A THEORY OF AUCTIONS AND COMPETITIVE BIDDING, II

by

Paul R. Milgrom and Robert J. Weber

Forward (April, 1999): This paper was written in early 1982. On the basis of results from previous research, we had informally conjectured that sequential auctions of the types treated here would yield upward-drifting price sequences at equilibrium. The 1981 RCA transponder lease auction described in the introduction, and then-contemporary comments of Ashenfelter concerning the "afternoon effect" in wine auctions, challenged this conjecture. Clearly, a proof was needed, and we generated one for the case of sequential first-price auctions with winning-bid announcements after each round. We wrote out the equilibrium-characterization and price-sequence proofs, and then sketched out what we thought might be analogous results for a variety of other auction formats.

Proofs for the sequential first-price and second-price auctions without price announcements refused to come together, and a bit of further thought indicated a possible reason, which we here add as a bracketed comment just before the "Summary of Results" section of the paper. Without the proofs we sought, the paper languished in informal distribution, until the editor of this volume requested permission to print the paper, and graciously agreed to the inclusion of this forward.

Introduction

Much of the existing literature on competitive bidding is focused on the sale of a single object at auction. Yet rarely does one actually observe sales involving only one object. Either a number of similar objects — stamps, books, cattle, paintings, antiques, mineral-rights leases, or the like — are put on the block in the course of a sale, or multiple units of essentially identical objects, such as U.S. Treasury bills, are offered for sale. In this paper, we offer some first steps in the extension of the theory of single-object auctions presented in [1] to a theory of multiple-object auctions.

An example of the type of sale we will consider here is provided by the November 1981 sale at Sotheby's (New York) of leases on RCA-owned satellite-based telecommunications transponders. At the time, RCA had announced plans to launch a satellite carrying 24 transponders, each capable of relaying a cable television signal from a single ground-based station to the nation's network of local cable television distributors. There was substantial interest among the major cable television companies in obtaining the right to use one of these transponders, since there was a general belief that many cable franchises would eventually choose to distribute to their customers the signals relayed by this particular satellite. Noting this interest, RCA decided to abandon previous plans to set a price and ration transponder leases among qualified customers at that price, and instead announced that seven of the leases would be sold at public auction. Sotheby's was given the responsibility of conducting the sale. They chose to sell the leases sequentially by English auction, so that the price resulting from each round was publicly known prior to the start of the next. The resulting auction set several records for Sotheby's. Each of the leases was purchased for substantially more than had ever before been paid for any single item sold at any branch of Sotheby's. In addition, the total revenue generated by the sale, slightly over $90 million, was nearly as great as the total revenue ever before generated by a sale Sotheby's of related items.

Two features of the sale just described are worthy of note. First, the objects being sold were identical, in the sense that a bidder would assess the same value for each of them. Second, it is likely (and the actual sales results bear this out) that no bidder wished to acquire more than one of the leases. This paper will be restricted to the consideration of multiple-item auctions that possess both of these features. Two questions raised by the RCA sale warrant discussion: How would we expect the bids made by a particular bidder in successive rounds to be related? What properties would we expect the sequence of prices generated by the auction to have?

In 1982, the U.S. Treasury will sell at auction approximately $1 trillion worth of bills and notes in order to finance the various operations of the federal government (including the rollover of previous debt). It is obviously not feasible for them to conduct their sales in such a manner as to sequentially sell individual bills, and indeed, they do use auction procedures that simultaneously allocate among the bidders all of the bills offered in a particular sale. Again, we note that the objects being sold at a particular sale are identical, in the sense that bidders value all of the bills of a particular issue identically. Of course, some bidders can be expected to seek to purchase a different quantity of bills than other bidders. We will consider several simultaneous-sale procedures in this paper, under the restriction that all winning bidders receive a single item; in the case of the sale of Treasury bills, we could view the items to be blocks of bills of equal total face value. Despite the lack of conformance of this model to the actual situation, we believe that the results we will offer shed some light on the long-standing debate over which of the available procedures should be used. One question we will discuss is how these procedures can be ranked in terms of their revenue-generating properties.

We mention here briefly another type of sale that our model fits to some degree. Almost every business day, some firm "goes public" for the first time, i.e., shares of stock in the firm are offered to the general public for purchase. In practice, these sales are conducted by underwriters who elicit price/quantity bids from stock exchange firms representing their clients. A highly-simplified way of modeling the trading which takes place during the first few days after the date of issue is as a sequential auction, in which blocks of equal size are sold, one after another, with the price of each sale revealed publicly (on the ticker tape) prior to the next sale. Again, our restriction of bidders to the purchase of single blocks is not very realistic; still, we hope that our analysis of sequential auctions will provide some insight into the manner in which the sequence of prices can be expected to vary over time.

Types of Auction Procedures

The procedures we will consider can first be characterized according to whether the items being sold are allocated among the bidders simultaneously or sequentially. We treat three simultaneous allocation procedures: the discriminatory, uniform-price, and English auctions.

In a *discriminatory auction*, the bidders submit sealed bids. The highest bidders each receive one of the objects, and each pays the amount of his own bid. This procedure generalizes the single-object "first-price" auction, and is the procedure most commonly used for the sale of U.S. Treasury bills.

