Discussion Paper #496

MULTIPLE-OBJECT AUCTIONS

by

Robert J. Weber

August 1981

J.L. Kellogg Graduate School of Management

Northwestern University

Evanston, Illinois 60201

This research was supported in part by the U.S. Office of Naval Research, the Xerox Corporation, and the Center for Advanced Studies in Managerial Economics at Northwestern University.

MULTIPLE-OBJECT AUCTIONS

by

Robert J. Weber

1. Introduction

The received theory of auctions has a principal focus on the sale of a single object. However, a great number of observed auction sales involve more than one object, sold either simultaneously or sequentially. In this paper, we take a brief look at some of the issues which arise in the study of multiple-object auctions; a number of these issues simply do not arise when the sale of only one object is being considered.

We classify the auctions we will discuss into three categories. A simultaneous dependent auction is one in which the bidders are called upon to each take a single action, subsequent to which the objects will be distributed among the bidders and payments will be made. An example of such a sale is the weekly auctioning of U.S. Treasury bills. A special case of such a sale, which we choose to view separately, is a simultaneous independent auction, in which the bidders must simultaneously act in several different auctions of individual items, and in which the outcome of each sale is independent of the outcomes of the others. The sale of mineral rights on federal land by the U.S. Department of the Interior frequently takes this form. It should be noted that such a sale may force a bidder to expose himself to the risk of obtaining more items than he desires. Finally, a sequential auction is just what the name suggests: the sale of one item at a time, perhaps with the public release of information concerning the outcome of one round prior to the beginning of the next. Examples of such sales are estate auctions at which a collection of objects - stamps, coins, antiques, or the like - are sold.

2. Bidders with Limited Consumption Capacity

Perhaps the most elementary situation occurs when a number of identical objects are to be sold and each bidder desires only one of them. But even in this setting, much of the richness of the multiple-object environment appears.

2.1 The Independent Private Values Model

We initially assume that each bidder knows the (identical) values of the objects to himself (this is the private values assumption), and that the values of the different bidders are independent observations of a nonnegative random variable X with a commonly-known continuous distribution. Let there be a total of n risk-neutral bidders, and let $X_1,\ldots,X_n$ be their valuations (which we will also refer to as their types). Finally, assume that a total of k objects are to be sold. For the sake of notational ease, but with no real loss of generality, we restrict our consideration to the case $k < n$.

Let $X_{(1)} \geq X_{(2)} \geq \ldots \geq X_{(n)}$ be the order statistics of the n types. For any particular bidder, the order statistics of the opposing types are central to his decision problem. Therefore, taking the perspective of bidder 1, we define $Y_1,\ldots,Y_{n-1}$ to be the order statistics of $X_2,\ldots,X_n$.

First-price and second-price auctions are two well-known procedures for the sale of a single object. In each, the bidders submit nonnegative sealed bids; the highest bidder receives the object, and pays either the amount of his own bid (first-price) or the highest of the opposing bids (second-price). Both can be generalized to simultaneous dependent sealed-bid auctions in which each of the k highest bidders receives an item. In the discriminatory auction, each of the k highest bidders pays the amount he bid. In the uniform-price auction, each pays the amount of the highest rejected ((k+1)-st highest) bid.

A strategy for a bidder in either of the auctions under consideration is a function which associates a bid with each of his possible types. When discussing symmetric models, we are primarily concerned with symmetric equilibria. The following result is due to Vickrey [1962] and Ortega-Reichert [1968].

Theorem 1: In the independent private values model:

(a) Let $b^D(x) = E[Y_k | Y_k < x]$. Then $(b^D,\ldots,b^D)$ is the unique symmetric equilibrium of the discriminatory auction.

(b) Let $b^U(x) = x = E[Y_k | Y_k = x]$. Then $(b^U,\ldots,b^U)$ is the unique symmetric equilibrium of the uniform-price auction.

(c) The total expected revenue of the seller is the same in both the discriminatory and uniform-price auctions, and equals $k \cdot E[X_{(k+1)}]$.

Recent studies have set the "revenue equivalence" result of Theorem 1(c) in a broad context. Consider any auction mechanism which delivers at most one object to each bidder. Given fixed strategies for $n-1$ of the bidders, the selection of any bid by the remaining bidder determines a particular probability p of his obtaining an item, and a particular expected payment e. Since the types are independent, the (p,e)-pair associated with any bid does not depend on his own type. The bidder's decision problem, when his type is x, is to choose a bid b which maximizes $x \cdot p(b) - e(b)$.

