NBER WORKING PAPER SERIES

AUCTION THEORY: A SUMMARY WITH APPLICATIONS TO TREASURY MARKETS

Sanjiv Ranjan Das
Rangarajan K. Sundaram

Working Paper 5873

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 1997

Forthcoming in *Financial Markets, Institutions and Instruments*. The second author would like to thank the National Science Foundation for support under grant no. 94-10485. This paper is part of NBER's research program in Asset Pricing. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

Auction Theory: A Summary with Applications to Treasury Markets
Sanjiv Ranjan Das and Rangarajan K. Sundaram
NBER Working Paper No. 5873
January 1997
JEL Nos. C72, G14
Asset Pricing

ABSTRACT

This review paper describes basic auction concepts, and provides a summary of the theory in this area, particularly as it relates to Treasury auctions.

Sanjiv Ranjan Das
Department of Finance
Harvard University
Soldiers Field
Boston, MA 02163
and NBER
sdas@hbs.edu

Rangarajan K. Sundaram
Department of Finance
New York University
New York, NY 10012
rsundara@stern.nyu.edu

1 Introduction

Auctions are among the oldest market mechanisms for price discovery, dating back at least two millenia. The list of commodities sold by auctions today is long; it includes, among other things, artwork, agricultural produce, antiques, and mineral rights. Our focus in this paper is on one of the largest and most important auction markets in the world, that for US Treasury bills (T–bills).

US Treasury Auctions

For several decades now, the US Treasury has used auctions as a means of testing the pulse of the short term interest rate market. A large number of auctions are conducted every year. Three– and six–month T–bills are auctioned every Monday, while one-year T–bills are auctioned every four weeks. Longer-term securities such as T–notes and T–bonds are also auctioned: for instance, two–year and five–year T–notes are auctioned every month, while other maturities (such as three–year, ten–year, and thirty–year securities) are auctioned quarterly.[1] The total annual volume of Treasury securities auctioned has increased steadily from a level of $670 billion in 1981 to $1.70 trillion in 1991, and reached $2.0 trillion in 1995.[2]

From an analytical standpoint, Treasury auctions share a number of features in common with other auction markets. The asymmetry of information between buyer and seller is, for instance, central; without this, the auctioneer could simply sell to the buyer with the highest valuation. However, Treasury auctions are also characterized by a number of institutional features that distinguish them from other auction markets, and that may render many of the standard analyses inapplicable.

First, Treasury auctions are preceded by *forward* trading among potential buyers in the security to be auctioned. Trading in this forward market commences on the date the auction is announced, with settlement taking place on the date the securities are issued. The presence of this forward market (known as the "when-issued" market) implies that traders could enter the auction with prior short or long positions. In turn, this could affect their bidding strategies and, thereby, the outcome of the auction.

Second, Treasury auctions are followed by trading in an active resale (or "secondary") market. As Bikhchandani and Huang [2] argue, this creates an important information linkage that affects bidding behavior and prices in the auction. How does this linkage work? The majority of bidders in the auction are primary dealers and large financial institutions whose information about interest rates is typically better than that available to investors in the secondary market. However, the public information available to buyers in the resale market includes information from the Treasury on the bids submitted in the auction. To the extent that these bids are relective of the bidders' private information, the resale price will be affected by the bids. This now creates an incentive for participants in the auction to signal their private information to secondary market participants through their bidding strategies.

[1] Although our discussion in this paper is couched in terms of T–bills, it also applies equally to these longer-term instruments.

[2] Treasury bills are also auctioned in other economies, but in some cases they are issued at a fixed discount. Without exception, however, T-bills are issued for 3 and 6 month maturities. (To be precise, in the US dollar markets they are issued for 91 days and 182 days respectively.)

Third, academic theory has mostly focussed on single-unit auctions, or on multiple-unit auctions in which each bidder is entitled to at most one unit of the commodity being auctioned. As pointed out by Wilson [33] and Back and Zender [1], it is not at all clear to what extent the intuition gained in this case can be be profitably transferred to the study of divisible-good auctions, such as the treasury auction.

Fourth, T–bill auction markets are characterized by the presence of two different classes of bidders, called "noncompetitive" and "competitive" bidders, respectively. Noncompetitive bidders, who are usually individual investors, submit sealed bids that specify only the quantity they desire. Competitive bidders, who are designated primary dealers and large financial institutions, submit sealed bids that are price–quantity pairs. Since multiple bids are permitted, each competitive bidder effectively submits a demand schedule. Noncompetitive bidders are guaranteed the quantity they bid for. The total quantity submitted by the noncompetitive bidders is subtracted from the amount offered by the treasury, and the remainder is distributed to the competitive bidders on the basis of their bids. The price charged the noncompetitive bidders is the quantity-weighted average of the price paid by the competitive bidders.

The Main Questions

The behavior of bidders in an auction is obviously affected by the rules governing the auction (the auction "form" or "mechanism"). A principal aim of auction theory is to identify the auction mechanism that maximizes the expected reveue of the seller. This problem is usually stated in a more constrained form. As with optimal contracts in principal-agent models, optimal auction forms sometimes have very complex and unintuitive forms. Therefore, the aim is often only to identify the best of a given set of auction forms.

In the case of Treasury auctions, two auction forms are usually considered: *discriminatory auctions* (DA's) and *uniform-price auctions* (UPA's). In a DA, winning bids are filled at the bid price; that is, the demands of the bidders are met by starting with the highest-price bidder down, until the entire quantity is exahausted. In a uniform-price auction, winning bidders pay a flat price, called the ***stop-out price*** for each unit they receive. The stop-out price is simply the lowest winning price, i.e., the maximum price at which the aggregate demand[3] equals or exceeds the available supply of T–bills. Since the 1970's, the Treasury has relied almost entirely on the DA format; however, begining in 1992, it also began experimenting with the UPA format in the sale of two-year and five-year notes.

Since bidders will alter their behavior depending on whether they are facing a DA or a UPA, it is not immediately apparent which auction form will generate higher expected revenue for the Treasury. Going back at least to Friedman [6], one strand of academic opinion has argued strongly in favor of the UPA format. The early (informal) arguments offered in support of this position were that discriminatory auctions discourage relatively uninformed bidders because the "winner's curse" becomes more severe.[4] Bidding therefore becomes concentrated among a few large bidders; in turn, this makes collusion easier and more profitable.

[3] Recall that each of the competetive bidders submit a demand schedule. The aggregate demand schedule is simply the sum of these schedules plus the total quantity demanded by the noncompetitive bidders.

[4] See section 2 for a description of the winner's curse.

In the early 1980's, game-theoretic models of auctions provided formal support to this position. Milgrom and Weber [17] showed that in the auction of an indivisible good, the expected revenues of the seller were higher in a UPA than in a DA. The Milgrom-Weber model has subsequently been extended in many directions but the basic result appears quite robust. For example, Bikhchandani and Huang [2] show that the same result obtains even if a secondary market, with a non-trivial information linkage to the primary market, is appended to the model.

Recently, however, Back and Zender [1] have shown that the assumption of indivisibility of the good being auctioned is critical to the Milgrom–Weber result. Back and Zender analyze a very general divisible good auction model and demonstrate that uniform-price auctions can have very undesirable equilibria; and that as a consequence, the UPA is no longer unambiguously superior (theoretically speaking) to a DA.[5] Building on this work, Wang and Zender [32] show that a ranking of the two auction forms in divisible-good auctions is, in general, impossible: they consider a class of divisible-good auctions, and prove that in this class, there always exist equilibria of the DA that dominate some equilibria of the UPA, and vice versa. Empirical studies of the debate have also been ambiguous in their conclusions; some studies of Treasury (and foreign exchange) auctions have indicated that the UPA may be superior to the DA, but this is contradicted by other studies.

In summary, the one clear conclusion to have come out of the recent theoretical studies is that no useful lessons on Treasury auction format can be gained from the study of auctions of indivisible goods. Much work, however, remains to be done. The identification of the "better" auction format is obviously harder than was originally realized, but may not be as hopeless as the Back–Zender and Wang–Zender results indicate. For one thing, not all of the equilibria in these papers are equally plausible. Thus, one can ask if there exists a strategically salient class of equilibria that could be used as the basis for comparison of the two auction forms. Second, the auction market is only one in a chain of linked markets; as we have explained above, it is preceded by the when-issued market, and followed by the secondary market. From a theoretical standpoint, the importance of the when-issued market and its effect on bidding behavior in the auction market is little understood.[6] Finally, the presence of noncompetitive bidders in the auction process creates "noise" that affects the actions taken by the competitive bidders; it is yet to be analyzed if this noise is a good thing, from the Treasury's viewpoint.

The Structure of this Paper

The remainder of this paper is organized as follows. In section 2, we describe the auction framework involving a single unit of an indivisible good, and provide a detailed summary of the basic results in this area. There are at least two good reasons for begining with this model. First, many of the modelling issues and analytical techniques involved in Treasury auctions also arise in indivisible-good auctions; however, the latter framework is a simpler one and, therefore, offers a better setting for understanding these. Second, the results on the ranking of different auction forms are unambiguous in the case of indivisible-good

[5] That divisible and indivisible good auctions are very different from the seller's viewpoint was pointed out as early as 1979 by Wilson [33]. However, Wilson was not concerned with Treasury auctions, and did not analyze the UPA vs. DA issue in a divisible-good auction framework.

[6] However, see Vishwanathan and Wang [30] for a recent contribution in this direction.

auctions. To fully appreciate the difference between divisible and indivisible good auctions, it is necessary to understand why this is the case.

In section 3, we examine the extension of the basic model undertaken by Bikhchandani and Huang [2]. The most important generalization Bikhchandani and Huang consider is the addition of a secondary market. Their main result (from our perspective) is that this leaves unaltered the ranking between UPA's and DA's.

Section 4 turns to an examination of divisible-good auctions. The material presented here focusses on two papers: Wilson [33], who first pointed out the difference between divisible and indivisible good auctions, and Back and Zender [1].