In a *uniform-price auction*, the bidders similarly submit sealed bids. However, the winning bidders are all charged the amount of the highest rejected bid. This procedure generalizes the single-object "second-price" (or "Vickrey") auction.

An *English auction*, as we model it here, begins with a "price clock" set at a very low level, and an "active-bidders" clock which indicates the number of bidders present at the start of the auction. Each bidder controls a button, which is initially depressed. The price clock is continuously increased, and bidders who wish to remain active merely keep their buttons depressed. A bidder leaves the auction by releasing his button; when he does this, the button locks (so he cannot re-enter the auction at a later point), and the active-bidders clock decreases by one (so at each instant, it registers the number of buttons still depressed). When the active-bidders clock first registers a number equal to the number of items being sold, the price clock stops, and the remaining bidders each receive an item, at the price then displayed. We are not aware of any institutions that use this procedure in this precise form, but we view this as the natural generalization of the single-object English auction treated in [1].

In specifying a sequential auction procedure, one important feature is the amount of information concerning the result of one round which is made available to the remaining bidders prior to the next. For both first-price and second-price sequential actions, we will consider two extreme informational assumptions.

In a *sequential first-price auction without price announcements*, the bidding proceeds in rounds. In each round, the remaining bidders submit sealed bids. The seller examines the submitted bids, determines the identity of the bidder who submitted the highest bid, and sends that bidder from the room. That bidder subsequently receives one of the items being sold, and pays the amount of his bid. The remaining bidders learn only that there was one bidder who outbid them all, and that he is no longer present. The next round is then conducted.

A *sequential first-price auction with price announcements* proceeds in much the same manner, except that the seller also reveals the amount of the winning bid prior to commencement of the next round.

A *sequential second-price auction* proceeds in essentially the same manner as a sequential first-price auction. The only difference is that the high bidder in each round, after leaving the room, is required to pay the amount of the second-highest bid submitted in the round he won. No price announcement is made to the remaining bidders.

Notice that in the three procedures just described, no specific information concerning the bids of any of the remaining bidders is revealed to them. A sequential second-price auction with price announcements would not have this property; indeed, the situation facing the remaining bidders in the next round would be highly asymmetric, in that (an unidentified) one of the remaining bidders would have his exact bid known to the others. In order to retain a bit of symmetry in the successive rounds, we choose instead to study the *sequential English auction* as the case polar to the sequential second-price auction without price announcements. In this procedure, an English auction is conducted among the remaining bidders in each round. Consequently, in the following round, all of the remaining bidders' bids will be a matter of public record. We will find below that this informationally rich procedure differs little in theory from the simultaneous English auction procedure.

[Here is where the problem occurs. In both the first-price and second-price auctions without price announcements, each remaining bidder enters the second round of the auction knowing that the winning bid in the first round exceeded his own. To the extent that having “better” information has value in subsequent rounds, a bidder might choose to bid a bit higher in the first round in order to have a better estimate of the winning bid, should he lose. This upward pressure is tempered, of course, by the risk of winning, and paying a higher price than would have been paid in the first or a subsequent round had the bid not been raised. In retrospect, it is not at all clear that conjectures in the following sections concerning the “without price announcements” auctions adequately deal with this issue.]

Summary of Results

After laying out the general environment (the “general symmetric model,” with risk-neutral bidders, and fewer items than bidders) in which our analysis takes place, we will treat three major sets of issues.

First, we determine the symmetric equilibrium strategies of the bidders in each of the seven auctions defined above. Two results follow immediately from our equilibrium characterizations. In the sequential first-price auctions, and the sequential second-price auction, the sequence of bids entered by a bidder will (at equilibrium) be strictly increasing. An intuitive explanation of this fact is that in each successive round the ratio of items remaining to bidders remaining decreases, and this decrease of supply relative to demand forces the bidders to become continually more aggressive in their bidding. In any round of a sequential English auction, as soon as the number of active bidders drops to the number of remaining objects, all of those bidders will immediately drop out (in this case, we assume that one of them is selected, at random, to become the winning bidder). At equilibrium, all of the bidders bid to the same level in every round in which they participate; consequently, all of the items sell at exactly the same price, and the result of the auction is identical to the result of the simultaneous English auction.

Second, we consider the equilibrium sequences of prices arising from the sequential first-price auctions and the sequential second-price auction. In all cases, we find the sequence is upward-drifting, i.e., conditioned on the price in any round, the expected price in the next round is at least as great as (and generally greater than) that price. To provide an intuitive explanation for this result, we can argue that the equilibrium sequence of prices cannot be downward drifting. If it were, then consider the purported equilibrium bid of a bidder in the first round. If he makes that bid, he exposes himself to the chance of winning the object in the first round. If instead, he were to bid so low as to have no chance of winning in the first round, he would later find himself in a position where he would expect to have more information than he had originally, and expect the selling price to be lower than the initial price. Both of these expectations are favorable to him; hence the strategy of abstaining in the first round and then seeking an object later in the auction would be superior to the purported equilibrium strategy.

