IMF WORKING PAPER

WP/95/47

INTERNATIONAL MONETARY FUND

Research Department

Auction Format Matters: Evidence on Bidding Behavior and Seller Revenue

Prepared by Robert A. Feldman and Vincent Reinhart 1/

Authorized for Distribution by Michael Mussa

May 1995

Abstract

This paper evaluates the importance of auction format on bidding behavior and seller revenue, focusing on differences in performance under uniform-price and discriminatory-price formats. The analysis is based on a standard benchmark model from which empirically-testable hypotheses are derived on the optimal amount of bid shading that generates revenue equivalence between the two formats. Applying this model to data from the IMF gold auctions run in 1976-80, we find evidence of statistically significant shading in excess of the theoretically-derived optimum under the discriminatory format. This evidence suggests greater seller revenue under the uniform-price format.

JEL Classification Numbers:
C15, D44

1/ International Monetary Fund and Board of Governors of the Federal Reserve System, respectively. The views expressed are our own and do not necessarily reflect the views of our institutions. We wish to thank Leonardo Bartolini and Paul Masson for helpful comments and Greg Belzer for excellent research assistance.

Table of Contents

Summary

This paper evaluates empirically the importance of auction format for bidding behavior and seller revenue, focusing in particular on differences in performance under uniform-price and discriminatory-price formats. The analysis is based on a standard benchmark model from which empirically testable hypotheses are derived on the optimal amount of bid shading that generates revenue equivalence between the two formats. More specifically, we apply Vickrey's (1961) model of shading, following closely the presentation of McAfee and McMillan (1987). Statistically significant shading in excess of the theoretically derived optimum under the discriminatory format suggests greater seller revenue under the uniform-price format.

The model is applied to a data set based on IMF press releases issued after each of the IMF's gold auctions in 1976-80. These auctions, which were run using both uniform- and discriminatory-price formats, represent a distinct experiment that has apparently escaped rigorous study in the literature. Thirty-five of the 45 auctions--and 10 of the first 20--followed a discriminatory-price format, while the other 10 followed a uniform-price format.

The appropriate choice of auction format is a matter of great practical concern. For example, in the United States, the Treasury currently sells some government securities in uniform-price auctions and others in discriminatory-price auctions, in an attempt to determine which technique provides higher revenues (lower interest costs) to the U.S. Government in auctioning its debt. Mexico has apparently shared those concerns over the comparative performance of discriminatory- and uniform-price formats, as evidenced by a switch from a discriminatory-price to a uniform-price approach in 1990 for its treasury bill auctions, and then, in 1993, by a return to the discriminatory approach. More generally, the uncertainty with regard to the "best" auction technique is readily apparent in the prevalence of both uniform- and discriminatory-price setups for auctioning similar items; practical advice on auction choice from empirical study is rather limited.

From the paper's findings, it is concluded that Vickrey's benchmark model offers useful and empirically valid insights into bidding behavior. What is particularly relevant and apparent in this regard is that auction participants do, in fact, shade their bids under a discriminatory-price format, as the basic model would suggest. At least as important, the paper also provides statistically significant evidence that the extent of this bid shading is, if anything, even larger than this model would indicate, pointing to the superior revenue-generating properties of uniform-price over discriminatory-price auctions.

I. Introduction

A good deal of the trade in goods, services, and assets over time, both within and across countries, involves auctions, particularly when a government or international organization is on one side of the transaction. While the particulars vary, an auction allows public access, monitoring, and equal treatment for participants at a level not normally secured when parties negotiate amongst themselves. Despite their importance to the public purse and the large body of work on bidding theory, practical advice on auction choice is limited. The complexity of the relevant theory is not the hurdle, since there are many accessible summaries available, including Feldman and Mehra (1993), Milgrom (1989), McAfee and McMillan (1987), and Reinhart (1992). Rather, the problem appears to be that there are few simple, applied examples of the importance of auction format beyond the classroom experiments represented by Kagel et al. (1989). In actual practice, the auction format chosen by the seller is implicitly based on a prior (and unobservable) assessment of bidder behavior, as explained in Hansen (1985). Also, real-world experiments with award technique--as in the U.S. Treasury's sale of bonds in the early 1970s--are usually accompanied by several changes in selling technique and not simply a change in auction format alone.

This paper evaluates a distinct experiment that, to our knowledge, has escaped rigorous study in the literature. From 1976 to 1980, the International Monetary Fund (IMF) sold one-fifth of its gold stock at 45 sealed-bid auctions to create a fund to assist developing countries. In general, in sealed-bid auctions awards are made either at the price that was bid (discriminatory-price format) or at a single, market-clearing price (uniform-price format). 1/ This choice, in turn, affects the behavior of bidders in ways that have occupied the attention of auction theorists for years. Importantly, 10 of the first 20 IMF auctions were varied in an announced and predictable manner between uniform-price and discriminatory-price formats, providing a controlled experiment in auction technique; the remainder were based on discriminatory awards.

The IMF press releases at the time provided a wealth of information about the specifics of delivery, and, after the auctions, the names of the participants and a detailed histogram of bids--that is, the quantities bid, sorted by price. This unique dataset facilitates empirical tests of various aspects of bidding behavior and the revenue performance of the two major types of sealed-bid auctions.

1/ Discriminatory- and uniform-price formats refer to auctions of multiple units. When a single item is being auctioned, the corresponding terminology is to first- and second-price formats, respectively. Reinhart (1992) and Feldman and Mehra (1993) discuss the confusing terminology applied to different auction formats in the academic literature and in the popular financial press.

Our strategy is to examine in some detail the benchmark predicted by theory when bidders' valuations of the item sold are independent of each other. While obviously not the case in the setting of a world gold market, this assumption of independent and private values has two advantages: as is familiar in the literature, revenue to the seller is equal under the two major auction formats; and optimal bidders reveal their true valuations in single-price auctions. Employing the latter result directly, we use the data on the ten uniform-price auctions to estimate the distribution of bidder valuations, which we assume holds in the 35 discriminatory-price auctions. We then compare the optimally-shaded valuations, as predicted by bidding theory, with the actual results of the discriminatory-price auctions.

By and large, we observe lower bids--that is, more bid shading--in the discriminatory-price auctions than predicted by the simple theory. Because bids, on average, appear to be lower than predicted by theory, then revenue must also have fallen short of neutral. In related work (Feldman and Reinhart (1995)), we have applied more flexible estimation techniques to this same dataset and reached the same conclusion: Revenue per unit of gold sold was higher in the ten uniform-price auctions than in the 35 discriminatory-price auctions.

The next section describes both the data set and some relevant details of the IMF gold auctions. Section III then reviews some theoretical characteristics underlying the two auction formats to provide the underpinnings for later empirical tests. Before performing these tests, section IV takes a broad-brush look at the performance of the auctions under each of these two formats by examining various summary statistics. Section V tests more rigorously hypotheses on the differences in actual bidding behavior between uniform-price and discriminatory-price formats, consistent with what would be expected on the basis of theoretically-derived optimal bid shading for discriminatory auctions. Concluding remarks are in section VI.

The appropriate choice of auction format is a matter of great practical concern. In the United States, the Treasury currently sells some government securities in uniform-price auctions and others in discriminatory-price auctions, in an attempt to determine which technique provides higher revenues (lower interest costs) to the U.S. government in auctioning its debt. Mexico has also apparently shared those concerns over the comparative performance of discriminatory- and uniform-price formats, as evidenced by the switch from a discriminatory-price to a uniform-price approach in 1990 for its Treasury bill auctions, and then, in 1993, by the return to the discriminatory approach. 1/ In a similar vein, Italy, in 1991, switched from uniform to discriminatory pricing in its local-currency Treasury bill

1/ Umlauf (1994) provides an interesting evaluation of that first shift. His results suggest a higher average selling price (lower borrowing cost) for these securities by switching to the uniform-price setup.

auctions, while Treasury bonds and ECU-denominated bills have been auctioned on a uniform-price basis. The uncertainty with regard to the "best" auction technique is readily apparent in the prevalence of both uniform- and discriminatory-price setups for auctioning similar items, a topic addressed in Rothkopf et al. (1990). 1/

In many other parts of the world, countries are trying to establish and improve market-oriented institutions. This effort is evident, for example, in the transforming economies of eastern and central Europe and the newly emerging states of the former Soviet Union. In some of these transition countries, auctions can, and in some cases already do, offer a useful way of setting prices and allocating resources in a market-oriented setting. Empirical evidence on the implications of using discriminatory- or uniform-price formats in running auctions could be helpful. 2/

II. The Data

The IMF's sales of part of its gold holdings were spread over 45 separate auctions that began in June 1976 and ended in May 1980. In total, the IMF sold 23 1/2 million ounces of gold. These sales were conducted against the backdrop of a large and active secondary market for gold, where the price per ounce ranged from US$110 to US$710 over the course of the IMF's sales (Figure 1).

