is no assurance that these bidding strategies would, in fact, be used.

Comparisons Among Auctions

The emphasis of much of recent bidding theory has been on ranking auctions on the basis of the expected receipts they generate. Sometimes this approach is taken to the extreme of determining the institutions that maximize expected receipts, on the grounds that such institutions will be the ones chosen by the auctioneer/seller. The results of these maximization problems are, for all but the simplest environments, auctions of outlandishly complicated forms involving payments by the seller to losers, required side bets among the bidders, and so on. That such forms are not observed in practice indicates that the "optimal auctions" theory in which the auctioneer can tailor a specific institution to each environment may be a poor way to explain actual institutions. The common auction institutions are all simple and robust, working well in a variety of environments, used by desperate sellers as well as by those with market power bordering on a monopoly, and usually leading to a tolerably efficient allocation of the items being sold. Comparisons of robustness, efficiency, transaction costs, and immunity to cheating offer an important alternative to the revenue-based approaches for explaining the popularity of specific auction institutions.

We have already seen that the various models differ in their conclusions about the efficiency of the allocations resulting from various kinds of auctions. The Hansen endogenous quantities model assumes symmetry among the bidders before the auction and predicts that the outcome of the first-bid auction is more efficient. When the quantity traded is fixed and the bidders are not symmetrical, the second-bid auction always assigns the goods to the proper bidder, but the first-bid auction may fail to do so.

Of course, the final allocation itself is only one aspect of the efficiency of auctions. Another important aspect is the cost of preparing a bid. Complicated auctions and those that provide large returns to information gathering are likely to increase bid preparation costs. The English auction system, in which a bidder's optimal bidding strategy does not depend on how his competitors bid, economizes on information gathering and bid preparation costs.

In summary, at least for fixed quantity environments, the English auction possesses a variety of characteristics that help to explain its popularity. It generates more receipts on average than the Dutch/sealed-bid auction. It leads to efficient outcomes in a wider range of environments. And, it economizes on information gathering and bid preparation costs.

English auctions also have some characteristic disadvantages, however. First, being open outcry auctions, they require the actual presence of the bidders, which may be expensive. One might think that this disadvantage could be overcome by substituting the equivalent second-bid auction, which is a sealed-bid auction and so

does not require the presence of the bidders. However, when the auctioneer opens the bids in a second-bid auction and learns what the bidders are willing to pay, what is to prevent him from inserting a false bid to drive up the price? This susceptibility to manipulation may account for the unpopularity of second-bid auctions. (Interestingly, a second-bid auction in which the auctioneer is free to insert extra bids after opening the sealed-bids is virtually identical to a first-bid auction, because the highest bidder wins and winds up paying a price approximately equal to his bid.)

A second disadvantage of English auctions is their easy susceptibility to *rings* of bidders.[9] A ring is a group of bidders who agree to re-auction the items they purchase among themselves, dividing the proceeds among the bidders. One member, representing the ring, bids up to the ring's reservation price (the highest of the members' reservation prices) in competition with other bidders, while the ring members refrain from bidding. In an English auction, no member of a ring can successfully exploit the ring agreement to gain a bargain for himself. For if a ring member were to have an anonymous associate bid aggressively on his behalf, the ring representative would continue to bid up to the ring's reservation price, and the selling price would be just as high as if no ring had been formed.[10] Rings of bidders are comparatively much less effective in the first-bid auction, in which a single defector from a ring agreement could reap a substantial gain by bidding slightly more than the agreed price. This observation favors the selection by auctioneers of the sealed-bid auction when the threat of rings is great, as when the bidders are well known to each other. It supplements Hansen's argument, that first-bid auctions lead to more efficient allocations when purchase quantities are endogenous, to explain the use of first-bid auctions for industrial procurement.

Comparisons with Non-Auction and Hybrid Institutions

So far, we have compared only alternative auction institutions for selling a single item. However, most trades are not made using auctions. There are goods for which stores post prices and others over which people haggle. High volume securities trading is conducted on organized exchanges using a complicated set of auction-like rules. These alternatives only scratch the surface in describing the huge variety of terms and institutions that govern trade in the modern world. What is it about the circumstances in each case that make one or another trading institution most appropriate? Unfortunately, that question has received less attention from researchers than the others we have asked. Nevertheless, some informed guesses are possible.

Posted prices are commonly used for standardized, inexpensive items sold in stores. Often, these are manufactured items for which there is no need to compare

---

[9]This discussion of the problem of rings is based largely on Graham and Marshall (1987).