It is interesting to note the sequence of prices generated in the RCA transponder lease auction. From first round to seventh, the prices were \$14.4, \$14.1, \$13.7, \$13.5, \$12.5, \$10.7, and \$11.2 million. This certainly does not appear to be the realization of an upward-drifting sequence! Yet the informational conditions in Sotheby's auction seem to place it somewhere between the sequential second-price and English auctions, and hence we would expect upward drift at equilibrium.

There is an explanation for this discrepancy between theory and observation that seems confirmed to some degree by subsequent events. It is not implausible that many of the firms involved in the auction would have studied the likely value to them of a lease, and on that basis given their bidding agents instructions of the form: "If you can obtain a lease for at most \$x million, do so." Implicit in these instructions is the fact that an agent can fail his client in only two ways: by purchasing an item at more than the authorized maximum price, or by failing to obtain an item, when at least one of the items sells for less than that authorized maximum. Only one strategy guarantees that the agent will not fail his client: he must in every round remain active as long as the price clock is below the authorized maximum. And if all agents do this, then the one with the highest authorization level will win in the first round, at a price equal to the second-highest authorization level; the one with the second-highest authorization level will win in the second round, at the third-highest authorization level, and so on. The resulting sequence of prices will be strictly decreasing!

Of course, if a firm anticipates that the others will be bidding in this fashion, it will benefit from abstaining in the first few rounds. The presence of several such clever firms would create price up-ticks near the end of the sale, just as was observed in the final round. At equilibrium, every agent should plan on bidding to higher and higher levels in each successive round, starting with a first-round maximum bid well below his authorization level, and bidding all the way up to the authorization level only in the last round. It does not appear that such strategies were in wide use.

If our explanation is correct, then several of the firms that won in early rounds paid substantially more than would have been called for in any round at equilibrium. Given a subsequent realization of their strategic error, and an opportunity to withdraw their original bids and purchase leases at a fixed (non-competitive) price, they would choose not to pay an amount near the amount they had originally bid. Such an opportunity presented itself when the FCC (in January of 1982) threw out the results of the auction, ruling that the use of a procedure which sold identical leases at different prices constituted discriminatory pricing, and violated existing common carrier regulations. RCA subsequently offered the leases for sale on a first-come, firstserved basis, at a fixed price of \$13 million. And several of the firms that had bid above \$13 million in the November sale declined in March to purchase leases at that lower price!

Finally, we consider the revenue-generating properties of the auctions under study. Two principal results emerge. The discriminatory auction generally yields lower expected revenues than the uniform-price auction, which in turn generally yields lower expected revenues than the English auction. These results generalize results presented in [1], and contribute to the debate over the choice between a discriminatory or uniform-price procedure for the sale of U.S. Treasury bills. Furthermore, the sequential first-price auctions and the sequential second-price auction generally yield greater expected revenues than their simultaneous counterparts (the discriminatory and uniform-price auctions, respectively). An explanation, related to the "linkage effect" found in [1] (the benefit to the seller of bringing information into the public domain prior to a sale), is that the use of sequential procedures forces the bidders to gradually bring some information into the public domain on their own.

One additional point is worthy of note. Over the past decade, a rich literature has developed reporting the results of auctions conducted in experimental settings. Many of the predictions of equilibrium theory are, however, extremely difficult to test, because of the need to compare results across experiments, and the near impossibility of replicating experimental conditions. Many of the results reported in this paper concern the relationships among statistics observed in the course of a single sale; therefore, they should be highly amenable to experimental testing.

The General Symmetric Model

Consider an auction in which n bidders compete for the possession of one of $k < n$ identical objects. Each bidder possesses some information concerning the objects being sold; let $X = (X_1, ..., X_n)$ be a vector, the components of which are the real-valued informational variables, or *types*, of the individual bidders. Let $S = (S_1, ..., S_m)$ be a vector of additional real-valued variables, unobserved by the bidders, which influence the value of each of the objects to the bidders. The actual value of an object to bidder i — which may, of course, depend on variables not observed by him at the time of the auction — will be denoted by $V_i = u_i(S, X)$. We make the following assumptions (which are identical to the assumptions made for the single object case in [1]:

Assumption 1: There is a function u on R^{m+n} such that for all i, $u_i(S, X) = u(S, X_i, \{X_j\}_{j \neq i})$. Consequently, all of the bidders' valuations depend on S in the same manner, and each bidder's valuation is a symmetric function of the other bidders' types.

Assumption 2: The function u is nonnegative, and is continuous and nondecreasing in its variables.

Assumption 3: For each i, $E[V_i] < \infty$.

Unless otherwise stated, we assume that the bidders' valuations are in monetary units, and that the bidders are neutral in their attitudes towards risk. Hence, if bidder i receives one of the objects and pays an amount b, his payoff is simply $V_i - b$.

Let $f(s,x)$ be the joint probability distribution of the random variables (S, X). We make two assumptions about the joint distribution of S and X:

Assumption 4: f is symmetric in its last n arguments.

Assumption 5: The variables $S_1, \ldots, S_m, X_1, \ldots, X_n$ are affiliated.

