CHAPTER 9

The economics of competitive bidding: a selective survey

Paul R. Milgrom

1 Introduction

In Western economies, a wide variety of institutions have emerged for determining prices and conducting trade. In retail stores, the price of each good is usually posted by the seller, and individual buyers can do little to influence that price. When costly inputs are sold to manufacturing firms, the price is often negotiated – so the buyer and seller both take an active part in setting the price. A third common institution for setting prices is the auction, in which the price of goods is determined by competition among potential buyers. Hybrid trading institutions can also be found; for example, on some securities exchanges, specialists bidding against one another determine the market bid and ask prices, while public customers act as price takers.

Each of these various arrangements for conducting trade has some unique merits.[1] Detailed negotiations provide a flexible way to determine prices, product features, financing terms, and so on. But negotiations are costly and time consuming. In contrast, posted prices are inexpensive to administer, but they are also relatively inflexible and unresponsive to the wants and needs of particular consumers and to short-run fluctuations in demand.

Auctions have properties that place them somewhere between negotiations and posted prices. Billions of dollars of U.S. Treasury bills are sold each week using a sealed-tender auction, and while that causes resources to be spent by potential buyers in preparing bids, it also allows T-bill

[1] Harris and Raviv [1981, 1982], Maskin, Riley [1981], and Riley and Zeckhauser [1983] all consider various arrangements for selling goods from the point of view of profit maximization in an attempt to explain why, out of all conceivable selling arrangements, posted prices and competitive bidding are so popular. Their analyses focus on production and selling costs, rather than such factors as shifting demand, buyer transactions costs, uniqueness of the goods, and the need to adapt to newly arriving information.

Reprinted from
Social goals and social organization
Essays in memory of Elisha Pazner
Edited by Leonid Hurwicz, David Schmeidler, & Hugo Sonnenschein

prices to reflect current demand conditions and to respond to new information about variables like the federal deficit or the money supply. Antiques are commonly sold at auction, and although bringing together the potential buyers at one place and time entails moderate costs, there are compensating benefits: the auction procedure leads to prices that reflect the unique features and appeal of each piece, rather than setting a single price for all "Early American Rockers."

There are many kinds of auctions, but we shall limit our attention primarily to the most popular forms: the Dutch, English, discriminatory, and Vickrey auctions. In the *Dutch auction*, the auctioneer begins by calling a relatively high price for the goods being sold, and then gradually reduces the price until some bidder claims the goods for the current price. The *English auction* is an increasing price auction; the auctioneer gradually raises the price until only one bidder is still active. That bidder is then awarded the goods for the specified price. (The early Roman auctions were probably also of this kind, for the word "auction" is derived from the Latin root "auctus," meaning "an increase" [Cassady, 1967].)

The discriminatory and Vickrey auctions are sealed-bid tender auctions that can be used to sell units of a homogeneous good. In the *discriminatory auction*, each bidder submits one or more sealed bids, and the available units are awarded to the highest bidders according to the terms of their bids. Thus, each unit may be sold at a different price. The discriminatory auction is currently used for the sale of U.S. Treasury bills to major buyers. (Individual customers are permitted to buy a small number of bills for the average price paid by the major buyers.)

The *Vickrey auction* is a sealed-bid tender auction that, like the discriminatory auction, awards the units to the highest bidders. The price paid by a bidder for the j^{th} unit he acquires is equal to the j^{th} highest rejected bid from among his competitors. Vickrey [1961] pointed out that the prices a bidder must pay under this procedure do not depend on his own bids; that is, each bidder is a price taker. Consequently (assuming there are no income effects), a bidder can maximize his expected payoff by submitting bids equal to his marginal willingness to pay for each unit. Notice that this bidding strategy is optimal regardless of the bids made by one's competitors; it is a dominant strategy. Vickrey argued that an advantage of this procedure is that it eliminates the bidders' incentives to gather strategic information about their competitors and so reduces bid preparation costs. Moreover, if every bidder adopts his dominant strategy, the allocation of goods resulting from the auction will be Pareto optimal, because each unit is assigned to the bidder for whom its marginal value is highest.

Closely related to the Vickrey auction is the *uniform-price* auction in which the units are all sold at a uniform price equal to the highest rejected

bid. Most formal analyses of auctions focus on the case where each bidder desires only one unit, and in that case the uniform price auction and the Vickrey auction are identical. Moreover, in that case, the uniform price is a competitive price: it is the lowest price at which supply (the number of units being offered) equals demand.

A storm of controversy arose in the 1960s when it was suggested that a uniform-price auction be used in place of the discriminatory auction for selling U.S. Treasury bills (Carson [1959], Friedman [1960], Brimmer [1962], Goldstein [1962], Rieber [1964], Smith [1966]). Proponents of the move argued that the uniform-price auction would lead to more efficient allocations than the discriminatory auction. Moreover, because uniform-price auctions are strategically simpler, the change would reduce bid preparation costs and encourage more bidders to participate. Opponents retorted that efficiency is but one objective of government; getting favorable terms on financing is another. In a discriminatory auction, it was argued, the government captures substantial advantages from price discrimination; more anxious buyers will pay higher prices to raise the likelihood of acquiring a unit. Experiments conducted by Cox, Roberson, and Smith [in press] lend support to this view.

As the following simple example shows, both the increased efficiency argument made by the proponents of the uniform price procedure and the increased revenue argument made by its opponents are subject to important qualifications. Let there be four bidders named A, B, C, and D, each of whom wants only one unit. Let their reservation prices be 22, 20, 15, and 12, respectively. Suppose that two units are available for sale. We model the auction as a noncooperative game among these bidders. If bidder A is awarded an object and pays a price p, his payoff is $22 - p$. If he loses, his payoff is zero. The payoffs of the other bidders are similarly determined. Each bidder selects a nonnegative real number to be his bid, and the outcome is then determined according to the rules of the auction.

In the uniform price auction, each player has a dominant strategy: A bids 22, B bids 20, C bids 15, and D bids 12. The units are awarded to A and B for a price of 15, and the seller's revenue is 30.

In the discriminatory auction, no bidder has a dominant strategy, but the bidders do have dominated strategies: they should never bid higher than their respective reservation prices. Because there is no dominant-strategy equilibrium, we shall seek a Nash equilibrium in undominated strategies. At every such equilibrium, A and B both bid 15 and C adopts some randomized strategy. For example, at one equilibrium, C selects his bid at random from the interval (14,15), using a uniform distribution, and D makes some undominated bid. At each equilibrium, the winners are A and B. Both pay a price of 15, and the seller's revenue is 30, just as with the uniform-price auction.

In this example, then, there are neither differences in efficiency nor differences in revenues generated between these two auctions. It seems clear that to reach any general conclusions, a more careful analysis is needed.

2 Equilibrium models of competitive bidding

The theoretical literature on competitive bidding consists mostly of studies of noncooperative game models of auctions in which a single item is offered for sale. To specify a game model, one must identify (i) the players, (ii) the information known by the players, (iii) the actions available to them, and (iv) how payoffs are determined. In addition, one must specify a solution concept that predicts how players in a game will behave.

In a typical auction game, the players are the bidders. All of the bidders share certain common knowledge about their environment – the rules of auction, certain characteristics of the item being sold, and so on. In addition, each bidder may have private information concerning his own tastes or the characteristics of the item being sold. To capture the idea that this information is private (and therefore not known by the other bidders), each player's information is modeled by a random variable that he alone observes. The bidders' uncertainty about one another's information is then represented by a joint-probability distribution over all these random variables. This distribution is itself assumed to be common knowledge among the bidders. Thus, but for the differences in their information, the bidders would share identical beliefs about the item being offered and about their competitors.

The auctions available to individual bidders depend on the particular auction game being played. In a sealed-bid auction, each bidder selects a single number, representing his bid. The payoff to the winning bidder will most often be the excess of the value of the item to the bidder over the price he pays. In most bidding games, the losing bidders' payoffs is set equal to zero.

A *strategy* for a bidder specifies what action to take as a function of what he knows. For example, in a sealed-bid auction, a bidder who estimates the value of the item to himself at x might bid an amount $\beta(x)$. Then, the function β is his strategy.

All formal auction models developed to date have assumed that the bidders do not collude or cooperate with each other. The noncooperative solution most often used in game theory is the *Nash equilibrium* or one of its variants. A set of strategies (one for each player) is a Nash equilibrium if every player, believing that his competitors will use their equilibrium strategies, maximizes his own payoff by playing his equilibrium strategy. Thus, if it were common knowledge that all planned to use their equilibrium strategies, no bidder would have an incentive to alter his plan.

There have been many formal equilibrium models of competitive bidding. The assumptions of these models differ along many dimensions, but two of these dimensions are of preeminent importance for our analysis: the determinants of the bidders' payoffs (*value assumptions*) and the determinants of their beliefs about their competitors (*distributional assumptions*).

