

J.L. KELLOGG GRADUATE SCHOOL OF MANAGEMENT

Strategic Decision-Making

DECS-452
Week #8

Professor Bob Weber

1. Be prepared to discuss the RCA transponder lease case a week from today.
2. Read the attached material.
3. Continue working on the project.



J.L. KELLOGG GRADUATE SCHOOL OF MANAGEMENT
Northwestern University

Strategic Decision-Making

Professor Bob Weber

The RCA Transponder-Lease Auction

"In a hushed Sotheby Park Bernet auction hall, as tense cable television executives signaled with paddles, a sale of telecommunications satellite facilities for their booming industry brought \$90.1 million in winning bids yesterday." Thus did *The New York Times* describe the November 9, 1981 sale of seven-year leases on seven transponders, all to be carried aboard the same satellite and each capable of relaying cable-television programming from an originating ground station to the nation's network of independent cable-television distributors. The combined sale set a record for both single-item and single-auction sales at Sotheby's, and generated nearly \$20 million more than RCA Americom had expected to receive under a previously-announced price schedule.

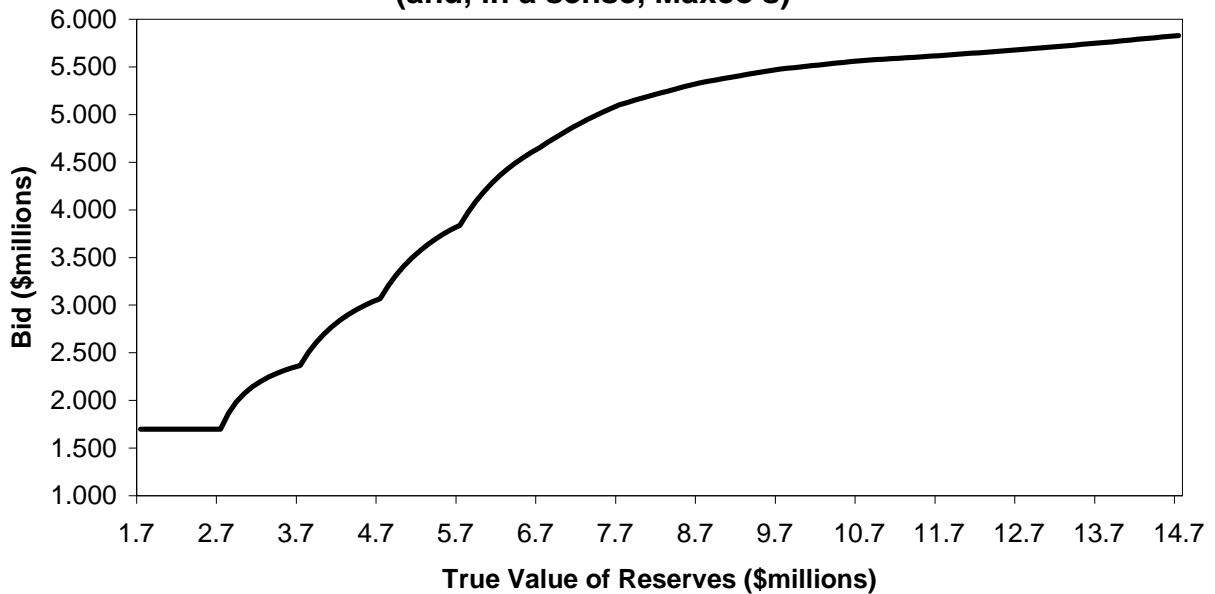
Fifty competing firms and individuals participated in the auction. The seven identical leases were sold in sequence, with an English auction conducted for each. The winners, and final bids, in the seven rounds are given below in the order of sale:

Transponder Leasing Company	\$14,400,000
Billy H. Batts (Faith Broadcasting)	\$14,100,000
Warner Amex	\$13,700,000
RCTV (Rockefeller Center)	\$13,500,000
Home Box Office	\$12,500,000
Inner City Broadcasting Corporation	\$10,700,000
UTV Cable	\$11,200,000

Contrary to theoretical predictions, this sequence of prices does *not* appear to be the realization of an upward-drifting random sequence. What happened?

- 1) Imagine yourself as (a strategically-naive) chief executive officer of a cable television firm, preparing to bid in the sale. You have decided that you are willing to pay as much as \$13.8 million (but no more) in order to obtain a lease. What instructions do you give your bidding agent, who will represent you at the auction?
- 2) Now, imagine yourself instead as the bidding agent who has just received these instructions. How can you eventually be viewed as having failed to carry them out? How might you bid, in order to avoid such failure?
- 3) If all firms (differing, of course, in their valuations of a lease) act in the manner you've just described, what will the sequence of prices look like?
- 4) If the other firms all act as in (3), how should *you* act?

Gambit's Equilibrium Bidding Strategy (and, in a sense, Maxco's)



At equilibrium:

Gambit observes the true value of the reserves, \$ x .7 million, then randomly chooses a point on the horizontal axis between \$ x .7 million and \$ $(x+1)$.7 million and makes the corresponding bid.

Maxco, not knowing the true value, picks a hypothetical true value according to the prior distribution, and then implements the same (randomized) strategy.