The equilibrium assumption - that each bidder follows a strategy which is optimal given the strategies of the others - yields for each bidder i a differential condition relating the p_i and e_i functions he faces. Consequently, at equilibrium the n functions e_i, and hence the seller's expected revenue, are fully determined by the n functions p_i and n boundary conditions. A convenient set of boundary conditions is provided by the expected payment made by each of the bidders when he is of his lowest possible type. The preceding discussion motivates the following theorem, which can be proved rigorously using arguments presented in Myerson [1981] or in Milgrom and Weber [1981].

Theorem 2: Consider any auction mechanism in the k-object, n-bidder, independent private values setting. Assume that an equilibrium point is given, such that the bidders with the k highest types are certain to receive items, and such that the bidder with the lowest type has an expected payment of zero. Then the seller's expected revenue at this equilibrium is $k \cdot E[X_{(k+1)}]$.

A sequential auction may be viewed as a simultaneous dependent auction, albeit one in which a bidder's "action" may be very complex (since it must specify his intended action at every stage, conditioned on the information revealed in previous stages). Consequently, Theorem 2 applies to sequential sales. Here we consider two such sales, the sequential first-price auction and the sequential second-price auction, each based on the corresponding single-object procedure. In each round, the bidders who remain (those who have not yet been awarded an item) submit sealed bids. The highest bidder is awarded an item and pays according to the corresponding single-object pricing rule. A public announcement is made concerning the outcome of that round, and the procedure is then repeated (until all k objects have been sold).

We shall consider two types of announcements. The first is simply: "An object has been sold." This provides each bidder only with information which must be available to him at the start of the next round. The second type of announcement is: "An object has been sold at the price p," (where p is indeed the price at which the sale occurred). In the sequential first-price procedure, this announcement permits each remaining bidder to draw certain inferences about the distribution of types of the other remaining bidders. In the sequential second-price auction, this announcement actually discloses the bid made by one of the remaining bidders. A remarkable fact is that symmetric equilibrium strategies for either auction are the same under either announcement policy.

Theorem 3: In the independent private values model:

(a) Let $b^F_\ell(x) = E[Y_k | Y_\ell < x < Y_{\ell-1}] = E[X_{(k+1)} | X_{(\ell)} = x]$ for $\ell = 1,\ldots,k$, and let $b^F = (b^F_1,\ldots,b^F_k)$. Then $(b^F,\ldots,b^F)$ is the unique symmetric equilibrium of the sequential first-price auction.

(b) Let $b^S_\ell(x) = E[Y_k | Y_\ell = x] = E[X_{(k+1)} | X_{(\ell+1)} = x]$ for $\ell = 1,\ldots,k$, and let $b^S = (b^S_1,\ldots,b^S_k)$. Then $(b^S,\ldots,b^S)$ is the unique symmetric equilibrium of the sequential second-price auction.

In both cases, $b_\ell(x)$ is the bid made by a bidder of type x in the ℓ-th round, if he has not yet received an item. In Theorem 3(a), Y_0 is interpreted to be infinite.

Both parts of the theorem can proved recursively, by working back from the final stage of the auction. A crucial consequence of the independence assumption is that

$$E[X_{(k+1)} | X_{(\ell)}, X_{(\ell-1)}] = E[X_{(k+1)} | X_{(\ell)}] \quad \text{for} \quad \ell = 1,\ldots,k \quad .$$

It is this which makes the equilibrium strategies independent of the type of announcement.

From Theorem 3, it is not difficult to see that the (unconditional) expected selling price at each stage is $E[X_{(k+1)}]$. But how does the sequence of actual prices behave? In the sequential first-price auction, the expected price in the ℓ-th stage, given that the $(\ell-1)$-st price is p, is:

$$E[b_\ell(X_{(\ell)}) | b_{\ell-1}(X_{(\ell-1)}) = p]$$

$$= E[E[X_{(k+1)} | X_{(\ell)}] | X_{(\ell-1)} = b_{\ell-1}^{-1}(p)]$$

$$= E[E[X_{(k+1)} | X_{(\ell)}, X_{(\ell-1)}] | X_{(\ell-1)} = b_{\ell-1}^{-1}(p)]$$

$$= E[X_{(k+1)} | X_{(\ell-1)} = b_{\ell-1}^{-1}(p)]$$

$$= b_{\ell-1}(b_{\ell-1}^{-1}(p)) = p \quad .$$

Hence, the sequence of prices is a martingale, i.e., on average, prices drift neither up nor down over time. A similar result can be established in the same manner for the sequential second-price auction.