[10]This is not to say that an auctioneer is helpless when faced by a ring. He can set high reserve prices and take bids from the chandelier to force some competition when he suspects that a ring is at work. Still, these devices are not complete remedies, so auctions that make secret and profitable violations of ring agreements possible have obvious advantages in cases where rings are a real problem.

competing bids: All comers can be served. Even when the supplies are limited, unlike organized auctions, the competing buyers are not all present at the store at any given time. If it is too expensive to gather the competing buyers together or if the item is storable and the timing of buyer demands varies, auctions are not a practicable selling institution. The alternative to posted prices, then, is individual bargaining. Bargaining, however, is a costly way to determine prices in societies where time is especially valuable. Also, when the salesman is not an owner, the lack of a fixed price makes it easy for the salesman to steal or take kickbacks from the buyer. These defects of bargaining are especially great when the opportunity to gain from price discrimination in individual sales is low. These observations explain why sales of consumer goods in developed economies so often utilize the posted price institution.

When goods are not standardized or when the market clearing prices are highly unstable, posted prices work poorly, and auctions are usually preferred. Thus, livestock are sold at auction because the value of individual animals varies, and needs to be determined separately. Similarly, the assets of bankrupt firms are valued depending on their age, condition, location, and so on. Fresh fish, being perishable, is sold at auction so that prices can be responsive to daily variations in the catch and demand.

Bargaining is a trading institution that is best avoided when there is enough competition for auctions to be used, Coase's theorem notwithstanding. It is well-known, for example, that in bargaining to divide a fixed surplus with some or all of the surplus being lost if agreement is delayed, there can be a problem of indeterminacy, with the outcome of bargaining depending on the bargainers' expectations (Roth and Schoumaker, 1983). If one bargainer demands 75 percent of the surplus, the other can do no better than to settle for 25 percent. Each, knowing this, may be tempted to play a game of brinkmanship, demanding the lion's share in the hope that the other party will back down. Disagreement and inefficiency can, and sometimes does, result. In contrast, when a seller employs an English auction to sell an item worth $100 to himself to one of a pair of potential buyers with reservation values of $170 and $200, the equilibrium theory predicts the sale will occur at a price of $170.[11] Not only is the result efficient, but the seller gets a good price: By bargaining singly with the $200 evaluator, the seller can at best hope to split the gains, getting a price of $150, and he may lose some of those gains if the parties fail to reach agreement.

As compared to bargaining, auctions have the additional advantage of being institutions whose conduct can be delegated to an unsupervised agent. Public auctions offer fewer opportunities for kickbacks and behind-the-scenes agreements between the seller's agent and a single buyer than do negotiated agreements. In the early New England textile trade, established merchants sponsored laws against auction sales, thus indicating their awareness of how effectively auctions narrow their margins and prevent them from extracting better terms from the cotton farmers.

---

[11] Indeed, he may get more if the $170 evaluator considers that he may be able to resell the item to the $200 party, dividing the $30 surplus equally. When resale is possible, the item is worth about $185 to this party, and the initial auction may lead to a price of $185. See Milgrom (1987) for more on auctions with resale.

These simple comparisons of the pure forms of auctions and bargaining tell only part of the story. Often, it is necessary to combine bargaining and bidding to support efficient trade. For example, in the problem of soliciting bids for a contract to build a bridge, to install a cable television service, or to perform R&D on a new weapons system, it would be foolish to invite bids from all comers and just take the lowest bid. Before the final bid can be made, there may be a round of negotiations over the specifications and then another round to determine which of the potential bidders are qualified to produce a product meeting the specifications. Then, the evaluation of the bid may take into account design or quality differences, service capabilities, the ability to deliver on time, and perhaps the need to maintain multiple sources of supply, as well as price. These appendages to the basic auction institution require a substantial element of managerial judgment and open again the possibilities for influence, favoritism, and bribes.

Recent Empirical Studies of the Winner's Curse

In a pair of recent papers, Hendricks, Porter and Boudreau (1987) and Hendricks and Porter (1987) have studied empirically the prediction of auction theory concerning price and profit for federal sale of leases on the Outer Continental Shelf. In lease sales, the government divides a map of an area with hydrocarbon potential into square areas called *tracts*. Periodically, the government conducts an auction sale of a group of tracts. Rights to the individual tracts are sold separately, by first-bid auction.

Both empirical studies focus on the bidding for *drainage tracts*, which are tracts adjacent to one on which deposits of oil and gas have already been found. A characteristic feature of these tracts is that the owner of the adjacent tract has better information about the underlying geologic structure than do any of the competitors. The theoretical analysis of competitive bidding under such conditions was initiated by Wilson (1967) and developed further by Weverburgh (1979), Engelbrecht-Wiggans, Milgrom, and Weber (1983), and Milgrom and Weber (1982a). The distinguishing assumption of drainage tract models is that all the bidders but one have access only to public records. A number of testable predictions are derived, including these: (1) The best informed firm wins at least half the auctions (exactly half if there are no economies of scope in developing adjacent tracts); (2) given the informed bid, the uninformed bids are uncorrelated with the actual value of the tract; and (3) the average profits of uninformed bidders are zero, being negative on the tracts where the informed bidders fail to bid (that is, where the reserve price exceeds the informed's value estimate) and positive on the tracts where the informed do bid.