The original Vickrey [1961, 1962] studies of competitive bidding as well as many more recent studies (Griesmer, Levitan, and Shubik [1967], Ortega-Reichert [1968], Matthews [1979], Holt [1980], Riley and Samuelson [1981], Harris and Raviv [1982]) adopted the *private-values* assumption. According to this assumption, a bidder's payoffs can depend only on (i) what he knows, (ii) whether he wins, and (iii) how much he pays. Value is treated as a purely personal matter: each bidder knows what the goods are worth to himself, and no bidder cares what they are worth to others, except possibly as strategic information to be used in choosing a bid. Thus, the object being sold is not a painting that may eventually be resold for a price depending on the tastes of others; it is not the rights to minerals or timber on federal territory, whose value depends on the unknown amounts and composition of the recoverable minerals or timber; and it is not a financial security, whose value depends on future prices, dividends, interest rates, and the like.

In extreme contrast to the private-values assumption is the *common-value* assumption, used most often in analyzing auctions for mineral rights (Wilson [1967], Ortega-Reichert [1968], Rothkopf [1969], Reece [1978], Engelbrecht-Wiggans, Milgrom, and Weber [1981], Milgrom and Weber [1982b]). According to the common-value assumption, the actual value (V) of the mineral rights is the same to all bidders, but the value will not be known until the resource is extracted. Also, the bidders may presently differ in their estimates of V.

Let X_i be bidder i's estimate of V, and suppose that it is an unbiased estimate; that is, $E[X_i|V] = V$. Other things being equal, the winning bidder will be the one with the highest estimate. The estimate of the winning bidder is biased upward; that is, $E[\max X_i|V] > \max E[X_i|V] = V$ (because "max" is a convex function). Intuitively, a bidder wins often when he overestimates V but wins only rarely when he underestimates V. Consequently, even if his estimates are unbiased, his estimates will be systematically high in those cases where he wins. The phenomenon is called the *winner's curse*.

Oil companies bidding for tracts of land in unexplored territory are allegedly among those accursed (Capen, Clapp, and Campbell [1971]), and other examples of the curse are easy to find. The new entrant in a local construction industry who is inexperienced in making competitive cost estimates may find that his bid wins the job only when he has underestimated the construction cost – he suffers from the winner's curse (cf. Brown [1975]). The university that makes the highest salary offer to a young assistant professor and the first-time homebuyer who outbids all rivals may both suffer the curse.

The analysis of common-value models focuses on the winner's curse and related issues. How does one bid to alleviate the curse? (Cautiously!) What are the incentives to acquire information? (Private information alleviates the curse for the informed bidder and intensifies the curse for his competitors.) How efficiently do prices aggregate information? How should the seller manage any information to which he may be privy?

A few papers develop bidding models that accommodate both private-values and common-value models, as well as a range of intermediate models (Wilson [1977], Milgrom [1979a,b, 1981a], Milgrom and Weber [1982a,b]). The most general models allow the value of the goods to any one bidder to depend on his tastes, the tastes of other bidders, the preferences of nonparticipants, and various unobserved qualities of goods. All of the issues described can be treated in this framework.

The second important way in which bidding models differ is in their distributional assumptions. The value estimate that a bidder makes – or his reservation price in a private-values model – is represented by a real-valued random variable, called the bidder's *type*. One common assumption is that the types are statistically independent. In an auction for a painting, the independent types assumption rules out the possibility that a bidder, finding the painting to be quite beautiful, might expect others to admire it as well. In the sale of mineral rights, it rules out the possibility that a bidder, upon receiving a discouraging geological report, may expect his competitors to receive discouraging reports. In Section 4, we consider a model in which the bidders' types are positively correlated or, more precisely, *affiliated*. The concept of affiliation is defined in that section.

The *independent private-values* model combines the independent types and private-values assumptions. We begin the formal analysis of auctions by studying that model. Henceforth, we make the simplifying assumption that each bidder wants only one unit of the good.

3 The independent private-values model

To penetrate to the heart of a bidder's decision problem, it is useful to abstract from the particular rules of the auction and focus directly on

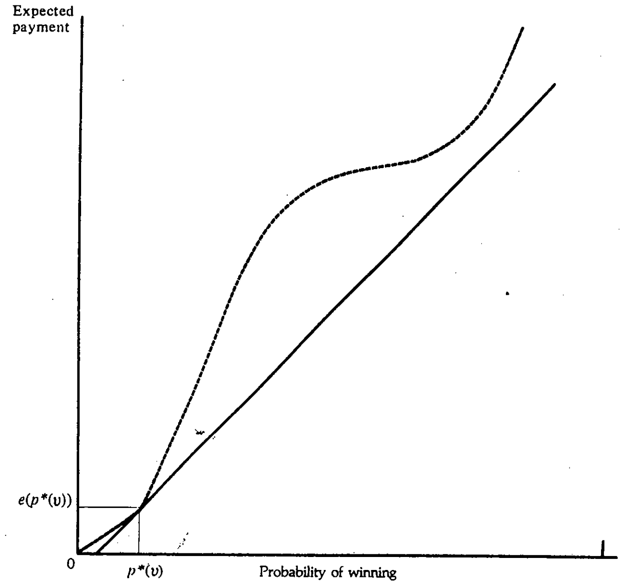

Figure 1

matters of concern to a bidder. Given the private-values assumption, if the bidder is risk neutral, his choice of bids affects his payoff only to the extent that it affects either his probability of winning a unit or the amount he expects to pay. Figure 1 displays the bidder's abstract choice problem.

In the figure, the curve $(p, e(p))$ represents the menu of (probability, expected payment) pairs available to the bidder. If the bidder values a unit at v, his expected payoff u from the point $(p, e(p))$ is $u = p \cdot v - e(p)$, so his indifference curves are lines with slope v. The intercept of the indifference line with the vertical axis is $-u$, the negative of the expected payoff.

Let $(p^*, e(p^*))$ denote the most preferred point for the bidder; then we may write $p^* = p^*(v)$.[2] One can see from the figure that the p^* function must be increasing and that if $e(\cdot)$ is differentiable at $p^*(v)$, then $e'(p^*(v)) = v$. Even if e is not differentiable, one can establish the following result.

[2] If there are several optimal points, p^* may be chosen arbitrarily from among them. Note that p^* is essentially the dual functional of e.

Lemma 1 (Myerson [1981], Riley and Samuelson [1981]):

$$e(p^*(x)) = e(p^*(0)) + \int_0^x s\, dp^*(s).$$

In most treatments of the independent private-values model, it is assumed that the seller sets some minimum price r (the "reserve" price) and that bidders' types are drawn from some known distribution F with a positive density f. For the Vickrey and discriminatory auctions, a strategy is a function β that maps a bidder's type x into his nonnegative bid $\beta(x)$. Because the rules of the auctions treat bidders symmetrically and because we have made symmetric assumptions about the bidders, we shall look for a symmetric equilibrium point.

For the Vickrey auction, the analysis of equilibrium is relatively easy. Let there be n bidders ($n > k$) and let their types be $\mathbf{X}_1, \ldots, \mathbf{X}_n$. Let $\mathbf{Y}_1, \ldots, \mathbf{Y}_{n-1}$ denote the maximum, ..., minimum among $\mathbf{X}_2, \ldots, \mathbf{X}_n$. A typical bidder, say bidder 1, has a dominant strategy. If he observes $\mathbf{X}_1 = x$, he bids $\beta(x) = x$. If the others also adopt this strategy, bidder 1 of type x will be awarded a unit of the good if $x > \max(r, \mathbf{Y}_k)$, and the price charged will be $\mathbf{Y} = \max(r, \mathbf{Y}_k)$. Bidder 1's probability of winning at equilibrium with type x is zero if $x < r$ and is $P\{\mathbf{Y} < x\}$ otherwise. In summary, we have $p^*(x) = P\{\mathbf{Y} \leqslant x\}$ and $e(p^*(x)) = P\{\mathbf{Y} \leqslant x\} \cdot E[\mathbf{Y}|\mathbf{Y} \leqslant x]$.

In the discriminatory auction, the required monotonicity of p^* implies that the symmetric equilibrium strategy β must be nondecreasing, and it isn't hard to show that β must be increasing. Hence, for this auction, $p^*(x) = P\{\mathbf{Y} \leqslant x\}$, just as it was for the Vickrey auction. Also, as with the Vickrey auction, $e(p^*(0)) = e(0) = 0$. Hence, by Lemma 1, the bidders must have the same expected payment function as in the Vickrey auction. Then since $e(p^*(x)) = p^*(x) \cdot \beta(x)$, we have the following result.

Proposition 1 (Vickrey [1961], Ortega-Reichert [1968]): *The symmetric equilibrium strategy of the discriminatory auction game is given by*

$$\beta(x) = \begin{cases} E[\mathbf{Y}|\mathbf{Y} < x] & \text{if } x \geqslant r \\ 0 & \text{if } x < r. \end{cases}$$

The expected price, the expected number of sales, and the seller's expected revenue from the discriminatory auction are the same as from the Vickrey auction.