Gambit's expected payoff is:

Value	Expected Payoff
\$1.7	\$0.00
\$2.7	\$0.03
\$3.7	\$0.12
\$4.7	\$0.31
\$5.7	\$0.67
\$6.7	\$1.31
\$7.7	\$2.13
\$8.7	\$3.03
\$9.7	\$3.97
\$10.7	\$4.93
\$11.7	\$5.90
\$12.7	\$6.88
\$13.7	\$7.87

Maxco's expected payoff is \$0.

Strategic Decision-Making (DS-452 Course Outline)

Auctions and competitive bidding

1. History
 - a) Herodotus and Babylonian wife auctions
 - b) The Roman legions: "Caveat emptor"
 - c) British textile "dumping"
2. Justification for use
 - a) Efficiency (non-economic)
 - b) Efficiency (economic)
 - c) Basing of price on information not held by seller
3. Procedures
 - a) Sealed-bid: First- and second-price
 - b) Dynamic: Ascending (open, oral, English) and descending (Dutch clock). General equivalence of Dutch to first-price.
 - c) Multiple-object: Simultaneous (OCS leasing) vs. sequential, uniform-price (Treasury bonds) vs. discriminatory (T-bills)
 - d) Marginal variations: Entry fees and reserve prices
4. Private-value auctions
 - a) The "revenue-equivalence theorem": With risk-neutral bidders chosen at random from the population of prospective bidders, first-price and second-price auctions should generate the same expected revenue for the seller. (The expected price is the lowest price at which supply (one unit) equals demand).
 - b) If bidders are risk-averse, the first-price auction should generate higher revenues, on average.
 - c) Entry fees or reserve prices enhance revenues.
5. Common-value auctions
 - a) Maxco-Gambit: The Winner's Curse
 - 1) Maxco can't afford to bid more than \$5.83 million, so...
 - 2) Gambit needn't bid more than \$5.83 million, no matter what the value of the lease is.
 - b) General oil-lease auctions
 - 1) Large amounts of money left on the table, ...
 - a) Explained by the requirement that firms participate simultaneously in several different auctions.
 - c) The Linkage Principle
 - 1) Rough intuition: Lessening the effect of the Winner's Curse puts the bidders in more direct competition with one another.
 - 2) When bidders are risk-neutral, second-price auctions should generate higher expected revenues than first-price auctions. Progressive "English" auctions should do better yet.
 - 3) Bringing information into the public domain can increase seller's expected revenues, so ...
 - a) "Honesty is the best policy."

Selections from:

Current Issues in Economics and Finance

The Treasury Auction Process: Objectives, Structure, and Recent Adaptations

February 2005, Volume 11, Number 2

Kenneth D. Garbade and Jeffrey F. Ingber

The When-Issued Market

Immediately following the announcement of a forthcoming auction, dealers and other market participants begin to trade the new security on a when-issued basis. Secondary market transactions in outstanding Treasury securities typically settle on the business day after the trade date, when sellers deliver securities to buyers and receive payment. When-issued transactions, by contrast, settle on the issue date of the new security (which can be as much as a week or more after a trade is negotiated) because the security is not available for delivery at any earlier date.

When-issued trading enables market participants to contract for the purchase and sale of a new security before the security has been auctioned. This type of trading is important because public dissemination of the yield at which a new note is trading, or the discount rate at which a new bill is trading, provides valuable information about the market's appraisal of the prospective value of the security.

Auction Structure

An auction market participant can submit either (a) one or more competitive bids, each specifying a minimum yield at which the participant is prepared to buy a specified quantity of notes (or a minimum discount rate at which it is prepared to buy a specified quantity of bills), or (b) a single noncompetitive bid, specifying the quantity of securities that it is prepared to buy at whatever price is paid by successful competitive bidders.

The Treasury conducts note auctions in a "single-price" format. After the close of bidding, it subtracts the noncompetitive bids from the total quantity of securities offered and then accepts competitive bids, in order of increasing yield, until it has exhausted the offering. The highest accepted yield is called the "stop." Bids specifying yields below the stop are filled in full, bids above the stop are rejected, and bids at the stop are filled on a pro rata basis.

All auction awards are made at a single price, computed from the yield at which the auction stopped. The 1990s witnessed several important changes in auction structure, including a change from a multiple-price format (in which successful competitive bidders paid prices computed from their own bid yields rather than from the stop), a significant

elaboration of restrictions on auction awards, and an increase in the arithmetic precision with which bids are expressed. Each of the changes was motivated by the Treasury's interest in promoting competitive bidding and liquid secondary markets.

The best-known feature of a Treasury auction—the single-price format—was introduced in 1992. The change from the earlier multiple-price format was part of a major overhaul of the auction process that followed several violations of auction rules in 1991. The Treasury first adopted the multiple-price format when it initiated bill auctions in 1929 and it continued to use that format when it introduced auctions of coupon-bearing securities in the early 1970s. However, when the auction process came under scrutiny in 1991, public officials became interested in alternative formats that might appeal to more investors and that might lead to lower financing costs. Academics had suggested earlier that single-price auctions might reduce financing costs. In a single-price auction, a participant can bid its actual reservation yield for a new security, that is, the minimum yield at which it is willing to buy the security. The bidder certainly has no reason to bid a lower yield, but if the auction stops at a higher yield it will get the full benefit of buying at that higher yield. In contrast, the multiple-price format encourages a participant to bid higher than its reservation yield in hopes of getting the security on more favorable terms. Whether the Treasury would be better off selling securities in a single-price format or a multiple-price format was a matter that could only be resolved by empirical analysis.

