

How the U.S. Treasury Should Auction Its Debt

V. V. Chari
Adviser
Research Department
Federal Reserve Bank of Minneapolis
and Harold H. Hines Professor of Risk Management
J. L. Kellogg Graduate School
of Management
Northwestern University

Robert J. Weber
Professor of Decision Sciences
J. L. Kellogg Graduate School
of Management
Northwestern University

The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

Auctions have been around for more than 2,000 years. The Babylonians arranged marriages by auction. The Roman legions sold booty at auction, and on one notable occasion, the Praetorian Guard killed the emperor and put up the whole empire for auction. Today, members of the general public sell at auction such diverse things as tobacco, fish, cut flowers, works of art, thoroughbred horses, and used cars. The U.S. government sells natural resources by auction and may soon take bids on radio airwaves and pollution rights. And in the largest auctions in recorded human history, the U.S. Treasury each year sells roughly \$2.5 trillion worth of debt. With such large amounts at stake, even small improvements in the Treasury's auction procedure can lead to large gains for taxpayers. In this paper, we review what economic theory tells us about ways to improve this procedure.

The Treasury's current procedure is what is known as a *multiple-price, sealed-bid* auction. Roughly a week before each of its more than 150 annual auctions, the Treasury announces the amount of debt it plans to sell. Eligible dealers and brokers submit competitive sealed bids which specify the price they are willing to pay for a particular quantity of debt. Investors may also submit so-called noncompetitive bids up to a fairly low quantity ceiling without specifying a price if they are willing to accept whatever will turn out to be the average accepted-bid price. Once all bids are in, the Treasury first adds up the quantity of noncompetitive bids and subtracts that from the total debt it plans to sell. Then, starting at the highest price bid and moving down, the Treasury adds up the competitive quantities bid until it hits its total. Each competitive bidder who has won (or, in the Treasury's jargon, has been "awarded" the bid) pays the price stated in his or her sealed bid; thus, each winning bidder may pay a different price. Noncompetitive bidders, again, pay the average of the awarded competitive bids.

This multiple-price, sealed-bid procedure, of course, is not the only way to design an auction. Indeed, most economists agree that it is not the best one for the Treasury. We argue here, based on economic theory, that the Treasury should switch to a *uniform-price, sealed-bid* auction. Under this procedure, with bids ordered by price, from the highest to the lowest, the Treasury would still accept quantities up to the amount it planned to sell, but the price winning bidders paid wouldn't vary. Instead, all bidders would pay the same price, that of the highest bid not accepted—the price that just clears the market.

The main reason to make this change is that the current auction procedure provides incentives for bidders to acquire more information than is socially desirable. In the current procedure, again, bidders pay the amount of their bids if they win. Therefore, bidders have an incentive to shade their bids below the maximum amount they are willing to pay in order to try to obtain the securities at a lower price. But bid-shading carries with it the risk that the bid is so low that the bidder is not awarded any securities. In selecting a bid price, therefore, bidders want to balance the gain from a lower winning bid against the risk of not winning. Thus, they have incentives to learn what others plan to bid.

In a uniform-price auction, by contrast, the price paid by a winning bidder does not depend on that bidder's bid. Therefore, bid-shading is less extreme than in multiple-price auctions, and the incentives to acquire information about what others plan to bid is smaller. Information about

how bidders plan to bid is of no value to society as a whole since such information merely ends up redistributing payments from uninformed to informed bidders. But acquiring this information is costly. The loser from the resources expended in information acquisition is the Treasury (and, of course, ultimately, the taxpayers). A uniform-price, sealed-bid auction will therefore yield more revenue to the Treasury.

Uniform-price auctions are also likely to be less susceptible to market manipulation. In 1990, Salomon Brothers Inc. violated Treasury rules designed to protect against market manipulation, and in 1991 the market was allegedly manipulated twice more. We argue that episodes of this kind are less likely under a uniform-price auction.

The Treasury did, in fact, experiment with such an auction briefly in the 1970s, but abandoned the experiment as largely inconclusive. As will become obvious, we think the experiment was abandoned too hastily. In any event, more recently, the Treasury embarked on a review of its auction procedures in collaboration with the Securities and Exchange Commission and the Board of Governors of the Federal Reserve System. (See U.S. Department of the Treasury et al. 1992.) Therefore, a review of what economic theory tells us about Treasury auctions seems particularly desirable. (For recent reviews of auctions in general, see McAfee and McMillan 1987, Mester 1988, and Milgrom 1989.)