One can analyze a wide variety of auctions by the same technique and always reach the same conclusion: the equilibrium allocation of goods,

the expected bidder profits, and the expected seller revenues are the same as in the Vickrey auction. A general statement of this equivalence principle, based on the lemma, goes as follows.

Proposition 2 (Revenue equivalence principle): *Consider the independent private-values symmetric auction model with risk-neutral bidders. If, at equilibrium, units of the good are awarded to just those bidders whose types exceed* $\mathbf{Y}$, *and if bidders whose types are less than* r *pay nothing, then the expected payment of a bidder of type* x *is equal to* $P\{\mathbf{Y} \leqslant x\} \cdot E[\mathbf{Y}|\mathbf{Y} < x]$.

Proposition 2 applies to a large number of auction games including both simultaneous auctions like the discriminatory and Vickrey auctions and some *sequential* auctions – auctions where units are sold in sequence, one at a time. The proposition can be used to derive the equilibrium strategies in some sequential auction games. For example, suppose k objects are to be sold by a sequence of k discriminatory auctions. Suppose that the winning bid is announced after each round, and let $b_1, b_2, \ldots, b_k$ denote these winning bids. A strategy is a collection of functions $\beta_1(\cdot), \beta_2(\cdot; b_1), \ldots, \beta_k(\cdot; b_1, \ldots, b_{k-1})$ that specify how to bid at each round, as a function of the bidder's type and the winning bids at earlier rounds.

Let us hypothesize that each $\beta_j(\cdot; b_1, \ldots, b_{j-1})$ is increasing, so that the winner at the jth round will be the jth highest type. By Proposition 2, at equilibrium, the expected payment by a bidder of type x ($x \geqslant r$) who has not yet won an object after the jth round is $P\{\mathbf{Y} \leqslant x|\mathbf{Y}_j > x\} \cdot E[\mathbf{Y}|\mathbf{Y} < x, \mathbf{Y}_j > x]$. Starting with $j = k$ and working backward, one can deduce that a bidder's expected payment at stage j is $P\{\mathbf{Y}_{j-1} > x > \mathbf{Y}_j\} \cdot E[\mathbf{Y}|\mathbf{Y}_{j-1} > x > \mathbf{Y}_j]$, and from that one can guess the equilibrium strategy.

Proposition 3: *The symmetric equilibrium strategy in the sequential discriminatory auction game is given by*

$$\beta_j(x; b_1, \ldots, b_{j-1}) = \begin{cases} E[\mathbf{Y}|\mathbf{Y}_{j-1} > x > \mathbf{Y}_j] & \text{if } x \geqslant r \\ 0 & \text{if } x < r \end{cases}$$

for $j = 1, \ldots, k$.

One can similarly analyze the sequential Vickrey auction. That auction has no dominant-strategy equilibrium, because at each round (except the last) the bidder must weigh the value of submitting a low bid at the current round and possibly acquiring a unit at a later round against the value of acquiring a unit immediately. It is clear that at the last round, the dominant strategy is $\beta_k(x; b_1, \ldots, b_{k-1}) = x$. At the next to last round, the expected payment of a player of type x who has not yet won is $P\{\mathbf{Y} \leq x | \mathbf{Y}_{k-2} < x\} \cdot E[\mathbf{Y}|\mathbf{Y}_k < x < \mathbf{Y}_{k-2}]$ by Proposition 2. If he wins at the last round, his expected payment is $E[\mathbf{Y}|\mathbf{Y}_k < x < \mathbf{Y}_{k-1}]$, so his conditional expected payment at the next to last round, conditional on winning then, must be $E[\mathbf{Y}|\mathbf{Y}_{k-1} < x < \mathbf{Y}_{k-2}]$. This must be equal to $E[\beta_{k-1}(\mathbf{Y}_{k-1}; b_1, \ldots, b_{k-2})| \mathbf{Y}_{k-1} < x < \mathbf{Y}_{k-2}.]$ So, $\beta_{k-1}(x; b_1, \ldots, b_{k-2}) = E[\mathbf{Y}|\mathbf{Y} < x < \mathbf{Y}_{k-1}]$. Arguing inductively leads to the following result.

Proposition 4: *The symmetric equilibrium strategy in the sequential Vickrey auction game is given by*

$$\beta_j(x; b_1, \ldots, b_{j-1}) = \begin{cases} E[\mathbf{Y}|\mathbf{Y}_{j+1} < x < \mathbf{Y}_j] & \text{if } x \geq r \\ 0 & \text{if } x < r \end{cases}$$

for $j = 1, \ldots, k$.

Several authors have considered the problem of designing an auction for a single object that maximizes the seller's expected revenue (an "optimal auction"). The intent of this theory is both positive (to explain existing selling arrangements) and normative (to determine an appropriate auction design for the sale of T-bills or government-owned properties, or for purchasing supplies or services). The various papers in the "optimal auctions" literature vary in their distributional assumptions: Myerson [1981] and Riley and Samuelson [1981] assume that the bidders' types are continuously distributed on a convex set, while Harris and Raviv [1982] and Maskin and Riley [1981] work with discrete distributions. By adopting Vickrey's [1961, 1962] original approach to modeling uncertainty in auctions, one can unify these diverse results. Indeed, using this approach, we extend them to the case where more than one item is to be sold, and more generally to the case where the kth item can be produced and delivered at a cost c_k where the sequence $\{c_k\}$ is nondecreasing. However, we retain the independence and private-value assumptions along with the assumptions that bidders are risk neutral and that each bidder desires only one item.

Vickrey's approach was this: Let each bidder's type be distributed uniformly on the interval (0,1). A bidder of type x has a reservation price $\bar{F}(x)$, where $\bar{F}$ is some nondecreasing function. Thus, a bidder of type x who acquires a unit for the price b has a payoff of $\bar{F}(x) - b$. Essentially,

$\bar{F}$ is the inverse of the distribution F of bidder reservation prices: for y in the support of F, $\bar{F}(F(y)) = y$. In general, for any given distribution of reservation prices F, $\bar{F}(x) = \inf\{t|F(t) > x\}$.

If, faced with a single buyer, the seller sets a price of $\bar{F}(t)$, he will sell the goods with probability $(1 - t)$, so his expected revenue is $(1 - t)\bar{F}(t)$. If this expected revenue function is not concave, then for some values of t there exist types $t' > t > t''$ and $s \in [0,1]$ such that

$$st' + (1 - s)t'' = t,$$

and

$$s(1 - t')\,\bar{F}(t') + (1 - s)(1 - t'')\,\bar{F}(t'') > (1 - t)\,\bar{F}(t).$$

In that case, the seller could randomize among the prices $\bar{F}(t')$ and $\bar{F}(t'')$ with probabilities s and $(1 - s)$, sell the object with probability $1 - t$, and gain more expected revenue than would be possible by fixing a price of $\bar{F}(t)$. Hence, the maximal expected revenue to a seller among all schemes that involve picking a price at random, and for which the probability of sale is exactly $1 - t$ is $H(t)$, where H is the minimum concave function such that for all t, $H(t) \geq (1 - t)\bar{F}(t)$.

As was noted earlier, in any auction game, the probability of acquiring a unit p^* must be a nondecreasing function of the bidder's type, and the seller's expected revenue depends only on that probability function p^* and $e(p^*(0))$. The seller can implement any nondecreasing p^* by randomizing over the price he demands. Then $-e(p^*(0))$ can be interpreted as a lump sum the seller chooses to pay buyers just for participating in the auction. In an expected-revenue maximizing auction, one must certainly have that $e(p^*(0)) = 0$. Therefore, *$H(t)$ is the maximal expected revenue to a seller among all selling schemes that sell with a probability of precisely* $1 - t$.[3] If the cost of supplying the item is c, then the seller's optimal selling scheme can be determined by maximizing $H(t) - (1 - t)c$. If the optimum value t^* is not zero or one, we must have $H'(t^*) = -c$. Generally, to allow that H' may not be differentiable at t^*, let $t^* = \inf\{t|H'(t) < -c\}$.

The general optimal selling problem for the independent private-values model with n risk-neutral buyers can be formulated by allowing the cost c_k of supplying the kth unit to vary with k. Consider the case where $c_1 \leq c_2 \leq \ldots$. By setting $c_1 = 0$ and $c_j = +\infty$ for all $j \geq 2$, one obtains the optimal selling strategy for a single costless object. Except for Harris and Raviv [1981], previous analysis have limited attention to this special case.

[3] Roger Myerson suggested this interpretation of H to me.