Section 5 presents a brief review of the empirical work that has been done on Treasury auctions. Section 6 concludes.

2 A Brief Review of Auction Theory

This section provides a brief review of the main concepts and results of auction theory. Our exposition here draws on the work of McAfee and McMillan [14], and Milgrom and Weber [17]. Although T–bill auctions are more complex in some important ways than the standard models we examine here, there are sufficiently many similarities in the modelling and analytical issues involved that this constitutes an excellent starting point for further study.

Two assumptions will be maintained throughout this section. First, we will treat the auction as a mechanism for *selling.* Second, we will assume that the object being auctioned is a single unit of an indivisible good. The first assumption is made solely for expositional convenience; in practice, auctions are also widely used as mechanisms for *buying*, and with minor (mostly notational) changes, the results reported here apply equally to this case also. The second assumption is, on the other hand, a much more serious one. Its principal virtue is that it simplifies analysis considerably. However, it also means that the seller cannot auction "shares" in the good; in particular, it leaves open the question of how much of the intuition gained here can be legitimately transferred to the study of T–bill auctions. We will return to this issue again in a later section.

Auction theory is concerned primarily with answering two questions. First, why is an auction used instead of some other selling or buying procedure? Second, there are a great many varieties of auctions that one sees in practice. What determines which auction form is chosen in a given situation? It is these questions that we shall be concerned with in the sequel. Our analysis proceeds in several steps.

We begin in subsection 2.1 with a discussion of the key feature of auctions—the ability of the seller to pre-commit to a set of rules and thereby to extract maximal gains from trade—that make them attractive mechanisms from the seller's viewpoint. Following this, in subsection 2.2, we provide a categorization of the most commonly used auction forms.

The remaining subsections are devoted to the second of the two questions raised above: what determines the specific auction form used in a given circumstance? Of course, the answer to this question will depend on the maintained background assumptions concerning the environment in which buyers and seller interact. In subsection 2.3, we list the alternatives facing the modeller in this situation. Subsection 2.4 sets the ball rolling by considering a specific set of assumptions that we call the benchmark model. The subsections that follow then examine the impact of weakening or otherwise altering these assumptions.

2.1 Pre-Commitment and Informational Asymmetry

Classical economic theory regards as indeterminate outcomes in markets in which there is monopoly on one side of the market (a single seller) and oligopoly on the other (a small set of potential buyers). Auction theory, which is concerned with the same setting, resolves the indeterminacy by viewing the seller as essentially a Stackelberg leader who has the ability to *pre-commit* to a set of policies (i.e., the choice of auction form, rules of the auction, etc.). Such pre-commitment is a key aspect of auction markets. Credible pre-commitment gives the seller a first-mover advantage; in particular, it enables the adoption of procedures that induce the bidders to bid in desirable ways.

The ability to pre-commit does not, however, imply that the seller extracts all possible gains from trade. The bargaining ability of the seller is limited by a second key aspect of auction markets, viz., the fact that the seller does not know the buyers' true valuations of the object being auctioned. If such *informational asymmetries* were not present, the seller could simply offer it to the bidder with the highest valuation at a price just below that bidder's valuation, and threaten to refuse to sell if this offer is rejected. Pre-commitment would make this threat credible, so the seller would realize virtually all of the gains from trade. When information is asymmetric, however, the seller's ability to exploit competition amongst the buyers is more limited. Certainly, the seller cannot always drive the price up to the highest valuation, because he is not aware of this valuation. The question of which auction form is "optimal" from the seller's viewpoint is, therefore, non-trivial, and the central one that auction theory attempts to answer.

2.2 Categories of Auctions

There are four basic types of auctions:

1. English Auctions: In an English auction, the price is raised successively until only one bidder remains. Also sometimes called the oral or ascending-bid auction, the English auction is perhaps the most commonly used auction form in practice.

2. Dutch Auctions: Dutch auctions are the converse of English auctions. In a Dutch auction, the auctioneer calls out an initial high price and then lowers the bid successively until some bidder accepts the current bid. Dutch auctions are used in various markets around the world, but are not as popular as English auctions.

3. First-Price Sealed Bid: Potential buyers in a first-price sealed bid auction submit sealed bids. The item being auctioned is awarded to the buyer who submits the highest bid at the price bid by him. First-price sealed bid auctions are commonly used by governments in awarding procurement contracts.

4. Second-Price Sealed Bid: Once again, buyers submit sealed bids, and the item is awarded to the buyer who submits the highest bid. However, the price the winner pays is equal not to his own bid, but to the second-highest bid. Second-price sealed bid auctions have useful theoretical properties, but are not widely-used in practice.

A number of variations on these basic forms are commonly employed. For example, the seller may impose a "reserve price" and discard all bids below this price. The seller may also charge the bidders an entry fee for the right to participate in the auction. The price to be paid by the winning bidder may depend not only on the bids received in the auction, but also on something correlated with the true value of the item (such as royalties).

2.3 Modelling Issues

Whether a particular auction form emerges as optimal will depend at least in part on how uncertainty is modelled, and there are at least four separate decisions that the modeller is faced with in this context.

First, there is the question of the various players' attitudes to risk. Should the buyers be modelled as risk-neutral or risk-averse? What about the seller?

The second decision concerns the relationship between different buyers' valuations of the object. Is it reasonable to model them as being independent, or as being correlated in some degree? Two polar extremes can be identified:

- *The Independent Private-Values Model:* In this model, it is assumed that bidder i's valuation v_i is a draw from a probability distribution F_i. Only bidder i observes the value of v_i, but the distribution F_i is itself common knowledge among all the players. Any one bidder's valuation is statistically independent of any other bidder's valuation. This model is a good approximation of situations in which bidders are buying for their own use and not for resale (for example, antiques or artwork).

- *The Common-Value Model:* In the common-value model, it is assumed that the item being bid for has a single objective, but unknown, value V. The bidders' valuations v_i are independent draws from a probability distribution $G(\cdot|V)$, which is presumed to be common knowledge. The common-value model may be applied to markets in which the commodity is being purchased primarily for resale (for example, the auction of mineral rights or Treasury bills).

Of course, many real-world auctions may consist of aspects of both models simultaneously. To handle such possibilities, Milgrom and Weber (1982) develop a general model of which the independent private-values model and the common-value model are both special cases. Their model, which uses the notion of *affiliated distributions*, is described in a later subsection.

Thirdly, there is the question of whether bidders may be modelled as essentially identical (up to informational differences), or whether one should admit the possibility of of different "types" of bidders. The first case, the case of *symmetric bidders*, is significantly easier to handle analytically. The second, that of *asymmetric bidders*, allows for a richer scenario such as the existence of systematic cost differences between bidders.

Finally, there is the issue of whether payment by the winning bidder to the seller should be made to depend on variables other than the bids themselves. In some situations where the only observable variables may be the bids themselves, this is not a relevant issue. In others, it is an important consideration. In auctions of publishing rights, for example, the final payment typically depends on the winning bid, as well as royalties based on the actual sales of the book.

2.4 The Benchmark Model

Our analysis of the optimality of different auction forms begins in this subsection with what is perhaps the simplest framework for the analysis of this question. It is based on the following four assumptions:

1. All the bidders are risk-neutral, as is the seller.
2. The independent private-values model holds: Bidder i's valuation v_i is a draw from a distribution F_i, and the draws are statistically independent.
3. The bidders are symmetric, so $F_i = F_j$ for all i, j. Denote the common distribution by F.
4. The final payment from the winning bidder to the seller depends on the bids alone.

Following McAfee and McMillan (1987), we will refer to this model as the "benchmark model." Subsequent to the analysis of the benchmark model, we will examine the impact of weakening or otherwise replacing each of the four assumptions in turn.

Throughout this paper, we use the concept of a *Bayes–Nash equilibrium* to analyze any given auction model. That is, it is presumed that all of the following are common knowledge among the seller and the bidders: the rules of the auction chosen by the seller, the number of bidders n, the probability distribution governing the valuation v_i of each player i, and the attitude to risk of each player. In addition, each player also knows his own true valuation. Based on his information, each player i chooses a "strategy" σ_i, i.e., a rule that decides the amount $\sigma_i(v_i)$ that should be bid by player i as a function of his valuation v_i. An equilibrium is a vector of strategies $(\sigma_1, \ldots, \sigma_n)$ such that for all i, given that players $j \neq i$ have adopted the strategy σ_j, player i can do no better than to adopt σ_i.

One final observation is useful in simplifying the material that follows. Under all circumstances (i.e., regardless of assumptions concerning risk-aversion or the relationship between different players' valuations, etc.) the outcome under a Dutch auction must necessarily be identical to that under a first-price sealed bid auction. This follows simply because bidders in either of these auction forms must choose how high to bid without knowing the others' decisions; in either case, moreover, the winning bidder pays an amount equal to his bid. In the formal analysis that follows, therefore, we ignore the Dutch auction altogether.

Returning to the main question: Which auction form—English, Dutch, first-price, or second-price—should a seller use if the conditions of the benchmark model are met? The surprising answer is: It does not matter. Under some weak technical conditions, all four auction forms have the same revenue implications! This result is called the *Revenue Equivalence Theorem*:

Theorem 2.1 (Revenue Equivalence Theorem) *Suppose that the assumptions of the benchmark model hold. Then, under some technical conditions on the distribution F, the expected revenue to the seller is the same under all four auction forms.*

Sketch of Proof The following notation will come in useful. Recall that F denotes the common distribution from which the n bidders' valuations are drawn. Let f denote the density of F. Let $v_1, \ldots, v_n$ represent the actual vector of valuations, and let $v_1^*, \ldots, v_n^*$ denote the valuations arranged in a non-increasing fashion.[7] In statistical terminology, v_i^* is called the *i-th order statistic*, $i = 1, \ldots, n$. Finally, define

$$J(v_1^*) \;=\; v_1^* - \frac{1 - F(v_1^*)}{f(v_1^*)} \qquad (2.1)$$

It will be assumed in the proof that $J(\cdot)$ is a strictly increasing function (the meaning of this assumption will become clear shortly).[8] The Revenue Equivalence Theorem will be proved by showing that under all four auction forms, the expected revenue to the seller is precisely the expectation of $J(v_1^*)$ with respect to the distribution of v_1^*.