Following these auctions, the IMF issued press releases designed, in part, to provide data of interest to analysts and students of the gold market. These releases reported data both on prices bid (expressed in intervals of one dollar) as well as on the number of bids and the total amount bid for (expressed in thousands of ounces) in each price interval. 3/ The press releases also provided data on the total number of successful bidders, the total number of bids accepted, the total number of bidders, and the total number of valid bids. Preceding these auctions, the IMF issued invitations to bid that detailed the terms and conditions for

---

1/ A study by Bartolini and Cottarelli (1994), which details the techniques used by various governments to auction government liabilities, finds that in auctioning such securities the discriminatory-price format is in much greater use than the uniform-price format. Kovanen (1994), in summarizing the various auction arrangements that have been employed by developing countries to allocate foreign exchange, finds that both uniform-price and discriminatory-price formats have been prevalent. Tenorio (1993) provides useful analysis of the switch between these two formats in the foreign exchange auctions in Zambia. His conclusions suggest that uniform-price auctions yielded significantly greater revenue.

2/ Discussions of the pros and cons of central bank credit auctions in transition economies are found in Mathieson and Haas (1994) and Saal and Zamalloa (1994).

3/ Participants were allowed to submit multiple bids.

each of the various auctions. Each IMF gold auction was conducted in one round, with bidders entering sealed tenders. Table 1 details main characteristics of the auctions.

Within the broad literature on auctions and related bidding strategy, the Fund's gold auctions fall into a special class, because the items that were sold had close substitutes trading actively in secondary markets as well as on a forward-delivery basis on organized futures markets. In effect, bidders at each auction were attempting to guess what other bidders were guessing--and what other markets had simultaneously priced--to be the common value of the gold being sold. Thus, in technical terms, these were common-value auctions. In such auctions, demand depends importantly on the information available to individual bidders and how they use it to try and outguess each other as to the common value of the auctioned gold. To the extent that the choice of a particular auction format itself serves to reveal information about the bidding intentions of the auction participants, it also influences their expectations of common value and, through that channel, total demand.

It is crucial to this paper that the IMF gold auctions were run using both uniform- and discriminatory-price formats. Thirty five of the 45 auctions--and 10 of the first 20--followed a discriminatory-price format in which bidding took place in private through sealed bids and awards were made at the highest prices covering the total auction size. Meanwhile, the remaining 10 of the first 20 auctions followed a uniform-price format, in which the IMF collected sealed bids, arranged them by price, and selected the highest single price that just placed the amount of gold up for auction. 1/

III. Theoretical Considerations

Vickrey (1961) established that the major auction formats provide equal expected proceeds to the seller when individual valuations are independent--that is, when the subjective worth of a single item on the

---

1/ Other forums for selling gold, such as the London gold fixing (described in O'Callaghan (1993)), employ multiple rounds of bidding, with participants in a single room bidding in public, or connected by phone. Awards can be made at prices that are progressively lowered until the fixed amount of gold to be auctioned is sold--a descending-price (Dutch) auction; alternatively, prices can be progressively increased until arriving at a single price that just exhausts the fixed amount to be auctioned--an ascending-price (English) auction. This terminology follows the pioneering work of Vickrey (1961), and along with the sealed-bid, uniform-price (also termed second-price when a single item is being auctioned) and discriminatory-price (first-price for a single item) auction formats completes the four primary types of auctions distinguished in Vickrey's work.

block is unrelated across bidders. In those circumstances, known as the private-values case in the auction literature, the expectation of the actions of the other bidders only influences the way that a given bidder shades her bid relative to her underlying valuation, not the valuation itself. Revenue in the private-values case is equal across auction types because the act of shading exactly offsets any attempt by the auctioneer in the award technique to extract surplus from bidders. Obviously, the IMF gold auctions violate Vickrey's original setup, as the value agents place on the common value of the auctioned gold reflects an estimate, however imperfect, of the common price in subsequent market trading--where they might want to resell the gold bought from the IMF.

In such a common-values case, any knowledge of other bidders' behavior would influence a given bidder's valuation of the item(s) to be sold, as well as the extent to which he or she shades. Because the auction format itself may convey information about bidders' valuations, the way that the auctioneer sells the item(s) can affect, in the aggregate, underlying valuations and, therefore, have consequences for revenue. Much of the progress of auction theory in the 1980s was related to making these consequences for revenue more specific.

Still, the original Vickrey model provides a useful benchmark for bidder behavior and revenue implications under alternative auction formats. Essentially, this model provides the benchmark for revenue neutrality and the standard from which we judge whether actual bid shading was greater (or less) than what theory would imply, and what would translate directly into less (or more) revenue for the auctioneer. Specifically, we will apply Vickrey's model of shading, following closely the presentation of McAfee and McMillan (1987), to the 35 discriminating-price auctions run by the IMF. 1/ As in that work, we rely on four main assumptions: bidders are risk neutral; individual valuations are independent (the private-values case); bidders are symmetric--that is, they use the same distribution function to estimate their valuations; and payment is a function of bids alone. The benchmark model derived from these assumptions has the advantage of being easy to analyze, allowing us to derive empirically-testable hypotheses under the maintained model. Importantly, we will consider the auction of a single unit, rather than multiple units, for the sake of tractability. 2/

We proceed by deriving an expression for how much an individual bidder would optimally shade his bid below his true valuation under the discriminatory-price format. The expected gain (π) from participating in the auction depends on the product of two terms representing, respectively, (i) the probability that the bid (b) will win, requiring that it be above the stop-out price or the lowest-priced winning bid, and (ii) the excess of

---

1/ Also see Gordy (1994) for a similar presentation.

2/ Some of the complications introduced by the more accurate assumption that there are multiple units for sale are discussed in Weber (1983).

the participant's assessment of the value of the gold (υ) over the award price (a). In simple terms,

$$\pi = \text{Prob}\ (b > \text{stop-out}) * (\upsilon - a) \qquad (1)$$

As to valuation, assume that each of the n bidders in the auction draws their valuation from a fixed probability density function f(•) that has a corresponding cumulative density of F(•). The symmetry of bidders implies that each bidder shares the same f(•), while independence implies one bidder's draw does not influence another's.

The goal is to arrive at a bidding function that describes bidder i's action solely as a function of subjective valuation, with the function perhaps varying by auction type. 1/ In a discriminatory-price auction, the bidding function $b = B(\upsilon)$ is complicated because there is an important tradeoff between the probability of winning and the winner's surplus (υ-a, from above). A high bid heightens the chance of winning the auction but lessens the value to having won. A rational bidder trades between the two terms, lowering b toward the market consensus so as to increase profits from winning the auction while accepting that this action reduces the probability of winning. To win the auction, a bidder's tender of b must beat the other (n-1) agents' bids, an event that, because of the symmetry of agents, only occurs when agent i's valuation is greater than all the other (n-1) valuations. That probability is simply $F(\upsilon_i)^{n-1}$, owing to the independence of valuations.

The expected surplus, noting that the award price (a) is the price bid (b), can be written as

$$(\upsilon_i - b_i)\ F(\upsilon_i)^{n-1}$$

It is helpful to work in terms of the bidding function, which can be inverted as

$$\upsilon_i = B^{-1}(b_i)$$

in order to write the expected surplus as

$$(\upsilon_i - b_i)\ F[B^{-1}(b_i)]^{n-1}. \qquad (2)$$

The optimizing bidder chooses b_i to satisfy the standard interior condition that the first derivative of the expected surplus with respect to the bid should equal zero, or

1/ And because bidders are symmetric, the same function applies across all n participants so that we do not have to index by individual.

$$B'(\upsilon_i) = \frac{F'(\upsilon_i)^{n-1}}{F(\upsilon_i)^{n-1}} \cdot (\upsilon_i - b_i). \tag{3}$$

This is a first-order differential equation that can be integrated forward from the boundary condition that at the lowest possible valuation the bidder earns no surplus:

$$B(\upsilon_i) = \upsilon_i - \frac{\int_0^i F(u)^{n-1} du}{F(\upsilon_i)^{n-1}} \tag{4}$$

The first term is simply the bidder's valuation. The second term is the degree to which he or she shades that valuation depending on the distribution of valuations. 1/

Deriving a bidding function is an easier goal for a uniform-price auction, which in the auction of a single item requires that the award be made at the price of the highest losing bid, p (which is why it is sometimes called a second-price auction). In this circumstance, the amount of the bid itself does not affect the bidder's surplus (υ-p), only the probability of winning. In other words, in terms of the expected gain to participating in the auction (π), uniform-price awards separate the probability of winning from the profits from having won. This is because the award price is not the same as the amount bid, as is the case for a discriminatory-price auction, but rather the highest single price (p) that just places the amount of gold up for auction. As a result, participants bid their true valuation of the gold being auctioned; and there should be no bid shading. The optimal bidding rule follows as

$$b = \upsilon. \tag{5}$$

1/ This issue is discussed rigorously in J. Smith (1981).

This rule implies that bids in a uniform-price auction reveal underlying valuations, a result that we will apply to our advantage in the empirical work that follows.

There are many reasons to suspect that the actual amount of bid shading can differ from that implied by equation (4), all related to one violation or another of the assumptions of the independent private values case. As already mentioned, the assumption most likely to be inappropriate in describing IMF gold auctions concerns the use of information. In fact, there was likely a high degree of correlation among the individual valuations because all participants could look at a common signal, the secondary market price of gold.

Allowing some commonality in valuations implies that the choice of auction technique can influence expectations formation. Intuitively, with all participants guessing about the same price (for where the gold will trade after the auction), a high bid signals a heightened probability of subsequent loss for that bidder. Planning on winning in a discriminatory auction therefore requires that the bidder trim back further his actual bid from his initial reading of his expected valuation of the gold for auction--that is, υ in equation (1) is lower after the bidder finishes the thought exercise of asking "What do I learn from winning the auction that I did not know before?". This is the "winner's curse" and causes aggressive bidders to rein in their enthusiasm by moving bids down toward the perceived market consensus.