One particularly transparent case to analyze that suggests the general solution is the case of constant unit production costs c. In that case, the optimal selling problem decomposes into n individual problems. Applying the previous analysis leads to the conclusion that an optimal scheme is to sell to buyer i if $H_i'(x_i) \leqslant -c$, at a price determined as in the one-buyer problem. The generalizable attributes of this solution turn out to be the following ones: (i) the number of items sold is the largest k such that there are at least k buyers with $H_i'(x_i) \leqslant -c_k$ (by convention, $c_o = 0$), (ii) the buyers awarded items are the k buyers with the smallest values of $H_i'(x_i)$ (ties can be broken in any fashion), and (iii) the price is determined as in the one-buyer problem.

Combining this result with the revenue equivalence theorem leads to the following result, which generalizes several previous optimal auction results.

Proposition 5: *A symmetric auction game with production cost c_k for the kth item maximizes seller profits if at its equilibrium, (i) the number of items sold is the largest k such that at least k bidders have $H'(x_i) \leqslant -c_k$, (ii) the items are awarded to the k buyers for which $H'(x_i)$ is lowest, with ties broken at random, and (iii) the expected payment by each bidder not awarded an item is zero.*

A great deal of attention has been focused on the case where exactly k items are offered for sale and the seller's personal reservation price is c. One question that is sometimes asked is: How do the standard auction mechanisms perform in this setting?

The answer hinges on the concavity of the function $(1 - t)\bar{F}(t)$. When that function is not concave, the standard auctions do not maximize expected seller revenues; but when it is concave (as it is when F is the cumulative distribution function of a normal, exponential, or uniform distribution), then the discriminatory, Vickrey, sequential discriminatory, and sequential Vickrey auctions with a reserve price of $\bar{F}(t^*)$ (where $H'(t^*) = c$) are all expected-revenue–maximizing auctions. This result relies on the independent types, private values, symmetry, and risk-neutrality assumptions. It does not generalize to the case of risk-averse bidders.

When bidders are risk averse, it is still a dominant strategy in the Vickrey auction for a bidder of type x to bid x. In the discriminatory auction, however, risk-averse bidders will tend to bid higher than is required to maximize expected profits. To see why this is so, notice that a small increase Δb in the bid from the expected-profit–maximizing level reduces profits on the order of $(\Delta b)^2$ (that follows from the first-order maximality condition), but the increase Δb reduces the riskiness of the lottery a bidder faces by raising his chance of winning (on the order of Δb) and reducing his expected profit conditional on winning.

Implicitly, by raising their bids slightly, risk-averse bidders can buy "partial insurance," on actuarially fair terms, against losing, and this opportunity is seized in the equilibrium strategies of the discriminatory, sequential discriminatory, and sequential Vickrey auctions. That fact leads to this result.

Proposition 6 (Matthews [1979], Holt [1980], Maskin and Riley [1980], Harris and Raviv [1982], Milgrom and Weber [1982a]): *In the independent private-values model with risk-averse bidders, the expected payment by a winning bidder of any type x (conditional on winning) is larger in the discriminatory auction than in the Vickrey auction. In particular, the expected revenue is larger in the discriminatory auction.*

Proposition 6 provides some formal justification for the argument that price discrimination in T-bill auctions raises the government's expected revenue from the sale. However, we shall see in Section 4 that statistical dependence among the bidders' types favors Vickrey auctions over discriminatory auctions, so the matter remains ambiguous.

Two common auction forms that we have not yet discussed in this section are the English and the Dutch. These auctions are most commonly used to sell one object at a time. In the English auction for a single object, each bidder has a dominant strategy: he should remain active until the price called by the auctioneer exceeds his reservation price. If this strategy is universally adopted, the object will be awarded to the bidder who values it most highly for a price equal to the second highest valuation. This outcome is identical to the outcome in the Vickrey auction, and for that reason Vickrey considered the two auctions equivalent. Actually, this equivalence hinges on the private-values assumption: when we drop that assumption, quite a different conclusion will emerge.

In the Dutch auction, a bidder of type x must decide as the price falls whether to stop the auction and claim the prize or whether to let the price continue to fall. Given any strategy that the bidder may adopt, there is a highest price $b = \beta(x)$ at which he will stop the auction. Thus, a strategy can be described as a function from types into nonnegative real numbers, which we may call "bids." The bidder choosing the highest "bid" wins and is awarded the object for a price equal to his "bid."

Notice that the Dutch auction game is identical to the discriminatory auction game. In both games, the bidder selects a "bid" as a function of his type, the high bid wins, and the winning bidder pays his bid. In short, the Dutch and discriminatory auctions are strategically equivalent, and this equivalence, unlike that of the Vickrey and English auctions, does not depend on any value or distributional assumption. However, experimental evidence obtained by Cox, Roberson, and Smith [in press] tends to refute this conclusion.

4 Affiliated types and monotone values

The model that we study in this section is more general in both its value assumptions and its distributional assumption than the independent private-values model. Specifically, let $\mathbf{X}_1, \ldots, \mathbf{X}_n$ be the types of the n bidders, and let $\mathbf{S}_1, \ldots, \mathbf{S}_m$ be any other variables that may influence the value of the goods to the bidders. The value of a unit to bidder i is designated by $\mathbf{V}_i$ where

$$\mathbf{V}_i = V(\mathbf{X}_i, \{\mathbf{X}_j\}_{j \neq i}, \mathbf{S})$$

where $\mathbf{S} = (\mathbf{S}_1, \ldots, \mathbf{S}_m)$. The expression $\{\mathbf{X}_j\}_{j \neq i}$ is designed to emphasize the assumed symmetry of the valuation function; bidder, i's payoffs may depend on the preferences of the competing bidders, but only in a symmetric fashion. The *monotone values* assumption asserts that the valuation function V is nondecreasing and $E[\mathbf{V}_i]$ is finite. The private-values assumption is the special case of the monotone-values assumption in which $\mathbf{V}_i = \mathbf{X}_i$. The common-value assumption is the special case in which $\mathbf{V}_i = \mathbf{S}_1$. Intermediate cases, in which the bidders' valuations are partly a matter of personal preference and partly dependent on observed qualities, can also be accommodated by the monotone-values assumption. Like the common-value model, these intermediate models include aspects of a winner's curse.

The equilibrium strategies in an auction game will implicitly reflect an adjustment for the winner's curse. Thus, consider bidder 1's problem in a first-price auction. Let $\mathbf{W}_1$ denote the highest bid among the $n - 1$ opposing bidders; it is a random variable from bidder 1's point of view. The bidder's expected payoff from a bid of b is

$$E[(\mathbf{V}_1 - b)\, 1_{\{\mathbf{W}_1 < b\}} | \mathbf{X}_1] = P\{\mathbf{W}_1 < b \mid \mathbf{X}_1\}\, E[\mathbf{V}_1 - b | \mathbf{X}_1, \{\mathbf{W}_1 < b\}],$$

where $1_{\{\mathbf{W}_1 < b\}}$ designates an indicator function that is one if $W_1 < b$ and zero otherwise. In words, the bidder computes his expected payoff as his probability of winning (given the information $\mathbf{X}_1$) times his conditional expected winnings $\mathbf{V}_1 - b$ *given both on his actual information* $\mathbf{X}_1$ *and the hypothesis that* $\mathbf{W}_1 < b$; the bidder anticipates the winner's curse in choosing his bid.

In addition to the monotone-values assumption, assume that all of the random elements of the model have positive partial correlations, and further that these positive correlations are preserved conditional on arbitrary restrictions on the ranges of the individual variables. Such restrictions arise in bidding models when one bidder learns or conjectures something about his competitors' bids.

The restriction on distributions just described is not vacuous. Indeed, it is identical to the assumption that the exogenous random elements of the bidding model are affiliated. A random vector $\mathbf{Z} = (\mathbf{Z}_1, \ldots, \mathbf{Z}_l)$ is *affiliated* if for every z and z' in R_l, (*) $f(z)\, f(z') \leq f(z \vee z')\, f(z \wedge z')$, where

$$z \vee z' = (\max(z_1, z_1'), \ldots, \max(z_l, z_l'))$$

$$z \wedge z' = (\min(z_1, z_1'), \ldots, \min(z_l, z_l')).$$

The inequality (*) is known as the *affiliation inequality* (and also as the "FKG inequality" and the "MTP_2 property"). The general theory of affiliation has been developed by Milgrom and Weber [1982a].[4]

Notice that if $\mathbf{Z}_1, \ldots, \mathbf{Z}_l$ are independent, then the affiliation inequality holds an equality. Also, in common-value models, it is usual to specify that the bidder's types $\mathbf{X}_1, \ldots, \mathbf{X}_n$ are independent estimates of $\mathbf{S}_1$ drawn from some common distribution with density $f(\cdot|s_1)$, where the family of densities $\{f(\cdot|s_1)\}$ is lognormal with mean s_1, or exponential with mean s_1, or uniform on $(0, s_1)$, or some other family with the monotone likelihood ratio property. In all such cases, regardless of the prior distribution for $\mathbf{S}_1$, the vector $(\mathbf{S}_1, \mathbf{X}_1, \ldots, \mathbf{X}_n)$ is affiliated.