In September 1992, the Treasury announced that, in an experiment, it would begin to auction two-year notes and five-year notes in the single-price format. It subsequently produced two empirical studies analyzing the results of the experiment. Although the evidence was not unambiguous, the Treasury decided in October 1998 that it justified extending the single-price format to all auction offerings.

Text of June 3, 1992 opening address at Fed/Treasury "Forum on Change"
Robert J. Weber

Thank you, Mr. Secretary. Good morning, ladies and gentlemen.

Auctions have been in use for more than two thousand years. The Babylonians arranged marriages by auction. The Roman legions sold plundered booty at auction. Today, tobacco, fish, cut flowers, works of art, thoroughbred horses, and used cars are sold at auction. Developers auction homes. Contracts are let by competitive bid. The federal government sells natural resources at auction. Pollution rights, and the radio airwaves, are soon to be up for bid. And, in the largest auctions in the history of mankind, the U.S. Treasury sells debt. It is this last auction market which we are here to discuss today.

My goal this morning is to set a general stage for the more specific discussions to follow. So let me begin by talking about the various types of auctions that are in common use. I'll begin with a discussion of single-item auctions, and later generalize what we know to the sale of multiple units, such as Treasury securities.

The oldest of auction procedures is the ascending-bid auction, in which bidders indicate a willingness to pay higher and higher prices, until only one remains. Such auctions predate the Roman Empire. Indeed, the Latin phrase *caveat emptor*, often mistranslated as "let the buyer beware," more literally means "let the highest bidder beware: All sales are final."

The ascending-bid auction solves a simple problem: It establishes a market-clearing price, in cases where the seller lacks the information needed to set that price directly. For when only one bidder remains, it can be assumed that the bidding has reached a level where only one unit of demand remains, matching the single unit of supply.

There are, of course, many prices at which the market might clear, ranging from the maximum amount that the second-highest bidder was willing to pay, on up to the maximum that the winner would have agreed to pay. The ascending-bid auction guarantees the seller no more than the lowest price in this range; the seller's inability to extract more revenue from the sale is the price he must pay for his lack of prior information about the true shape of the demand curve.

The standard sealed-bid auction, in which the submitter of the highest bid is declared the winner and pays the amount of his bid, evolved quite naturally from the ascending-bid auction. Bidders, knowing of a sale but unable to attend in person, would transmit a bid to the auctioneer in advance of the auction. The auctioneer would then enter this bid on the submitter's behalf. If no "active, present" bidder was willing to top that bid, the absentee bidder would be declared the winner, and would (quite naturally) be charged the amount of his submitted bid.

But was this really the "natural" way for the absentee bidder to be represented? The bidder would, of course, have preferred to submit instructions of the form, "Keep me in the auction at minimum price increments, until my submitted limit is passed, or until the bidding stops." But the auctioneer was typically either the seller himself, or an agent of the seller. The auctioneer's incentive was to obtain the highest possible price. And the absentee bidder, unable to monitor the bidding directly, was forced to accept the auctioneer's word about when the competition had finally dropped out. The cleanest resolution to this incentive problem was simply to agree that the submitted bid would be entered directly into the bidding at the specified level; then, at least, no charge of auctioneer dishonesty need ever be made.

This understanding forced an absentee bidder to think somewhat strategically. One profits from a transaction by buying at a price below one's actual valuation of the object being sold. Bidding in person provided a profit, as long as the bidding of others ceased somewhere below the valuation of the eventual winner: The appropriate strategy for any present bidder was to remain active until the bidding of others ceased, or his own valuation was passed. But an absentee bidder could only hope to profit by submitting a bid strictly below his true valuation.

And, once absentee bidding was allowed, the existence of even a single absentee bid raised the specter of sale taking place at a non-market-clearing price. If the item at auction was sold to an active bidder at a price only slightly above a submitted bid, there remained the chance that the absentee bidder truly valued the item at auction at more than the final price. Thus was born the notion of "loser's regret", a notion not relevant to pure ascending-bid auctions: A bidder could learn that he had lost to a bid which he would, if present, have been willing to top. Another way of looking at this is that, even if all of the bidders bid rationally, the price determined by a traditional sealed-bid auction is not necessarily a market-clearing price: There might be residual demand at that price.

Of course, "winner's regret" was also a new phenomenon, as the submitter of a winning bid always faced the possibility that a lower bid would still have won for him, and would have yielded a lower price.

Still, the submission of sealed bids became common, to the point that, for some types of sales, the ascending-bid phase of the auction was abandoned completely. The growing efficiency of communications infrastructures made it possible for the specifics of a sale to be widely distributed, and for the submission of bids via postal or electronic communications to be inexpensive and reliable.

But these sealed-bid auctions - called "first-price" auctions, since the highest of all bids established the price - were economic aberrations, arising from an inability on the part of bidders to trust that the seller would carry out more detailed instructions. The economic rationale of these procedures was incomplete, since they did not always establish market-clearing prices. Yet the discriminatory-price auctions currently used to sell Treasury securities are based directly on the general notion of charging each winning bidder the amount of his bid.