The empirical investigations of these hypotheses involves the study of 11 drainage tracts adjacent to a single producing tract auctioned in the period 1954-69. There are various empirical and theoretical issues to be faced in attributing profit levels to firms over the producing period of the wells, during which world oil prices jumped sharply.

Still, the results are encouraging, leading Hendricks and Porter to conclude that their analysis "has provided strong, but not necessarily definitive, support" for the model.

Further Reading

The interested reader can turn to several recent survey papers covering auction theory for additional details. McAfee and McMillan (1987) thoroughly survey both the theoretical and empirical literatures on auctions. Milgrom (1987) explores the connections between auctions and bargaining and gives a precise mathematical account of the Linkage Principle. Wilson's (1987) survey pursues the links between auction research and research into market mechanisms when there are many buyers and sellers of a homogenous good. Much of the work in experimental economics has been focused on tests of auction theory. Useful surveys of that field can be found in Plott (1982) and Smith (1982). The influential "optimal auctions" theory is not fully developed in any of the surveys cited above. A highly readable introductory account of it has been given by Riley and Samuelson (1981), and a new account that analyzes the problem using familiar economic tools is given by Bulow and Roberts (1988). Finally, any serious student of auction theory must read the paper by Vickrey (1961). His analysis, which compared auctions on the basis of expected receipts, allocational efficiency, bid preparation costs, and vulnerability to cheating by the auctioneer and others, and which extended the analysis to procurement problems where the government needs to use multiple suppliers, provided the outline for much of what has been reported in this paper.

References

Bulow, Jeremy, and John Roberts, "The Simple Economics of Optimal Auctions," *Journal of Political Economy*, October 1989, forthcoming.

Capen, E. C., R. B. Clapp, and W. M. Campbell, "Competitive Bidding in High Risk Situations," *Journal of Petroleum Technology*, June 1971, *23*, 641–53.

Cox, James, Vernon Smith, and James Walker, "A Test that Discriminates Between Two Models of the Dutch-First Auction Non-Isomorphism," *Journal of Economic Behavior and Organization*, 1983, *4*, 205–19.

Engelbrecht-Wiggans, Richard, Paul Milgrom, and Robert Weber, "Competitive Bidding and Proprietary Information," *Journal of Mathematical Economics*, June 1983, *11*, 161–9.

Friedman, Lawrence, "A Competitive Bidding Strategy," *Operations Research*, 1956, *4*, 104–12.

Graham, Daniel, and Robert Marshall, "Collusive Bidder Behavior at Single-Object Second-Price and English Auctions," *Journal of Political Economy*, December 1987, *95*, 1217–39.

Hansen, Robert, "Auctions with Endogenous Quantity," *Rand Journal of Economics*, Spring 1988.

Hendricks, Kenneth, and Robert Porter, "An Empirical Study of An Auction with Asymmetric Information," *American Economic Review*, December 1988, *78*, 865–83.

Hendricks, Kenneth, Robert Porter, and Bryan Boudreau, "Information, Returns and Bidding Behavior in OCS Auctions: 1954–1969," *Journal of Industrial Economics*, June 1987, *35*, 517–42.

Kagel, John, and Dan Levin, "The Winner's Curse and Public Information in Common Value Auctions," *American Economic Review*, December 1986, *76*, 894–920.

McAfee, R. Preston, and John McMillan, "Auctions and Bidding," *Journal of Economic Litera-*ture, June 1987, *25*, 699–738.

Milgrom, Paul, and Robert Weber, "A Theory of Auctions and Competitive Bidding," *Econometrica*, November 1982, *50*, 1089–1122.

Milgrom, Paul, and Robert Weber, "The Value of Information in a Sealed-Bid Auction," *Journal of Mathematical Economics*, June 1982a, *10*, 105–14.

Milgrom, Paul, "Auction Theory." In Bewley, Truman, ed., *Advances in Economic Theory, 1985: Fifth World Congress*. London: Cambridge University Press, 1987, pp. 1–32.

Plott, Charles, "Industrial Organization Theory and Experimental Economics," *Journal of Economic Literature*, December 1982, *20*, 1485–1527.

Riley, John, and William Samuelson, "Optimal Auctions," *American Economic Review*, June 1981, *71*, 381–92.

Roth, Alvin, and Francoise Schoumaker, "Expectations and Reputations in Bargaining: An Experimental Study," *American Economic Review*, June 1983, *73*, 362–72.

Smith, Vernon, "Microeconomic Systems as Experimental Science," *American Economic Review*, December 1982, *72*, 923–55.

Vickrey, William, "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, March 1961, *16*, 8–37.

Weverburgh, M., "Competitive Bidding with Asymmetric Information Reanalyzed," *Management Sciences*, March 1979, *25*, 291–4.

Wilson, Robert, "Competitive Bidding with Asymmetric Information," *Management Science*, July 1967, *13*, 816–20.

Wilson, Robert, "Auction Theory." In Eatwell, J., M. Milgate, and P. Newman, eds., *The New Palgrave: A Dictionary of Economic Theory and Doctrine*. London: MacMillan, 1987.