The key property of affiliation for our analysis is this:

Proposition 7 (Milgrom and Weber [1982a]): *If* $(\mathbf{Z}_1, \ldots, \mathbf{Z}_l)$ *is affiliated and* g *is a nondecreasing function, then the function*

$$h(a_1, \ldots, a_l; b_1, \ldots, b_l) = E[g(\mathbf{Z}) \mid a_1 \leq \mathbf{Z}_1 \leq b_1, \ldots, a_l \leq \mathbf{Z}_l \leq b_l]$$

is nondecreasing.

To describe the equilibria of various auction games, it is useful to define the random variables $\mathbf{Y}_1, \ldots, \mathbf{Y}_{n-1}$ to be the maximum, ..., minimum from among $\mathbf{X}_2, \ldots, \mathbf{X}_n$. It has been shown (Milgrom and Weber [1982a]) that $\mathbf{S}, \mathbf{X}_1, \mathbf{Y}_1, \ldots, \mathbf{Y}_{n-1}$ are affiliated.

Let $F_j(y|x) = P\{\mathbf{Y}_j \leq y | \mathbf{X}_1 = x\}$ and let $f_j(y|x)$ be the corresponding density. Let $v_j(x, y) = E[\mathbf{V}_1 | \mathbf{X}_1 = x, \mathbf{Y}_j = y]$. In view of Proposition 7, v_j is nondecreasing.

[4] A survey of the theory of affiliated variables is also given by Karlin and Rinott [1980].

The model with monotone values and affiliated types is called the *general symmetrical model*. The equilibrium strategies for the Vickrey and discriminatory auctions for this model are given by the next two propositions, where $x^* = \inf\{x | E[\mathbf{V}_k | \mathbf{X}_1 = x, \mathbf{Y}_k \leqslant x] \geqslant r\}$; x^* is called the *screening level* corresponding to r.

Proposition 8 (Milgrom [1981]): *The symmetric equilibrium strategy in the Vickrey auction game is given by*

$$\beta^V(x) = \begin{cases} v_k(x, x) & \text{if } x \geqslant x^* \\ 0 & \text{if } x < x^*. \end{cases}$$

Proposition 9 (Milgrom and Weber [1982a,b], Wilson [1977]): *The symmetric equilibrium strategy in the discriminatory auction game is given by*[5]

$$\beta^D(x) = \begin{cases} r \cdot L_k(x^*|x) + \int_{x^*}^{x} v_k(t, t)\, dL_k(t|x) & \text{if } x \geqslant x^* \\ 0 & \text{if } x < x^* \end{cases}$$

where

$$L_k(\alpha|x) = \exp\left[- \int_{\alpha}^{x} \frac{f_k(t|t)}{F_k(t|t)}\, dt \right].$$

Notice that with a private-values assumption, the equilibrium strategy in the Vickrey auction becomes $\beta^V(x) = x$. With independent types, $L_k(t|x) = F_k(t) / F_k(x)$, and the discriminatory auction equilibrium strategy becomes $\beta^D(x) = E[\max(r, v_k(\mathbf{Y}_k,\mathbf{Y}_k)) \mid \mathbf{Y}_k < x]$, which further simplifies under private values to $\beta^D(x) = E[\max(r,\mathbf{Y}_k) | \mathbf{Y}_k < x]$, in accordance with Proposition 1.

Now, motivated by the Treasury bill controversy, we may ask which of the two auctions, Vickrey or discriminatory, leads to greater revenues for the seller in the general symmetric model.

Proposition 10 (Milgrom and Weber [1982a,b]): *The expected price paid by a winning bidder of type x in the Vickrey auction is as high as, or higher than, that paid in the discriminatory auction:*

$$\beta^D(x) \leqslant E[\max(r, \beta^V(\mathbf{Y}_k)) \mid \mathbf{X}_1 = x, \mathbf{Y}_k < x].$$

Consequently, the expected revenue to the seller is higher for the Vickrey auction than for the discriminatory auction.

[5] For s not in the support of $\mathbf{X}_1$, $f_k(s|s)/F_k(s|s)$ is taken to be zero.

Results reported by Tsao and Vignola [1977] tend to confirm this conclusion using data drawn from the weekly T-bill auction.

An important general intuition lies behind Proposition 10 that will shortly enable us to generate a host of similar comparisons. The key insight is best seen by abstracting from the particular rules of the auction and representing a bidder's choice problem in a new way. In each auction game that we have studied, the bidders' strategies were increasing functions of their types, and by bidding $\beta(x)$, any bidder, regardless of his actual type, could arrange to win whenever $\max(x^*,\mathbf{Y}_k) < x$. We may therefore think of a bidder as choosing x, rather than as choosing particular bids.

Consider the problem faced by bidder 1 when his type is z in some auction game "A". If he chooses some $x \geqslant x^*$ and wins, then conditional on winning the expected value received is $R(x, z) = E[V_1 | X_1 = z, \mathbf{Y}_k < x]$ and the expected payment is some amount $W^A(x, z)$. The two sides of the inequality in Proposition 10 represent $W^A(x, x)$ for the discriminatory and Vickrey auctions, respectively.

When bidder 1 chooses x, his expected payoff is $[R(x, z) - W^A(x, z)] F_k(x|z)$. At a symmetric equilibrium, it must be optimal for the bidder to choose $x = z$, so the first-order necessary condition is

$$0 = [R(z, z) - W^A(z, z)]f_k(z|z) + [R_1(z, z) - W_1^A(z, z)]F_k(z|z),$$

where subscripts on R and W^A denote partial derivatives. Solving for W_1^A, one can compute the total derivative of W^A.

$$\frac{d}{dz} W^A(z, z) = W_1^A(z, z) + W_2^A(z, z)$$
$$= R_1(z, z) + [R(z, z) - W^A(z, z)]\,\frac{f_k(z, z)}{F_k(z|z)} + W_2^A(z, z). \qquad (**)$$

In each of the auctions studied, $W^A(x^*, x^*) = r$, and that boundary condition plus the differential equation (**) completely determine the $W^A(z, z)$ function. The differential equations corresponding to different auctions are identical, except for the W_2^A term.

It now follows that the winner's expected payment across different auctions can be compared by comparing the partial derivatives W_2^A. When the bidders' types are statistically independent, W_2^A is necessarily zero for all auction games. This fact can be used to derive Proposition 2, the revenue equivalence principle. More generally, different auction games yield different average revenues. The key to comparing games is the *linkage principle*, which is a consequence of equation (**).

The linkage principle: Let A and B be two auction games with symmetric equilibrium at which (1) units are awarded to all bidders with types $z > \max(x^*, \mathbf{X}_{(k+1)})$ and (2) bidders with types $z \leq x^*$ have expected payoff zero. If $W_2^A(z, z) \geq W_2^B(z, z)$ for all $z \geq x^*$, then $W^A(z, z) \geq W^B(z, z)$ for all $z \geq x^*$.

The function W_2^A summarizes the effect of the bidder's unobserved type on the amount he expects to pay. That effect arises only when the variables of the model are statistically dependent. As the winning bidder's actual type z rises, he expects the types of others to rise as well. If the price he pays depends on those types, then he expects the price to rise; that is, $W_2^A > 0$.

In the discriminatory auction, $W^D(x, z) = \beta^D(x)$ and $W_2^D = 0$: the equilibrium price is determined exclusively by the winner's type, so there is no linkage at all. In the Vickrey auction, $W^V(x, z) = E[\max(r, v_k(\mathbf{Y}_k, \mathbf{Y}_k))|X_1 = z, \mathbf{Y}_k < x]$. The Vickrey price depends on $\mathbf{Y}_k$ which is statistically linked to $\mathbf{X}_1$. By Proposition 6, $W_2^V \geq 0$. These observations and the linkage principle establish Proposition 10.

A second application of the linkage principle arises in connection with analyzing whether the seller, if he has information $\mathbf{X}_0$ of his own, should establish a policy of revealing that information, or whether he would be better off concealing it. A key word here is "policy": the idea is that the seller must commit himself to revealing information according to some rule. For example, the U.S. Congress could instruct the Department of Interior to conduct geologic surveys before offering the mineral rights on any piece of property. It could also order the Department to adopt a policy of always reporting the survey in detail, or summarizing it, or reporting it only if the information is favorable. In this last case, however, the bidders would "hear the silence" – withholding the report would be a sure sign of an unfavorable survey.

When $\mathbf{X}_0$ is reported, the set of bidders willing to bid at least r will be changed. However, the seller can choose to use the survey data to set an appropriate reserve price. We shall say that the seller has adopted a *fixed screening-level* policy at x^* if, upon reporting his information $\mathbf{X}_0$, he sets the reserve price $r = r(\mathbf{X}_0)$ to attract exactly those bidders whose types are at least x^*. When the seller reports all of his information and adopts a fixed screening-level policy at x^*, it can be shown that the conditions of the linkage principle are satisfied. That leads to the following result.