Not until the late 1950's did the proper sealed-bid analog of an ascending-bid auction come under discussion. This alternative procedure is known as a second-price auction. The submitter of the highest bid is awarded the item being sold, but is charged only a price equal to the second-highest of the submitted bids. This second-price procedure accurately emulates an ascending bid auction, in which the instructions of absentee bidders are to keep them in the auction at minimum bid increments until the bidding stops or their submission level is reached.

A naive reaction to a proposal to replace a first-price auction with a second-price auction would be to expect that the seller would reap smaller revenues. But of course bidders could be expected to bid somewhat more aggressively, if they knew that the winner's payment would be only the amount of the second-highest bid, so revenue implications actually are not immediately obvious. In order to discuss the impact of auction rules on the seller's revenues (and consequently, on the bidders' potential profits from participation, since the two are inversely related), it is necessary to specify an economic context in which the auction takes place. Two extreme contexts, at opposite ends of the economic spectrum, have received substantial attention.

One is the independent private-values context, in which each bidder knows his personal valuation of the object being sold, and these valuations differ from bidder to bidder. The independent private-values context is descriptive of the market for a pure consumption good, or of a market involving bidders who

are acting as agents for their customers, with predetermined prices waiting for them should they happen to win the auction.

Of course, this context is not completely descriptive of the sale of Treasury securities, a sale which involves bidders acting both as agents for their customers and in their own behalves, with an active secondary market following the sale. But what we know about auctions in the private-values setting can provide a benchmark for discussion of the more relevant "common-value" context which I'll talk about in a few minutes.

In the independent private-values setting, an analysis of the strategic elements of a sealed-bid second-price auction is trivial: With no further assumptions concerning differences between the bidders in valuations, or differences in their attitudes towards risk, or even concerning the level of sophistication of one's competitors, we find that one strategy is dominant: Bid precisely your own private valuation. If all of the others bid less than you, you win and turn a profit. A lower bid would not have yielded any more profit, since the price you pay if you win is unaffected by a reduction in your own bid. The reduction would only expose you to some chance of losing the auction, and making nothing. Similarly, the only change which could result from bidding more than your own valuation is that you pass a higher bid, and win. But the price you will face will be equal to the bid you passed, and hence you'll lose money. Bidding your own value is the only sensible action.

And what of the seller's revenue? If all bidders act sensibly, the seller will collect an amount equal to the second-highest of all outstanding valuations, which, as noted before in our discussion of the ascending-bid auction, is the lowest price at which the market clears.

A strategic analysis of sealed-bid first-price auctions is somewhat more complicated, and is dependent on assumptions concerning the distribution of private valuations among the bidders, and about their level of aversion to risk. It would not be useful to go into details in this forum, other than to note that, obviously, bidders should bid at a discount from their true valuations. But one striking fact results from the detailed analysis: If the private valuations are independent draws from a fixed distribution of "tastes", and the bidders are all risk-neutral over the range of relevant payoffs, then, on average, the seller's revenues will be precisely the same as from a second-price auction: The winning bid will, on average, just equal the second-highest of all valuations! This result, first noted by Professor Vickrey in the late 50's, has come to be known to economists as the "Revenue Equivalence Theorem." (Note that this is an "on-average" result. It does not assert that auction outcomes will be the same under both sets of rules. It only says that the seller, prior to the auction, has no reason to expect either type of procedure to generate higher revenues than the other.)

About 15 years ago, the Revenue Equivalence Theorem was extended to cover all other methods by which a sale might be organized. One might imagine other types of rules. The high bidder might be charged the average of all of the submitted bids. Or all bidders might simply send money to the seller, with the sender of the highest amount obtaining the object being auctioned, and the seller keeping all of the submitted amounts. (One might cynically view this as a model of "legislative lobbying".)

But, as long as the bidders know their own valuations, are risk-neutral, act rationally, and the sale eventually goes to the bidder who values the object most highly, the seller always will obtain, on average, an amount equal to the second-highest valuation.

A natural question, then, is whether the seller should care which procedure is used. And the answer is YES, a seller SHOULD care. This theoretical result should be viewed only as a benchmark, against which variations can be compared. Consider three variations in the assumptions:

First, the revenue-equivalence result assumes sensible, rational competitive behavior from the bidders. This assumption breaks down somewhat more frequently in first-price auctions, where detailed strategic considerations are quite complex, than in second-price auctions. The breakdown can yield differences in revenues.

The result also assumes risk-neutral bidders. It is not difficult to see that, when bidders are risk-averse, the first-price auction will extract higher revenues than the second-price auction. In a first-price auction, the "loser's regret" phenomenon - the possibility of losing the auction, when a somewhat higher bid might have profitably won - will lead to higher bids from risk-averse bidders, while in the second-price auction, the "bid-your-own-valuation" strategy remains dominant. So, in the presence of risk aversion, the seller benefits from the first-price procedure.

