

Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market

Ali Hortaçsu ¹

February 1, 2002

¹Department of Economics, University of Chicago, 1126 E. 59th Street, Chicago, IL 60613. E-mail: hortacsu@uchicago.edu. I would like to thank my thesis advisors Patrick Bajari, Lanier Benkard, Timothy Bresnahan, and Robert Wilson for their generous encouragement and support. Bradley Efron, Philip Haile, Jonathan Levin, Preston McAfee, Leonardo Rezende, and Frank Wolak made important contributions to this research with their valuable comments and suggestions. I would like to thank Berna Bayazitoğlu for letting me use her dataset, Erdal Yılmaz of the Central Bank of Turkey, Volkan Taşkın of the Turkish Ministry of Treasury, Ahmet Arzan and Derya Tamerler for helpful discussions about the Turkish Treasury auction system. Financial support for this research was provided by a John M. Olin Dissertation Fellowship from SIEPR and the Olin Program in Law and Economics.

Abstract

An important question in the theory of divisible good auctions is to determine whether a discriminatory auction yields higher revenues to the auctioneer than a uniform price or Vickrey auction. This question can be answered empirically using estimates of the bidders' true marginal valuations for the good being auctioned. Based on the share auction model of Wilson (1979), I propose an estimation method to reconstruct the distribution of bidders' marginal valuations using data on individual bids. I apply the method to data from the Turkish Treasury auction market. Using the estimated marginal valuations of the bidders to calculate what the revenue would have been under a uniform price or a Vickrey auction. I find that, ex-post, the discriminatory auction yielded more revenue to the Turkish Treasury than the alternatives. However, in ex-ante expected revenue comparisons, I fail to reject revenue equivalence between the discriminatory auction and the "best-case" outcome of the uniform price auction, in which bidders reveal their valuations truthfully. I investigate the robustness of my results to the presence of a common-value component in bidders' utilities when comparing the revenue performance of the discriminatory and Vickrey mechanisms. I also show that the estimation framework can be generalized to allow for deterministic vs. uncertain supply, asymmetries, and stochastic participation by bidders.

Keywords: Multi-unit auctions, divisible good or share auctions, Treasury auctions, structural econometrics, nonparametric identification and estimation

1 Introduction

What is the most effective way for a Treasury to sell government securities? Since governments sell about \$4 trillion dollars worth of securities every year (Bartolini and Cottarelli (1997)), many economists have tried to answer this question, interpreting “effectiveness” from both the revenue maximization and efficiency standpoints. At least since Friedman (1960), the unanimous suggestion of the profession has been to conduct the sale through an auction. Unfortunately there is much less of a consensus among economists as to what the optimal auction mechanism should be, giving way to a lively debate in the auction theory literature.¹

This paper explores what we can learn about revenue-maximizing mechanism choice using a structural econometric model of strategic bidding that is estimated using data on individual bids. Before I explain how this paper contributes to the debate on the optimal auction mechanism, however, I will briefly review the alternative auction mechanisms that have been considered in the literature.

Cross-country studies of Treasury practices reveal that practitioners overwhelmingly prefer one mechanism over others: 39 out of 42 countries surveyed by Bartolini and Cottarelli (1997) use the *discriminatory auction* mechanism, also known as the “pay-as-bid” or “multiple-price” mechanism. In this auction format, bidders may submit multiple price-quantity pairs as their bids, which trace out *bid functions* on the price-quantity plane. Treasury officials aggregate individual bid functions and find where the *aggregate bid function* meets supply, as in Figure 1(a). The revenue of the auctioneer is the area under the aggregate bid function up to the market clearing price.

An alternative mechanism that has been the favorite of many economists, including Friedman (1960), is the *uniform price auction*. Here, winning bids are determined in the same manner as in the discriminatory auction, but bidders pay the market clearing price for all the units they purchase. In this case, the auctioneer’s revenue is the rectangle defined by the total quantity being sold and the market clearing price, as shown in Figure 1(b).

To understand the revenue tradeoff between these two mechanisms, a brief look at an individual bidder’s problem is in order. In Figure 2, I draw a typical bid function in one of these auctions. Observe that this is a step function due to the finite number of price-quantity pairs. I also plot the *residual supply function* for this particular bidder, calculated by subtracting the aggregate bid function of all

¹See the surveys by Klemperer (2000), Bikhchandani and Huang (1993), Nandi (1997), Chari and Weber (1992) and Malvey and Archibald (1998).

other bidders from the total supply. The market clearing price of the auction is at the point where the individual bid function intersects the residual supply function. This point of intersection also defines the total quantity won by the bidder.