The notion of affiliated random variables is developed in [1]. Rather than present a formal definition here, we simply note that affiliation means (roughly) that large values for some of the variables make the others more likely to be large than small, a not unreasonable assumption for the models we have in mind.

Let $X_{(1)} > \ldots > X_{(n)}$ be the order statistics of the bidders' types. Of central importance to any given bidder is the distribution of the types of the others. Therefore, many of our arguments will focus on bidder 1, and we will let $Y_1 > \ldots > Y_{n-1}$ denote the order statistics of $X_2, \ldots, X_n$. We denote the distribution of Y_ℓ by F_{Y_ℓ}, and the associated density by f_{Y_ℓ}. We make two further definitions here:

$$
\begin{aligned}
v_\ell(x,y) &= E[V_1 \mid X_1 = x, Y_\ell = y], \\
v_\ell(x;y_1,\ldots,y_\ell) &= E[V_1 \mid X_1 = x, Y_1 = y_1,\ldots,Y_\ell = y_\ell].
\end{aligned}
$$

Three consequences of the affiliation assumption, which we shall require in subsequent sections, are stated in the following lemma (proved in [1]).

Lemma 1:

(a) $v_\ell(x,y)$ and $v(x;y_1,\ldots,y_\ell)$ are monotone increasing in all of their arguments.

(b) $f_{Y_\ell}(x \mid X_1 = z; y_1,\ldots,y_\ell) / F_{Y_\ell}(x \mid X_1 = z; y_1,\ldots,y_\ell)$ is increasing in z.

(c) The conditional distribution of Y_ℓ, given X_1, $Y_1,\ldots,Y_{\ell-1}$, and given that $Y_\ell \le X_1$, is stochastically increasing in X_1.

In order to simplify some of the ensuing arguments, but without any real loss of generality, we make the nondegeneracy assumption that (both versions of) v_ℓ are strictly increasing in x.

We shall also require the following simple lemma, which is an immediate consequence of the mean value theorem.

Lemma 2: Let g and h be differentiable functions for which (i) $g(\underline{x}) \geq h(\underline{x})$, and (ii) $g(x) < h(x)$ implies $g'(x) \geq h'(x)$. Then for all $x \geq \underline{x}$, $g(x) \geq h(x)$.

Equilibrium Characterizations

We begin our analysis by characterizing symmetric equilibrium strategies for each of the seven auction procedures under consideration. An equilibrium characterization result is usually developed in two stages. Initially, "first-order" and "boundary" conditions are determined, which must be satisfied by any symmetric equilibrium strategy. Then, it is shown that a solution to those conditions is indeed a best response, when all other bidders act according to that solution.

In [1], first-order conditions were determined by examining the function E(x,z), the expected payoff to a bidder when his type is z, he acts (according to a strategy $b(\cdot)$) as if it were x, and the other bidders act according to b. The partial derivative of E with respect to x must be zero when $x = z$, in order that $b(z)$ be a best response when the bidder's type is z; this typically yields a differential equation which b must satisfy at equilibrium. Here, we take an approach that, while formally equivalent to the one just described, sheds additional insight into the form of the first-order conditions.

Consider, for example, the sequential first-price auction with the winning bid announced after each stage. Let $b = (b_1,...,b_n)$ be a symmetric equilibrium strategy for this auction; $b_\ell(x;y_1,...,y_{\ell-1})$ denotes the bid made in the ℓ-th stage by a bidder with type x, if he has not yet won an object and the previously-announced prices lead him to infer that the types of the previous winners are $y_1,...,y_{\ell-1}$.

Consider the decision problem faced by a bidder with type x, in the ℓ-th stage. He must choose whether to bid (according to b_ℓ) as if his type were x, or as if it were some other type. If he acts as if his type were slightly greater than x, say $x + dx$, there are two possible consequences. If he was already going to win, i.e., if Y_ℓ is less than x, then he still wins, and pays slightly more: His unconditional expected payment increases by approximately

$$b'_\ell(x;y_1,...,y_{\ell-1}) \cdot dx \cdot F_{Y_\ell}(x \mid x;y_1,...,y_{\ell-1}).$$

The other consequence arises if his more-aggressive action leads him to win in this round, when he was otherwise going to lose. In this case, it must be that the highest of the opposing types lies between x and $x + dx$; therefore, he was almost certain to win an object in the next stage of bidding. His unconditional expected gain is approximately

$$\left(\left[v_\ell(x;y_1,...,y_{\ell-1},x) - b_\ell(x;y_1,...,y_{\ell-1})\right] - \left[v_\ell(x;y_1,...,y_{\ell-1},x) - b_{\ell+1}(x;y_1,...,y_{\ell-1},x)\right]\right) \cdot$$
$$f_{Y_\ell}(x \mid x;y_1,...,y_{\ell-1}) \cdot dx.$$