The independent private values model can be generalized through the introduction of uncertainty concerning the number of bidders. Assume that the number N of bidders is a random variable that is independent of the bidders' types and is almost surely finite. Let the types of the bidders present at the auction be $X_1,\ldots,X_N$. Extend this finite list of types by adjoining an infinite number of zeroes, and define the sequences $X_{(1)}, X_{(2)}, \ldots$ and $Y_1, Y_2, \ldots$ of order statistics as before. Then Theorems 1, 2, and 3 remain valid as stated, and the martingale property again holds for both the first-price and second-price sequential auctions.

Another generalization of the independent private values model arises when the private values assumption is relaxed. Let the value of an object to bidder i be V_i. Assume that there is a real-valued function u of n variables, such that for every i,

$E[V_i|X_1,\ldots,X_n] = u(X_i,\{X_j\}_{j\neq i})$. (That is, assume that each bidder has the same "valuation function," and that the other bidders' types enter his function in a symmetric manner.) Further assume that u is strictly increasing in its first argument, and monotone increasing in all of its arguments. Define $v_k(x,y) = E[V_1|X_1 = x, Y_k = y]$. Then Theorems 1, 2, and 3 are valid for this "independent types, symmetric values" model, when $v_k(Y_k,Y_k)$ (or $v_k(X_{(k+1)},X_{(k+1)})$) is substituted for Y_k (or $X_{(k+1)}$) before the conditioning signs in the various expectations.

2.2 The General Symmetric Model

Clearly, the assumptions that drove the results of the previous section were symmetry and type independence. In this section, we relax the independence assumption. The following two paragraphs present what we call the general symmetric model.

Let $S_1,\ldots,S_m,X_1,\ldots,X_n$ be real-valued random variables. The first m variables represent qualities of the (identical) objects not directly observable by the bidders; the remaining variables are the bidders' private estimates. Let V_i be the value of any object to bidder i. We assume that there is a nonnegative function u on R^{m+n}, such that for each i,

$$E[V_i|S_1,\ldots,S_m,X_1,\ldots,X_n] = u(X_i,\{X_j\}_{j\neq i};S_1,\ldots,S_m) \quad .$$

Consequently, the m state variables enter all of the bidders valuations in the same manner. We further assume that u is increasing in its first argument, and nondecreasing in all arguments.

Let $f(s,x)$ denote the joint density of the $m + n$ random elements of the model. Our final symmetry assumption is that f is symmetric in its last n arguments. We also assume that the elements of the model have a positive statistical linkage. Let z and z' be points in R^{m+n}, let $z \vee z'$ denote the coordinate-wise maximum, and $z \wedge z'$ the coordinate-wise minimum. Then f is affiliated if for all z and z', $f(z \vee z')f(z \wedge z') \geq f(z)f(z')$. Roughly, this condition states that higher values of some variables make higher values of the others more likely, a not unreasonable assumption in the kinds of situations we wish to study. A formal development of affiliation is presented in Milgrom and Weber [1981].

In this general setting, the revenue-equivalence result of the previous section fails to hold. (The argument given before Theorem 2 breaks down due to the dependence of the functions p_i and e_i on bidder i's estimate.) Indeed, we have the following result.

Theorem 4: In the general symmetric model, the uniform-price auction generally yields greater expected revenues than the discriminatory auction.

The proof of this theorem is based on the idea that, when the price paid by a bidder reflects the estimates of other bidders, that price is more closely linked to his own estimate, even when his choice of a bid is held fixed. This extra linkage leads to a steeper expected payment (a function of his estimate) for each bidder. Since the expected payment functions associated with different auction procedures are the same (at equilibrium) when a bidder has the lowest possible estimate, the extra linkage generates higher expected revenues.

There is a colorful history of debate surrounding the format used for the weekly U.S. Treasury bill auctions. Tradition has been to use a discriminatory auction, although a number of authors have argued for a change to uniform pricing. Many of the arguments presented deal with questions of bidder collusion, or with the degree to which one format or the other will attract a greater number of small-volume risk-averse bidders. (Indeed, current practice is to allow small-volume bidders to enter "noncompetitive" bids, and to award them bills at the average of the accepted "competitive" bids.) It is natural to assume that the bidders' estimates of economic trends are statistically linked. Therefore, to the extent that the assumption that bidders desire equal quantities of bills is valid, Theorem 4 provides a new dimension to the debate. It should be noted, however, that the introduction of risk-aversion complicates matters.