In the discriminatory auction, the bidder pays the area under his bid function up to the quantity he wins, area $A + B + C$ in Figure 2. Therefore, for any given quantity, a rational bidder would bid a price that is lower than his true (marginal) valuation for that quantity, i.e. “shade” his bid. The amount by which a bidder shades his bid relative to his valuation, however, depends on where he believes the market clearing price, or equivalently, the residual supply function will lie. If competing bidders have private information about their marginal valuations, the residual supply functions for each bidder will be random. Hence, what is relevant for a strategic bidder trying to make an optimal decision under uncertainty is the distribution of residual supply functions that he is going to face.

The incentive for bid-shading is not as strong in the uniform price auction, since the bidder only pays area $A + B$ in Figure 2, the rectangle defined by the market clearing price and the total quantity he wins. In fact, as Vickrey (1961) argued, if the auctioneer charges the bidder only the area under the residual supply curve (area A), then the bidder’s optimal response will be to bid his marginal valuation. Since bidders do not shade their bids as much, the market clearing price in a uniform price or Vickrey auction will be higher than in a discriminatory auction. Hence, the auctioneer extracts higher revenues from marginal units. However, in the discriminatory price auction, the auctioneer extracts more revenue from inframarginal units. Therefore, the revenue tradeoff between a discriminatory vs. a uniform price or Vickrey mechanism depends on the amount of bid-shading each bidder decides to undertake in the discriminatory auction.

In the case where bidders have demand for only a single unit of the good being sold, the celebrated revenue equivalence theorem of Vickrey (1961) and Myerson (1981) and the revenue ranking results of Milgrom and Weber (1982a) can be used to compare the expected revenues of the different mechanisms. However, in the case where bidders have demand for more than a single unit of the good, these results can not be applied even in the independent private values case, as recent results by Ausubel and Cramton (1997) and Engelbrecht-Wiggans and Kahn (1998) suggest that the three mechanisms are not, in general, allocationally equivalent.² In the case of common or affiliated

²The revenue equivalence theorem requires two compared mechanisms to be allocationally equivalent. In the case of independent private values, the Vickrey mechanism is efficient. Ausubel and Cramton (1997) show that the uniform

values, a general theoretical comparison becomes even more difficult, and hence the existing theoretical literature is usually dependent on specific functional form assumptions or considerable simplifications of the economic environment.³

This paper attempts to fill the gap left by auction theory by letting the data decide on the relative ranking between mechanisms. This is how I proceed: building on the seminal work of Wilson (1979), I model bidder behavior in a discriminatory auction as an incomplete information game. The model yields a first-order necessary condition in which the price bid for a given quantity is determined as the bidder's marginal value less an oligopsonistic "mark-down" that depends on the inverse elasticity of the residual supply function each bidder expects to face. By inverting this mark-down rule, I can calculate the unobserved marginal valuations of the bidders, and to conduct counterfactual calculations in which the revenue performance of a uniform price or Vickrey auction can be compared to the performance of a discriminatory auction. I apply the framework to data from the Turkish Treasury, covering 3-month T-bill auctions between 1991 and 1993. I find that the discriminatory mechanism yielded higher ex-post revenues to the Turkish Treasury than a uniform price or Vickrey mechanism would have. However, in ex-ante expected revenue comparisons, I fail to reject revenue equivalence between the discriminatory auction and the "best-case" outcome of the uniform price auction, in which bidders reveal their valuations truthfully.

To my knowledge this is the first empirical study of mechanism choice in divisible good auctions whose underlying model and thus its counterfactual predictions are consistent with a *strategic* divisible good auction framework. Previous empirical research addressing the question of selecting the correct mechanism for Treasury auctions has almost exclusively focused on "policy experiments" in which different auction formats have been used in different time periods or in the sale of securities of different maturities. These studies have compared the differential between the auction price and the resale or forward ("when-issued") market price of the security across separate samples of discriminatory and uniform price auctions.⁴

The assumption implicit in this line of research is that bidders' true valuations for the security are better reflected in resale or forward markets than in the auction. Hence, the claim is that smaller price mechanism is generically inefficient. Ausubel and Cramton (1997) and Engelbrecht-Wiggans and Kahn (1998) also construct examples in which the discriminatory mechanism is inefficient.

³See Bikchandani and Huang (1989), Back and Zender (1993), Noussair (1995).