(In the k-th stage, the second bracketed term must be replaced by zero, since there will be no subsequent chance to obtain an object.) At equilibrium, these potential losses and gains from deviation must just offset one another. Consequently, the equilibrium strategy b must satisfy:

$$b'_{\ell}(x;y_1,\ldots,y_{\ell-1}) = \left(b_{\ell+1}(x;y_1,\ldots,y_{\ell-1},x) - b_{\ell}(x;y_1,\ldots,y_{\ell-1})\right)\cdot\frac{f_{Y_\ell}(x\,|\,x;y_1,\ldots,y_{\ell-1})}{F_{Y_\ell}(x\,|\,x;y_1,\ldots,y_{\ell-1})}$$

Furthermore, a bidder with the lowest possible type $\underline{x}$ must at each stage have an expected profit of zero, conditioned on the event of his winning at that stage. (If it were negative, he would prefer to abstain from the bidding; if it were positive, then bidding slightly more would yield a greater unconditional expected profit.) Hence, it must be that for every $1 \le \ell \le k$, $b_{\ell}(\underline{x};y_1,\ldots,y_{\ell-1}) = v_{\ell}(\underline{x};y_1,\ldots,y_{\ell-1},\underline{x})$. These are the boundary conditions mentioned earlier.

Arguments analogous to that just given yield the characterizations stated in the following theorem. The proof of the theorem shows that the strategies satisfying these conditions are indeed symmetric equilibrium strategies.

Theorem:

(a) A symmetric equilibrium strategy in the discriminatory auction is determined by

$$b'(x) = \left(v_k(x,x) - b(x)\right)\frac{f_{Y_k}(x\,|\,x)}{F_{Y_k}(x\,|\,x)}, \text{ and } b(\underline{x}) = v_k(\underline{x},\underline{x}).$$

(b) A symmetric equilibrium strategy in the uniform-price auction is determined by

$$b(x) = v_k(x,x).$$

(c) A symmetric equilibrium strategy in the English auction is determined by

$$d_m(x;y_{m+1},\ldots,y_{n-1}) = \\ E[v(x;Y_1,\ldots,Y_{m-1},x,y_{m+1},\ldots,y_{n-1})\,|\,X_1 = x, Y_k = x,\ldots, Y_m = x, (Y_{m+1},\ldots,Y_{n-1}) = (y_{m+1},\ldots,y_{n-1})]$$

for $k \le m \le n-1$. The quantity $d_m(x;y_{m+1},\ldots,y_{n-1})$ denotes the level at which a bidder will drop out of the auction if his type is x and, when the price reaches this level, n-m-1 bidders have already dropped out, indicating by the levels at which they drop out that their types are $y_{m+1},\ldots,y_{n-1}$.

(d) A symmetric equilibrium strategy in the sequential first-price auction without price announcements is determined by

$$b'_\ell(x) = \left(b_{\ell+1}(x) - b_\ell(x)\right)\frac{f_{Y_\ell}(x \mid X_1 = x, Y_{\ell-1} > x)}{F_{Y_\ell}(x \mid X_1 = x, Y_{\ell-1} > x)}, \text{ and } b_\ell(\underline{x}) = v_k(\underline{x}, \underline{x})$$

for $1 \le \ell \le k$, where $b_{k+1}(x) \equiv v_k(x,x)$.

(e) A symmetric equilibrium strategy in the sequential first-price auction with price announcements is determined by

$$b'_\ell(x; y_1,\ldots,y_{\ell-1}) = \left(b_{\ell+1}(x; y_1,\ldots,y_{\ell-1}, x) - b_\ell(x; y_1,\ldots,y_{\ell-1})\right)\frac{f_{Y_\ell}(x \mid x; y_1,\ldots,y_{\ell-1})}{F_{Y_\ell}(x \mid x; y_1,\ldots,y_{\ell-1})},$$

$$b_\ell(\underline{x}; y_1,\ldots,y_{\ell-1}) = v_\ell(\underline{x}; y_1,\ldots,y_{\ell-1}, \underline{x})$$

for $1 \le \ell \le k$, where $b_{k+1}(x; y_1,\ldots,y_k) \equiv v_k(x; y_1,\ldots,y_k)$.

(f) A symmetric equilibrium strategy in the sequential second-price auction without price announcements is determined by

$$b_\ell(x) = E\left[b_{\ell+1}(Y_{\ell+1}) \mid X_1 = x, Y_\ell = x\right]$$

for $1 \le \ell \le k-1$, and $b_k(x) = v_k(x,x)$.

(g) A symmetric equilibrium strategy in the sequential English auction is determined by $d_{\ell,m}(x; y_{m+1},\ldots,y_{n-1}) = d_m(x; y_{m+1},\ldots,y_{n-1})$, as defined in (c), where $d_{\ell,m}$ is the dropout level in the ℓ-th round when m bidders (and the bidder following this strategy) remain active.

Proof: The proofs that the strategies specified in (a), (b), and (c) are symmetric equilibrium strategies are direct analogues of the proofs given for single-object auctions in [1]. The proofs for (d), (e), and (f) have similar outlines. Therefore, we shall give only the proof for (e). [In keeping with previous comments concerning this version of the paper, parts (d) and (f) should be regarded as being in doubt.]