Proposition 11 (Milgrom and Weber [1982a]): *In a discriminatory auction, a policy of revealing $\mathbf{X}_0$ and fixing the screening level at x^* results in greater expected revenue than withholding the information and setting the screening level at x^*.*

Adopting the prescribed policy links the price to $\mathbf{X}_0$, and $\mathbf{X}_0$ is affiliated with the winning bidder's estimate. That linkage results in W_2^A being positive for the auction with $\mathbf{X}_0$ announced. Proposition 11 follows as a consequence.

If the seller reports only a summary statement about $\mathbf{X}_0$ – for example if he reports only that $\mathbf{X}_0$ lies in some set T_0 – then conditional on that information, the variables $\mathbf{X}_0, \mathbf{X}_1, \ldots, \mathbf{X}_n$ are still affiliated. Hence, by Proposition 11, reporting the remaining details of $\mathbf{X}_0$ and fixing the screening level would further raise expected revenues. This observation leads to Proposition 12.

Proposition 12 (Milgrom and Weber [1982a]): *In a discriminatory auction, the expected-revenue maximizing policy for revealing information and setting a reserve price involves reporting $\mathbf{X}_0$ precisely and in full detail.*

In the Vickrey auction, there is already some linkage of the price to variables other than the winner's type, but one can sometimes raise revenues further by introducing additional linkages. For example, the seller can raise revenues by reporting his information $\mathbf{X}_0$.

Proposition 13 (Milgrom and Weber [1982a]): *In a Vickrey auction, a policy of revealing $\mathbf{X}_0$ and fixing the screening level at x^* results in greater expected revenues than withholding the information and setting the screening level at x^*.*

Proposition 14 (Milgrom and Weber [1982a]): *In a Vickrey auction, the expected-revenue maximizing policy for revealing information and setting a reserve price involves reporting $\mathbf{X}_0$ precisely and in full detail.*

In a discriminatory auction, revealing information has two effects. First, if $\mathbf{X}_0, \ldots, \mathbf{X}_n$ are strictly affiliated, then the seller's information tells each bidder something about his competitor's type. As a result, the bidders with lower types will, on average, raise their assessments of their competitors' bids. They will then bid more aggressively, and that will tend to make the higher types raise their bids, too. This effect is present even in a private-values model if types are affiliated. The next effect appears only in models involving the winner's curse. In such models, bidders tend to shade their bids to avoid the curse, and the lower types, who are overly pessimistic, shade their bids excessively. When the seller provides public information, he alleviates the winner's curse, allowing lower types to bid more aggressively on average, which in turn causes everyone to bid more aggressively. This raises revenues.

In contrast to the discriminatory auction, revealing information has no effect in the Vickrey auction game with private values, since the bidders all follow their dominant strategy ($\beta^V(x) = x$) regardless of the content of the seller's information. But if the private-values assumption does not hold and a winner's curse effect is present, then revealing affiliated information alleviates the winner's curse and causes the average price to rise: the inequality in Proposition 13 becomes strict.

Recall that in any private-values model, the Vickrey and English auctions are equivalent. In the general model, this equivalence may not hold. The equivalence or lack thereof depends on how much of the bidding behavior of the $n - 2$ lowest types can be observed during the auction by the two highest types. If none of their bidding can be observed, then the Vickrey and English auctions are strategically equivalent. Milgrom and Weber [1982a] have computed equilibrium strategies for the case where all of the bidding behavior can be observed. Observing those bids passes information to the last two bidders, with much the same effect (at equilibrium) as if the seller had revealed his information: Prices become linked monotonically to the types of $n - 2$ lowest bidders as well as to the types of the second highest bidder. Then, three new propositions follow, using the linkage principle.

Proposition 15 (Milgrom and Weber [1982a]): *The expected price paid by a winning bidder in the English auction is as high as, or higher than, that paid in the Vickrey auction.*

Proposition 16 (Milgrom and Weber [1982a]): *In an English auction, a policy of revealing $\mathbf{X}_0$ and fixing the screening level at x^* results in greater expected revenue than withholding the information and setting the screening level at x^*.*

Proposition 17 (Milgrom and Weber [1982a]): *In an English auction, the expected-revenue maximizing policy for revealing information and setting a reserve price involves reporting $\mathbf{X}_0$ precisely and in full detail.*

Some work has been done studying the incentives of bidders to gather information in a common value model. In a detailed example, Schweizer and von Ungern-Sternberg [1980] show that it may not pay a bidder to acquire information if his competitors will learn that he has done so. Lee [1982] notes that the seller can, by providing information, discourage bidders from gathering their own information (if such is costly). That policy tends to raise the expected price in a discriminatory auction.

Milgrom [1981] has studied the bidders' incentives to acquire information in a Vickrey auction and has proved the following result.

Proposition 18: *In a common-value model for a single object, let bidders 1, . . . ,n adopt the equilibrium strategies given in Proposition 7. Let there be an $(n + 1)$st bidder whose information $\mathbf{X}_{n+1}$ is a* garbling *of that of bidder 1 (i.e., the joint distribution of $\mathbf{S}$, $\mathbf{X}_0, \ldots, \mathbf{X}_n$ given both $\mathbf{X}_1$ and $\mathbf{X}_{n+1}$ does not depend on the value of $\mathbf{X}_{n+1}$). Then there is no strategy for bidder $n + 1$ that yields a positive expected payoff. Hence, an equilibrium of $(n + 1)$-bidder game ensues if bidders 1, . . . ,n use β^V and bidder $n + 1$ always bids zero.*

Following a different line of thought, Wilson [1977] and Milgrom [1979a,b] have studied how the price resulting from a first-price (discriminatory) auction aggregates the information of the many bidders. It might seem that the price could not reflect more information than was available to the winning bidder, since his bid sets the price. However, this reasoning is not correct. The winning bidder's type is a maximum order statistic from a (possibly large) sample, and such a statistic can sometimes reveal quite a lot of information.

In one version of the Wilson–Milgrom model, the bidder's types are independent estimates of the common value $\mathbf{S}_1$, drawn from a distribution with density $f(\cdot|\mathbf{S}_1)$, where the family of densities $\{f(\cdot|s)\}$ has the monotone likelihood ratio property. The upshot of their investigations is the following theorem.

Proposition 19: *Let a subscript of n denote the number of bidders in an auction. The following three statements are equivalent.*

(1) *The winning bid $\beta_n^D(\mathbf{X}_{(1)})$ is a consistent estimator[6] of $\mathbf{S}_1$.*
(2) *There exists some functions g_n such that $g_n(\mathbf{X}_{(1)})$ is a consistent estimator of $\mathbf{S}_1$.*
(3) *For any $s < s'$, $\inf_x[f(x|s)/f(x|s')] = 0$.*

The equivalence of the first two statements means that the first-price auction generates consistent estimates of $\mathbf{S}_1$ whenever any consistent estimator based on the maximum order statistic exists. The third statement makes it easy to check which distributions lead to consistency; in particular, if the types are normally distributed with mean $\mathbf{S}_1$ and fixed variance, or if they are uniformly distributed on $[0, \mathbf{S}_1]$, consistency does

[6] In other words, if $\mathbf{W}_n = \beta_n^D(\mathbf{X}_{(1)})$ denotes the winning bid in an auction with n bidders, then $\mathbf{W}_n$ converges in probability to $\mathbf{S}_1$.

follow, but if they are exponentially distributed with mean $\mathbf{S}_1$, it does not.

When the winning bid is not a consistent estimator of $\mathbf{S}_1$ in this model, there are two other possibilities. The first is that there is a degeneracy: there are two values s' and s'' such that for all $s \in (s', s'')$, $f(\cdot|s) \equiv f(\cdot|s')$. This case can be eliminated by reformulating the problem, replacing $\mathbf{S}_1$ by its expected value conditional on the whole sequence $\{\mathbf{X}_k\}$. The resulting game has the same normal form as the original game, and the degeneracy noted previously does not arise.

If the degeneracy described previously has been ruled out and the winning bid still does not consistently estimate $\mathbf{S}_1$, then there is often a nonnegligible difference between the highest and second highest bidders' "estimates" of $\mathbf{S}_1$, where the bidder's estimate is made conditional on his observed $\mathbf{X}_i$ and the hypothesis that $\mathbf{X}_i$ is the maximum observation. In that case, the bidder will attempt to earn nonnegligible positive profits. He will often succeed. Indeed, it can be shown for this case that as the number of bidders grows large, the bidders' total expected profits remain bounded away from zero and the seller's expected revenue remains bounded away from $E[\mathbf{S}_1]$. The asymptotic expected profits depend only on the asymptotic likelihood ratios: $\inf_x f(x|s)/f(x|s')$.