The plan for our review is as follows. We begin by laying out a general framework for analyzing bidder behavior in auctions. We apply this framework to two models in which only one unit of an object is being sold at auction: a simple model called the *independent private-values* model and an extension called the *correlated-values* model. We then discuss more complicated auctions like the Treasury's, auctions in which more than one unit is sold. We argue that the incentives to acquire information are smaller with uniform-price auctions than with multiple-price auctions and that uniform-price auctions are less susceptible to manipulation.

The General Framework: Game Theory

Let's start by describing how economists generally think about bidder behavior under any type of auction procedure. The framework we use is *game theory*. This is a way to analyze how rational decisions are made by competitors in uncertain conditions. In auctions, of course, the competitors are primarily the bidders.

A seller faced with the problem of choosing among auction procedures must predict how bidders will act. Each bidder in a given auction, in turn, must predict how other bidders will act. These actions depend both on how much the bidder values the object being sold and on guesses about how others will bid on it. Each bidder's valuation of the object depends on his or her information about the object. For example, a bidder on an oil tract may know something about oil or neighboring tracts. Or in an art auction, bidders may know how valuable a painting will be to them. Successful bidding at an auction, therefore, involves successful guesses about other bidders' information and successful guesses about how these others will guess about each other's information.

This is an apparently intractable problem, but the language of noncooperative game theory offers a neat way around it. The way is to shift attention from bids to *bidding strategies*. Formally, a (pure) *strategy* for a bidder is

a description of the relationship between what is known—the information of the bidder and the history of the auction—and what should occur—for each bidder’s information and each stage of the auction, the appropriate decision for the bidder to make. Of course, in practice, bidders simply choose their bids rather than their strategies. A strategy for a particular bidder is simply a way of describing how other bidders imagine the particular bidder will act under various circumstances. A *Nash equilibrium* is a collection of strategies, one for each bidder, such that given the strategies of the other bidders, no one prefers to change his or her own strategy.

Single-Object Auctions

From this perspective, the nature of the information possessed by each bidder is critical in determining the outcome of a given auction procedure. The seller’s problem is simply to compare equilibrium outcomes across auction procedures and pick the one that does best for him or her. (Of course, the chosen auction procedure may alter bidders’ incentives to acquire information. We return to this theme later.) Here we consider two models of the information possessed by bidders: the independent private-values model and the correlated-values model. In both, we assume only one object is being sold.

The Independent Private-Values Model

Suppose a painting is being auctioned. Each buyer knows how valuable this painting is to him- or herself but is uncertain about its value to other bidders. The seller is also uncertain about its value to the bidders. No bidder plans to resell the painting. This assumption of no resale means that each bidder cares about the value of the painting to others only insofar as it affects how others will bid. Put differently, even if bidders knew each other’s values, no bidder would change the maximum amount he or she would be willing to pay. In this model, that is, bidders have *independent private values*.

Assume the model has N bidders. Let v_i denote the value of the painting to bidder i . That is, v_i is the maximum amount bidder i is willing to pay for the painting. We model the uncertainty about other bidders’ values by assuming that bidder i ’s value is a random variable drawn from a distribution $F_i(\cdot)$ on $[0, \bar{V}]$. We assume that bidders are risk-neutral, so that a bidder who pays m and receives the painting has a payoff of $(v_i - m)$. The seller is also risk-neutral. Bidders and the seller care about expected payoffs.

Each auction procedure can be described as a set of rules for bidders, describing at each stage what bidders can do as a function of the history of the auction. Given these rules, a *strategy* for a bidder prescribes what the bidder should do at each stage of the auction as a function of the history up to that stage and as a function of that bidder’s private valuation v_i . A collection of such strategies for each of the bidders together with the rules of the auction procedure determines the outcome of the auction. This outcome should be thought of as who gets the object and how much each bidder pays. This outcome determines the *payoffs* of each bidder. Again, a *Nash equilibrium* is a collection of such prescriptions or strategies, such that given the strategies of the other bidders, none strictly prefers to change his or her strategy.

Now, this may sound numbingly complex, and it is. Fortunately, an insight due to Myerson (1979) and Harris

and Townsend (1981) allows us to simplify the problem considerably.

Consider replacing a complicated auction procedure by the following mechanism. All bidders, privately and confidentially, report their valuation, v_i , to an impartial computer. The computer is programmed with the equilibrium strategies of the complicated auction and uses them by, in effect, running through the entire auction, doing what the bidders would have done, and producing an outcome. This outcome, of course, depends on the valuations of all the bidders. With this computerized mechanism, the decision problem of an individual is simply what value to enter into the computer. This mechanism is called a *revelation mechanism* since each individual reveals his or her private information to the computer.