Consider the English auction first. In this case, the second-to-last bidder drops out of the auction as soon as the item exceeds his own valuation. Consequently, the highest-valuation bidder receives the item, but pays an amount equal to the second-highest valuation. Thus, the expected revenue of the seller is equal to the expectation of v_2^*. This expectation can also be expressed in terms of v_1^* as follows. The winning bidder earns a rent of $(v_1^* - v_2^*)$. It can be shown (see McAfee and McMillan [14] for references) that the expected value of this rent is given by the expectation, with respect to the distribution of v_1^*, of

$$\frac{1 - F(v_1^*)}{f(v_1^*)}.$$

The amount the seller receives is, by definition, the valuation of the winning bidder minus the rent of the winning bidder. This difference is precisely expression (2.1).

Now consider the second-price sealed bid auction. We claim that in this case it is a dominant strategy for each bidder to simply bid his true valuation. To see this, note that the amount bid only affects the probability with which the bidder wins; the amount the winning bidder pays in a second-price auction is beyond his control. If bidder i were to bid *less* than his true valuation, then the auction outcome is affected only if the new bid is lower than the bid submitted by another bidder, say j, and as a result j wins the auction. Since rents from winning are non-negative, lowering his bid below his valuation clearly cannot make bidder i better off. If bidder i were to bid *more* than his true valuation, then the outcome of the auction changes only if there is another bidder (say j) whose bid is above i's valuation, but below the new bid. In this case, the new higher bid causes i to win, but now he also has to pay more than the item is worth to him. Consequently, this also cannot make bidder i better off than simply bidding his true valuation. Of course, if all bidders bid their true valuation, then—for the same reasons as in the English auction—the expected revenue of the seller is again given by the expectation of (2.1).

The proof that the seller's expected revenue under a first-price sealed bid auction is also equal to (2.1) is more complex. Unlike the case with English and second-price auctions, Nash equilibria in a first-price auction are not dominant strategy equilibria, and are therefore more difficult to identify. Let $B(\cdot)$ be the (common) bidding strategy used by all bidders $j \neq i$;

[7]That is, v_1^* is the highest valuation, v_2^* is the second highest, and so on.

[8]The technical conditions referred to in the statement of the theorem are precisely that F admits a density f, and that $J(\cdot)$ is a strictly increasing function. For sufficient conditions on F that will result in $J(\cdot)$ being a strictly increasing function, see McAfee and McMillan [14].

that is, bidder j bids $B(v_j)$ is his true valuation is v_j. Assume that $B(\cdot)$ is a monotonically increasing function (i.e., that bids increase if the valuation is higher). Consider bidder i's best-response. If i bids b_i, then the probability of winning is given by

$$\text{Prob}\,[b_i > B(v_j),\ j \neq i] \;=\; [F(B^{-1}(b_i))]^{n-1}$$

Thus, bidder i's expected surplus from bidding b_i is given by

$$\pi_i(v_i) \;=\; (v_i - b_i)[F(B^{-1}(b_i))]^{n-1} \tag{2.2}$$

Bidder i's optimal bid b_i must therefore satisfy $\partial\pi_i(v_i)/\partial b_i \;=\; 0$. It follows that when we differentiate π_i with respect to v_i, we obtain

$$\frac{d\pi_i(v_i)}{dv_i} \;=\; \frac{\partial\pi_i(v_i)}{\partial v_i} + \frac{\partial\pi_i(v_i)}{\partial b_i}\frac{db_i}{dv_i} \;=\; \frac{\partial\pi_i(v_i)}{\partial v_i} \;=\; [F(B^{-1}(b_i))]^{n-1} \tag{2.3}$$

At a symmetric Nash equilibrium, player i's optimal strategy will be the same as the strategy $B(\cdot)$ chosen by $j \neq i$. If $B(\cdot)$ represents a Nash equilibrium strategy, therefore, we must have $b_i = B(v_i)$. Substituting this in (2.3), we obtain the condition that

$$\frac{d\pi_i(v_i)}{dv_i} \;=\; [F(v_i)]^{n-1}\,. \tag{2.4}$$

Let v_l denote the lowest possible valuation. A bidder who has the valuation v_l must earn zero surplus (i.e., we must have $\pi_i(v_l) = 0$). Using this boundary condition, and integrating in (2.4), we obtain:

$$\pi_i(v_i) \;=\; \pi_i(v_l) + \int_{v_l}^{v_i}[F(x)]^{n-1}dx \;=\; \int_{v_l}^{v_i}[F(x)]^{n-1}dx \tag{2.5}$$

By equating (2.5) with the definition (2.2) of π, and invoking the Nash equilibrium condition $b_i = B(v_i)$, we finally obtain:

$$B(v_i) \;=\; v_i - \frac{\int_{v_l}^{v_i}[F(x)]^{n-1}dx}{[F(v_i)]^{n-1}}. \tag{2.6}$$

[9]

Note that $B(\cdot)$ is indeed an increasing function as was assumed earlier in the proof.

Since $B(\cdot)$ is increasing, the bidder with the highest valuation v_1^* wins the auction. From the point of view of the seller, thereore, the expected revenue from the auction is equal to the expected value of $B(v_1^*)$. It can be shown (see McAfee and McMillan [14] for references) that for any v, $B(v)$, as given by (2.6) is actually the expected value of the second-order statistic, conditional on the first-order statistic being v. Therefore, the expected revenue to the seller is simply the expected value of the second-order statistic, which, as we have already mentioned, is simply the expectation of the expression (2.1).

Since the Dutch auction is strategically equivalent to the second-price auction, the proof of the Revenue Equivalence Theorem is complete. □

It is very important to be clear about what the Revenue Equivalence Theorem asserts. The theorem claims only that the ex-ante expected revenue to the seller is the same in all four auction forms. It does not suggest that the realizations themselves always coincide. In the notation of the theorem, the revenue of the seller under the English or second-price auctions is equal to v_2^*, whereas under the first-price auction, it is equal to $B(v_1^*)$. Of course, $B(v_1^*)$ need not coincide with v_2^*. However, the expectation of $B(v_1^*)$ equals that of v_2^*, so that *on average* the seller's revenue is the same.

Building on this point, Vickrey [29] shows that the variance of the seller's revenue is smaller in an English or second-price auction than in a first-price or Dutch auction. Therefore, if the seller were risk-averse (rather than risk-neutral as we have assumed), he would prefer the English or second-price auctions to the first-price or Dutch.

Several other remarks are in order before moving to extensions of the Benchmark model:

1. The Revenue Equivalence Theorem is devoid of any empirical content, since it asserts essentially that "anything is optimal." We will see in the subsection following that this is no longer the case when we alter the assumptions of the benchmark model.

2. Auction outcomes in the benchmark model are always Pareto-efficient, since the bidder with the highest valuation receives the object.

3. An increase in the number of bidders increases the revenue of the seller on average. This is intuitive: when the number of bidders increases, the second-highest valuation also increases on average. Indeed, it can be shown that as n becomes unboundedly large, the price approaches the highest possible valuation.

2.5 The Optimal Auction in the Benchmark Model: A Comment

The Revenue Equivalence Theorem only compares the expected revenue across the four standard auction forms. It is a natural question to ask whether there are other (perhaps more complicated) auction forms that yield a higher expected revenue under the assumptions of the Benchmark Model. The answer is a qualified yes. Using the Revelation Principle, it can be shown (see, e.g., Harris and Raviv [9]) that if the seller's own valuation of the object is given by v_S, then the auction that produces the highest expected revenue has the following implications:

> If every bidder's valuation v is such that satisfies $J(v) < v_S$,[10] then the item is not sold to anybody. Otherwise the item is sold to the bidder with the highest valuation v at a price of $B(v)$.

Thus, the optimal auction involves the seller effectively setting a reserve price of $J^{-1}(v_S)$. Since $J(v) < v$, this reservation price raises the possibility that the outcome of the optimal auction could be Pareto-inefficient: it is possible that for all i we have $J(v_i) < v_S$, but there exists j such that $v_j > v_S$. Note, however, that if the reservation price is not binding, then the optimal auction has the same outcome on average as all the auction forms of the previous section, and is also therefore, Pareto-efficient.

[10] $J(\cdot)$ is the function defined in (2.1).

2.6 Beyond the Benchmark Model

In this section, we examine the effect of dropping in turn three of the Benchmark Model's four assumptions. The last assumption—that of independent private-values—is the subject of the next subsection. It is presumed throughout this subsection, that when the modification of one assumption is being discussed, the Benchmark Model's remaining three assumptions continue to hold. We begin with the assumption of risk-neutrality of the buyers.

Risk-Averse Buyers

Regardless of his attitude to risk, it is optimal for a buyer in an English auction to remain in the bidding until the price exceeds his personal valuation. Since the seller can always choose to sell using an English auction, it follows that, *ceteris paribus*, the seller is no worse off if the buyers are risk-averse rather than risk-neutral.

In fact, the seller can do strictly better. It can be shown that, with risk-averse bidders, the first-price sealed bid auction generates higher revenues on average than the English or second price auctions. The reason is not far to seek. Under the strategies described in the Benchmark Model, a typical bidder earns a positive rent if he wins, and zero rent if he does not. By increasing the size of his bid by a "small" amount, he can increase the probability of winning. Although this will decrease his rent if he wins, a risk-averse bidder will nonetheless find the trade-off worthwhile since it smooths his utility.