And finally, the revenue equivalence result only bounds the seller's revenues if the object being auctioned is guaranteed to be sold. But, just as in negotiations, sellers can benefit from threatening to not close the deal. For example, by setting a reserve price (that is, a bid of his own), a seller threatens to not sell if bids are too low. And it can be shown that there is always some reserve price which will, on average, increase the seller's revenues over those obtained from the use of the same auction rule, without a reserve price. (Indeed, while making more money per sale, the seller also gets to sometimes retain the item being sold!)

Still, these three variations are of little relevance to us today. The magnitude of the issues of Treasury securities makes it likely that participants will not regularly engage in unsophisticated bidding behavior. At the same time, the financial size of most participants and the frequency of sales means that risk aversion will not be a major factor in any single auction. And finally, current fiscal policy requires the complete placement of every issue, so the use of a reserve price is not considered a viable option.

But a fourth violation of our assumptions, which again leads to invalidation of the revenue-equivalence result, is real, and of extreme importance. This is the fact that Treasury securities are not privately valued - that Treasuries are durable goods, for which there exists an active secondary market. Uncertainties about future prices in that market create uncertainties for each bidder, at the time of the auction, about the true value of the securities being sold.

So let's consider an economic setting polar to the independent private-values context. Consider the common-value context, in which the item being sold has the same value to all bidders, yet no bidder knows that value for certain. Instead, each bidder holds a private estimate of what that value is. A new, striking phenomenon arises in this common-values setting - a phenomenon known as the "Winner's Curse."

The common-value setting is roughly descriptive of the federal government's leasing at auction of petroleum extraction rights on a tract of land: There's a fixed amount of oil there, but no-one knows how much, or what the extraction costs will be, or what the market price of crude will be at the time of extraction. Yet all of these economic factors are approximately equal for all bidders, each of whom holds a private estimate of the value of the lease based on their own geological estimates and economic forecasts. If the estimation processes are unbiased, then it is likely that some of the estimates are too high, and others too low. Now, imagine that your firm, facing a first-price auction, has submitted a sealed bid of \$50 million for a tract that you estimate to be worth \$60 million. And the telephone rings, and you're informed that you've won the auction. As you set down the phone, what goes through your mind? Well, you've just learned something that you didn't know before: That every competitor bid less than you. If \$50 million was truly the appropriate strategic bid for a firm holding a private estimate of \$60 million for the value of the tract, then you can now infer that the other competing firms all held estimates somewhat less than yours: The estimate of \$60 million was an extreme outlier. Assuming the other firms to be as

good as yours at making estimates, this new news is bad news, and forces you to downgrade your original estimate.

Now, the Winner's Curse should not be misinterpreted. It certainly does not imply that the winning bidder will lose money. If, in the example just given, the "bad" news that every other competitor held a lower estimate than you leads you to lower your original estimate from \$60 million to \$55 million, you still are happy to have won the auction, for you still have an expected profit of \$5 million. Exposure to the Winner's Curse simply forces bidders to scale back their bids somewhat, to protect themselves in case they do, in fact, win the auction.

And thus, in a rational market, the Winner's Curse does not truly curse the bidders: It curses the seller! For this scaling back of bids reduces the expected price obtained by the seller from the auction.

How can a seller fight the revenue-reduction caused by the bidders' reaction to the Winner's Curse? The natural approach - the only approach - is to reduce the exposure of the bidders to the Winner's Curse. For example, if the seller knew the true value of the object being sold, and could convincingly reveal that information to all of the bidders, one could expect competition to push the price up to that true value, maximizing the seller's revenues in the process of clearing the market at the unique market-clearing price.

Of course, in the sale of mineral extraction rights, or for that matter in the sale of Treasury securities, the seller - the government - does not have perfect information. Still, some information is held by the government at the time of sale. And it has been shown (in theory), and should be expected to be true in practice, that a policy of full and accurate public revelation of all knowledge held by the seller is the policy which maximizes revenues.

Note that I used the word "policy" in the last statement. Of course, suppressing bad news, or even distorting it into good news, can be advantageous to a seller in the short run. A major art auction house, holding incontrovertible evidence that a particular work is a forgery, could tell bidders that it believed the work to be genuine. It would make more money from the sale of that work. But its reputation - if not for honesty, then at least for the ability to make accurate appraisals - would be damaged, leading to lower revenues from future auctions. We have here one of the few instances in which economic theory justifies a well-known adage, namely, that "Honesty is the best policy."

Honest announcement of appraisals is one approach that a seller can take in reducing the exposure of individual bidders to the Winner's Curse. Another approach is to let the market generate relevant public information on its own, prior to the moment of sale. Again referring to the leasing of mineral rights, note that the holders of rights on tracts of land already leased must file extraction reports with the government, and some of the information contained in these reports becomes a matter of public record. This information is of value to all of the bidders involved in subsequent auctions of leases on nearby tracts.

The when-issued market plays much the same role prior to the sale at auction of Treasury securities. It provides a publicly-available composite view of market participants concerning the true eventual value of the securities to be auctioned. In doing so, it reduces the exposure of individual bidders to the Winner's Curse, and consequently elicits somewhat higher bids, on average, at auction than would be obtained if when-issued trading were not allowed.