Other things being equal, as the number of bidders increases, it is prudent to bid more conservatively. This is because the range of the distribution of bids, and thus the highest bid, is likely to expand with the number of bidders, reinforcing the winner's curse and thereby creating a greater shading of bids below the bidder's true valuation. Furthermore, the gap between the highest bid and the "true" value of the auctioned gold should decrease as the amount of information available increases. The winner's curse is therefore muted by increasing information about the value of what is being auctioned. It is also to be noted that avoiding the winner's curve may lead to a pooling of bids, as a group of investors is more likely to have a clearer view of the market consensus. Also, in this simple model we assumed that only one item was being sold and that the number of bidders was fixed. Intuitively, auctioning multiple items allows a bidder to shade more in a discriminatory-price auction, because to get one of the k items for sale, he or she does not have to beat the n-1 other bidders, only n-k-1 of them. Offsetting this, if other bidders can enter or exit the competition based on their expectation of the likely outcome, one would presume that the scope for an auction award well below valuation is narrowed.

In light of the discussion above, a comparison of the theoretically derived bids predicted by the simple independent values case with the data on actual bids can shed some light on the importance and direction of the net effect of these other factors in determining bidding behavior.

IV. Summary Statistics

Some of the theoretical characterizations discussed above are consistent with the summary statistics on the gold auctions reported in Table 2, which provide information on prices bid, weighted by the volume of bids. In particular, as shown in the first two columns, bidding, on average, appeared to be somewhat more aggressive for uniform-price auctions. The averages of all bids and winning bids, at, respectively, 99.76 and 100.68 percent of the secondary market price posted before the auction, were greater than that recorded either at the first ten or at all discriminatory-price auctions. 1/ More aggressive bidding in the uniform-price auctions, of course, need not translate into higher revenue for the seller, because awards are made at one market-clearing price. The last column of the table presents revenue (relative to the secondary market price) at these auctions. The 0.75 difference between revenue and the mean of winning bids at uniform-price auctions represents the consumer surplus that the auctioneer deliberately chooses not to claim. In contrast, with awards made at the price bid, the auctioneer seized all the surplus shown at the discriminatory-price auctions -- the volume-weighted prices of winning bids equals revenue. On net, more aggressive bidding more than compensated for not seizing consumer surplus and revenue was higher in the uniform-price auctions than in the discriminatory-price auctions. The difference, in revenue, however, is quite small. The variance of winning bids is also significantly higher for uniform-price auctions, consistent with the view that bidders shade less under that auction format. This same sense of more aggressive bidding in uniform-price auctions can be seen in Figure 2, which displays the range of winning bids according to auction format. Quite clearly, maximum bids tended to be higher under the uniform-price format.

The theoretical discussion suggests that shading should make bids under the discriminatory-price format cluster closer to the market consensus than under the uniform-priced format. It would therefore seem reasonable to expect that bids under the discriminatory auctions would be more tightly compressed than for the uniform auctions, and that the variance of the bids would also be lower, especially because of the effect of less aggressive bidding in the range of winning bids (near the consensus). This is in fact the case: The range of winning bids tended to be considerably lower when the discriminatory-price format was used (see Figure 2); and the variance was lower as well (see Table 2). When all bids are considered, the variance is higher under the discriminating-price format. Inspection of the lower panel of Figure 2 would suggest that minimum bids at times fell well below

---

1/ In all the empirical work that follows, U.S. dollar bids at the auction will be converted into relative prices by dividing by the secondary market price of gold from the previous day. This helps to control the substantial variation on gold prices over the period. Given the substantial upward trend to gold prices over the period, we compare the uniform price auctions, which were concentrated early in the sample, with both the first 10 and remaining 25 discriminatory-price auctions.

those for the uniform-price auctions--that is, participants at discriminatory-price auctions often placed quite low, off-market, bids presumably in the hope of catching a bargain. In the event, no such bargains were evident, and the average variance of winning bids was markedly below that of all bids in discriminatory price auctions.

All told, the data appear to be broadly consistent with a number of theoretical priors, although the analysis at this stage is only meant to be suggestive. The next sections contains more rigorous empirical examination.

V. Testing Bidding Behavior

Putting theory into practice can be difficult, as closed-form representations for bidding in a discriminatory auction are not generally available, except when the distribution of valuations take very simple forms. 1/ However, the numerical solution to equation (4) is trivial for any distribution of valuations. This property has been used quite profitably by Laffont and Vuoung (1993), Paarsch (1992), and Gordy (1994) to test bidding theory in several applications using simulated nonlinear least squares. Such techniques are quite data intensive and so we, instead, adopt an indirect test of the independent-values bidding paradigm given our limited sample.

Note that in the independent (also known as private) values case, actual bids in uniform-price auctions uniquely identify the distribution of valuations. That is, by equation (2), b and υ are equal so that the empirical distribution of bids can be used to estimate F(•), which we call $\hat{F}$(•). Further, if the distribution of valuations is similar across auctions, the estimate $\hat{F}$(•) can be used to calculate the appropriate degree of bid shading in discriminatory-price auctions, given by equation (4). This theoretical construct can, in principle, be compared to the actual results of the discriminatory-price auctions. However, in our application, the 35 discriminatory-price auctions were spaced over four years and so it is probably inappropriate to assume that F(•) was unchanged over time. Fortunately, we can easily allow F(•) to vary in a mean-preserving manner to capture the observed variance of bids in the discriminatory-price auctions, introducing a single parameter to be estimated auction by auction. The resulting distribution of optimally-shaded bids can be compared to the actual distribution of bids, which completes our indirect test of bidding theory.

To repeat this strategy in four explicit steps:

1/ An example when F(•) follows a uniform distribution is provided in V. Smith (1989).

1. Assume that the independent values case holds and that bidders act optimally.

2. Assume that unobserved valuations at the 35 discriminatory-price auctions were proportional to those at the 10 uniform-price auctions, i.e., $F[k_i(\upsilon-\hat{\mu}) + \hat{\mu}]$ where μ is the estimated mean calculated from the bids for the uniform-price auctions and k_i is the unknown scaler multiple.

3. Estimate the proportionality factors, k_i, for the 35 separate auctions by an iterative technique that sets the estimated variances of optimally-shaded bids by repeatedly applying equation (4) to those of the actual bids. Note that the actual number of independent bidders varies across auctions.

4. Compare the distributions of optimally-shaded bids with actual bids for the 35 discriminatory price auctions by parametric and nonparametric methods. The parametric method is a simple one-tailed t-test for equality of means, while the nonparametric technique is the Kolmogorov-Smirnov test statistic.

By following this strategy, we nest two sets of assumptions (items 1 and 2 above) and employ two separate estimators (items 3 and 4 above). To implement this procedure, we had to cope with two consequences of the major limitation to our dataset--that we do not have individual bids, but only the aggregate volume of bids falling in $1 price ranges. First, this implies that the number of observations that we have are the total number of bins reported for each auction, which is what we used for sample size in the statistical tests that follow. Second, because of the way that the data are reported, we cannot determine if a bidder placed multiple bids. In calculating the appropriate degree of bid shading, we assume that bids, n, are equivalent to participants.

The three panels of Figure 3 apply this process for all bids pooled across the 35 discriminatory-price auctions. The upper panel plots the cumulative density functions for relative bids as a share of all bids for the 10 single-price auctions, estimated by a kernel-smoothing algorithm. 1/ The middle panel compares the optimal strategies at the discriminatory-price auctions (the dashed line) and a single-price auction (the solid line), which were provided in equations (4) and (5), when $n = 18.6$ and $k = 2.03$. The former value is simply the average number of participants in the 35 discriminatory-price auctions, while the latter is the multiplicative term that adjusts the variances of the optimally-shaded

---

1/ As has already been noted, dollar prices have been converted into relative prices by dividing by the secondary market price of gold from the previous day. When pooling across auctions (where the amount sold varied from 444,000 ounces to 780,000 ounces), we divided the quantity bid by the auction stock.

distribution in the upper panel so that the variances of the two distributions in the lower panel are equal. In that lower panel, the solid line is the actual average distribution of bids and the dashed line is the outcome of optimal bidding.

The important result in this comparison is that the distribution of optimally-shaded bids lies to the right of that of actual bids. Bidders, in practice, tended to shade more than the optimal strategy predicted in the independent values case. 1/ This is important because the independent values case generates revenue neutrality for the auctioneer. If bidding is less aggressive in a discriminatory-price auction than that benchmark, then the auctioneer must have received less revenue than would have been the case under single-price awards.

The inset to the lower panel provides two statistical tests of the equality of these distributions. Having chosen k so that the variances of actual and optimal bidding are equal, we can still test if they have equal means. As is evident from the t-test of the equality of means, the sample average that is actually observed is significantly lower than the average optimal bid. Similarly, the Kolmogorov-Smirnov statistic, which measures the widest spread between the two density functions and is a useful nonparametric test of equality of distributions, would not likely be observed if the distributions, in fact, were identical.