Even the first-price auction is not the optimal auction in this case. It is possible to design auctions with even greater expected revenues. Unfortunately, the optimal auction with risk-averse bidders lacks the simple form that obtains when bidders are risk-neutral; instead it involves a complex scheme which requires subsidizing high bidders who lose and penalizing all low bidders (see Maskin and Riley [12] for details).

Asymmetric Bidders

In the Benchmark Model, we assumed that bidders were symmetric, so that we had $F_i = F$ for all i.[11] If we allow for asymmetric bidders, then the F_i's need no longer coincide; as a consequence, it turns out once again that revenue equivalence also breaks down.

To understand why this is the case, note that the effect of asymmetry on bidders in an English auction is minimal. It remains optimal for a bidder to stay in the bidding until the price exceeds his private valuation. Thus, the revenue the seller obtains is equal to the second-highest valuation;[12] moreover, the bidder with the highest valuation receives the object, so that outcomes in an English auction remain Pareto-efficient.

In a first-price sealed bid auction, it continues to remain an equilibrium for each player to bid his estimate of the second-highest valuation conditional on his information (which is his own valuation). However, the situation is no longer symmetric: bidder i's estimate uses the distribution functions F_j for $j \neq i$. Therefore, two bidders who have the same valuation v could differ in their estimate of the gap between their valuation and the second-highest

[11] Recall that F_i denotes the distribution from which bidder i's valuation is drawn in the independent private-values model.

[12] Of course, the expectation of this second-highest valuation is more difficult to compute when the F_i's differ.

valuation. In turn, this implies that the winner in a first-price sealed bid auction need not necessarily be the bidder with the highest valuation, so outcomes could be Pareto-inefficient in this case. It also implies that the revenue implications of the first-price auction could differ from that of the English auction.

Unlike the earlier cases, however, no ordering between these auctions is possible in general. There are examples in the literature in which the expected revenue under the English auction is higher than under a first-price auction, and examples in which it is lower. Nor are any of the basic auction forms optimal. Myerson [18] has shown that the optimal auction with multiple types of bidders is a discriminatory auction in which the seller sets a different reserve price for each type of bidder.

Payments that do not depend only on the bids

There are many real-world auctions (such as book publishing, the music industry, and in some mineral-rights auctions) in which payment by the winning bidder to the seller depends (usually via a royalty), on information about the value of the good that is revealed after the auction. Notationally, we can represent the total payment by the winning bidder as

$$p = b + r\tilde{v}, \tag{2.7}$$

where b is the winning bid, r is the royalty rate, and $\tilde{v}$ the value of the commodity that is unknown at the time of the uaction, but that is revealed subsequently.

If the distribution of $\tilde{v}$ is exogeneous to the model (in particular, if it cannot be affected by actions of the seller or the winning bidder), then it can be shown (see, e.g., McAfee and McMillan [13]) that the expected revenue of the seller is an increasing function of the royalty rate r. The *optimal* royalty rate is therefore 100%. In reality, however, the distribution of $\tilde{v}$ is not likely to be exogeneous, but to be influenced by the winning bidder's actions. (Book sales, for instance, depend on the publicity and other promotion provided by the publisher.) In this case, the optimal royalty rate is less than 100%.

2.7 Affiliation and the Common-Values Model

What happens if the assumption of independent private values in the Benchmark Model is replaced by the other polar extreme, the common-value model? As in any auction model, bidders in a common-value model base their bids on their estimate of the item's value. This raises the paradoxical possibility that winning could be bad news: a bidder wins if and only if every other bidder estimated the common value to be lower. This phenomenon has been labelled the *winner's curse.*[13]

A rational bidder in a common-values auction anticipates the winner's curse effect, and takes this into account in deciding his bidding strategy. Milgrom and Weber [17] describe the equilibrium strategies for a common values model under each of the four basic auction mechanisms.[14] They show that of the four auction forms, the English auction now provides

[13]To express the winner's curse in statistical terms, suppose that the signal x_i about the common value v has the property that a larger signal x_i implies a larger true value of v. (This would be the case, for instance, if the distribution of signals $f(x_i|v)$ possessed the monotone likelihood ratio property.) Let x_1^* denote the largest of the n signals $x_1, \ldots, x_n$. Then, it is the case that $E(v|x_1^*) \geq E(v|x_1, \ldots, x_n)$.

[14]We do not describe equilibria of the common-values model in this section, since the next section (on the Bikhchandani–Huang paper) is concerned with a similar, but more general, model.

the highest expected revenue, followed by the second-price auction, and then by the Dutch and first-price auctions. (The last two, of course, continue to remain equivalent.)

These results are not unintuitive. The process of bidding in an English auction, for instance, provides information to the bidders, and the incorporation of this additional information into the bidding strategy reduces the impact of the winner's curse. There is no corresponding effect in the other three auction forms.

The full model considered by Milgrom and Weber is actually significantly more general than either the independent private-values model, or the common value model; indeed, it contains both as special cases. Since their model forms much of the basis of modern auction theory, we shall describe it in a little more detail here.

The valuation v_i of bidder i in the Milgrom–Weber model is a function of n "information variables" $x = (x_1, \ldots, x_n)$ and m other variables $s = (s_1, \ldots s_m)$:

$$v_i = v_i(s, x).$$

Bidder i observes the realization of x_i, but may not observe the realization of x_j for $j \neq i$, or the realization of any of the variables $s_1, \ldots, s_m$. Thus, it is possible that player i does not know with certainty the true worth (to him) of the object prior to the auction.

All bidders are assumed to be risk-neutral. In addition, bidders are *symmetric*: there exists a function $v: \mathbb{R}^{>+\kappa} \to \mathbb{R}_+$ such that

$$v_i(s, x) = v(s, x_i, (x_j)_{j \neq i}).$$

Finally, let $z = (s, x)$, and let G denote the joint distribution of z. The most important assumption of the Milgrom–Weber model lies in the assumption that the random variables z are *affiliated*. In mathematical terms, affiliation means that the density g of G satisfies the condition that[15]

$$g(z \vee z')g(z \wedge z') \geq g(z)g(z') \qquad \text{for all } z, z'. \tag{2.8}$$

Affiliation implies that valuations are positively correlated, i.e., that a high valuation by one bidder makes high valuations by other bidders more likely. Densities exhibiting affiliation possess a number of strong properties. For instance:

Theorem 2.2 (Properties of Affiliation) *Let $y_1, \ldots, y_k$ be a vector of affiliated random variables.*

1. *Let $h_1, \ldots, h_k$ be non-decreasing functions. Then, the variables $h_1(y_1), \ldots, h_k(y_k)$ are also affiliated.*

2. *Let $H: \mathbb{R}^{\rceil} \to \mathbb{R}$ be non-decreasing. Let constants (a_i, b_i) be given for $i = 1, \ldots, k$, where $a_i \leq b_i$ for all i. Then, the following function h^* is itself a non-decreasing function of its arguments:*

$$h^*(a_1, b_1; \ldots; a_k, b_k) = E[H(y_1, \ldots, y_k) \mid a_i \leq y_i \leq b_i \text{ for all } i]$$

[15] The notation $\vee$ and $\wedge$ stand for coordinate-wise maximum, and coordinate-wise minimum, respectively.

For a list of the many other properties of affiliated random variables, we refer the reader to Milgrom and Weber [17].

It is not very hard to see that both the independent private-values model and the common value model are actually special cases of the Milgrom–Weber framework. To obtain the former, we take

1. $m = 0$.

2. $v_i(s, x) = x_i$ for each i.

3. $G(z) = F(x_1) \times \cdots \times F(x_n)$.

(The last condition is the statement of independence in values.) It is not very hard to check that the distribution G, thus defined, is affiliated.

Similarly, to obtain the common-value model, we define

1. $m = 1$.

2. $v_i(s_1, x_1, \ldots, x_n) = s_1$ for all i.

Thus, the common value v is given by s_1. Sufficient conditions for the joint density of $(s_1, x_1, \ldots x_n)$ to meet the requirement of affiliation are that (i) conditional on s_1, the signals $x_1, \ldots, x_n$ be independent, and (ii) the conditional density $g(x_i|s_1)$ have the monotone likelihood ratio property.

Two results that Milgrom and Weber [17] prove in this model are of particular interest for us. First, they show that the ranking

English $\geq$ Second-Price $\geq$ First-Price = Dutch

continues to hold in this more general setting. The inequalities are typically, though not always, strict: we already know that equality obtains throughout for the special case of the independent private-values model. Second, no matter what the auction form chosen, the seller's expected revenue increases if he releases any private information he possesses about the item's true valuation (assuming that this can be done in a credible fashion). The basis for both results is the effect on the winner's curse. If buyers are uncertain about their true valuation, then any information they obtain from observing the bids of others will reduce the winner's curse; thus, the English auction does better than the others. The credible release of private information by the seller also reduces the winner's curse; this encourages bidders to bid more aggresively, increasing the seller's revenues on average.

2.8 Lessons for T–Bill Auctions?

Treasury Bill auctions share several features in common with the auction forms discussed above. They are common value auctions, in which (to a first approximation) bidders can be regarded as symmetric and risk-neutral. Moreover, payments to the seller in Treasury Bill auctions depend only on the winning bids.

However, there are also differences. As we mentioned in the Introduction, the T–Bill auction market is actually the middle of three linked markets; it is preceded by the when-issued market in which potential bidders can trade forward in the security that will be auctioned, and is followed by an active secondary market. Moreover, even when viewed in isolation, T–Bill auctions are essentially divisible good auctions, whereas the material of this section has focussed on auctions of an indivisible good. Bids in a T-bill auction are not just prices, but are *price-quantity pairs*; each bidder specified the amount he is willing to buy at different prices. Thus, effectively, bidders in a T–Bill auction submit *demand curves.*

The chief question in the study of T–Bill auctions is whether Uniform-Price Auctions (UPA) generate a greater expected revenue for the seller than Discriminatory Auctions (DA). In a UPA, each bidder pays the same price for the units he buys; this price, called the *stop-out price*, is the maximum price at which total demand exceeds supply. In a DA, each bid is filled at the bid price for that unit; the seller begins with the highest price bidder and works down the demand curves till the entire quantity is exhausted. In notational terms, if $b_i(p)$ denotes the quantity demanded by bidder i at the price p, and p^* denotes the stop-out price, then the bidder in a discriminatory auction pays

$$p^*b_i(p^*) + \int_{p^*}^{\infty} b_i(p)dp,$$

while the bidder in a uniform-price auction pays $p^*b_i(p^*)$.