There is yet another approach a seller can take, in fighting the revenue-reducing effects of the Winner's Curse. That is to use an auction procedure which lets the price paid by a winning bidder depend on information other than just his own. In a second-price auction, the winning bidder expects to pay a relatively high price only if some other bidder holds an estimate nearly as high as his own. If all other

bidders hold significantly smaller estimates, and make significantly smaller bids than he, the price he pays will be low. This correlation between price and the estimate of another bidder lessens the effect of the Winner's Curse on each individual bidder, and therefore benefits the seller: In a common-value setting, the Revenue-Equivalence result breaks down, and sellers should prefer to use second-price auctions.

An ascending-bid auction cashes in on both of the previously-discussed methods of reducing the exposure of the bidders to the Winner's Curse. Information about the estimates of others, inferred from the level of competition as the price climbs, has the same effect as information revealed by the seller. Then the price finally paid by the highest bidder is close to the point at which the second-highest bidder chooses to drop out of the competition, yielding an additional benefit similar to that gained from a second-price auction. The combination of these two effects leads to the natural conclusion: The seller's revenue from the use of an ascending-bid auction will, on average, be greater than from the use of a second-price auction, which in turn will be greater than from the use of a first-price auction.

What of strategic issues? In a second-price or ascending-bid auction, strategic issues remain relatively simple even in a common-values setting. Take all of the publicly-available information into account, and then ask yourself: "If I knew that my estimate was really the highest, and that the second-highest estimate was just marginally below mine, what would I then revise my estimate to be?" Bid that amount! In game-theoretic language, the market will resolve itself "in equilibrium" if all bidders follow this strategy.

Just as in the independent private-values setting, second-price and ascending-bid auctions have the "no-regret" feature, that even when the auction is over, no bidder will find himself wishing that he'd bid differently. The assumption on which the winner's bid is based, that the highest rejected bid is as high as his own, is overly optimistic. The assumption on which the price-determining, next-to-highest bid is based - that the winning bid is no higher - is overly pessimistic. Therefore, the winning bidder is always happy to have won at the price he must pay, while the losing bidder would not, even after the fact, have wanted to have entered a bid higher than the winner's. The auction will always establish a market-clearing price, even in the common-value setting.

In summary, this is the state of economic theory as it currently stands, with respect to seller's revenues: If the aversion of individual bidders to risk is a small factor, relative to the existence of common uncertainty about the value of the item being sold, and if the market is reasonably symmetric, in the sense that no bidders hold substantial informational advantages over the others, then sellers benefit from using second-price auctions in place of first-price auctions, and benefit even more from the use of ascending-bid auctions.

How do these results carry over to the sale of multiple units, such as the auctioning of Treasury securities?

The discriminatory pricing procedure currently in use is, of course, a direct extension of the first-price sealed-bid auction, in which each winning bidder pays the amount of his bid.

The second-price auction naturally generalizes to a uniform-price auction, where all bidders pay an equal price corresponding to the highest rejected bid. (Note that this is not quite the same as the proposed uniform-price auction, in which the lowest accepted bid, rather than the highest rejected, determines the uniform price. But with the large number of units available at Treasury auctions, it doesn't hurt to view the two pricing rules as equivalent.)

And finally, the ascending-bid auction corresponds closely to the procedure to be laid out in more detail by Dr. Reinhart in this afternoon's session.

The revenue effects found in the single-item setting all carry over directly to the multiple-unit setting: Ascending-bid and uniform-price auctions should generate higher revenues - which means, for the Treasury, lower financing costs - than the currently-used discriminatory-price auctions.

Next, let's consider the incentives of individual bidders to obtain informational advantages over their competitors. How might a change in procedures affect current participants, and the structure of the pre-auction and post-auction markets?

Consider the value of private information. For first-price and discriminatory-price auctions, the situation is simple: Without private information, or an advantage in terms of attitude towards risk, one can't expect to make any money. In fact, in the presence of several bidders, each strictly better-informed than you, stay out of the auction. The Winner's Curse lashes out at uninformed bidders with a vengeance.

This provides substantial incentive for some bidders to stake out positions providing themselves with extreme informational advantages. One shouldn't be surprised to find a structure very similar to that in the current market for U.S. Treasury securities. A modest number of firms (such as a number of the primary dealers) make a substantial investment in information-gathering technology. Once they achieve an informational advantage, it makes little sense for less-well-informed firms to attempt to participate directly in the auctions. And at the same time, the marginal returns to becoming as well-informed as those who came before become smaller and smaller, until finally no new firms can justify putting information-collection departments of their own in place. Instead, they either participate only indirectly, as customers of the primary dealers, or they do most of their trading on the secondary market. Competition in the auction thins out, which benefits the remaining bidders at the expense of the seller. Thinner competition in the auction market also provides additional opportunities to those still in to take manipulative actions.

In this regard, note that many institutions, finding themselves at an informational disadvantage, currently enter non-competitive bids for the maximum allowed quantity. Indeed, the Treasury has at times been forced to actions specifically designed to stop employees of firms from bidding noncompetitively on their employer's behalf, and to stop bank holding companies from having every branch submit a separate, maximum-quantity noncompetitive bid.

In a uniform-price auction, private information is still of value. But the value is lessened. Private information allows a firm to correlate the times it wins with the times that favorable prices result, but most of the correlation between the amount bid and the price paid disappears. Informationally-disadvantaged firms find it much more attractive to participate, because the price they will pay, if they win, is a price which incorporates much of the information held by their better-informed counterparts.