No doubt, there was a great deal of heterogeneity among the IMF gold auctions, spread as they were across four years and a price of gold that ranged from US$110 to US$710 per ounce. The two panels of Figure 4 repeat a comparison of actual and optimal bidding, this time for the pooled data for the first ten discriminatory-price auctions (the upper panel) and the last twenty-five auctions (the lower panel). The first ten discriminatory auctions were conducted within the overall plan to experiment with auction technique, and the price of gold varied in a relatively narrow range. After auction 20, the experiment was dropped so that only discriminatory-price auctions followed. Meanwhile, the price of gold sky-rocketed on world markets. The split in the sample offered in Figure 4, then, appears quite natural.

In both cases, the distribution of optimally shaded bids lies to the right of actual experience--that is, bidders shaded more than called for by the independent private values case. Again, our results suggest that less revenue was received than if the gold offerings were sold at single prices.

1/ It might help the geometric intuition to think of the distribution functions in the upper and lower panels as demand curves plotted with the axes reversed and the quantity scale inverted. In other words, the corresponding demand curves plotted in the (q,p) quadrant would look like those panels if the origin were given at the upper left of each panel--turn Figure 3 counter clockwise for one quarter rotation. Viewed that way, we see that the actual demand curve lies inside of the optimally chosen one.

The statistics offered in the inset to the charts indicate that these are significant differences.

Because this procedure is not particularly data intensive, we examined bidding behavior on an auction-by-auction basis, scaling the benchmark shading derived by applying equation (4) to the average result from the ten uniform-price auctions to match the observed variability in each of the 35 discriminatory-price auctions and testing for equality of means and of distributions. Table 3 records summary statistics for this procedure when applied to each of the 35 discriminatory price auctions. These results suggest that theory holds up well, in the sense that bidders do, in fact, shade their bids in discriminatory-price auctions. However, they tend, and by statistically significant amounts, to shade more than would be suggested by the basic theoretical model. In 24 of the 35 discriminatory auctions, and 7 of the first 10 of them, the means of actual bids were significantly below what theory dictated. Because the benchmark model makes the best case for revenue equivalence, bid shading in excess of what the theoretical model implies would suggest that these discriminatory-price auctions resulted in less revenue to the seller than uniform-price auctions would have produced.

Of course, a main reason for rejecting the hypotheses derived from the benchmark model could be that the simplifying assumptions of private values and symmetric information do not in practice hold. One way to proceed would therefore be to relax these assumption and test alternative hypotheses derived from less stylized and more realistic models. Unfortunately, this task would be complicated and data intensive. Importantly, we would have the very difficult task of modeling alternative information structures.

Further analysis of this issue is the focus of another paper by Feldman and Reinhart (1995), which goes back to an older strand of the literature and fits separate (reduced-form) demand curves to the data for the uniform- and discriminatory-price auctions, respectively. The main issues examined in that paper are whether the demand curves under each of these auctions formats are in fact different, as theory would suggest; and what can then be said about whether one auction format or the other produces higher revenues.

VI. Concluding Remarks

We conclude from our findings that the benchmark model offers useful and empirically-valid insights into bidding behavior. What is particularly relevant and apparent in this regard is that auction participants do, in fact, shade their bids under a discriminatory-price format, as the basic model would suggest. At least as important, we also provide evidence that the extent of bid shading, if anything, is even larger than this model would indicate, pointing to the superior revenue-generating properties of uniform-price over discriminatory-price auctions. This was the result from our comparison of actual and optimal bidding on the basis of pooling all 35 discriminatory-price auctions, as well as on the basis of splitting the sample after the first ten of them. The results from analyzing each of the

35 discriminatory-price auctions individually also confirm the hypothesis of bid shading under this format, and by amounts in a significant number of auctions that were larger than that derived from the benchmark model; in only one of the 35 discriminatory auctions was bid shading absent.
It is also relevant to note that the summary statistics discussed earlier in the paper provide further evidence of differential behavior under the two auction formats studied, even if the analysis is undertaken in a less rigorous manner.

While our results have important implications for the revenue performance under the two auction formats, we do not explore this subject in more detail in this paper, which has focused directly on differences in bidding behavior under the alternative auction formats; nor do we attempt to "identify" which particular violations of the assumptions of the benchmark model might have led to our empirical findings of bid shading in excess of theoretically-derived amounts. Clearly, however, the assumption of independent private values breaks down in the common values situation of the gold auctions and the winner's curse is an important factor; a breakdown in the assumption of symmetric information is also an obvious candidate. Ongoing and future research should continue to strive to develop more realistic models for the case of common values and asymmetric information, which could allow for more structurally revealing statistical tests of differentiated bidding behavior and corresponding analysis of revenue performance under different auction formats. Another avenue is to explore further the differences in (reduced-form) demand curves under the two auction formats and their revenue-generating properties, an approach we take in a subsequent paper using various statistical methods.

Figure 1

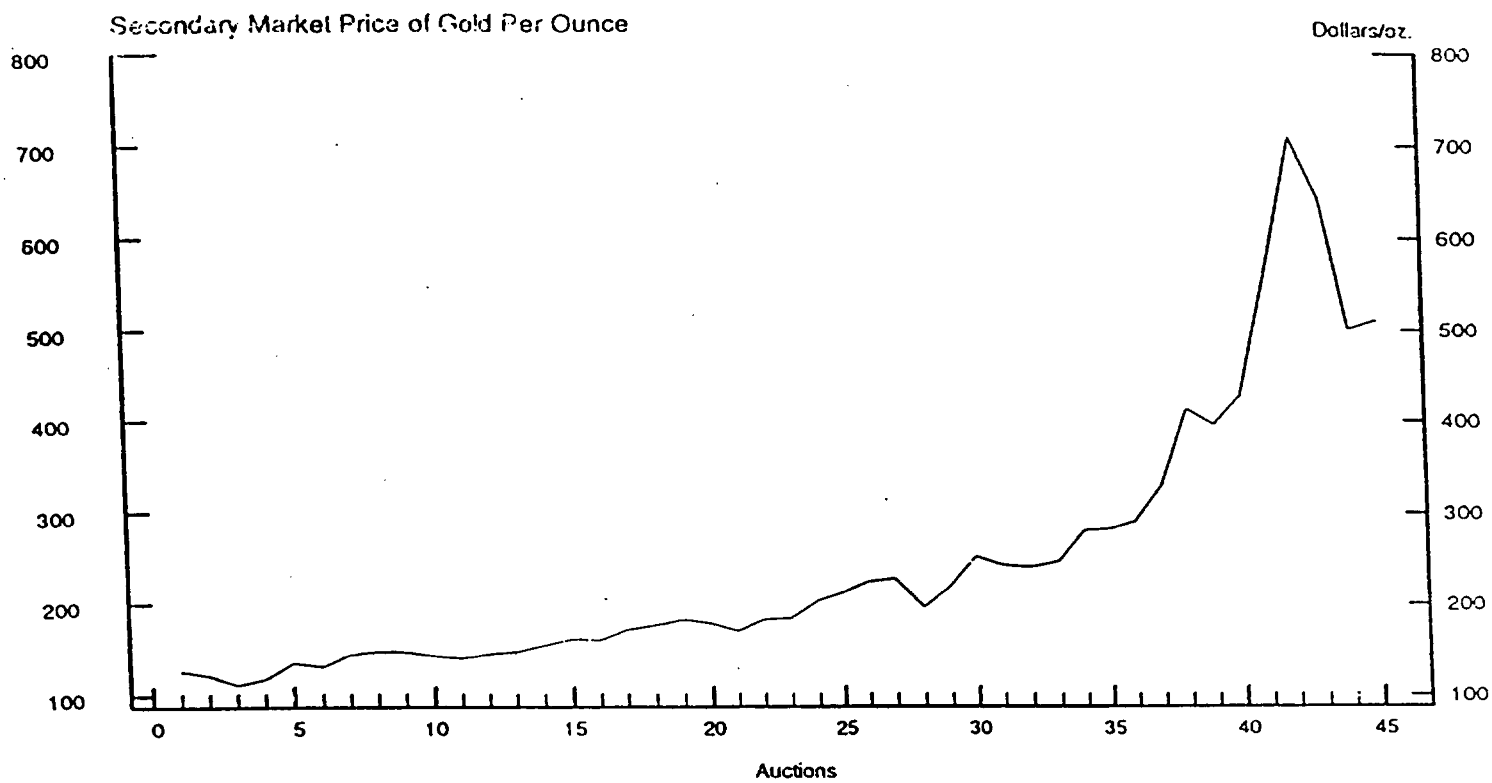
Secondary Market Price of Gold Per Ounce
Dollars/oz.
800
700
600
500
400
300
200
100
0
5
10
15
20
25
30
35
40
45
Auctions