A little reflection shows that the indivisible-good analog of the UPA is the second-price auction, while that of the DA is the first-price auction. As we have seen above, the second-price auction does strictly better than a first-price auction in common value auctions of indivisible goods, since it reduces the impact of the winner's curse. Thus, if we ignored the differences mentioned above, we would conclude that a second-price auction (i.e., the UPA) would generate greater revenues for the Treasury than the first-price auction (i.e., the DA). Indeed, invoking precisely this reasoning, a number of authors (e.g., McAfee and McMillan [14], Bikhchandani and Huang [2], or Smith [24]) have recommended in recent years that the Treasury switch from discriminatory to uniform-price auctions

The validity of this result clearly depends on the legitimacy of identifying T–Bill auctions with the auction models studied here. In at least one direction, the difference may not matter. Bikhchandani and Huang [2] show that when a secondary market is appended to the auction market of this section, then—even in the presence of information linkages between the markets that affect auction behavior in a non-trivial way—it is the case that the UPA dominates the DA in terms of expected revenue for the seller.[16] We review their results in the next section.

Unfortunately, the decision to treat a divisible-good auction as an indivisible-good auction is not a legitimate one. That this could be problematic was pointed out a number of years ago by Wilson [33] in his analysis of the auctions of shares. Wilson's work has been reinforced and extended recently by Back and Zender [1]. The main thesis of the Back–Zender paper is that with divisible goods, the use of uniform-price auctions gives rise to the

[16] Bikhchandani and Huang actually consider the auction of k units of an indivisible good rather than one unit as we have assumed here. This difference is not, however, very significant, since they assume that each bidder is entitled to at most one unit of the good.

possibility of implicit collusion among the bidders, and could lead to a drastic reduction in revenue from the discriminatory auction scenario. We return to this issue in Section 4.

3 Auctions with Resale Markets

Bikhchandani and Huang [2] study a common-value model in which a secondary (resale) market is appended to the primary auction market. They assume that bidders in the primary market have better information on average than investors in the secondary market.[17] Prices in the secondary market will, in such a situation, be affected by the information revealed in the process of bidding in the auction, and this leads to interesting trade-offs facing both the bidders and the seller in the auction market.

Consider the bidders first. The information linkage between the markets prompts bidders to use their bids as signalling mechanisms with a view to increasing the secondary market price. Thus, *ceteris paribus*, they would like to increase their bids (which increases seller revenue). However, higher bids also exacerbate the winner's curse effect, which could lower expected profits for the bidders.

The seller faces a similar problem. From the material of the previous section, we have seen that it is in the seller's interest to reveal his private information concerning the item's true value, since this reduces the impact of the winner's curse and therefore leads to higher bids. However, when there is a secondary market whose price is affected by the information contained in primary market behavior, the release of information by the seller reduces the buyers' incentive to send "high" signals and thus drives bids down.

The presence of these trade-offs makes the auction problem with resale markets materially different from those without such markets. In particular, the issue of which auction form is superior from the seller's viewpoint needs to be visited afresh; nor is it apparent whether the seller should release any private information he has about the good's true value. Under some technical conditions, Bikhchandani and Huang [2] establish the following results:

1. In both the discriminatory (first-price) and uniform-price (second-price) auctions, the seller's revenues are higher in the presence of the information linkage with the secondary market than when this information linkage is absent. Therefore, ignoring the information linkage can cause an underestimation of the seller's expected revenues, regardless of the auction mechanism used.

2. The uniform-price auction generates greater expected revenues than the discriminatory auction even in the presence of the information linkage between the primary and secondary markets.

3. The release of private information by the seller could sometimes decrease expected revenues. However, there are sufficient conditions under which expected revenues increase with the release of private information.

In the subsections that follow, we turn to a more detailed look at the first two results.

[17] This is a plausible assumption in the Treasury auction scenario, where auction market bidders are primary dealers and large financial institutions.

3.1 Structure of the Model

The Bikhchandani–Huang (henceforth, BH) model has n symmetric bidders and k identical items for sale. It is assumed that $n > k \geq 1$. Although multiple units may be available for sale, it is assumed that each buyer is entitled to at most one unit of the good. Thus, bids continue to be prices, rather than demand schedules which specify the quantity demanded at any price. The common-values model is operative. Finally, it is assumed that items are for 100 percent resale.

Let v denote the true value per unit of the good. Seller i's private information signal concerning v will be denoted x_i, $i = 1, \ldots, n$. Let $X = (x_1, \ldots, x_n)$. It is assumed that x_i lies in some interval $[x_L, x_H]$ for each i. After the auction market is over, but before the resale market, additional partial information may become available on the good's true value. Let p denote this information. Finally, let f denote the joint density function of (p, v, X). Two assumptions are made in this context:

1. The density f is affiliated in (p, v, X), and is symmetric in its last n arguments.
2. $E(v|p) \neq E(v|p, X)$. That is, given p, the bidders' private information has value.[18]

The *strategy* for a generic bidder (say, i) is, as earlier, a function B_i mapping i's private information into a bid $B_i(v_i)$. Since there are k items for sale, the k highest bidders all win. Bidders are risk-neutral: they choose their strategies to maximize their expected profits from the auction. A symmetric Bayesian-Nash equilibrium is a vector of strategies $(B_1, \ldots, B_n)$ such that for each bidder i, given that bidder $j \neq i$ is using the strategy B_j, it is optimal for i to adopt the strategy B_i.

To make the notion of an equilibrium more formal, we must define players' payoff functions in more detail. To accomplish this, we must first describe (a) the auction form under consideration, and (b) the price formation process in the secondary market.

Concerning (a), we focus on two auction forms: *first-price* or *discriminatory* auctions in which each winning bidder pays the amount he bid, and *second-price* or *uniform-price* auctions in which all winning bidders pay an amount equal to the highest losing bid. Concerning (b), it is assumed that following the auction, the seller releases information about the k winning bids and the highest losing bid.

Some additional notation will be useful in what follows. Since the model is symmetric, it makes sense to focus on symmetric equilibria, i.e., equilibria in which all bidders use the common strategy B. We sharpen this further by restricting attention to strategies B that are strictly increasing and differentiable. To identify and analyze such equilibria, we examine the optimization problem faced by a typical bidder, say bidder 1.

Let y_j denote the j-th order statistic of the signals $x_2, \ldots, x_n$ of the remaining $(n-1)$ bidders. Suppose all bidders apart from bidder 1 are following the strategy B. Suppose bidder 1 gets the signal $x_1 = x$, submits a bid b, and wins. If investors in the secondary market believe that i is also using the strategy B, the secondary market price will be

$$\begin{aligned} r(B^{-1}(b), y_1, \ldots, y_k, p) &= E[v|x_1 = B^{-1}(b), B^{-1}(B(y_1)), \ldots, B^{-1}(B(y_k)), p] \\ &= E[v|x_1 = B^{-1}(b), y_1, \ldots, y_k, p] \end{aligned} \tag{3.1}$$

[18]This condition would be violated, for instance, if $v = p$, i.e., if observing p also revealed the true value of v. If this condition is violated, the BH model reduces simply to the common value model of Milgrom and Weber.

Finally, let $r^*(x', x, y)$ denote the expected resale price conditional on $x_1 = x$ and $y_k = y$, given that secondary market buyers believe bidder 1's signal was actually x':

$$r^*(x', x, y) \;=\; E[r(x', y_1, \ldots, y_k, p) \mid x_1 = x, y_k = y]. \tag{3.2}$$

Note that, by affiliation, both r and r^* are increasing in their arguments (see the second part of Theorem 2.2).

3.2 Equilibrium in Discriminatory Auctions

The central result that Bikhchandani and Huang prove about discriminatory auctions is the following:

Theorem 3.1 *Under some technical conditions, there is a symmetric equilibrium in increasing, differentiable strategies of the discriminatory auction. At each possible vector of signals, the equilibrium bids are larger than they would have been in a standard common-value auction.*

Before describing the proof of the theorem, it is necessary to clarify the "technical conditions" under which the result holds. The random variables $z_1, \ldots, z_m$ are said to be *information complements* with respect to another random variable w if it is the case that

$$\frac{\partial^2 \phi}{\partial z_i \partial z_j}(z_1, \ldots, z_m) \geq 0, \qquad i \neq j,\, i, j = 1, \ldots, m,$$

where $\phi(z_1, \ldots, z_m) = E[w|z_1, \ldots, z_m]$. The proof of the theorem requires that $x_1, \ldots, x_n, p$ be information complements with respect to v. What does this condition achieve? The responsiveness of the secondary market price to player 1's submitted bid is measured by $r_1^*(x', x, y)$, the partial of r^* with respect to its first argument. The condition of information complementarity states that this responsiveness increases with a bidder's information realization. Thus, a bidder's incentive to signal increases with his realization under this circumstance.