The remarkable result, due to Myerson (1979) and Harris and Townsend (1981), is that the equilibrium outcome of any auction procedure can be reproduced as a truth-telling equilibrium of the revelation mechanism. The reason is simple. Since the original strategies constituted an equilibrium and the computer is going to play those strategies anyway, no bidder could do better by reporting a different value than the true one; all that a different report would do is make the computer choose a different—and, hence, less desirable—course of action.

From the perspective of bidder i , the revelation mechanism induces three outcome functions. Each of these is a function of the report, \hat{v}_i , of bidder i . These functions are the probability of winning the object, $p_i(\cdot)$; the expected payment conditional on winning the object, $w_i(\cdot)$; and the expected payment conditional on losing the object, $l_i(\cdot)$.

The payoff to a bidder who reports a value \hat{v}_i and whose true valuation is v_i is, then, given by

$$(1) \quad \pi_i(v_i, \hat{v}_i) = p_i(\hat{v}_i)[v_i - w_i(\hat{v}_i)] - [1 - p_i(\hat{v}_i)]l_i(\hat{v}_i).$$

At a truth-telling equilibrium, we have, for all i ,

$$(2) \quad \pi_i(v_i, v_i) \geq \pi_i(v_i, \hat{v}_i).$$

Alternatively, if the expected payoff is differentiable in \hat{v}_i , we have

$$(3) \quad \partial \pi_i(v_i, v_i) / \partial \hat{v}_i = 0.$$

Condition (2) or (3) can be used with some additional assumptions to establish a remarkable result known as the *revenue equivalence theorem*. This theorem requires that we specialize the model further. Assume that the bidders’ valuations are symmetric; that is, the distribution functions are the same for all bidders. Denote this common distribution function by F . We will say that an auction procedure is *efficient* if it allocates the object to the bidder with the highest value. Assume also that the lowest valuation bidders receive zero expected payoff. The theorem asserts that all auction procedures with these properties have the same expected payoff to the seller. Formally, we have this:

PROPOSITION (Revenue Equivalence Theorem). *Every efficient auction with symmetric, risk-neutral, independent private-values bidders which assigns zero expected payoffs to bidders with the lowest values yields the same expected revenues to the seller.*

Proof. Note that, with symmetric bidders, the expected payoff functions of all bidders are the same. Denote these common expected payoffs at the truth-telling equilibrium by $\pi^*(v_i)$. Using the envelope theorem in (3), we have

$$(4) \quad \pi^*(v_i) = p_i(v_i).$$

Integrating (4) and using the hypothesis that $\pi^*(0) = 0$, we have

$$(5) \quad \pi^*(v_i) = \int_0^{v_i} p_i(x) dx.$$

For an efficient auction, $p_i(x)$ is simply the probability that the highest bidder's valuation is x . Thus, the expected payoffs of bidders are entirely determined from the distribution function $F(\cdot)$. Therefore, all auction procedures satisfying the hypotheses of the proposition yield the same expected payoffs to bidders. Since all auctions generate the same total surplus, the expected payoff of the seller is the same. Q.E.D.

To understand the relevance of this result, consider some examples of specific types of auctions when only one object is being sold. In *sealed-bid* auctions, each bidder silently submits a bid. In a *first-price, sealed-bid* auction, each bidder submits a bid and the object is awarded to the highest bidder at that bidder's price. In a *second-price, sealed-bid* auction, each bidder submits a bid and the object is awarded to the highest bidder at the price bid by the next-highest bidder. In *open-outcry* auctions, an auctioneer calls out prices to all bidders. In a *descending-price, open-outcry* auction (also called a *Dutch* auction since it was used in Holland to sell tulips), the auctioneer starts the price high and lowers it until some bidder claims the object. In an *ascending-price, open-outcry* auction (also called an *English* auction), the auctioneer starts the price low and raises it, stopping when only one bidder remains. With symmetry, all of these auctions are efficient and give zero payoffs to the lowest bidder. Thus, they all yield the same expected revenue.