Table 1

Summary of IMF Gold Auctions

Auction Number	Date	Type	Press Release #	Size (troy ounces)	Total Bids (troy ounces)	Minimum Bid (troy ounces)	Deposit	Payment (Days) 1/	Delivery Place	Delivery Time (Days) 2/
1	6/02/76	U	76/34	780,000	2,320,000	2,000	$50,000	28	FRBNY	30
2	7/14/76	U	76/46	780,000	2,114,000	1,200	$50,000	28	FRBNY	42
3	9/15/76	D	76/70	780,000	3,662,400	1,200	$50,000	28	FRBNY	42
4	10/27/76	D	76/80	779,200 3/	4,214,400	1,200	$50,000	28	FRBNY	42
5	12/08/76	U	76/94	780,000	4,307,200	1,200	$50,000	21	BOE	30
6	1/26/77	U	77/2	780,000	2,003,200	1,200	$50,000	14	FRBNY	28
7	3/02/77	D	77/12	524,400	1,632,800	1,200	$50,000	14	FRBNY	28
8	4/06/77	D	77/16	524,800	1,278,000	1,200	$50,000	14	FRBNY	28
9	5/04/77	D	77/24	524,800	1,316,400	1,200	$50,000	14	FRBNY	28
10	6/01/77	U	77/40	524,800	1,014,000	1,200	A	9	FRBNY	23
11	7/06/77	U	77/47	524,800	1,358,400	1,200	A	9	BOF	28
12	8/03/77	U	77/53	524,800	1,439,200	1,200	A	9	BOE	23
13	9/07/77	D	77/62	524,800	1,084,400	1,200	A	9	FRBNY	23
14	10/05/77	D	77/74	524,800	971,200	1,200	A	9	FRBNY	23
15	11/02/77	D	77/83	524,800	1,356,400	1,200	A	9	BOE	23
16	12/07/77	U	77/88	524,800	1,133,600	1,200	A	9	FRBNY	23

Table 1
(Continued)

Summary of IMF Gold Auctions

Auction Number	Date	Type	Press Release #	Size (troy ounces)	Total Bids (troy ounces)	Minimum Bid (troy ounces)	Deposit	Payment (Days) 1/	Delivery Place	Delivery Time (Days) 2/
17	1/04/78	U	77/94	524,800	984,800	1,200	A	9	FRBNY	23
18	2/01/78	U	78/3	524,800	598,400	1,200	A	9	BOF	27
19	3/01/78	D	78/13	524,800	1,418,000	1,200	A	9	FRBNY	23
20	4/05/78	D	78/17	524,800	1,367,600	1,200	A	9	FRBNY	23
21	5/03/78	D	78/24	524,800	3,104,000	1,200	A	9	BOE	23
22	6/07/78	D	78/31	470,000	1,072,400	1,200	A	9	FRBNY	23
23	7/05/78	D	78/42	470,000	797,200	1,200	A	9	FRBNY	23
24	8/02/78	D	78/53	470,000	1,467,600	1,200	A	9	FRBNY	23
25	9/06/78	D	78/62	470,000	773,200	1,200	A	9	FRBNY	23
26	10/04/78	D	78/66	470,000	805,600	1,200	A	9	BOE	23
27	11/01/78	D	78/74	470,000	689,600	1,200	A	9	FRBNY	23
28	12/06/78	D	78/83	470,000	1,965,200	1,200	A	9	BOF	23
29	1/03/79	D	78/92	470,000	1,479,600	1,200	A	9	FRBNY	23
30	2/07/79	D	79/6	470,000	1,489,600	1,200	A	9	FRBNY	23
31	3/07/79	D	79/12	470,000	1,534,400	1,200	A	9	BOE	23
32	4/04/79	D	79/18	470,000	1,186,800	1,200	A	9	FRBNY	23
33	5/02/79	D	79/28	470,000	1,514,800	1,200	A	9	FRBNY	23

Table 1
(Concluded)

Summary of IMF Gold Auctions

Auction Number	Date	Type	Press Release #	Size (troy ounces)	Total Bids (troy ounces)	Minimum Bid (troy ounces)	Deposit	Payment (Days) 1/	Delivery Place	Delivery Time (Days) 2/
34	6/06/79	D	79/36	444,000	1,452,400	1,200	A	9	FRBNY	23
35	7/03/79	D	79/44	444,000	1,518,800	1,200	A	10	FRBNY	24
36	8/01/79	D	79/52	444,000	1,138,800	1,200	A	9	FRBNY	23
37	9/05/79	D	79/62	444,000	1,646,000	1,200	A	9	BOF	23
38	10/10/79	D	79/72	444,000	665,600	1,200	A	9	FRBNY	23
39	11/07/79	D	79/81	444,000	1,798,400	1,200	A	9	FRBNY	23
40	12/05/79	D	79/95	444,000	1,746,000	1,200	A	9	BOE	23
41	1/02/80	D	79/101	444,000	1,342,400	1,200	A	9	FRBNY	23
42	2/06/80	D	80/8	444,000	1,939,600	1,200	A	9	FRBNY	23
43	3/05/80	D	80/15	444,000	1,412,400	1,200	B	9	FRBNY	23
44	4/02/80	D	80/24	444,000	802,800	1,200	B	9	FRBNY	23
45	5/07/80	D	80/31	443,200 3/	1,822,000	1,200	B	9	FRBNY	23

Notes: U = Uniform price auction.
D = Discriminatory price auction.
A = Greater of $25,000 or $10 per ounce of final bid.
B = $40 per ounce of final bid.
FRBNY = Federal Reserve Bank of New York.
BOE = Bank of England.
BOF = Bank of France.
1/ = Time from auction date to payment due date.
2/ = Time from auction date to delivery date.
3/ = 800 ounces could not be awarded because they did not reach the minimum award of 1,200 ounces under the terms and conditions of the auction.

Table 2

Summary Statistics Relative to the Secondary Market Price
(weighted by ounces)

	Prices Bid				
	Mean		Variance		
	All	Winning	All	Winning	Revenue
Uniform Price	99.76	100.68	1.58	0.50	99.93
Discriminating Price					
First 10	99.71	99.90	5.17	0.11	99.90
All	98.76	99.88	4.16	0.28	99.88