First, note that $v_k(x; y_1,\ldots,y_{k-1}, x)$ is strictly increasing in x; this is a consequence of affiliation and the nondegeneracy assumption. Therefore, it follows from Lemma 2 that $b'_k(x; y_1,\ldots,y_{\ell-1}) > 0$ for all $x > \underline{x}$. Similarly, it follows by induction that all components of b are strictly increasing, and therefore $b_{\ell+1}(x; y_1,\ldots,y_{\ell-1}, x) > b_\ell(x; y_1,\ldots,y_{\ell-1})$ for all $\underline{x} < x < y_{\ell-1}$.

Assume that the n - 1 bidders other than bidder 1 act according to the strategy b. Let H_ℓ denote the following statement: "At stage ℓ, if the previous winners' types were $y_1,\ldots,y_{\ell-1}$ and bidder 1's type is z, then bidder 1 can do no better than to bid $b_\ell(\min\{z, y_{\ell-1}\}; y_1,\ldots,y_{\ell-1})$." (Clearly, bidder 1 will not bid more than $b_\ell(y_{\ell-1}; y_1,\ldots,y_{\ell-1})$.) We next show that $H_{\ell+1}$ implies H_ℓ for $1 \le \ell \le k-1$. Let E(x,z) be the expected payoff to bidder 1 if his type is z, if he bids in stage ℓ as though his type were x > z, and if he acts optimally in all subsequent rounds. Then

$$E(x,z) = \int_{\underline{x}}^{x} [v_\ell(z; y_1,\ldots,y_{\ell-1},\alpha) - b_\ell(x; y_1,\ldots,y_{\ell-1})]\, f_{Y_\ell}(\alpha \mid z; y_1,\ldots,y_{\ell-1})\, d\alpha \, +$$

$$+ \int_{x}^{y_{\ell-1}} [v_\ell(z; y_1,\ldots,y_{\ell-1},\alpha) - b_{\ell+1}(\alpha; y_1,\ldots,y_{\ell-1},\alpha)]\, f_{Y_\ell}(\alpha \mid z; y_1,\ldots,y_{\ell-1})\, d\alpha,$$

where the second expression results from $H_{\ell+1}$. Consequently,

$$\frac{\partial}{\partial x} E(x,z) = \{[b_{\ell+1}(x; y_1,\ldots,y_{\ell-1},x) - b_\ell(x; y_1,\ldots,y_{\ell-1})] \cdot \frac{f_{Y_\ell}(x \mid z; y_1,\ldots,y_{\ell-1})}{F_{Y_\ell}(x \mid z; y_1,\ldots,y_{\ell-1})}$$

$$- b'_\ell(x; y_1,\ldots,y_{\ell-1})\} \cdot F_{Y_\ell}(x \mid z; y_1,\ldots,y_{\ell-1}).$$

The expression within the braces is zero when $x = z$. The bracketed expression is positive for all $\underline{x} < x < y_{\ell-1}$, and Lemma 1 implies that the associated fractional factor is increasing in z. Therefore, $\partial E/\partial x$ is negative when $x > z$. We can similarly show that $\partial E/\partial x$ is positive when $x < z$. It follows that E(x,z) is maximized when $x = \min\{z, y_{\ell-1}\}$.

Statement H_k can be shown to hold by slightly modifying the preceding argument. Therefore, H_1, ..., H_k are all true. It follows that bidder 1 has no better action at any stage than to bid as if his type were z, i.e., b is a symmetric equilibrium strategy.

Finally, (g) follows easily from the observation that the selling price resulting from the indicated strategy is the same in all k stages. Any deviation by a bidder from the indicated strategy either obtains for him an object with expected value less than the price he has to pay, or has no effect on subsequent rounds. **Q.E.D.**

The Independent Private-Values Model

In order to discover what properties the various auction procedures possess, we begin with an investigation of the independent private-values model, wherein each $V_i = X_i$ and $X_1,\ldots,X_n$ are independent, identically distributed, random variables. Even in this simple setting, several striking results are available.

Theorem:

(a) A symmetric equilibrium strategy in the discriminatory auction is determined by $b(x) = E[Y_k \mid Y_k < x]$.

(b) A symmetric equilibrium strategy in the uniform-price auction is determined by $b(x) = x = E[Y_k \mid Y_k = x]$.

(c) A symmetric equilibrium strategy in the English auction, either sequential or not, is determined by $d(x) = x$.

(d) A symmetric equilibrium strategy in the sequential first-price auction, either with or without price announcements, is determined by $b_\ell(x) = E[Y_k \mid Y_\ell < x < Y_{\ell-1}] = E\lfloor X_{(k+1)} \mid X_{(\ell)} = x, \text{for } 1 \le \ell \le k.\rfloor$

(e) A symmetric equilibrium strategy in the sequential second-price auction is determined by $b_\ell(x) = E[Y_k \mid Y_\ell = x] = E\lfloor X_{(k+1)} \mid X_{(\ell+1)} = x \rfloor$ for $1 \le \ell \le k$.

Of course, these results can be obtained as corollaries to the theorem of the preceding section. Several points are worthy of note. First, the seller's expected revenue at equilibrium from each of procedures is the same, and is equal to $k \cdot E\lfloor X_{(k+1)} \rfloor$. This is a consequence of a revenue-equivalence theorem analogous to Theorem 0 of [1].