Now, on the surface, these auctions seem quite different. In a second-price auction, for example, the best a bidder can do is submit his or her true valuation. Obviously, bidding higher than the true valuation would mean running the risk of paying more than the object is worth. Might a bidder want to bid less than the value of the object to that bidder? No, because all that such a strategy would do is reduce the chances of winning. It would have no effect on the price paid if the bidder wins. Thus, bidding one's valuation regardless of the actions of others is a *dominant strategy*. The seller's revenues are given by the value of the object to the second-highest bidder.

How could such an auction yield the same revenue as a first-price, sealed-bid auction? The result, due to Vickrey (1961), comes from the following reasoning. In a first-price auction, bidders shade their bids below their valuations. By doing so, they risk losing the object but pay less when they win. In equilibrium, each bidder's strategy is an increasing function of value. Thus, the object is assigned to the bidder with the highest valuation. The revenue equivalence theorem tells us that the bid-shading in the first-price auction results in exactly the same revenue as the second-price auction yielded. One can apply similar reasoning to other auctions [as Milgrom and Weber (1982) have done].

The Correlated-Values Model

The revenue equivalence theorem tells us that which type of auction the seller chooses doesn't matter much. But, of course, the theorem follows from assumptions which may or may not be relevant in actual applications. For the example of an auction for Treasury debt, one assumption is very questionable. Treasury securities are easy to resell in the active secondary market. Thus, the value of a particular Treasury security to a bidder depends on how much others are willing to pay for it. Thus, the assumption of independent private values seems very unlikely to hold here. We turn, therefore, to a model with correlated values.

Consider, again, the example of the painting being sold, but now assume that the painting can be bought and then sold to others. In this situation, a bidder's willingness to pay is affected both by the bidder's own valuation and by what the painting would fetch if it were resold in the secondary market. The price in the resale market, in turn, depends on the willingness to pay of others. That is, in determining a bid, each bidder must take into account all the bidders' values: bidder values are *correlated*. In the example of the painting, suppose that the winning bidder plans to keep the painting for some time and then to sell it. A bidder's willingness to pay now depends both on how much the painting is worth to that bidder while the bidder owns it and on how much it will fetch when sold. We will use the term *value estimate* to describe a bidder's maximum willingness to pay given that bidder's information.

In this situation, a phenomenon known as the *winner's curse* can emerge. Consider, for example, a first-price, sealed-bid auction for our well-worn painting. Imagine that you have submitted a bid and have just been called and told "Congratulations, you have won." Along with the thrill of winning comes a frightening thought. By winning, you have found out that your bid was higher than anybody else's; thus, others probably value the painting less than you do. Therefore, if you wanted to resell it, you would probably lose money. As a winner, you are cursed.

On more careful inspection, though, this phenomenon does not imply that winners should vow never again to attend an auction. Instead, it implies that bidders will optimally shade their bids, recognizing that, if they win, their bid was the highest.

The revenue equivalence theorem now does not necessarily hold. Milgrom and Weber (1982) have shown that the expected revenues of someone selling a single object in four different types of auctions can be ranked this way, from highest to lowest revenues:

1. The ascending-price, open-outcry (English) auction.
2. The second-price, sealed-bid auction.
3. Tied: The first-price, sealed-bid auction and the descending-price, open-outcry (Dutch) auction.

Rather than repeat Milgrom and Weber's (1982) formal results here, we provide some intuition. Recall that when values are correlated, the winner's curse causes bid-shading. A first-price auction awards the object to the highest bidder at the bid price. If other bidders value the object much less than the highest bidder does, then the object is worth much less than the bid price if the winner wants to resell it. Thus, all bidders fearing this kind of event end up shading their bids well below their own estimates. In contrast, in a second-price auction, the winner pays the price

bid by the next-highest bidder. Thus, bidders are induced to raise their bids above their first-price auction bids by the knowledge that they will not lose if other bidders estimate the value of the object to be very low. In fact, the equilibrium bidding strategies in a second-price auction turn out to be fairly simple. Each bidder tries to answer the following question: "If I knew that my estimate was the highest and that the second-highest estimate was just marginally below mine, what would I then revise my estimate to be?" The equilibrium strategy is to bid the revised estimate. So, while the equilibrium strategy is not quite as simple as it was in the independent private-values environment, it is still relatively easy.

In an ascending-price, open-outcry (English) auction, revenues are even higher than in the second-price auction. The reason is that, as this auction proceeds, it reveals information about the value estimates of other bidders. As the auctioneer raises prices, some bidders drop out. Other bidders gain information about the value estimates of the dropouts and thus are able to revise their own estimates. The availability of this information reduces the winner's curse and causes bidders to bid more aggressively. That raises the seller's revenues.