Figure 2

Bid Ranges Relative to Secondary Market Price

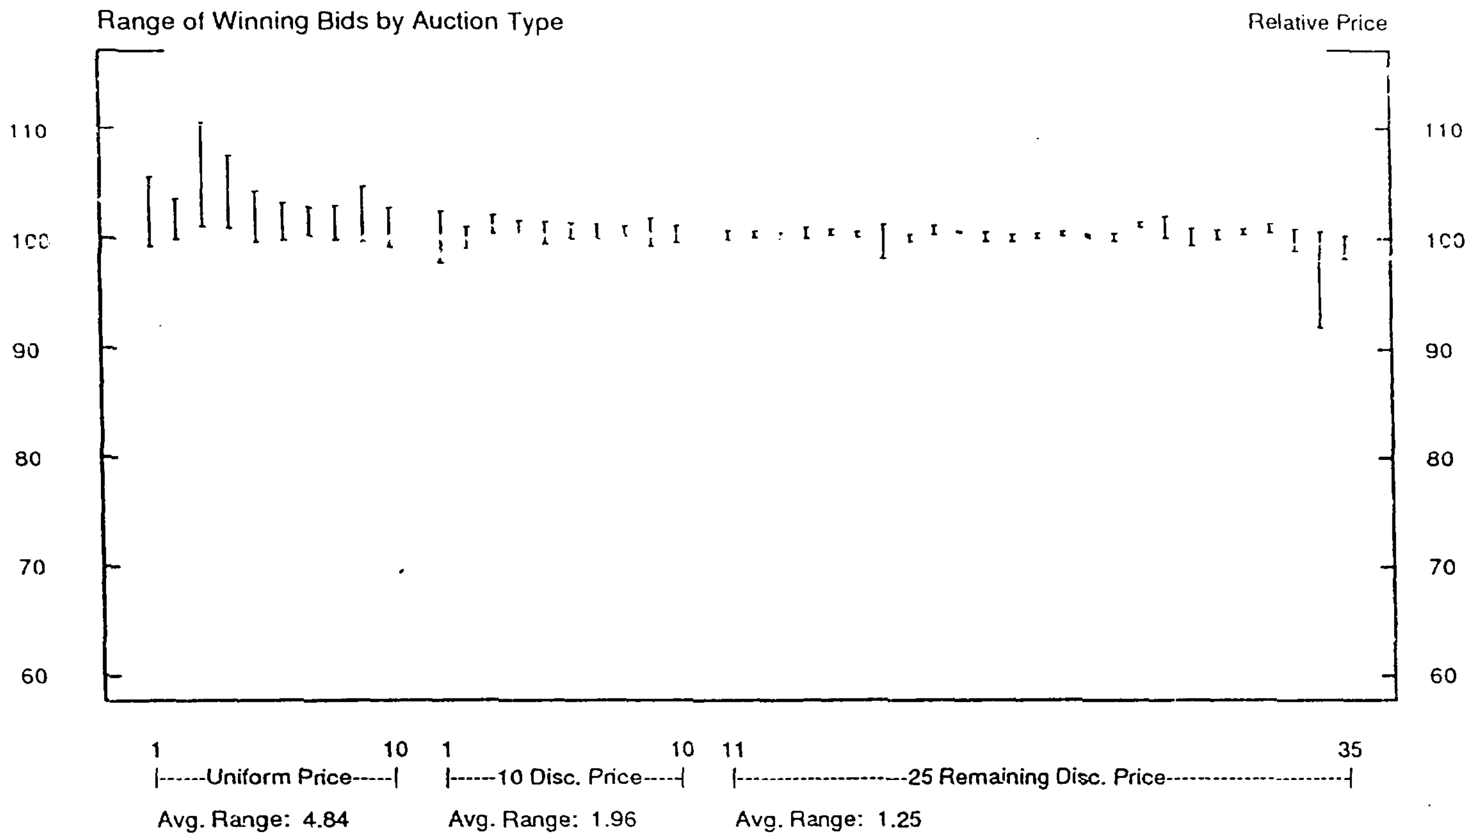

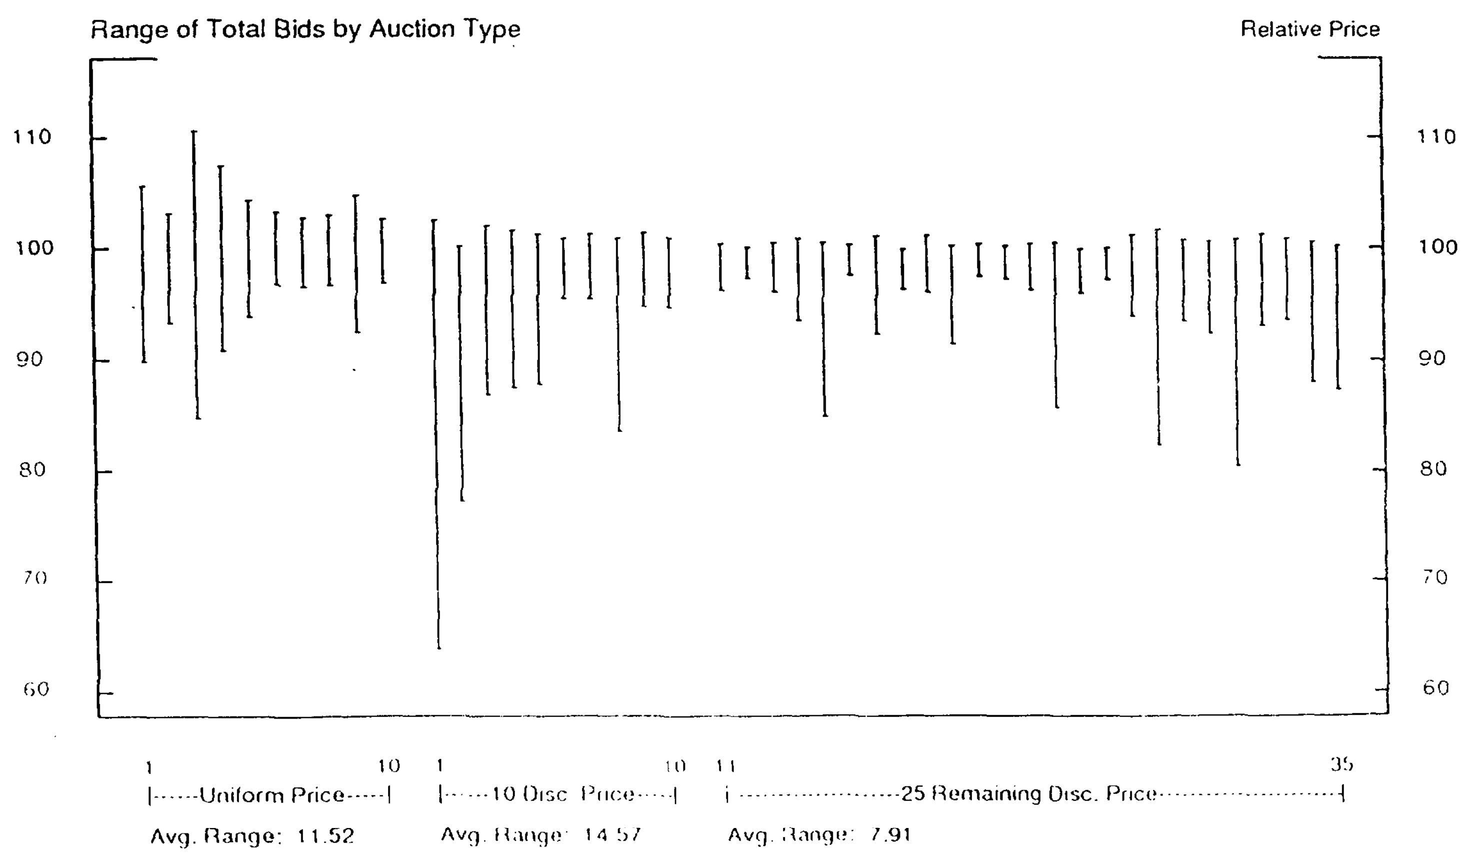

Figure 3

Optimal Shading for Discriminating Price Bids at All Auctions

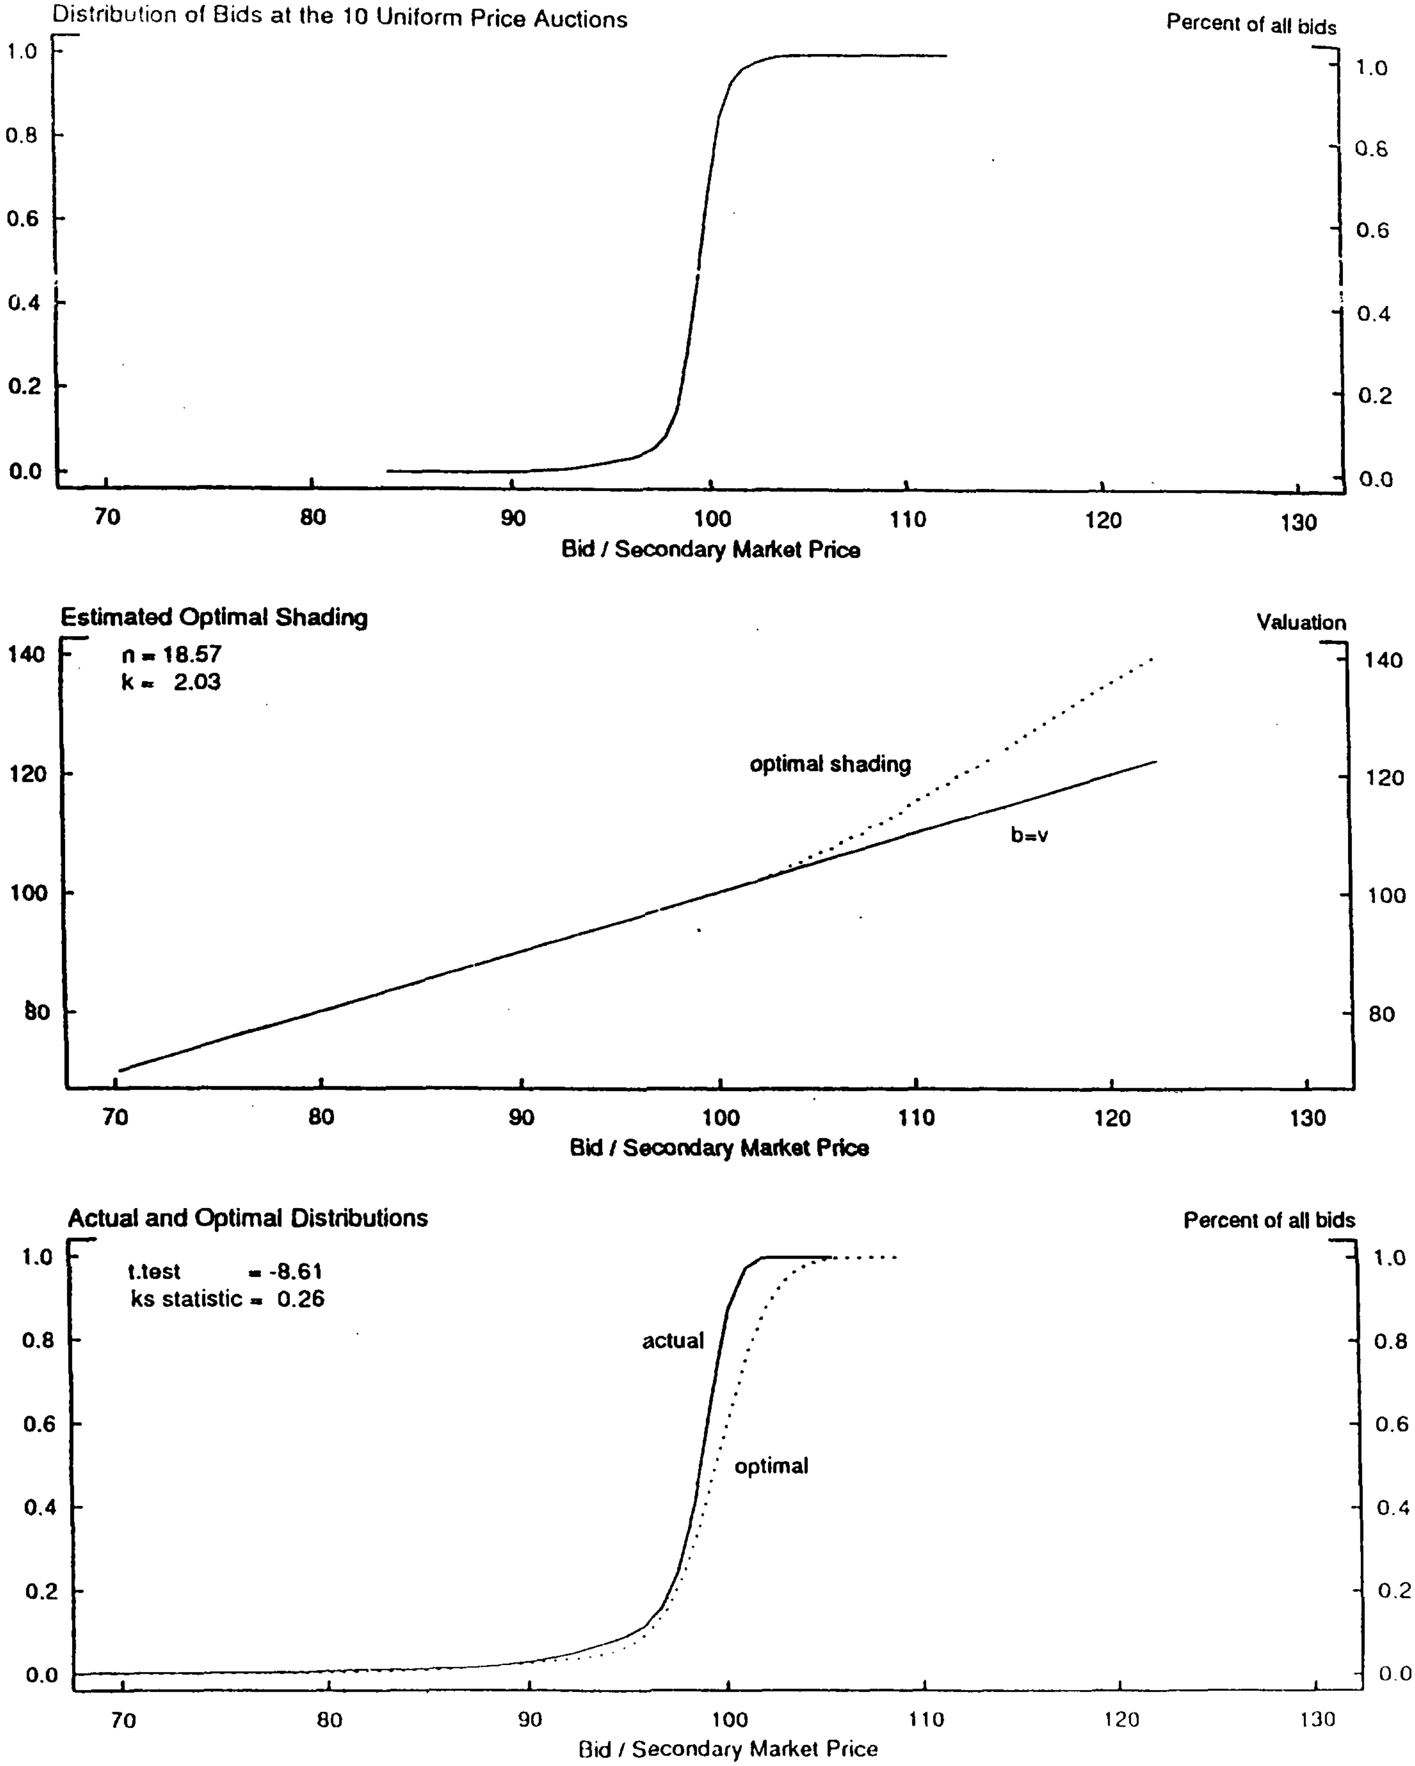

Figure 4

Optimal Shading for Discriminating Price Bids

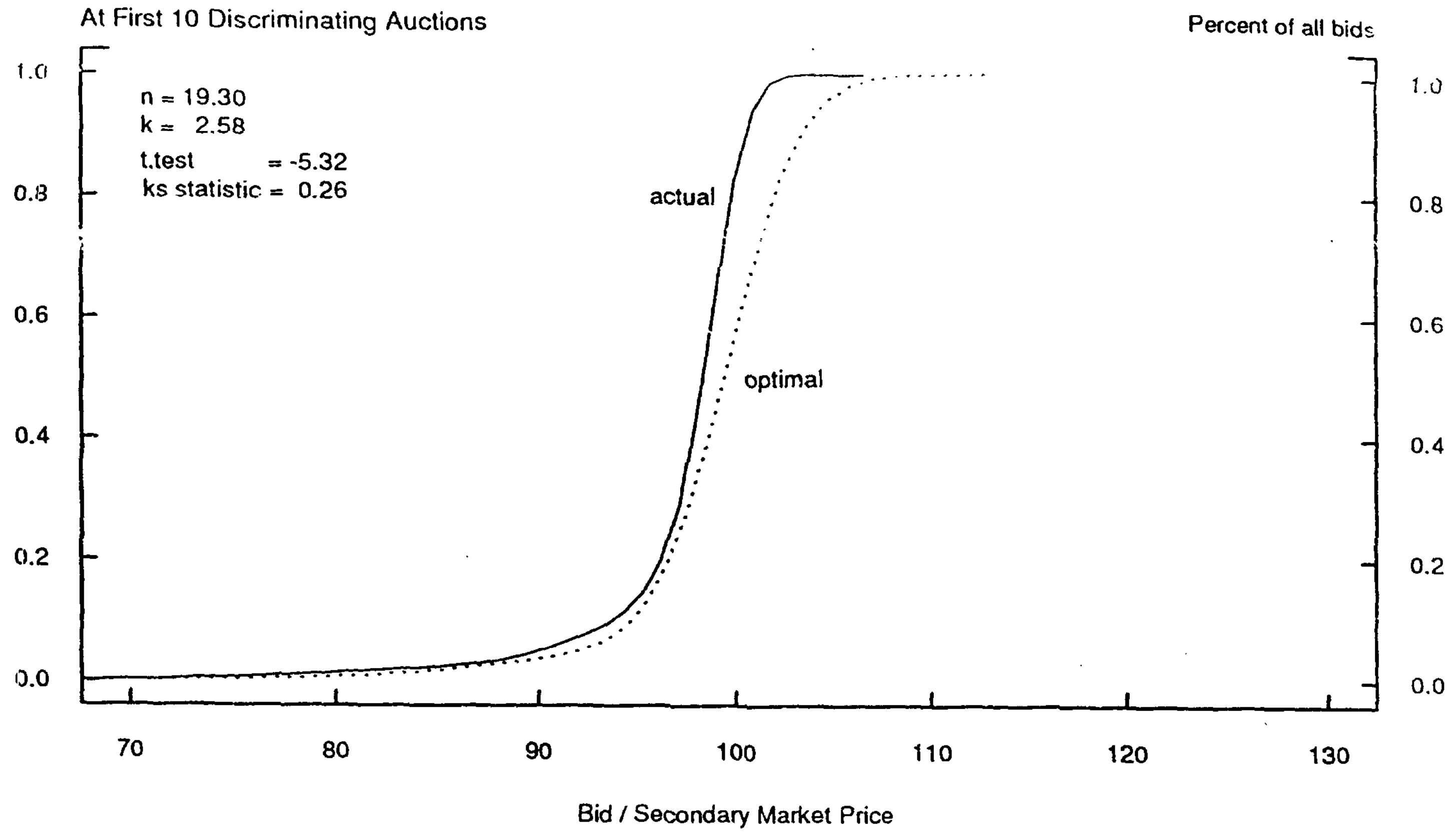

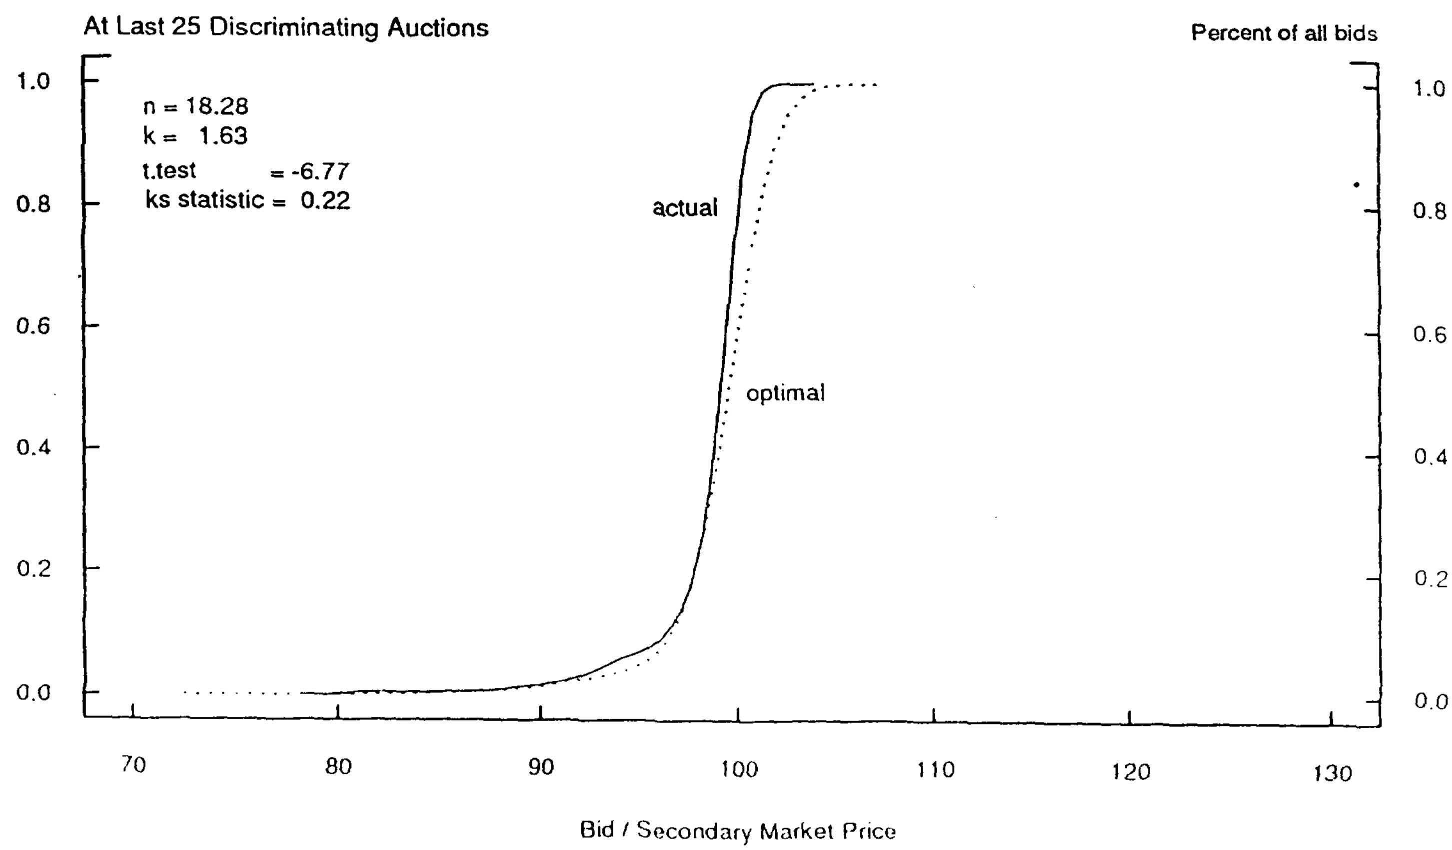

Table 3

Actual and Optimal Bidding Strategies

Auction number	Number of bidders	Number of bins	Variance of bids	Estimated k	Mean actual bids	Mean optimal bids	Test of Equality of :	
							means (t-statistic)	*distribution* (k-s statistic)
3	23	30	36.55	4.02	93.6	99.5	-5.38	0.39
4	24	21	8.19	1.90	97.6	99.6	-3.17	0.37
7	21	16	2.80	1.11	99.3	99.6	-0.74	0.12
8	18	12	1.70	0.87	99.5	99.6	-0.43	0.11
9	17	14	0.67	0.55	99.4	99.7	-1.29	0.19
13	15	8	0.27	0.35	99.7	99.7	0.14	0.02
14	17	10	0.28	0.36	99.9	99.7	1.39	0.25
15	18	10	0.44	0.44	99.8	99.7	0.50	0.06
19	19	13	0.46	0.45	99.2	99.7	-2.26	0.30
20	21	18	0.73	0.57	99.0	99.7	-3.12	0.32
21	24	8	1.33	0.76	98.6	99.7	-2.57	0.39
22	21	6	0.28	0.35	99.7	99.7	0.08	0.11
23	22	7	0.39	0.42	99.7	99.7	0.08	0.08
24	21	11	1.88	0.91	99.0	99.7	-1.46	0.20
25	20	9	0.93	0.64	100.0	99.7	0.96	0.33
26	18	7	0.24	0.33	99.8	99.7	0.57	0.11
27	14	17	3.12	1.20	97.6	99.6	-4.68	0.53
28	16	8	0.45	0.45	98.9	99.7	-3.09	0.47

Table 3
(Concluded)

Actual and Optimal Bidding Strategies

Auction number	Number of bidders	Number of bins	Variance of bids	Estimated k	Mean actual bids	Mean optimal bids	Test of Equality of :	
							means (t-statistic)	distribution (k-s statistic)
29	17	12	1.12	0.71	99.6	99.7	-0.29	0.06
30	19	13	0.80	0.60	99.8	99.7	0.41	0.21
31	18	8	0.37	0.40	99.2	99.7	-2.43	0.35
32	17	8	0.11	0.22	99.5	99.7	-1.39	0.25
33	20	9	0.19	0.29	99.6	99.7	-0.47	0.06
34	19	12	0.67	0.55	99.6	99.7	-0.09	0.09
35	20	7	0.09	0.20	99.8	99.7	0.81	0.13
36	20	9	0.53	0.49	99.2	99.7	-1.74	0.22
37	21	14	0.79	0.59	100.4	99.7	3.07	0.38
38	16	28	6.23	1.68	98.8	99.6	-1.75	0.27
39	16	30	2.31	1.02	98.1	99.6	-5.69	0.40
40	18	16	0.64	0.54	99.4	99.7	-1.30	0.18
41	10	23	40.54	4.37	97.6	99.2	-1.18	0.29