Theorem 5: In the independent private values model, if the bidders are equally averse to risk, then the discriminatory auction generally yields greater expected revenues than the uniform-price auction.

Consequently, if arguments in favor of one procedure or the other are to be based on their revenue-generating properties, one must decide whether the statistical linkage of the bidders' estimates or the aversion of the bidders to risk is the over-riding factor.

2.3 Simultaneous Independent Sales

From time to time, the U.S. Department of the Interior leases the mineral rights on various federal properties. In this section, we focus on drilling rights on offshore territory. These lease sales often involve more than a hundred tracts, all in the same area. Bidders are required to submit separate, non-retractable sealed bids for all tracts on which they wish to compete, and to submit a substantial down-payment with each bid. All bids on all tracts are unsealed on the same day, after the deadline for the submission of bids is reached. The high bidder on each tract receives the rights on that tract, and is required to complete his payment to the level of his bid. (Actually, the government reserves the right to withdraw a tract from sale if it finds the highest bid unsatisfactory. Also, in addition to his bid a winning bidder is required to pay a royalty on the petroleum he extracts. These complicating factors need not concern us here.)

Historically, the variance of the bids submitted on a tract is quite high. Studies of sales conducted in the late 1960's found the winning bid to typically be twice that of the second-highest bid (Capen, Clapp, and Campbell [1971]). Substantial sums of money are involved here: Examples abound in which tracts selling for around $100 million drew second-high bids of less than $30 million. Consequently, the spread between high and second-high bids (known in picturesque oil-industry parlance as "money left on the table") is of some concern to the competing bidders.

Some authors have cited the substantial uncertainty concerning the extractable resources present on a tract, as a factor which makes large bid spreads unavoidable. However, another factor which can lead to sizable spreads has been presented by Engelbrecht-Wiggans and Weber [1979].

When a firm prepares to bid on tracts in a particular area, it has two conflicting concerns. It does not wish to be "shut out" of the area: a certain number of leases must be obtained in order to protect its competitive position and to provide a use for its costly exploratory equipment. Yet it also doesn't wish to win too many tracts: there are many dangers inherent in an over-commitment of capital in a single area. Competing firms face similar concerns.

If it were possible, a firm might wish to submit contingent bids: "If we win less than k of the first ℓ tracts for which bids are opened, then our remaining bids stay in submission; otherwise, those bids are withdrawn." However, this is not allowed.

Assume that the supply of tracts is roughly equal to the demand for them. This is a plausible assumption, since the government can regulate the supply, and has as one of its goals to have tracts explored and developed expeditiously. Then one possible strategy for a firm is to bid quite aggressively on some tracts (in order to minimize the chance of being shut out), and to bid much less aggressively on others, expecting to win them only if the other firms also bid unaggressively, and then at a bargain price.

Indeed, in a relatively simple model in which all tracts are identical and there is no uncertainty concerning the values of the tracts being sold, Engelbrecht-Wiggans and Weber have shown that there is an equilibrium in such agressive/nonaggressive strategies, and that the distribution of the spread between winning and second-high bids at equilibrium is similar to that observed in practice. While it is unarguable that a certain amount of the variance in bids submitted in oil-lease auctions is the result of bidder uncertainties about the state of nature, the preceding discussion suggests that the strategic variance in bids forced by the use of a simultaneous independent auction procedure might be an important factor in explaining the large sums of money left on the table in these auctions.

3. Sequential Auctions

In Section 2.1, our results concerning sequential auctions in the independent private values model relied on an important property of the equilibrium strategies. To wit, a bidder never had to concern himself with how the other bidders' perceptions of him were affected by his actions, since their equilibrium strategies were independent of those perceptions.

When the bidders have dependent value estimates, the situation becomes much more complicated. If a bidder bids more conservatively in one round than his estimate would appear to warrant, he might lead others to believe that he has private information concerning unfavorable aspects of the objects being sold. The others would then bid more conservatively in future rounds, enabling him to "steal" an item at a bargain price. At equilibrium, of course, the temptation of deception must somehow be accounted for.