One simple way of thinking about this English auction (as in Milgrom and Weber 1982) is to consider a situation with only two bidders. Each bidder's strategy is described by a single number which specifies at what price that bidder will drop out. The price paid by the winning bidder is marginally higher than the dropout price of the losing bidder. From a strategic point of view, of course, this is the same as a second-price auction. Thus, the seller's revenues are the same in both types of auctions.

Now consider an auction which initially had N bidders but $N - 2$ have dropped out. The remaining two bidders know the prices at which the others have dropped out and have revised their value estimates accordingly. These two bidders now engage in a second-price auction with appropriately revised estimates. Why does this revelation of information during the auction cause bidders to bid more aggressively during the auction?

One way to think about the reasoning is as follows. Recall that in a second-price auction, the equilibrium strategy is to assume that one's own value estimate is the highest and that the next-highest bid is just slightly lower. Assume that three bidders are engaged in a second-price auction. Label these bidders 1, 2, and 3. Consider informing bidder 1, just before that bidder is to submit a bid, that his or her value estimate is the highest. Would this change bidder 1's strategy? The answer is that it would not since the bidder has effectively already assumed that his or her estimate is the highest. Suppose now that you informed bidder 1 of the opposite: that his or her estimate is the lowest. Clearly, this information would cause the bidder to bid more aggressively. In effect, the English auction reveals this information to bidder 1 as it proceeds. Thus, it leads to higher revenues for the seller.

One implication of the theory is that auctions should be conducted as ascending-price, open-outcry auctions. Indeed, most auctions are of this type. To maximize the seller's revenue, most others should use the second-price, sealed-bid procedure. Yet auctions often use the first-price, sealed-bid procedure. Why are some auctions of this apparently inferior type?

In some situations, sealed-bid auctions are simply more practical than open-outcry auctions. Obviously, open-out-

cry auctions require bidders or their trusted agents to be present during the auction, and that is not always possible. Now, in theory, bidders could effectively duplicate the ascending-price, open-outcry auctions by submitting written or electronic bid schedules, telling the auctioneer how they would bid as a function of the prices at which other bidders drop out. But such bid schedules would have to be so incredibly long and complicated that they are just not feasible.

The next-best procedure should be the second-price, sealed-bid auction. But this has a serious problem too—one it shares, in fact, with the open-outcry auction. Both of these procedures require that the auctioneer be completely trustworthy, a somewhat unrealistic condition in the private sector. Consider what happens, for example, in a second-price, sealed-bid auction if the auctioneer is not trustworthy. Once the bids are opened, the auctioneer has a great incentive to cheat: to insert bids just below the winning bid in order to extract higher revenue. Bidders, of course, recognize that fact before they bid and respond by treating second-price auctions as first-price auctions. In the private sector, therefore, the lower-revenue first-price auctions are common.

When a government agency is the auctioneer, however, cheating seems much less likely. Thus, unlike private sellers, the U.S. Treasury could use the higher-revenue second-price auction procedure.

Complications

While the theory developed thus far has dealt with single-object auctions, the results generalize relatively straightforwardly to more complicated situations, like auctions with more than one object for sale. The results do not generalize quite so straightforwardly if bidders are risk-averse or if they collude. However, we will argue that risk-neutrality and competitive behavior are reasonable assumptions for Treasury auctions.

Consider first a situation with N bidders, each of whom wishes to buy one unit, and where $M < N$ units are offered for sale. The multiple-price analog of a first-price auction is a *discriminatory* auction, where the M highest bidders are awarded the items at their bid prices. The analog of a second-price auction is a *uniform-price* auction, where each bidder pays the price bid by the highest rejected bidder. The theory can be extended to cover these situations, and the results are the same: the uniform-price (second-price) auction dominates the discriminatory (first-price) auction.

Matters are more complicated when bidders have demand schedules expressing the number of units they are willing to buy at various prices. While the theory has not been completely developed for that situation, the economic logic of the arguments for the single-object environment seem likely to carry over.

Thus far, we have assumed that bidders are risk-neutral. Now let's see what happens if they're risk-averse.

If bidders are risk-averse in the independent private-values context, then the seller's expected revenues are higher in a first-price (discriminatory) auction than in a second-price (uniform-price) auction. The reason is that submitting one's true valuation remains a dominant strategy in the second- (or uniform-) price auction. Risk-aversion implies a willingness to pay an actuarially unfair premium to avoid large losses. Thus, in a first-price (discriminatory) auction, risk-averse bidders are willing to pay

more than risk-neutral bidders to avoid the large loss from failing to win the object. (See Matthews 1983 for an analysis of auctions with risk-averse bidders.)