42	17	31	1.10	0.71	99.8	99.7	0.54	0.27
43	16	27	6.53	1.72	97.0	99.6	-5.22	0.38
44	16	33	10.24	2.15	94.6	99.6	-8.87	0.55
45	21	54	16.93	2.74	94.2	99.6	-9.53	0.47

References

Bartolini, Leonardo, and Carlo Cottarelli, "Treasury Bill Auctions: Issues and Uses," IMF Working Paper WP/94/135 (Washington: International Monetary Fund, November, 1994).

Feldman, Robert A., and Rajnish Mehra, "Auctions: Theory and Applications," Staff Papers, International Monetary Fund, Vol. 40 (September 1993), pp. 485-511.

Feldman, Robert A., and Vincent R. Reinhart, "Flexible Estimation of Demand Schedules and Revenue under Different Auction Formats," (1995, forthcoming)

Gordy, Michael, "Structural Analysis of Reservation Price Policy in a Treasury Bill Auction," Working Paper, MIT, January 1994.

Graham, Daniel A., and Robert C. Marshall, "Collusive Bidder Behavior at Single-Object, Second-Price and English Auctions," Journal of Political Economy, Vol. 95 (December 1987), pp. 1217-39.

Hansen, Robert G., "Empirical Testing of Auction Theory," American Economic Review, Vol. 75 (May 1985), pp. 156-59.

Joint Report on the Government Securities Market, Washington D.C.: U.S. Department of Treasury, Securities and Exchange Commission, and the Board of Governors of the Federal Reserve System, 1992.

Kagel, John H., and others, "First-Price Common Value Auctions: Bidder Behavior and the 'Winners Curse,'" Economic Inquiry, Vol. 27 (April 1989), pp. 241-58.

Kovanen, Arto, "Foreign Exchange Auctions and Fixings: A Review of Performance," IMF Working Paper WP/94/119 (Washington: International Monetary Fund, September 1994).

Laffont, Jean-Jacques, and Quang Vuoung, "Structural Econometric Analysis of Descending Auctions," European Economic Review, Vol. 37 (April 1993), pp. 329-41.

Mathieson, Donald J., and Richard D. Haas, "Establishing Monetary Control in Financial Systems with Insolvent Institutions," IMF Paper on Policy Analysis and Assessment PPAA/94/10 (Washington: International Monetary Fund, June 1994).

McAfee, R. Preston, and John McMillan, "Auctions and Bidding," Journal of Economic Literature, Vol. 25 (June 1987), pp. 699-738.