So, uniform-price auctions have two beneficial effects. They encourage more competition, which benefits the seller directly. At the same time, they lessen the incentives for individual firms to seek marked informational advantages on a continuing basis.

Of course, part of the impetus for today's forum is the fallout from Salomon's misadventures last year. How might a change in the rules lessen the chances of similar problems in the future?

Well, how can one strategically generate extra profits in the secondary market? As you are all aware, one of the simplest ways to do this is by taking control of an issue, and generating a short squeeze. Indeed, Salomon took precisely this approach last year. They were punished. But what really was their mistake? A cynical view is that they were simply wrong in trying to do it all by themselves. For there was nothing that they tried that could not have been done without falsified bids by three bidders working in concert.

Should the Treasury be concerned with short squeezes? A naive, short-term view is that an attempted squeeze generates extra revenues for the Treasury, since the attempt requires outbidding most of the other auction participants. But of course the concern is that exposure to squeezes will increase risk in the market for Treasury securities, and in the longer term will raise financing costs.

Here, note that the use of uniform-price auctions will make attempted squeezes more costly. Second-price auctions yield greater bid dispersion (and price variance) than do first-price auctions. That greater dispersion in bids carries over to uniform-price auctions, making it more costly to outbid other competitors. Indeed, this dispersion will be increased even more by the fact that there will be no practical limit on noncompetitive bids. For, in a uniform-price auction, a noncompetitive bid is equivalent to a bid at an extremely high price, and I can't conceive of any way to formally outlaw such bids. Indeed, I suspect that, under a uniform-pricing arrangement, the submission of multi-price bid schedules will become more common, and the dispersion in individual schedules will increase, as bidders take advantage of the opportunity to guarantee that a part of their demand will be filled, by bidding very aggressively for that part.

It is worth noting here that the current system, as well as any revisions that are implemented, should be supplemented with guidelines concerning the appropriateness of different types of pre-auction communication amongst intended bidders. Discussion of the likely stop-price must be distinguished and separated from discussion of precise bidding intentions. Issues concerning strategic manipulability of various auction procedures will be discussed in further detail by Professor Marshall this afternoon.

Much of what I said earlier was targeted at a single important point: The proposed changes in auction procedures will not increase the cost of financing the national debt. But how much lower can financing costs go? What additional revenues might the Treasury hope to capture by a change in rules?

Professor Sundarasan will discuss relationships between Treasury auctions and the when-issued and futures markets in some detail this afternoon. But I want to make a few broad comments up front.

Studies based on data from the 70's and 80's have suggested that auctioned securities sold on average at substantially lower prices than did the same securities on the secondary market. Estimates of losses to the Treasury of 4 or more basis points are the norm emerging from those studies. On the other hand, a recent Fed analysis of data from just the last two years suggests that the gap has narrowed substantially, to what appears to be less than a single basis point.

While a fraction of a basis point is still substantial when viewed across more than a trillion dollars in annual sales, the reduction in revenue loss at auction suggests that the market has evolved substantially over the past twenty years. Consider that evolution, and how it relates to the choice of a sales mechanism.

I've already noted that it is in the seller's best interest to have as much public information as possible available to the bidders prior to the time of an auction. Back when sales were less systematic, and volumes substantially lower, it was probably appropriate to provide an incentive for some individual competitors to gather information. The use of a discriminatory-pricing procedure provided that incentive, by offering extra profits to those who made the market.

But today, the market is much more complete. An active when-issued market, a complementary futures market, and greater predictability concerning the Treasury's long-term borrowing needs all help to lessen the need to encourage substantial investment in information-gathering from any individual bidder. With that encouragement no longer needed, it becomes appropriate to switch to a procedure which lessens the rewards which accrue to the informationally-advantaged.

As the Appendix to the January joint report indicates, several other nations already place a portion of their debt through the use of uniform-price auctions. In preparing for today's forum, I was surprised to hear, albeit very informally, that some of those nations are thinking of switching to discriminatory-price auctions. But on further reflection, I was able to reconcile that with the current push to move the U.S. towards the use of uniform-price auctions. Our secondary and ancillary markets are more mature - more complete - than those foreign markets. And therefore it might well be that, while they perceive a need to offer extraordinary compensation to central market participants, we have finally passed the point where we should do so.

Much of my discussion this morning has focused on uniform-price procedures, rather than on the further step to ascending-bid auctions. One reason for my emphasis is that the precise proposals concerning that second step will not be on the table until later today.

But another reason is that, from my personal point of view, the most important step is the first one, to uniform-price procedures. I can see no revenue-related issue, no market-related issue, no strategic issue which fails to favor this first step, and I consider it possible that the first step will rationalize the current market and address its problems sufficiently well to make the further step to ascending-bid auctions unnecessary.