If value estimates are correlated, however, the comparison for seller revenues across auction types becomes ambiguous. The theory, therefore, does not have much to say about the consequences for the Treasury if bidders are risk-averse. If risk-aversion is a major concern, the Treasury should not switch to a uniform-price auction. But we think risk-aversion should not be a major concern: no single Treasury auction is large relative to the wealth of actual and potential market participants, and it is not clear whether the Treasury does, or should have, attitudes toward risk that are substantially different from those of the participants. We thus think risk-aversion issues may be reasonably ignored for Treasury auctions.

The theory is also ambiguous if bidder valuations are drawn from different distributions. (See Milgrom 1989 for a nice example.) As a practical matter, market participants acquire information about eventual market prices in roughly the same way. Therefore, this issue too may be safely ignored.

What about collusive behavior among bidders? Let's answer that first for single-object auctions. In the independent private-values context, second-price auctions are more susceptible to collusive behavior than first-price auctions are. To see this, suppose a second-price auction has only two bidders, and they agree to tell each other their valuations and to adopt a strategy where the one with the lower value bids zero and the one with the higher value bids that value. For a promised side-payment, the lower-value bidder agrees to this arrangement and has an incentive to abide by it. Consider now what happens with a first-price auction. The only way the higher-value bidder can gain over the outcome without collusion is to bid less than the lower-value bidder's valuation. (Recall the revenue equivalence theorem.) But now the lower-value bidder has an incentive to defect.

In situations with multiple objects, recall, the analog of a second-price auction is a uniform-price auction and the analog of a first-price auction is a discriminatory auction. Thus, considerations of collusive behavior seem to suggest that discriminatory auctions should be favored for Treasury debt. However, two considerations militate against accepting this conclusion too quickly. First, the Treasury's current system has 39 primary dealers. Setting up, and enforcing, collusive arrangements among this large a group would be a formidable task. Second, as we shall see in the next section, uniform-price auctions stimulate entry into bidding, which is anticollusive.

Benefits of Uniform-Price Auctions

If the Treasury switched its auctions to the more-feasible of the two auction types that yield the highest revenue—the uniform-price auction—then the general public welfare would be improved in at least two ways.

Less Information Acquisition

The current system for auctioning Treasury debt creates large incentives to acquire information about other bidders' actions as well as the eventual state of market demand. These incentives would be much smaller under the uniform-price auction system.

To see this, consider a situation where one of the bidders—say, bidder 1—incurs a cost and acquires the esti-

mated values of all other bidders. Acquiring these estimates will cause bidder 1's estimate to be revised; but if the original estimate was unbiased, then the expected value of the revised estimate will be the same as the original estimate. For now, therefore, assume that acquiring this information causes no change in the estimated value for bidder 1. Assume for now also that other bidders do not change their strategies.

We want to focus on how bidder 1's bidding strategy changes after acquiring the information. Recall that bids are increasing in the estimated values. When bidder 1 does not have the information of other bidders' estimates, the bidder wins whenever his or her estimate is the highest.

Consider, first, the situation when the bidder's estimate, before the information was acquired, was the highest. In a second- or uniform-price auction, such a situation would not change the amount paid by the bidder (since we have assumed no change in the behavior of other bidders). In contrast, in a first-price or discriminatory auction, bidder 1 now shades his or her bid further down to just above the bid of the next-highest bidder. Thus, whenever the bidder would have won, a first-price or discriminatory auction yields a gain to information and a second- or uniform-price auction does not.

In situations where the bidder would have lost, matters are more complicated. Once bidder 1 acquires the information, that bidder is willing to pay any amount up to the estimate. Two possibilities must be considered. Either some bidders' bids are more than bidder 1's estimate, or some bidders' bids are less than bidder 1's estimate but more than bidder 1's bid without the information.

If some bidders are willing to pay more than bidder 1's estimate, the bidder drops out of the auction. Now, recall that bid-shading is more extreme with first-price (or discriminatory) auctions than with second- (or uniform-) price auctions. In both types of auctions, bidder 1 would have lost without the information and is happy to do so with the information. Since bid-shading is more extreme with first-price (discriminatory) auctions, the probability that some bidders will bid higher than bidder 1's estimate is smaller in those auctions than in second- (or uniform-) price auctions. The potential gains to changing the strategy are therefore higher.