Milgrom, Paul, "Auctions and Bidders: A Primer," Journal of Economic Perspectives, Vol. 3 (Summer 1989), pp. 3-22.

Milgrom, Paul, and Robert J. Weber, "A Theory of Auctions and Competitive Bidding," Econometrica, Vol. 50 (September 1982), pp. 1089-1122.

O'Callaghan, Gary, The Structure and Operation of the World Gold Market, Occasional Paper 105 (Washington: International Monetary Fund, September 1993).

Paarsch, Harry J., "Deciding Between the Common and Private Values in Empirical Models of Auctions," Journal of Econometrics, Vol. 51 (1992), pp. 191-215.

Reinhart, Vincent, "An Analysis of Alternative Treasury Auction Techniques," Federal Reserve Bulletin, Vol. 77 (June 1992), pp. 403-14.

Rothkopf, Michael H., Thomas Teisberg, and Edward P. Kahn, "Why Are Vickrey Auctions So Rare?," Journal of Political Economy, Vol. 98 (February 1990), pp. 94-109.

Saal, Matthew I., and Lorena M. Zamalloa, "Use of Central Bank Credit Auctions in Economies in Transition," IMF Paper on Policy Analysis and Assessment PPAA/94/11 (Washington: International Monetary Fund, June 1994).

Smith, James L., "Non-Aggressive Bidding Behavior and the 'Winner's Curse,'" Economic Inquiry, Vol. 19 (July 1981), pp. 380-88.

Smith, Vernon L., "Auctions," in The New Palgrave: A Dictionary of Economics, edited by John Eatwell, Murray Milgate, and Peter Newman (New York: Macmillan Press, 1987).

Tenorio, Rafael, "Revenue Equivalence and Bidding Behavior in a Multi-Unit Auction Market: An Empirical Analysis," Review of Economics and Statistics, Vol. 75 (May 1993), pp. 302-14.

Umlauf, Steven R., "An Empirical Study of the Mexican T-Bill Auction," Journal of Financial Economics, Vol. 33 (1993), pp. 313-40.

Vickrey, William, "Counterspeculation, Auctions, and Competitive Sealed Tenders," Journal of Finance, Vol. 16 (March 1961), pp. 8-37.

Weber, Robert J., "Multiple-Object Auctions," in Auctions, Bidding, and Contracting: Uses and Theory, edited by Richard Englebrecht-Wiggans, Martin Shubik, and Robert M. Stark (New York: New York University Press, 1983), pp. 165-91.