5 Incentives for gathering information

It is well known that in any formal decision problem, information has nonnegative value. Effectively, information enlarges a decision maker's set of strategies because it permits basing a decision on more variables. In multiperson settings, the issues become subtler, because *if a decision maker is known to have gathered additional information*, the other agents may choose to revise their strategies. Such changes may either benefit or harm the decision maker. It is always true, however, that if a decision maker can costlessly gather information without letting anyone else become aware of that fact (use "covertly gathered information"), then he or she would benefit from (or at least not be harmed by) doing so.

To study the incentives of the bidders and the seller in an auction game to gather information, we shall deal with a simple model of asymmetrically informed bidders. This model was introduced by Wilson [1967] in response to a description by Woods [1965] of a competition between two oil companies bidding for oil and gas rights on U.S. government-owned territory. One company owned the rights on an adjacent tract and had been able to explore the new tract by drilling at an angle. The other company had access only to publicly available geologic data.

To model this situation, let $\mathbf{V}$ be the common value of the rights. Assume that $\mathbf{V} > 0$ and $E[\mathbf{V}]$ is finite. Let the random variable $\mathbf{X}$ represent the private information of the better-informed bidder whom we call bidder A. No assumptions about $\mathbf{X}$ are necessary; its values may lie in any measurable space. We assume the reserve price set by seller is zero.

Bidder A's decision problem actually depends on $\mathbf{X}$ only through the resulting estimate of $\mathbf{V}$: let $\mathbf{H} = H(\mathbf{X}) = E[\mathbf{V}|\mathbf{X}]$ denote that estimate. Let $\mathbf{U}$ denote a random variable, independent of $(\mathbf{V}, \mathbf{X})$, that is uniformly distributed on (0,1). We assume that A observes $\mathbf{U}$ and uses it whenever he needs to randomize. A strategy for A is then a function $\beta: R_+ \times (0,1) \rightarrow R_+$, where $\beta(h, u)$ represents the bid made when $(\mathbf{H},\mathbf{U}) = (h,u)$. There is no loss of generality in restricting β to be nondecreasing in u.

A randomized strategy for the uninformed bidder B is simply a probability distribution function G on R_+. The auction is a discriminatory auction, and its equilibrium is given below.

Proposition 20 (Weverburgh [1979], Engelbrecht-Wiggans, Milgrom, and Weber [1981]): *In the asymmetric common-value model described previously, there is a unique Nash equilibrium point of the first-price auction game. The equilibrium strategies are*

$$\beta(h, u) = E[\mathbf{H}|\{\mathbf{H} < h \text{ or } (\mathbf{H} = h \text{ and } \mathbf{U} \leq u)\}]$$

$$G(b) = P\{\beta(\mathbf{H}, \mathbf{U}) \leq b\}.$$

To analyze how the seller's revenues and the bidders' profits depend on the better-informed bidder's information, let F denote the distribution of $\mathbf{H}$. Then, profits and revenues at the unique equilibrium are as follows.

Proposition 21 (Engelbrecht-Wiggans, Milgrom and Weber [1981]): *At equilibrium the expected profit of the better-informed bidder conditional on* $\mathbf{H} = h$ *is*

$$\int_0^h F(s)\, ds$$

and unconditionally it is

$$\int_0^\infty F(s)\,(1 - F(s))\, ds.$$

The expected profit of the worse-informed bidder is zero. The seller's expected revenue is

$$\int_0^\infty (1 - F(s))^2\, ds.$$

A host of conclusions flow from Proposition 21. First, consider what would happen to bidder B if he acquired some of A's information. If he could do so *covertly*, there would be no competitive response, and B would clearly benefit. If, however, it were common knowledge that B had gathered that information, Proposition 21 would apply, and B's profit would still be zero. Thus, as Milgrom and Weber [1982] have observed, B cannot gain by overtly collecting some of A's information.

If A could observe some additional variable **Z**, his new estimate would be $\mathbf{H}' = E[\mathbf{V}|\mathbf{X}, \mathbf{Z}]$. If B believed that A had observed only **X**, he would bid according to G. Then, if $\mathbf{H}' = h$, bidder A's maximal expected profit would be $\int_0^h F(s)\, ds$, by Proposition 21. (Notice that this profit depends only on B's strategy and on A's estimate; it does not matter whether A based his estimate on much information or on little.) If, on the other hand, B knew that A had observed both **X** and **Z**, then A's maximal expected profit would be $\int_0^h F'(s)\, ds$, where F' is the distribution of $\mathbf{H}'$. Which of these two scenarios does A prefer?

Proposition 22 (Milgrom and Weber [1982]): *For any realizations of* **X** *and* **Z**, *bidder A prefers that B know that A has observed both variables over having B believe that A has observed only* **X**; *that is, for every h,*

$$\int_0^h F'(s)\, ds > \int_0^h F(s)\, ds.$$

In summary, the better-informed bidder prefers to do his information gathering *overtly*. When B knows that A has observed **Z**, he realizes that he has become more vulnerable to the winner's curse, so he shades his bids, and A benefits from that response. On the other hand, if B can gather some of A's information, he would choose to do so *covertly*, because A's response to B's better information would be to bid more aggressively when the rights are relatively valuable, thereby depriving B of his best opportunity for earning a profit.

Notice that, in view of Proposition 21, the seller's interests and A's are strictly opposed. The sum of A's expected profit and the seller's expected revenue is $E[\mathbf{V}]$. Thus, if the seller has some information **Z**, it may pay him to report it to reduce A's informational advantage.

We consider two cases. First, suppose A already knows **Z**. For example, A might be better informed from his drillings on an adjacent tract and **Z** might be the royalty report that A filed for that tract.

Proposition 23 (Milgrom and Weber [1982]): *If A knows* (**X**, **Z**), *then a policy of announcing* **Z** *reduces bidder A's expected profit and raises the seller's expected revenue. Moreover, no policy of summarizing, garbling, or sometimes withholding* **Z** *results in greater expected revenue than the policy of reporting it precisely and in full detail.*

If A does not already know the seller's information, the matter becomes somewhat trickier. It is possible that the seller's information, though useless to B, is very useful to the better informed bidder A, much as half of a treasure map is most useful to the holder of the other half. A statistical example capturing this idea has been given by Milgrom and Weber [1981]. There is, however, an important class of models in which this problem can never arise.

Proposition 24 (Milgrom and Weber [1982]): *Let* **X** *and* **Z** *be real valued and suppose* (**V**, **X**, **Z**) *is affiliated. Then revealing* **Z** *raises the seller's expected revenue. Moreover, no policy of summarizing, garbling, or sometimes withholding* **Z** *results in greater expected revenue than the policy of reporting it precisely and in full detail.*

Proposition 24 is yet another consequence of the linkage principle.

6 Miscellaneous topics: collusion, sequencing, and bundling

The theory surveyed in the preceding sections deals with only a few of the many interesting and important questions concerning the conduct of auctions. Certainly, the assumption that bidders behave noncooperatively cannot be taken uncritically, especially in auctions like those for timber rights where the few buyers in each region may all be members of a single trade association (cf. Mead [1967]). Nor is the assumption that the bidders compete against one another once and for all a particularly appealing one. If bidders do bid against each other repeatedly, then they may attempt to infer something about their competitors' characteristics from their bidding history.

An interesting analysis that captures this learning feature has been given by Ortega-Reichert [1968]. In his model, two manufacturers compete for supply contracts in two periods. The cost of production of each manufacturer in each period is drawn from an exponential distribution with unknown mean **t**. The unknown mean has a gamma distribution.

In the first period, each bidder i offers a supply contract for a price $p_{i1} = p_{i1}(\mathbf{C}_{i1})$. The low bidder wins and earns a profit of $p_{i1} - \mathbf{C}_{i1}$. At the end of the first period, the bids are announced and the supply contract is awarded.

In the second period, the bids are again tendered, but this time each bidder has more information to use

$$p_{i2} = p_{i2}(\mathbf{C}_{i2} \mid \mathbf{C}_{i1}, p_{i1}, p_{j1})$$

where j denotes the other bidder. Ortega-Reichert found that, at equilibrium, p_{i2} is independent of $\mathbf{C}_{i1}$, but is increasing in its other arguments. In particular, by placing higher bids in the first period, a firm could induce its competitor to place higher bids in the second period. In effect, the competitor views the firm's bid as a statistic providing relevant information about the "technology" **t**.

The result of the bidders' attempts to influence their competitors' beliefs is that the prices bid at the first round are higher than they would otherwise be. In modern parlance, this is a "signaling equilibrium," because each firm tries to signal to its competitors that its future costs will be high. At equilibrium, however, each bidder recognizes the others' incentive to signal and no one is fooled.