Sketch of Proof Suppose all bidders apart from 1 are using the strategy B, where B is strictly increasing and differentiable. If $x_1 = x$ and bidder 1 submits a bid of b, his expected profits are

$$\pi(b|x) \;=\; E\left[\left(r(B^{-1}(b), y_1, \ldots, y_k, p) - b\right) 1_{b \geq B(y_k)} \mid x_1 = x\right]$$

where $1_{\mu \geq \nu}$ is the indicator function that takes on the value 1 if $\mu \geq \nu$, and zero otherwise. Let $F_k(y|x)$ and $f_k(y|x)$ denote, respectively, the distribution function and density of y_k given $x_1 = x$. Some manipulation shows that we have

$$\pi(b|x) \;=\; \int_{x_L}^{B^{-1}(b)} \left[r^*(B^{-1}(b), x, y) - b\right] f_k(y|x) dy. \tag{3.3}$$

For B to constitute a symmetric equilibrium strategy, it is necessary that the above function be maximized at $b = B(x)$. Taking the first-order condition with respect to b in (3.3) and evaluating it at $b = B(x)$, we obtain the following necessary condition for a symmetric equilibrium B:

$$\begin{aligned} 0 &= [r^*(x,x,x) - B(x)]f_x(x|x)[B'(x)]^{-1} \\ &\quad - F_k(x|x) + [B'(x)]^{-1}\int_{x_L}^{x} r_1^*(x,x,y)f_k(y|x)dy, \end{aligned} \tag{3.4}$$

where r_1^* denotes the partial of $r^*(x',x,y)$ with respect to its first argument. Rearranging this, we obtain the ordinary differential equation (ODE):

$$B'(x) = [r^*(x,x,x) - B(x)]\frac{f_k(x|x)}{F_k(x|x)} + \int_{x_L}^{x} r_1^*(x,x,y)\frac{f_k(y|x)}{F_k(y|x)}dy. \tag{3.5}$$

The boundary condition for (3.5) is $B(x_L) = r^*(x_L,x_L,x_L)$.[19] Using this to solve the ODE (3.5) , we obtain

$$B(x) = r^*(x,x,x) - \int_{x_L}^{x} L(u|x)dt(u) + \int_{x_L}^{x} \frac{h(u)}{f_k(u|u)}dL(u|x), \tag{3.6}$$

where

$$t(u) = r^*(u,u,u) \tag{3.7}$$

$$h(u) = \int_{x_L}^{u} r_1^*(u,u,y)f_k(y|u)dy \tag{3.8}$$

$$L(u|x) = \exp\left\{-\int_{u}^{x}\frac{f_k(s|s)}{F_k(s|s)}ds\right\}. \tag{3.9}$$

Of course, expressions (3.6)–(3.9) have been derived using *necessary* conditions for an optimum. To complete the proof, it remains to be shown that (i) the strategy B as defined by (3.6)–(3.9) does, in fact, maximize expected profits for plyer i when all the other players use B, and (ii) B is an increasing function. Indeed, there is also a third condition: for B to be an equilibrium bid, it must satisfy $B(x) \leq r^*(x,x,x)$ since the expected profit for a bidder must be non-negative in equilibrium. All three conditions can be verified using the condition of information complementarity. The details are lengthy and are omitted. The interested reader is referred to Theorem 1 of Bikhchandani and Huang [2].

Finally, it is shown in Milgrom and Weber [17] that, when t and L are defined exactly as in (3.7)–(3.8), the strategy

$$\beta(x) = r^*(x,x,x) - \int_{x_L}^{x} L(u|x)dt(u) \tag{3.10}$$

[19] Since bidders must earn positive expected profits, we must have $b(x) \leq r^*(x,x,x)$ at all x, and in particular at x_L. If $b(x_L) < r^*(x_L,x_L,x_L)$, then by raising the bid at x_L to $b(x_L)+\epsilon$ for ϵ sufficiently small, expected profit can be made strictly positive.

constitutes an equilibrium of the common-value model with affiliated private signals. It is immediate from a comparison of (3.6) and (3.10), that the former is strictly larger by the amount

$$\int_{x_L}^{x} \frac{h(u)}{f_k(u|u)} dL(u|x).$$

Note the intuitively appealing feature that the magnitude of this difference depends (through h) on r_1^*, the responsiveness of secondary market price to the submitted bid; and that if $r_1^* \equiv 0$, we have $\beta(\cdot) = B(\cdot)$.) □

Bikhchandani and Huang also show that a stronger result than Theorem 3.1 can be proved under strengthened hypotheses. Namely, if secondary market beliefs are *monotone* in the sense that they are increasing in the bids submitted in the primary market, then the strategies defined by (3.6)–(3.9) describe the *only* symmetric equilibrium in increasing strategies. Appealing to this result, we will refer in the sequel to "the" symmetric equilibrium of the discriminatory auction.

3.3 Equilibrium in Uniform-Price Auctions

In a uniform-price auctions, the k highest bidders win, and each pay an amount equal to the highest losing bid. Bikhcahandani and Huang show that at any symmetric equilibrium of the uniform-price auction, the seller's expected revenues are greater than at the symmetric equilibrium of the discriminatory auction. The intuition behind this result is simple. In the standard common value problem (i.e., one without a secondary market), we have already seen that the winner's curse is weakened in a uniform-price auction leading to more aggressive bidding and greater revenues for the seller. The presence of a secondary market only intensifies this effect: the uniform-price auction format makes it cheaper to submit high bids in order to influence the secondary market price.[20]

What do the equilibrium strategies in a uniform-price auction look like? Suppose all bidders but bidder 1 use the bidding strategy B. Continuing with the notation introduced at the top of this section, if bidder 1 receives the signal $x_1 = x$ and submits the bid b, his expected profit is

$$\xi(b|x) = \int_{x_L}^{B^{-1}(b)} \left[r^*(B^{-1}(b), x, y) - B(y)\right] f_k(y|x) dy. \tag{3.11}$$

For B to be a symmetric equilibrium strategy, it is necessary that $b = B(x)$ maximize (3.11). The first-order conditions for a maximum are

$$0 = [r^*(x, x, x) - B(x)] f_k(x|x) + \int_{x_L}^{x} r_1^*(x, x, y) f_k(y|x) dy. \tag{3.12}$$

[20]This last point raises an interesting issue. The more responsive secondary market price to primary market bids, the greater the incentive to bid high in a uniform-price auction. (The price a bidder pays in such an auction is beyond his control, but a large bid can increase the profits from winning.) Bikhchandani and Huang provide an example to show that equilibria may fail to exist altogether in uniform-price auctions because of this problem.

It is seen that the solution to (3.12) is given by

$$B(x) = r^*(x,x,x) + \frac{h(x)}{f_k(x|x)}. \tag{3.13}$$

where $h(x)$ is exactly as defined in (3.8). Therefore, *if* a symmetric equilibrium exists, it must have the form (3.13). It must be stressed that we are only working with necessary conditions. For the reasons mentioned in the last footnote, equilibria may fail to exist altogether.

3.4 Comparison of Auction Forms

Since any symmetric equilibrium of the UPA must be of the form (3.13), the necessary conditions suffice for the comparison of equilibrium revenues received by the seller under the two auction forms. Taking expectations in (3.13) and (3.6), Bikhchandani and Huang establish one of their central results (see their Theorem 4) that the uniform-price auction remains superior to the discriminatory auction even in the presence of the secondary market.

The Bikhchandani–Huang paper clearly establishes the importance of including secondary markets in any analysis of Treasury auctions. Nonetheless, their result on the superiority of uniform-price auctions is of limited value in the design of Treasury auctions. As we will see in the next section, when the item being auctioned is a divisible good (as in a Treasury auction) the intuition gained from the study of indivisible-good auctions is no longer applicable.

4 Auctions of Divisible Goods

Treasury-Bill auctions are essentially auctions of divisible goods, in which bidders compete for shares of the quantity being auctioned. It is a natural question to ask whether the analysis of indivisible-good auctions has any implications for this case. In particular, is it the case that *uniform-price auctions* (UPA's), which are the analogs of second-price auctions, always lead to greater expected revenue for the seller than *discriminatory auctions* (DA's) which are the analogs of the first-price auctions?

The answer, in a nutshell, is no. A significant difference between the second-price auction of an indivisible good and the uniform-price auction of a divisible good, is that the latter gives rise to the possibility of implicit collusion among the bidders, that could reduce the seller's revenues dramatically. This was originally pointed out by Wilson [33]. Building on Wilson's analysis, Back and Zender [1] have recently carried out an elaborate comparison of UPA's and DA's in general divisible-good auctions. We present a summary of their results here.

4.1 An Example

We begin with a simple, but elegant, example due to Wilson [33] that illustrates the possibility of implicit collusion in uniform-price auctions. The example compares a uniform-price auction for a divisible-good (Wilson calls this a "share auction") to an auction where the same good is treated as indivisible (a "unit auction" in Wilson's terminology). The main result is that the revenue obtained by the seller in the former case could be as low as one-half the revenue obtained in the latter case.

Example 4.1 Consider a common-value auction in which one unit of a good is being sold. Suppose further that all n bidders and the seller know that the true post-auction valuation of the good is $v = 1$. Finally, assume that if there are two or more bidders submitting the highest bid, then the winner is chosen randomly with all of the highest bidders being equiprobable.

Consider first the case where the item being auctioned is an indivisible good. It is easy to see that in any symmetric equilibrium of a second-price auction, the selling price must be $b = 1$.[21] Consequently, the seller receives the full value of the good.