Second, the strategies in (d) and (e) call for a bidder to make progressively higher bids in each successive stage of the auction. This is not surprising, since the ratio of supply to demand is decreasing from stage to stage.

Third, previous price announcements have no effect on a bidder's strategy in the sequential first-price auction. This is a bit surprising, since one might suspect that the most-recent price would provide some indication of how aggressively a bidder must compete in the current round. A formal demonstration that this is not the case follows from the characterization theorem of the previous section, once one notes that in the independent private values model, for all $x < y_{\ell-1}$, $f_{Y_\ell}(x \mid y_1,\dots,y_{\ell-1}) / F_{Y_\ell}(x \mid y_1,\dots,y_{\ell-1}) = f_{Y_\ell}(x) / F_{Y_\ell}(x)$.

Finally, the sequence of prices generated at equilibrium in any of the sequential auctions is drift-free, i.e., is a martingale. For example, in the sequential first-price auction,

$$
\begin{aligned}
&E\left[b_{\ell+1}(X_{(\ell\_+1)}) \mid b_\ell(X_{(\ell)}) = p\right] \\
&= E\left[E\left[X_{(k+1)} \mid X_{(\ell+1)}\right] \mid X_{(\ell)} = b_\ell^{-1}(p)\right] \\
&= E\left[E\left[X_{(k+1)} \mid X_{(\ell+1)}, X_{(\ell)}\right] \mid X_{(\ell)} = b_\ell^{-1}(p)\right] \\
&= E\left[X_{(k+1)} \mid X_{(\ell)} = b_\ell^{-1}(p)\right] \\
&= b_\ell(b_\ell^{-1}(p)) = p.
\end{aligned}
$$

In the next sections, we shall investigate what happens to these results in the more general context of the general symmetric model.

The Equilibrium Price Series

We begin by looking at the series of bids made by a bidder at equilibrium in a sequential auction. Consider the sequential first-price auction with price announcements. If $y_1,\ldots,y_\ell$ are the ℓ highest opposing types, and if $x < y_\ell$ (so the bidder is present in both the ℓ-th and $(\ell+1)$-st stages, then $b_{\ell+1}(x; y_1,\ldots,y_\ell) \geq b_{\ell+1}(x; y_1,\ldots,y_{\ell+1}, x) > b_\ell(x; y_1,\ldots,y_{\ell-1})$. Consequently, he bids increasingly more in successive stages. Similar results hold for the sequential first-price auction without price announcements, and in the sequential second-price auction. (In the sequential English auction, bidders bid the same amount in every round.)

In general, bidders bid different amounts in the sequential first-price auctions, depending on the announcement procedure. Indeed, the sequence of expected bids is steeper in the case of price announcements than when no announcements are made; in particular, the first-stage bid is lower, and the final-stage bid is higher in expectation, when announcements will be made than when they will not.

We prove here a remarkable property of the equilibrium price sequence.

Theorem: In a sequential first-price auction with price announcements, the series of prices generated at equilibrium is a submartingale.

<u>Proof</u>: Let $W(x,z;y_1,\ldots,y_{\ell-1})$ be the expected price in the $(\ell+1)$-st round, given (a) each bidder other than bidder 1 follows his equilibrium strategy, (b) the $\ell-1$ highest opposing types, as revealed through the earlier-round prices, are $y_1,\ldots.,y_{\ell-1}$, (c) bidder 1's type is z, (d) he bids as if it were x, and (e) he wins in the ℓ-th stage. Then

$$
W(x,z;y_1,\ldots,y_{\ell-1}) = \int_{\underline{x}}^{x} b_{\ell+1}(\alpha; y_1,\ldots,y_{\ell-1}, x) \frac{f_{Y_\ell}(\alpha \mid y_1,\ldots,y_{\ell-1}, z)}{F_{Y_\ell}(x \mid y_1,\ldots,y_{\ell-1}, z)} d\alpha .
$$

Note first that $W(\underline{x},\underline{x}); y_1,\ldots,y_{\ell-1}) = b_\ell(\underline{x}; y_1,\ldots,y_{\ell-1}) = v_\ell(\underline{x}; y_1,\ldots,y_{\ell-1},\underline{x})$. Next, observe that

$$
\begin{aligned}
W'(x,x;y_1,\ldots,y_{\ell-1}) &= W_1(x,x;y_1,\ldots,y_{\ell-1}) + W_2(x,x;y_1,\ldots,y_{\ell-1}) \\
&= b_{\ell+1}\frac{f_{Y_v}}{F_{Y_\ell}} - W\frac{\partial}{\partial x}\ell n\, F_{Y_\ell}(x \mid y_1,\ldots,y_{\ell-1},z|_{z=x} + \int_{\underline{x}}^{x} \frac{d}{dx} b_{\ell+1}\frac{f_{Y_\ell}}{F_{Y_\ell}} d\alpha + W_2 .
\end{aligned}
$$

Hence,

$$
(W - b_\ell)' = (W - b_\ell)\frac{\partial}{\partial x}\ell n\, F_{Y_\ell}|_{z=x} + C + L ,
$$

where C represents the "commonality effect," which is nonnegative because $b_{\ell+1}(\alpha; y_1,\ldots,y_{\ell-1},x)$ is nondecreasing in x, and L represents the "linkage effect," which is nonnegative due to the affiliation assumption. Consequently, $W - b < 0$ implies $(W - b)' > 0$; hence $W \geq b$ always, by Lemma 2. **Q.E.D.**

Note that the linkage effect is nontrivial whenever the types are not independent, and the commonality effect is nontrivial whenever the expected value of an object to bidder 1 varies with the types of the other bidders. Hence, we obtain strict upward drift except in the case of the independent private-values model.