Certainly, I can see no reason to wait until electronic bid submission is in place before beginning the uniform-price experiments. I am confident that the experiments will be deemed a success - after all, it appears that the series of experiments in the early 1970's were already successful, and were ended for political, rather than economic reasons. Indeed, I am so confident in the eventual success of the experiments that I hope they will be structured in a manner designed not only to collect data, but also to begin the evolutionary process of fully replacing discriminatory-price auctions. I would personally support a phased schedule which moves several issues over to uniform pricing on a long-term commitment basis, at several different points along the yield curve, e.g. the 26-week bill, and the two- and seven-year notes. And I would like to see a schedule which holds open the possibility of adding further issues to the uniform-price list as favorable evidence comes in, rather than seeing an experiment which moves individual issues back and forth between methods. Professor Chari will discuss the issue of the scheduling of experiments in more detail this afternoon.

My segment of this morning's program is nearly over. So let me give one final summary of what economic theory has to say: In settings with a fixed pool of bidders, uniform-price and ascending-bid auctions yield greater revenue for the seller than do discriminatory-price auctions. With respect to Treasury securities in particular, this translates to a lesser burden on the taxpayer to finance the national debt. Beyond this, uniform-price and ascending-bid auctions should generate a larger pool of active bidders by eliminating both winner's and loser's regret, and by reducing the risk of participation for bidders at an informational disadvantage. This should consequently reduce the overall cost to the economy of information-collection activities, by reducing the competitive advantages gained by a bidder from becoming appreciably better-informed than his competitors. And finally, by eliciting more competition, increasing the dispersion of bids, and permitting individual bidders to guarantee that a portion of their demand will be filled, uniform-price and ascending-bid auctions will make market manipulation much more difficult.

First-price auctions, and their discriminatory-price analogues, are economic aberrations which have their roots in a distrust in sellers. Second-price and uniform-price auctions are already beginning to supplant them in many circles. The Italian government has placed construction contracts using second-price auctions. There are individuals in today's audience who are actively involved in arranging second-price


auctions for major corporate clients. The Argentine government placed better than a billion shares in the national telephone system into private hands this past winter using a uniform-price auction.

The time is right to move the largest auction market of all onto a more rational economic foundation.

In 1959, Milton Friedman stood before a Congressional committee and recommended that Treasury securities be sold through the use of uniform-price auctions. He continued to press his recommendation on a number of subsequent occasions. Perhaps, as I mentioned earlier, the market had not yet evolved to the point where the transition was justified. Perhaps he was somewhat ahead of his time. But the right time has now arrived.


A couple of weeks ago, I spoke with Professor Friedman concerning this meeting. He wished us all luck, but also predicted that, once again, the political process would ultimately subvert any attempt to bring about the recommended changes. I'd like to see him proven wrong, and I guarantee you all that he, as well, would be happy to be wrong this time. A third of a century is long enough to wait. The market has matured to the point where the recommended changes are practical. The climate is right to restore confidence in the non-manipulability of an extraordinarily important market. The time to act is now!

Robert J. Weber
J.L. Kellogg Graduate School of Management



A Forum on Potential Changes to the Treasury's Auction Technique

Sponsored by the Board of Governors of the Federal Reserve System and the U.S. Department of the Treasury
June 3, 1992



Board of Governors of the Federal Reserve System
Martin Building, Terrace Level, Dining Room E

20th and C Streets, N.W.
Washington, D.C. 20551

Morning Session	Afternoon Session
<p>8:30 A.M. Coffee</p> <p>9:00 A.M. Introductory Remarks Secretary Nicholas F. Brady <i>U.S. Department of the Treasury</i></p> <p>9:10 A.M. Opening comments by moderator Vice Chairman David W. Mullins, Jr. <i>Board of Governors of the Federal Reserve System</i></p> <p>9:20 A.M. Survey of the auction literature Robert Weber <i>Kellogg School, Northwestern University</i></p> <p>10:10 A.M. Background on alternative auction techniques William Vickrey <i>Columbia University</i></p> <p>10:25 A.M. Break</p> <p>10:40 A.M. The view from the Treasury Deputy Assistant Secretary Deborah Danker <i>U.S. Department of the Treasury</i></p> <p>11:10 A.M. Comments on Presentations Robert Glauber <i>Kennedy School, Harvard University</i> Stephen Thieke <i>J.P. Morgan Securities</i> Richard Roll <i>Anderson Business School, U.C.L.A.</i></p> <p>11:40 A.M. Roundtable discussion</p> <p>12:40 P.M. Lunch</p>	<p>2:00 P.M. Opening comments by moderator Under Secretary Jerome H. Powell <i>U.S. Department of the Treasury</i></p> <p>2:10 P.M. Issues to be resolved in a single-price open-outcry auction system Vincent Reinhart <i>Board of Governors of the Federal Reserve System</i></p> <p>2:40 P.M. Comments on Presentation</p> <p style="padding-left: 20px;">How to set the initial price Steven Wunsch <i>AZX Securities</i> Suresh Sundaresan <i>Columbia University</i></p> <p style="padding-left: 20px;">The logistics of iterated bidding Kenneth deRegt <i>Morgan Stanley</i> Robert Marshall <i>Duke University</i></p> <p style="padding-left: 20px;">The design of the experiment Jefferson DeAngelis <i>Northwestern Mutual Life Insurance Company</i> V.V. Chari <i>FRB Minneapolis</i></p> <p>3:40 P.M. Roundtable discussion</p> <p>4:40 P.M. Closing remarks Joseph Stiglitz <i>Stanford University</i></p>