Next, consider the case when some bidders' bids are between bidder 1's estimate and bidder 1's bid without the information. In this situation, the theory is ambiguous about which auction provides greater incentives to acquire information. However, given the gains in the other two situations, the overall effect is likely to enhance the incentives to acquire information.

Of course, if other bidders recognize that bidder 1 has acquired information, they will modify their strategies as well. One way of modeling the change in other bidders' behavior is to assume that, by incurring a cost, bidder 1 is informed, with some small probability, of the valuations of other bidders. The other bidders do not know whether or not bidder 1 has acquired this information. If the probability of acquiring the information is sufficiently small, then the change in the bidding strategies of the other bidders will be small, and the analysis above applies. If this probability is 1, it can be shown that the expected payoff to the less-informed bidders is zero under both types of auctions. Thus, in this case, the incentives to acquire information are the same under both types. Therefore, we ar-

gue that the incentives to acquire information are generally higher with first-price (or discriminatory) auctions.

From this result comes the conclusion that first-price (discriminatory) auctions yield lower revenues to the Treasury and lead to larger amounts of resources devoted to gathering information than do second-price (uniform-price) auctions. Is this information-gathering a socially valuable activity? To the extent that it involves gathering information about how much other bidders are willing to pay, it merely redistributes payments from uninformed to informed bidders. This information has no value to society as a whole. Even worse, the existence of informed bidders drives relatively uninformed bidders away from the auction. Thus, auction procedures which provide large incentives to acquire information lead to fewer active uninformed bidders. This reduction in the number of bidders tends to reduce revenues to the Treasury.

We want to emphasize here that the true social cost of the current auction procedures is the excessive resources devoted to gathering information about potential bidders. Channeling these resources to other activities is likely to enhance welfare.

Less Market Manipulation

A switch to a uniform-price auction procedure would improve welfare in at least one other way. The Treasury's recent review of its auction procedures was spurred, in part, from violations of Treasury rules by Salomon Brothers in 1990 and from two instances of so-called short squeezes in 1991. These sorts of attempts to manipulate the market for Treasury securities should be less likely under a uniform-price auction procedure.

Let's briefly review the current structure of the market for Treasury securities. Approximately a week to 10 days before a Treasury auction, dealers and investors actively participate in a *when-issued market*. This is a market in forward contracts. Participants agree to deliver and accept delivery of specified quantities of a Treasury security when it is issued at a currently agreed-on price. Those who agree to deliver soon after the security is issued are known as the *shorts*; those who agree to wait for delivery, the *longs*. The market performs a *price-discovery* role; that is, it provides information to bidders about the likely state of market demand for the Treasury security when it is issued. This information benefits auction bidders who face uncertainty about the prices at which they will be able to resell Treasury securities.

To see the possibilities for market manipulation, consider the following scenario. A trader or group of traders commits to a forward contract for a large amount of a Treasury security, on the long side of the contract; they agree to accept delivery of the security when it is issued. The same person or group then purchases a large amount of the Treasury security at the auction. Now, those who have committed to deliver the security (the *shorts*) must acquire the security in the marketplace. But they find that most of the securities are held by those on the long side of the forward contract. Since the forward contract specifies delivery of that particular security, the shorts are squeezed. (See Sundaresan 1992 for a proposal to replace the when-issued market by a cash-settled futures market.) This possibility tends to reduce the volume of trade in the when-issued market and thus raise the costs of price discovery. The risks imposed on bidders are then passed on to the Treasury as lower revenues.

The Treasury recognizes this problem and is also sensitive to general concerns that particular traders may seek to corner a market. It therefore imposes limits on the amount that bidders can bid at the auction. These are the rules that traders at Salomon Brothers tried to circumvent in 1990 by submitting fraudulent bids in customers' names.

What are the likely consequences of our proposed reform on the when-issued market and on the prospect for short squeezes? First, a switch to a uniform-price auction procedure would reduce the role of the when-issued market. With uniform pricing, bidders would have less of an incentive to acquire information about other bidders' willingness to pay. We have already argued that this reduction is socially desirable. Second, to the extent that short traders fear the prospect of a squeeze, a uniform-price auction procedure would let these traders purchase the security at the auction more cheaply than they can under the current system. The reason is that, under the current system, a short trader must submit a bid at a high price and be willing to pay that price to guarantee not being squeezed. Under a uniform-price auction, short traders are unlikely to substantially affect the price they pay for the security by submitting a high price. Therefore, they can protect themselves better, and the prospects of market manipulation are reduced.