Similar arguments can be used to establish the following theorem.

Theorem: In a sequential first-price auction without price announcements, and in a sequential second-price auction, the series of prices generated at equilibrium is upward-drifting, i.e., given the price in any stage, the expected price in the next stage is at least as great.

Upward-drifting sequences are a bit more general than submartingales, and indeed, it is not necessarily true that for these auctions, the series of equilibrium prices is a submartingale.

Revenue Comparisons

In [1], it is shown that first-price auctions generally yield lower expected revenues than second-price auctions, which in turn generally yield lower expected revenues than English auctions. The following theorem generalizes this result, and is proved in the same manner as the single-object results.

Theorem: The discriminatory auction generally yields lower expected revenues yields than the uniform-price auction. The uniform-price auction generally yields lower expected revenues than the English auction.

As usual, we prove the following theorem only for the case of the sequential first-price auction with price announcements.

Theorem: The sequential first-price auctions generally yield greater expected revenues than the discriminatory auction. The sequential second-price and English auctions generally yield greater expected revenues than the uniform-price auction.

Proof: Conditional on winning an object, a bidder's expected payment in the discriminatory auction, when his type is z and he bids as if it were x, is $P^D(x,z) = b^D(x)$. The corresponding expected payment function in the sequential auction is

$$P^F(x,z) = E\left[b_L^F(x; Y_1, \dots, Y_{L-1}) \mid X_1 = z, Y_k < x\right],$$

where $L = \max\{\ell : Y_{\ell-1} > x\}$. P^D does not increase in z; P^F does (since L, conditioned on $Y_k < x$, is stochastically increasing in X_1, and $b_{\ell+1}(x; y_1, \dots, y_{\ell-1}, x) > b_\ell(x; y_1, \dots, y_{\ell-1})$). Using an argument analogous to that used to prove Theorem 15 of [1], it follows from Lemma 2 that $P^D(x,x) \le P^F(x,x)$ for all $x \ge \underline{x}$. **Q.E.D.**

Risk Aversion

It is natural to wonder which of the preceding results are dependent on the assumption that the bidders are neutral in their attitudes towards risk. Recall from [1] that the revenue ordering between the second-price and English auctions continues to hold when the bidders' preferences are characterized by constant absolute aversion to risk, but that risk aversion weakens the ordering between the first-price and second-price auctions, and can even reverse the ordering (as it always does, for example, in the independent private-values model). Analogous results hold for the discriminatory, uniform-price, and English auctions.

Risk-averse bidders will still bid progressively more in successive rounds of any of the sequential auctions (at equilibrium); however, the upward drift in the equilibrium price series is weakened, and can be reversed. For example, consider the sequential second-price auction in the presence of risk aversion. It is not difficult to show that the symmetric equilibrium strategy, when the bidders' utility functions are $u(\cdot)$, is characterized by

$$u(x - b_\ell(x)) = E\left[u(x - b_{\ell+1}(Y_{\ell+1})) \mid Y_\ell = x\right],$$

for $1 \le \ell \le k-1$, with $b_k(x) = x$. But then

$$\begin{aligned} b_\ell(x) &= x - u^{-1}\left(E\left[u(x - b_{\ell+1}(Y_{\ell+1})) \mid Y_\ell = x\right]\right) \\ &\ge x - E\left[x - b_{\ell+1}(Y_{\ell+1}) \mid y_\ell = x\right] = E\left[b_{\ell+1}(Y_{\ell+1}) \mid Y_\ell = x\right]. \end{aligned}$$

Therefore, if bidder 1 wins in the ℓ-th round, the price he expects to see in the next round will be less than or equal to (and generally less than) the price he pays, $b_\ell(Y_\ell)$.

Summary

We view the analysis presented here as only the first step in the treatment of auctions involving more than one object. A natural continuation would be to consider settings in which bidders desire more than one of the objects being sold. A number of new (and in some cases, surprising — see, for example, the latter sections of [2]) phenomena appear in these settings, and many of the analytical techniques used in this paper will require substantial extension in order to be applicable in the more general settings.

References

[1] P. R. Milgrom and R. J. Weber, "A Theory of Auctions and Competitive Bidding," *Econometrica* 50 (1982), 1089-1122.

[2] R. J. Weber, "Multiple-Object Auctions," *Auctions, Bidding, and Contracting: Uses and Theory* (ed. R. Engelbrecht-Wiggans, M. Shubik, and R.M. Stark), NYU Press, 1983.