Now suppose the item being auctioned is a divisible good, and the seller uses a uniform-price auction. Each bidder submits a demand schedule $B_i(p)$ which specifies the fraction of the total quantity he is willing to buy at the price p. The sale price p^* is then defined as the price p at which $\sum_{i=1}^{n} B_i(p) = 1$.[22]

We claim that it is a symmetric equilibrium for all n bidders to adopt the following strategy:

$$B(p) = \frac{1}{n-1} - \frac{2p}{n(n-1)}. \tag{4.1}$$

To see this, suppose all bidders but i adopt this strategy, and consider i's best response. If bidder i submits the schedule $B_i(p)$, the resulting price will be the value p^0 which satisfies

$$B_i(p^0) + (n-1)B(p^0) = 1.$$

The consequent profit earned by i is

$$(1-p^0)B_i(p^0) = (1-p^0)[1-(n-1)B_i(p^0)] = (1-p^0)\frac{2p^0}{n}.$$

A simple calculation shows that bidder i's profit is maximized at $p^0 = 1/2$. Another simple calculation shows that $p^0 = 1/2$ is precisely the price that arises if i also submits the schedule $B(p)$ given by (4.1). Therefore, the strategies in (4.1) describe an equilibrium; the revenue the seller receives in this equilibrium is $p^0 = 1/2$. □

At first sight, it may appear that the driving feature of Example 4.1 is the assumption of a degenerate distribution for the post-auction value v, but this is incorrect. Back and Zender [1] have shown recently that even in a much more general setting, there always exist equilibria in uniform-price auctions of divisible goods in which the equilibrium price is less

[21] Suppose some bidder (say, i) faces a situation where the highest bid from the remaining $(n-1)$ bidders is $b^* < 1$. If i submits a bid $b_i \leq b^*$, the *maximum* surplus that could arise is $(1-b^*)/2$. This can be improved on by bidding $b^* + \epsilon$ for ϵ sufficiently small. Incidentally, note that the selling price will be $b = 1$ in an English auction or a first-price auction also.

[22] This definition is incomplete. If demand curves are discontinuous, it is possible that total demand at any price below p^* *exceeds* the available supply of one unit, while total demand at any price above p^* is less than unity. Moreover, p^* may not be unique. These are important considerations in general, but are not relevant for this particular example, since the equilibrium we study has p^* uniquely defined.

than the lowest possible post-auction price. Even more importantly, they show that there may exist equilibria of the discriminatory auction which strictly dominate these equilibria of the uniform-price auction. Thus, unlike the case with indivisible goods, uniform-price auctions are *not* necessarily superior to discriminatory auctions from the seller's viewpoint. We describe the Back–Zender results in more detail below.

4.2 The Back–Zender Model

Back and Zender consider a model with the following structure. There are $n > 1$ buyers and a single seller. The auction involves a quantity Q (normalized to unity) of a perfectly divisible good. The post-auction value of the good is a random variable v that takes values in the interval $[v_L, v_H]$. Prior to the auction each bidder observes a signal x_i that is correlated with v. The joint distribution of v and $x = (x_1, \ldots, x_n)$ is common knowledge to all the players in the game.

The seller sets a reserve price $p_L \geq 0$. After observing his signal x_i, each bidder i submits a demand curve $B_i(\cdot|x_i)$, which specifies for each $p \geq p_L$, the quantity $B_i(p|x_i)$ demanded by i. A *strategy* for player i in the auction is, therefore, the specification of a demand curve for each possible value of x_i. Let $B(\cdot|x) = \sum_{i=1}^{n} B_i(p|x_i)$ denote the aggregate demand that arises under the vector of signals x.

The Stop-Out Price

The *stop-out price* p^* is defined as the maximum price at which demand equals or exceeds supply. Thus, given a vector of strategies $((B_i(\cdot|\cdot))$, the stop-out price that results from the vector of signals x is

$$p^*(x) = \begin{cases} p_L & \text{if } B(p|x) < 1 \text{ for all } p \geq p_L \\ \max\{p : B(p|x) \geq 1\} & \text{if } B(p|x) \geq 1 \text{ for some } p \geq p_L \end{cases}$$

Quantity Allocation

If total demand at the stop-out price $p^*(x)$ exactly equals the available supply of unity, then each bidder receives the quantity he demanded. It is possible, however, that because of "flats" in the demand curve, the total demand at p^* may exceed supply. In this case, the supply is distributed pro-rata among the bidders according to the following rules. Fix x, and for notational simplicity, let $p^* = p^*(x)$. Let

$$\Delta B_i(p^*|x_i) = B_i(p^*|x_i) - \lim_{p \downarrow p^*} B_i(p|x_i)$$

be the flat in bidder i's demand curve at p^*. The flat in the aggregate demand curve is

$$\Delta B(p|x) = \sum_{i=1}^{n} \Delta B_i(p|x_i).$$

The fraction of the flat in the aggregate demand curve that *cannot* be filled is

$$\lambda(x) = \max\left\{\frac{B(p^*|x) - 1}{\Delta B(p^*|x)}, 0\right\}.$$

Since the available amount is issued pro-rata, the amount received by bidder i is

$$B_i^*(x) = B_i(p|x_i) - \lambda(x)\Delta B_i(p|x_i).$$

Seller Revenues in UPA and DA

In a uniform-price auction, each bidder pays the stop-out price for each unit that he receives; thus bidder i pays a total of $p^*(x)B_i^*(x)$. In a discriminatory auction, each bidder pays the entire area under his demand curve out to p^*; thus, bidder i pays

$$p^*(x)B_i^*(x) + \int_{p^*(x)}^{\infty} B_i(p|x_i)dp.$$

4.2.1 Equilibria in the Uniform-Price Auction

The central result that Back and Zender derive concerning uniform-price auctions is that there are a continuum of symmetric Bayesian-Nash equilibria in which the stop-out price is less than v_L, although v_L is the lowest possible post-auction price. In order to appreciate fully the difference between divisible-good and indivisible-good auctions, it is necessary to understand intuitively how this result could be true.

The key lies in the best-response problem faced by bidders in uniform-price auctions of a divisible good. The marginal cost curve facing such any such bidder is *endogeneous*: it is determined by the residual supply curve after subtracting the total demand curve of the other bidders. Now, if the total demand curve submitted by the other bidders is suitably steep, then marginal cost escalates very rapidly for the last bidder. Thus, it is possible that the equilibria which result from such strategies could also result in the auction price being below the good's value. The following result is proved simply by formalizing this intuition. (See also the remark following the proof.)

Theorem 4.2 *Assume that the seller's reserve price p_L satisfies $p_L \leq v_L$. Then, for any $p^* \in [p_L, v_L]$, there is an equilibrium in which the stop-out price is p^* regardless of the value of x.*

Proof Pick any $p^* \in [p_L, v_L]$, and let p' be defined by $p' = [(n-1)v_H + p^*]/n$. Consider the following strategies:

$$B_j(p|x_j) = \begin{cases} \dfrac{1}{n-1}, & p_L \leq p \leq p^* \\ \dfrac{p'-p}{n(p'-p)+p-p^*}, & p^* < p < p' \\ 0, & p > p'. \end{cases} \tag{4.2}$$

Note that the right-hand side of this expression does not depend on x_j, so the strategies are effectively independent of x_j. Note also that $\lim_{p \downarrow p^*} B_j(p) = 1/n$, so the demand curve has a flat of $[1/(n-1)] - [1/n]$ at p^*.

Suppose all bidders $j \neq i$ adopt the strategy (4.2). The *residual supply curve* $s(p)$ is defined as the quantity bidder i would obtain, if his demand were to make the stop-out price equal to p. This curve is easily defined at all points except at p^*. At p^*, the demand curves specified by (4.2) have a flat, and so therefore, will the supply curve. Some messy computation using the pro-rata rule shows that bidder i can obtain any quantity $q \in [0, 1/n]$ at this price, and the residual supply curve therefore is:

$$\begin{aligned} s(p) &= 0, & p < p^* \\ s(p) &\in [0, 1/n], & p = p^* \\ s(p) &= \frac{(v_H - p^*)}{n(v_H - p)}, & p^* < p \leq p' \\ s(p) &= 1, & p > p' \end{aligned} \tag{4.3}$$

Now observe the following:

- Since the total quantity demanded by the remaining $n - 1$ bidders at the price p^* equals unity, no bid below p^* can be successful.
- At a price of stop-out price of p^*:
  1. The maximum quantity that i can obtain is $1/n$.
  2. Therefore, if p^* is the best price for bidder i, he will want this maximum quantity.
  3. In this case, his demand must satisfy $\lim_{p \downarrow p^*} B_i(p) = 1/n$.
  4. If the realized post-auction price turns out to be v, bidder i will receive a profit of $(v - p^*)/n$.
- At any price above p^*, no bidder $j \neq i$ has a flat. Thus, even if i's demand were to push total demand above unity, each bidder $j \neq i$ would receive his full requested amount (see the pro-rata rule). Therefore, i cannot gain by pushing demand above unity at any price above p^*.
- Finally, bidder i can obtain the entire quantity available by submitting a demand of unity at the price p'.[23] For any realization v of the post-auction price, this would result in a profit of $(v - p')$. A simple calculation using the definition of p' shows that, regardless of the value of $v \in [v_L, v_H]$, this profit is dominated by $(v - p^*)/n$. Therefore, by 2(d) above, it is suboptimal for i to make the stop-out price p' or higher.

Summing up, bidder i's best-response problem is to find the best stop-out price $p \in [p^*, p']$. The quantity obtained at the stop-out price $p \in (p^*, p']$ is

$$q_i(p) = \frac{1}{n}\left(\frac{v_H - p^*}{v_H - p}\right).$$

Therefore, if the realized post-auction price were v and bidder i chose the stop-out price p, the profit realized by i would be

$$(v - p)q_i(p) = \frac{1}{n}\left(\frac{(v - p)(v_H - p^*)}{v_H - p}\right).$$

[23] Clearly any bid above p' is senseless.

For $p^* < p$, it is easy to check that—regardless of the value of v—this quantity is decreasing in p. The quantity received by i as $p \downarrow p^*$ tends to $1/n$. Thus, the maximum profit i could receive by selecting $p \in (p^*, p')$ is $(v - p^*)/n$.

Since this is preceisely the profit received by i by submitting the demand curve (4.2) (again, see point 2 above), it follows that (4.2) is a best-response for bidder i, when all the other players are also using this strategy. Since this strategy does not depend on i's signal, by symmetry it is an equilibrium of the auction game. Of course, this implies that p^* is an equilibrium price.[24] □

4.2.2 Equilibria of Discriminatory Auctions

It is obvious that, *ceteris paribus*, a bidder in a discriminatory auction would wish to submit a flatter demand curve than in a uniform-price auction. Back and Zender show that, under relatively weak conditions, there is an equilibrium of the discriminatory auction in which each bidder submits a totally flat demand curve, with each bidder bidding for the entire quantity at a single price.