Although it would be risky to draw any general inferences from this two-period model, the nature of the analysis does at least suggest that equilibrium in repeated contests moves away from the competitive theory in the direction of a collusive theory. This seems to hold here even in a finitely repeated game model at a symmetric Nash equilibrium. The results reported by Kreps et al. [1982] suggest that frequent interaction among the bidders may quite generally enhance incentives to collude.

All of the auction theory that we have so far considered rests on the implicit assumption that preferences have a special additive structure so that auctions for single goods can be considered in isolation. This assumption is perhaps most clearly violated in sales of the assets of bankrupt of manufacturing firms, where the value of the land, plant, and equipment together may, for some buyers, exceed the sum of their individual values. To cope with that complementarity, an institution called "entirety bidding" has emerged. In one version, the entirety bids are made before the piecemeal auction begins, and the individual objects are awarded to the winning bidders only if their sum exceeds the highest entirety bid (Cassady [1967]). In another variation, entirety bidding occurs after the piecemeal auction. At this level, auction theory merges into the general theory of resource allocation.

7 Conclusion and postscript

Auctions represent an important institution used for conducting trade. The results that have been described in this survey give a good idea of why English auctions are more popular among sellers than sealed bids and

other auction designs, and why auction houses adopt the practice of revealing their own estimates of worth for the items being offered.

Still, the received theory is far from complete: it does not consider entirety bidding and problems of complementary goods; it does not deal adequately with collusion, or repeated bidding, or the nature of competition when there are many sellers as well as many buyers; it does not treat competitive bidding in connection with procurement, where suppliers often have differentiated products so that the evaluation of bids is not so straightforward. Finally, the received theory has little to say about when auctions are more appropriate than other arrangements for conducting trade. These open questions are part of the agenda for future research.

Since this survey was written in October 1981, a great deal more has been learned about the economics of auctions. Specific predictions of the theory have been tested (sometimes confirmed, sometimes not) using both empirical data and a series of laboratory experiments. The independent private-values model has been extended to models with multiple buyers and sellers and to situations where two parameters of individual preference – such as the bidder's valuation and relative risk aversion – are both unknown. The theory of optimal auctions has been modified to include the case of risk-averse buyers. And the affiliated–monotone-value model has been generalized to accommodate the sequential sale of several identical items.

It is interesting to contrast the research achievements in auction theory of the last few years with the research agenda I proposed in 1981. Little has been learned about collusion in auctions; that topic certainly deserves more attention. Less has been learned about comparisons of alternative modes of transaction: When should a purchaser bargain with individual suppliers and when should he seek bids? Too much recent research effort in auctions has been simply applying the latest techniques (principally "mechanism design") to ever more complicated models; too little has been devoted to the very real and important economic questions that auctions raise.

REFERENCES

Brimmer, A. [1962]. "Price Determination in the United States Treasury Bill Market." *Review of Economics and Statistics*, *44*:178–183.

Brown, K. [1975]. "A Note on Optimal Fixed-Price Bidding with Uncertain Production Cost." *Bell Journal of Economics*, *6*:695–697.

Capen, E. C., R. V. Clapp, and W. M. Campbell [1971]. "Competitive Bidding in High-Risk Situations." *Journal of Petroleum Technology*, *23*:641–653.

Carson, D. [1959]. "Treasury Open Market Operations." *Review of Economics and Statistics*, *41*.

Cassady, R., Jr. [1967]. *Auctions and Auctioneering*. Berkeley: University of California Press.
Cox, J., B. Roberson, and V. Smith [In press]. "Theory and Behavior of Single Object Auctions," in *Research in Experimental Economics*, vol 2, edited by V. Smith.
Engelbrecht-Wiggans, R., P. R. Milgrom, and R. J. Weber [Forthcoming]. "Competitive Bidding and Proprietary Information." *Journal of Mathematical Economics*.
Friedman, M. [1960]. *A Program for Monetary Stability*. New York: Fordham University Press.
Goldstein, H. [1962]. "The Friedman Proposal for Auctioning Treasury Bills." *Journal of Political Economy*, *70*:386–392.
Griesmer, J., R. Levitan, and M. Shubik [1967]. "Toward a Study of Bidding Processes, Part IV: Games with Unknown Costs." *Naval Research Logistics Quarterly*, *14*:415–433.
Harris, M., and A. Raviv [1981]. "A Theory of Monopoly Pricing Schemes with Demand Uncertainty." *American Economic Review*, 71:347–365.
Harris, M., and A. Raviv [1982]. "Allocation Mechanisms and the Design of Auctions," *Econometrica*, *49*:1477–1499.
Holt, C. A., Jr. [1980]. "Competitive Bidding for Contracts under Alternative Auction Procedures." *Journal of Political Economy*, *88*:433–445.
Kreps, D., P. Milgrom, J. Roberts, and R. Wilson [1982]. "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma." *Journal of Economic Theory*, *27*:245–252.
Lee, T. [1982]. "Resource Information Policy and Federal Resource Leasing." *Bell Journal of Economics*, *13*:561–568.
Maskin, E., and J. Riley [1980]. "Auctioning an Indivisible Object." JFK School of Government, Discussion Paper no. 87D, Harvard University, Cambridge, Mass.
Maskin, E., and J. Riley [1981]. "Multi-unit auctions, Price Discrimination and Bundling." Economics Discussion Paper no. 201, University of California at Los Angeles.
Matthews, S. [1979]. "Risk Aversion and the Efficiency of First- and Second-Price Auctions." CCBA Working Paper no. 586, University of Illinois.
Mead, W. J. [1967]. "Natural Resource Disposal Policy – Oral Auction Versus Sealed Bids." *Natural Resource Planning Journal*, *M*:194–224.
Milgrom, P. R. [1979a]. *The Structure of Information in Competitive Bidding*. New York: Garland.
Milgrom, P. R. [1979b]. "A Convergence Theorem for Competitive Bidding with Differential Information." *Econometrica*, *47*:679–688.
Milgrom, P. R. [1981]. "Rational Expectations, Information Acquisition, and Competitive Bidding." *Econometrica*, *49*:921–943.
Milgrom, P. R. [1981a]. "Good News and Bad News: Representation Theorems and Applications." *Bell Journal of Economics*, *12*:380–391.
Milgrom, P. R., and R. J. Weber [1982]. "The Value of Information in a Sealed-Bid Auction." *Journal of Mathematical Economics*, *10*:105–114.
Milgrom, P. R., and R. J. Weber [1982a]. "A Theory of Auctions and Competitive Bidding." *Econometrica*, *50*:1089–1122.
Milgrom, P. R., and R. J. Weber [1982b]. "Sequential Auctions." Unpublished notes, Northwestern University, Evanston, Ill.

Myerson, R. [1981]. "Optimal Auction Design." *Mathematics of Operations Research*, *6*:58–73.
Ortega-Reichert, A. [1968]. "Models for Competitive Bidding under Uncertainty." Department of Operations Research Technical Report no. 8 (Ph.D. dissertation), Stanford University, Stanford, Calif.
Reece, D. K. [1978]. "Competitive Bidding for Offshore Petroleum Leases." *Bell Journal of Economics*, *9*:369–384.
Rieber, M. [1964]. "Collusion in the Auction Market for Treasury Bills." *Journal of Political Economy*, *72*:509–512.
Riley, J., and R. Zeckhauser [1983]. "Optimal Selling Strategies: When to Haggle, When to Hold Firm." *Quarterly Journal of Economics*, *98*:267–290.
Riley, J., and W. Samuelson [1981]. "Optimal Auctions." *American Economic Review*, *71*:381–392.
Rothkopf, M. [1969]. "A Model of Rational Competitive Bidding." *Management Science*, *15*:362–373.
Schweizer, U., and T. von Ungern-Sternberg [1980]. "Sealed Bid Auctions and the Search for Better Information." Department of Economics, University of Bonn.
Smith, V. [1966]. "Bidding Theory and the Treasury Bill Auction: Does Price Discrimination Increase Bill Prices?" *Review of Economics and Statistics*, *48*:141–146.
Tsao, C., and A. Vignola [1977]. "Price Discrimination and the Demand for Treasury's Long Term Securities." Preliminary Report, U.S. Department of Treasury.
Weverburgh, M. [1979]. "Competitive Bidding with Asymmetric Information Reanalysed." *Management Science*, *25*:291–294.
Wilson, R. [1967]. "Competitive Bidding with Proprietary Information." *Management Science*, *13*:A816–A820.
Wilson, R. [1977]. "A Bidding Model of Perfect Competition." *Review of Economic Studies*, *4*:511–518.
Woods, D. [1965]. "Decision Making under Uncertainty in Hierarchical Organizations." D.B.A. dissertation, Harvard University, Cambridge, Mass.
Vickrey, W. [1961]. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance*, *16*:8–37.
Vickrey, W. [1962]. "Auctions and Bidding Games." In *Recent Advances in Game Theory* (conference proceedings), pp. 15–27. Princeton University, Princeton, N.J.