Concluding Remarks

A switch to either an ascending-price, open-outcry auction or a uniform-price auction for U.S. Treasury debt is likely to raise Treasury revenues and reduce excessive resources devoted to information-gathering. The ascending-price auction has the disadvantage of requiring physical presence at the auction. (Of course, this type of auction could be conducted with remote electronic terminals.) To the extent that such presence is costly, it raises entry barriers to the auction and is wasteful. Furthermore, it is more burdensome to bidders since the strategic calculations involved are more complicated. The uniform-price auction is strategically much less complicated. This feature also tends to reduce entry barriers.

Ultimately, the issue is relatively simple. The current organization of the Treasury market has primary dealers who purchase at the auction and resell to the public at large. With so many close substitutes and an efficient Treasury market, no reforms of the auction procedure will change prices to the ultimate holders very much, if at all. Entry is possible into the dealer/broker arena, and the market is competitive enough that, as a first approximation, such dealers make no more than the normal return on their investments. The only questions that remain are whether those investments are affected by the Treasury's auction procedure and whether they are at the socially optimal level.

The investments of Treasury dealer/brokers are in the form of a network of people who have learned to work with each other, ultimate buyers, and the Treasury. An important part of their activity is to acquire information about the behavior of actual and potential bidders. The when-issued market serves this role. We have argued that, given the auction procedure, this information is privately valuable and that market participants will rationally invest to acquire it. We have also argued that this information has dubious social value. With our proposed change in the auction procedure, the incentives to acquire this information would be lower, and over time, these investments

would not be replaced. Thus, over time, the returns going to these investments would accrue to the Treasury.

These arguments suggest that changes in the auction procedure will take time to yield gains in Treasury revenue. No experiment conducted over any period as short as even a year is likely to generate significant changes in Treasury revenue. Patience appears to be a must.

Another implication of our arguments is that if the reforms are implemented, the when-issued market will likely shrink. This market currently serves a variety of purposes, one of which is price-discovery: the market lets participants learn about each other's willingness to pay for Treasury securities. This role is extremely important for participants who want to reduce the risks of the Treasury auction. But the risks are largely due to the current form of the auction. The talents and resources now involved in the when-issued market are a rational response to the current auction procedure. With a uniform-price auction, some of these talents and resources could go to more socially productive activities.

We are certainly not the first to advocate a change in the Treasury auction procedure, or even this specific change. Milton Friedman advocated a uniform-price auction in testimony to the Joint Economic Committee in 1959 (excerpted in Friedman 1991). More recently, other economists have advocated this proposal—for example, Merton Miller in Henriques 1991. We think a large majority of economists support the proposal (which, admittedly, may be a popular argument against it). In this paper, we have tried to argue that everything we know from economic theory tends to support it.

References

- Friedman, Milton. 1991. How to sell government securities. *Wall Street Journal* (August 28): A8.
- Harris, Milton, and Townsend, Robert M. 1981. Resource allocation under asymmetric information. *Econometrica* 49 (January): 33–64.
- Henriques, Diana B. 1991. Treasury's troubled auctions. *New York Times* (September 15): 13.
- Matthews, Steven A. 1983. Selling to risk averse buyers with unobservable tastes. *Journal of Economic Theory* 30 (August): 370–400.
- McAfee, R. Preston, and McMillan, John. 1987. Auctions and bidding. *Journal of Economic Literature* 25 (June): 699–738.
- Mester, Loretta J. 1988. Going, going, gone: Setting prices with auctions. *Federal Reserve Bank of Philadelphia Business Review* (March/April): 3–13.
- Milgrom, Paul R. 1989. Auctions and bidding: A primer. *Journal of Economic Perspectives* 3 (Summer): 3–22.
- Milgrom, Paul R., and Weber, Robert J. 1982. A theory of auctions and competitive bidding. *Econometrica* 50 (September): 1089–1122.
- Myerson, Robert B. 1979. Incentive compatibility and the bargaining problem. *Econometrica* 47 (January): 61–73.
- Sundaresan, Suresh. 1992. Pre-auction markets and post-auction efficiency: The case for cash settled futures on on-the-run Treasuries. Manuscript. Columbia University.
- U. S. Department of the Treasury, Securities and Exchange Commission, and Board of Governors of the Federal Reserve System. 1992. *Joint report on the government securities market*. Washington, D.C.: U.S. Government Printing Office.
- Vickrey, William. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16 (March): 8–37.