Theorem 4.3 *Suppose that the first-price unit auction[25] has a (possibly mixed-strategy) equilibrium $(p_1, \ldots, p_n)$. Then, it is an equilibrium of the discriminatory auction for bidder i to demand nothing at any price above p_i, and to demand the entire quantity at any price below p_i.*

Proof See Theorem 2 of Back and Zender [1]. □

4.2.3 Comparison of the Auction Forms

Unlike indivisible-good auctions, divisible-good auction models have a large number of possible equilibria. In the case of uniform-price auctions, we described a continuum of equilibria (one for each $p^* \in [p_L, v_L]$), but it must be remembered that even these are only one class of equilibria; other equilibria may also exist. A similar remark applies to discriminatory auctions.

If the comparison is limited to the classes of equilibria shown to exist in the two auction forms, then (under some additional conditions concerning the signals) Back and Zender [1] provide a parametrized family of examples where the discriminatory auction never does worse, and sometimes does better, than the uniform-price auction.

There are at least two ways of interpreting these results. The less controversial is to treat the central message of the paper as pointing out that there are important differences between divisible-good and indivisible-good auctions, and that one should not use the latter framework to draw conclusions about the former. A stronger interpretation would be that discriminatory auctions are, in general, superior to uniform-price auctions in the case of divisible goods. Although some have drawn such a conclusion, this does not seem warranted. Even if it were true within the Back-Zender equilibria that discriminatory auctions

[24] A persual of the proof shows that i's best-response is forced to the point p^* precisely by the way the residual supply curve behaves for $p \geq p^*$. The steep slope in this supply curve increases i's marginal cost rapidly; consequently a lower price works out better.

[25] Recall that "unit auction" refers to the case where bidders are only allowed to bid for the entire quantity.

dominated uniform-price auctions, one is still faced with the problem of explaining why these equilibria are strategically salient. This is not an easy task; equilibria such as those in Theorem 4.2 appear to require a considerable amount of coordination among the bidders, most notably concerning the stop-out price p^*.[26] More work appears to be needed to resolve this issue.

4.3 Other Work

Wang and Zender [32] study a model similar to that of Back and Zender. Their focus is on equilibria in continuously differentiable strategies.[27] They show that a continuum of such equilibria exist under both the the UPA and the DA formats; however, all but one equilibrium of the discriminatory auction format disappear when the seller imposes a reserve price.

A comparison of the seller revenues under the two formats reveals an interesting picture. When all bidders are risk-neutral, it is shown that the seller's revenue in the unique surviving equilibrium of the DA dominates all but one of the continuum of equilibria of the UPA format. When bidders are risk-averse, however, the equilibrium under a DA dominates some of the equilibria under a UPA, but it may also be dominated by some of the latter. Thus, an unambiguous ranking of auction formats according to seller revenues is impossible.

5 Empirical Testing of Auction Models

Empirical examinations of Treasury auctions have tended to focus mostly on one of three issues: (i) the revenue generation ability of the auction format (seller's perspective), (ii) the examination of whether bidders demonstrate rational bidding strategies (the buyer's perspective), and (iii) the degree of manipulability of the auction, and the impact of auctions on other market traded instruments (the market's perspective). Our brief empirical summary follows this framework.

5.1 Uniform-Price vs. Discriminatory Auctions

Bolten [3] represents one of the earliest tests of revenue generation. His analysis also attempts to understand the effects of competitive and non-competitive bidding demand. Bolten tests the the effect of the non-competitive bidders moving to the secondary market. His results show that if the non-competitive bidders stay then revenue from the UPA increases, but if they leave, then revenue declines. Given that the non-competitive bidders are likely to stay, the practical implication of the test would appear to be that the UPA would do better than the DA.

[26]Of course, it could be argued that since Treasury auctions take place repeatedly, there is enough of a time element to enable bidders to coordinate their strategies by trial and error. This is an inadequate defence. If the time element is really important, then we should be analyzing the repeated game, not the static game, and now new, more salient, equilibria could arise.

[27]A recent paper by Viswanathan and Wang [30] studies a divisible-good auction model with when-issued and secondary markets. Unfortunately, the paper came to our attention too late for a summary to be included here.

The auction process drew renewed attention recently when the Salomon scandal broke in May 1991 (see Jegadeesh [10] and Jordan and Jordan [11] for a analysis of this episode). In 1992, this prompted the Treasury to experiment with the UPA as an alternative mechanism in the issue of Treasury securities. A number of recent studies have since focussed on the UPA vs DA question, but with mixed results.

Simon [22] argues that the UPA costs the Treasury money; he finds that the markup of auction yields over when-issued yields is much higher when the Treasury used the UPA rather than the DA. Umlauf [28] and Tenorio [27] study Mexican Treasury auctions and Zambian foreign exchange auctions, respectively. The Mexican Treasury switched from a discriminatory to a uniform-price format in 1990. Zambia's auction of US dollars followed a uniform-price format upto 1986, and a discriminatory format thereafter. Both studies report higher seller revenues under the UPA. (Mester [15] also suggests that the UPA may prove to be mildly better for revenue generation.) In addition, Umlauf asserts that the UPA format lowered bidder profits substantially as it hampered collusion. However, the lack of a competitive secondary market in either case, and the relative lack of sophistication in the two markets, make it difficult to draw any conclusions concerning the US Treasury auction format.

5.2 Tests of Bidder Efficiency

In one of the few tests of bidder efficiency, Scott and Wolf [21] empirically examine two dealers A and B. A is a small dealer in T-Bills and B is a big dealer. They assume that the dealers have quadratic utility functions for wealth. The data used is 74 consecutive weeks of forecasts and bids of the 2 dealers. Each dealer is asked to provide the following forecasts before each auction: (i) a discrete probability distribution of stop out prices, and (ii) point estimates of post-auction opening selling prices. Using this information, Scott and Wolf solve for the bids that maximize the dealers' utility. They also compute a mean-variance efficient frontier of bids. They then compare the profits made from points on the efficient frontier with actual profits. While the actual bidding was different from the computed efficient bid, it was not significantly so. The evidence in favor of or against bidder efficiency is quite inconclusive. Technically, a problem here is that both frontier and actual bids are drawn from the same subjective decision maker. It is therefore even more surprising that they are different. It is thus hard to justify either bidding efficiency or inefficiency.

Simon [23] undertakes a more recent study of bidder efficiency. Using intra-day quotes he examines the risks and rewards of Treasury coupon auctions for bidders who face different trade-offs between the winner's curse and quantity risk. The data indicates that markups of treasury auction average rates over bid when issued rates averages 3/8 of a basis point. Bidders often established long positions in the security in the when-issued market when they could have obtained the paper in the auction more cheaply, suggesting a degree of bidder inefficiency and also a degree of information asymmetry prior to the auction.

Cammack [5] undertakes a study of the information aggregation properties of auctions and bidder efficiency. She finds that bidders rarely agree on the value of the bill in the auction, making the common value format particularly interesting. Auction prices are lower than secondary market prices, reflecting information asymmetry in the absence of a full market and the lack of agreement on price amongst bidders.

5.3 Other Work

In a recent paper, Sundaresan [26] finds that the auction markets were very different in the period 1980-83 as against the behavior in 1984-91. This is possibly on account of the fact that the Fed had widely divergent methods of interest rate management during these periods. In the earlier period, they managed money supply, leading to a sharp increase in the volatility of interest rates. After 1983, they switched to managing the level of the short rate by targeting the Fed funds rate. This has led to far less volatile interest rates.

A salient feature of the data appears to be that the bid-cover ratio (or the ratio of total bids to the amount of the issue) bears an inverse relationship to the dispersion of winning bids. This is not, perhaps, surprising. The greater the number of bids, the smaller the percentage of bidders who win. This means that they are likely to fall under a smaller portion of the distribution of bids, and therefore be bunched together. It is also noticed that the dispersion of bids is related to the level of yields, which is, as is well known, related to the volatility of the interest rate. The bid-cover ratios tend to be higher for auctions of short maturity debt, and the percentage of the issue sold to competitive bidders increases as the maturity of the debt increases.

Sundaresan also finds a strong relationship between the repo markets and auction activity. Prior to the auction, substantial pre-trading of the issue occurs in the when-issued market. As a consequence of this trading, several short positions are established, and are often hard to fulfill on delivery date, resulting in short squeezes. This forces the shorts to pay high premia for the security which is reflected in the repo market rates. Sundaresan reports a significant jump in premia for borrowing securities in the 10 day window around the auction.

Finally, Wachtel and Young [31] investigate the impact of auction announcements on the interest rate markets, and finds little impact. Hence, the announcements of auctions contain little information of surprise value to the market regarding the supply of treasury securities.

6 Concluding Comments

In summary, the implications of the theoretical and empirical work so far seems to indicate that the following important issues are raised by the evidence from the auctions markets:

1. On a theoretical level, the modelling of T-bill auctions must account for several market stages: the when-issued market, the auction itself, the repo market and the post-auction secondary market. One model attempting to do so is the recent work of Viswanathan and Wang [30].

2. Models of the auction with indivisible units or single units seem to give results that are quite the opposite to those of multi-unit auctions. Hence, the only viable models should be those permitting the submission of price-quantity schedules as bids.

3. In understanding the winner's curse, a simple comparison of the prices pre-auction with those post-auction will not provide correct results. This is because of price distortions caused by different pre-auction positions, bidder risk heterogeneity, and squeezes.

4. Auctions are not completely 'common-value' in form. This is because the bidders in the game have different objectives and prior positions off which they trade in the auction. Understanding this aspect from a modelling view point is necessary, though probably impossible to account for in any framework of the Treasury markets.

5. Finally, the liquidity effects of auctions are an important effect on the markets. Auctions punctuate the time line of the debt markets by changing the mix of off-the-run and on-the-run bonds.