3.1 A Two-Stage Signalling Model

Important insights into the nature of equilibrium behavior in sequential actions were first provided by Ortega-Reichert [1968]. He considered a two-object, two-bidder, first-price sequential auction in a private values model. Assume that Nature chooses a state variable, unobserved by the bidders but distributed according to a commonly-known distribution. Conditional on this state variable, the private values of the first item to the two bidders are independently determined. Each learns his own value, from which he can draw inferences about the state of nature. These inferences affect his beliefs about the other's value, as well as his beliefs about the value of the second item to him. However, this second value, drawn from the same distribution as his first, will not be revealed to him until the sale of the first item is concluded and both bids have been revealed.

Incentives to deceive arise from the common uncertainty about the state variable. If bidder 1 could convince bidder 2 that his first-stage value was low, then bidder 2's second-stage beliefs about the state variable (which are conditioned on his two values and the information he thinks he possesses concerning bidder 1's first-stage value) would be falsely conservative. Bidder 2 would therefore bid less aggressively in the second-stage than he otherwise might, thinking that bidder 1's second-stage value was more likely to be low than it actually is. This would afford bidder 2 a chance to turn a large profit in the second stage.

Ortega-Reichert found that there is a symmetric pure-strategy equilibrium for the game in question. In the first stage, each bidder bids less than he would were the auction composed of only that stage. This is perhaps surprising. A pure strategy is invertible, in the sense that a bidder's value can be deduced with certainty from the bid he makes. But if his first-stage bid is fully revealing, what gain is there in not making a bid which maximizes his expected return in that stage? The answer lies in the inferences drawn by the other bidder. If bidder 1 makes a higher bid in the first stage, it is true that his first-stage expected profit will be greater than it is from his equilibrium bid. But bidder 2 will then falsely conclude that bidder 1's value was higher than it actually was, will have falsely high beliefs about the state variable when he prepares to bid in the second stage, and hence will bid aggressively, expecting bidder 1 to have another high value. At equilibrium, the cost to bidder 1 from stimulating this aggression in the second stage outweighs the gain to be had from the use of a non-equilibrium bid in the first stage.

Milgrom and Roberts [1980] have used a related two-stage model, which has an equilibrium of a similar nature, to give a cogent analysis of the phenomenon of limit-pricing.

3.2 An Asymmetric Model

Assume that k identical objects are to be sold to two bidders via a sequential first-price auction. There is initial uncertainty concerning the quality of the objects; they are either all of high value (1), or all of low value (0). The probability p that they are all of value 1 is known to both bidders. The values of the objects are additive, i.e., a collection of ℓ of them is worth either ℓ or 0.

Just prior to the sale, bidder A learns the true value of the objects. This fact is publicly known; in particular, bidder B is aware of it. How will this affect the auction?

The single-object version of this model has been analyzed by Engelbrecht-Wiggans, Milgrom, and Weber [1981]. They showed that at (the unique) equilibrium the expected profit of A is positive, and the expected profit of B is zero. In fact, this result holds no matter what the distribution of the object's value is, as long as there is <u>some</u> initial uncertainty. If both know the value of the object for certain, then at equilibrium both have expected profits of zero (the object sells at precisely its value).

When k is greater than one, the game has multiple equilibria (Engelbrecht-Wiggans and Weber [1981]). However, only one has the property that A always bids zero when that is the value of the objects. (At the other equilibria, he makes positive bids which are always beaten by B's bids. Formally, these equilibria are "imperfect.") At this equilibrium, B's expected profit is positive. Indeed, when k is sufficiently large, the expected profit of uninformed bidder B exceeds that of informed bidder A !

The explanation of this result is surprisingly straight-forward. Although A may know that the objects are of high value, he can only claim a profit from that knowledge in one stage. At equilibrium, B never bids zero (although as long as A has bid zero in all previous stages, B has a substantial probability of making only a nominal positive bid). Therefore, A cannot profit from his favorable knowledge until he makes a positive bid. But if ever he enters a positive bid, he reveals his knowledge. In all subsequent stages, both bidders will bid 1, and neither will profit. Consequently, A's expected profit is bounded above by p (that is, by 1 when the objects are valuable, and 0 otherwise). A's equilibrium choice of a time to act, when the objects are of value 1, is distributed rather evenly over the k stages of the auction. Therefore, B is quite likely to be able to claim a number of objects at low prices in the early stages, and the expected values of these objects at the times he claims them are not much less than p.

The phenomenon appearing here can be viewed from the perspective of bargaining under uncertainty, as developed by Myerson [1980]. Bidder A is of either of two "types": one type knows that the objects are valuable, and the other knows they are not. These types are involved in a sort of bargaining game (internal to A). The first type wants the second to occasionally make positive bids, so as to cloud the informational content of his own positive bids. But the first type can provide no incentives to the second for such actions (since the second wants nothing from the first), and hence is in a disadvantageous position. This internal conflict between types works to the detriment of A, and in the process, B reaps substantial benefits.

If the objects being sold can take more than two distinct values, it appears that equilibrium behavior involves several of A's "types" randomizing on overlapping intervals of bids in the early stages. Such behavior is different from that observed in any of the classical auction models.

4. Nonidentical Values

There are two basic manners in which the items for sale can have different values to a bidder. Either the items are identical, and the bidder's marginal value from an item varies with the number he possesses, or the items are truly different.

If the items are identical, the seller may wish to elicit a price-quantity function from each of the bidders. Several different schemes of this nature have been studied for the case in which the bidders' marginal values are decreasing. Perhaps the most simple generalizes the uniform-price procedure: Each bidder submits k bids, the highest k bids secure items, and a bidder who receives ℓ items is charged the sum of the ℓ highest rejected bids. If the bidders' valuations are independent, then a dominant strategy for each is to submit bids equal to his first k marginal values.

A similar, but more complex scheme can be used when the objects are different and there is no statistical linkage across bidders. Each bidder submits a bid for every subset of the objects. The set of objects is distributed among the bidders according to the partition of the set which draws a maximum total bid amount (summing over the high bids on the elements of the partition). Each bidder is charged the difference between the maximum total which would have occurred had his bids not been submitted, and the sum of the high bids placed on the subsets other than the one he receives in the actual maximizing partition. This procedure is a generalization of the one described in the preceding paragraph (which in turn generalizes the uniform-price auction), and again, a dominant strategy for each bidder is to bid his actual valuations. Both of these schemes are based on suggestions made by Vickrey [1961].

Other types of procedures have drawn little or no theoretical work. We offer as an example the bid-for-the-right-to-choose sequential auction, at each stage of which the bidders compete for the right to select one from among the as-yet-unclaimed objects. Even in the independent private values setting, the strategic issues associated with this procedure seem quite involved.

Another example is the "gymnasium" auction, in which bid sheets for the objects are posted in a central location. Bidders can observe the current high bid for each object, and enter a higher bid whenever they choose. Questions of timing seem important in this setting.

5. Prospects

Much remains to be learned about the auctioning of several objects. Should similar objects be sold individually, or in indivisible batches? How can contingent bids be handled? (Cassady [1967] presents a fascinating example of "entirety bidding," in which contingent claims play an important role.) What are the fundamental properties of auction-like two-sided markets?

To these and and many other questions, we can currently give few answers. But the recent explosion of interest in auctions holds the promise that our understanding will deepen substantially in the near future.

References

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

Cassady, R., Jr. [1967], Auctions and Auctioneering, University of California Press.

Engelbrecht-Wiggans, R., P. R. Milgrom, and R. J. Weber [1981], "Competitive Bidding and Proprietary Information," CMSEMS Discussion Paper No. 465, Northwestern University.

Engelbrecht-Wiggans, R., and R. J. Weber [1979], "An Example of a Multi-Object Auction Game," Management Science, 26, 119-142.

Engelbrecht-Wiggans, R., and R. J. Weber [1981], working notes.

Milgrom, P. R., and D. J. Roberts [1980], "Equilibrium Limit Pricing Doesn't Limit Entry," CMSEMS Discussion Paper No. 399, Northwestern University.

Milgrom, P. R., and R. J. Weber [1981], "A Theory of Auctions and Competitive Bidding," CMSEMS Discussion Paper No. 447R, Northwestern University.

Myerson, R. [1980], "Solutions for Two-Person Bargaining Problems with Incomplete Information," CMSEMS Discussion Paper No. 432, Northwestern University.

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, Stanford University.

Vickrey, W. [1961], "Counterspeculation, Auctions, and Competitive Sealed Tenders," Journal of Finance, 16, 8-37.

Vickrey, W. [1962], "Auctions and Bidding Games," Recent Advances in Game Theory (conference proceedings), Princeton University, 15-27.

Acknowledgement: I am indebted to Paul Milgrom for the insights into auction theory which I have derived from our collaboration. Many of the uncited results in Section 2 are drawn from joint working notes.