⁴See Umlauf (1993), Simon (1994), Nyborg and Sundaresan (1996) and Malvey and Archibald (1998).

differentials between the auction price and the transaction prices in secondary markets reflects better surplus extraction by the auctioneer. The validity of this comparison relies heavily on the assumption that one can control for factors that may have changed between the conclusion of the auction and the start of trade in the resale market, including the release of new information that might impact bidders' valuations of the security. The advantage of the structural approach taken in this paper is that estimation of model primitives allows me to construct counterfactual simulations in which I can hold the *ex-ante* information sets of the bidders constant. Furthermore, the policy experiment studies cited above have relied on aggregated price data from the auctions, rather than utilizing bidder level data. Therefore distributional aspects of policy intervention through a change of mechanism can not be analyzed. An important advantage of this work is that I have access to individual bid data. In many circumstances, data from such policy experiments is not available, restricting the applicability of the methodology.

The paper that is closest to this work is by Heller and Lengwiler (1998), who use a similar first-order condition to estimate the amount of (counterfactual) bid-shading in a discriminatory auction.⁵ However, their model is based on the Nautz (1995) model of bidding, which treats bidders as price-takers, and hence overlooks the strategic nature of the game. In particular, their model predicts that all bidders will shade their bids by the same amount, regardless of whether they possess very different degrees of market power. Heller and Lengwiler (1998) model bidders' expectation about the market clearing price distribution using the distribution of past market clearing prices, and hence do not make use of the information contained in individual bids to investigate the strategic effect of each bidder on the equilibrium market clearing price distribution.

Previous attempts to analyze Treasury auctions within a *strategic* equilibrium framework have been constrained to use models of single-unit auctions, since strategic bidding in divisible good auctions is difficult to analyze both analytically and computationally. However, none of these papers have estimated a full structural model to conduct the policy counterfactuals reported here.⁶

⁵Observe that the very related idea of using first-order conditions to reconstruct unobserved marginal cost functions from observed supply decisions has existed in the industrial organization literature at least since Rosse (1970) and Bresnahan (1981).

⁶For example, Nyborg, Rydqvist and Sundaresan (1997) motivate their empirical investigation with a model of bidding; however, they use a single-unit sealed-bid first-price auction model, and do not estimate structural parameters of their model. Gordy (1994), builds and estimates a structural model to quantify the effect of reservation price policy

The work presented here also offers several methodological advances to research on structural econometric modelling of auctions. The identification and estimation method employed in this essay extends the nonparametric identification and estimation framework of Elyakime, Laffont, Loisel and Vuong (1994) and Guerre, Perrigne and Vuong (2000) to divisible good auctions, and offers an empirical algorithm that has low computational demands and minimizes the impact of distributional assumptions on estimation results. The statistical properties of the estimation algorithm are also discussed.

The outline of the paper is as follows: In section 2, I present a model of strategic bidding in a discriminatory price divisible good auction. Section 3 discusses how the model developed in section 2 can be used as an empirical device to estimate the unobserved marginal valuations of the bidders. In section 4, I discuss the institutional setup of the Turkish Treasury auction market, and report descriptive statistics from my data, presenting evidence on inventory/liquidity concerns as a source of private information. In section 5, I use bidding data from the Turkish Treasury bill auctions between October 1991 and October 1993 to estimate the marginal valuations of the bidders. Using my estimates, I conduct comparisons between the discriminatory, uniform price and Vickrey mechanisms. In section 6, I discuss the robustness of my results in the presence of a common value component, supply uncertainty, asymmetries, and stochastic participation. Section 7 concludes. The appendix contains proofs of the propositions and provides further discussions of the empirical method.

2 Modelling and identification framework

2.1 The Model

I will follow the share auction setup of Wilson (1979) to model bidder behavior in the divisible good discriminatory auction. Each bidder's private information regarding his utility from winning a given quantity of the security is given by a signal, s_i , which can be a scalar or vector. Let $F(s)$ denote the joint distribution of signals across bidders, and $F_i(s)$ be the marginal distribution of bidder i 's signal. Bidder i 's marginal utility from winning q units of the security is given by the marginal valuation function, $v_i(q, s_i, s_{-i})$. This function can depend on the information of other bidders, as in

in Portuguese Treasury auctions. However, Gordy's structural model treats the quantity choice of a bidder as being exogenous, and does not account for multiple price-quantity pairs submitted by a single bidder, a defining feature of Treasury auctions.

the generalized mineral rights model of Milgrom and Weber (1982b).

A special case of this utility specification is the symmetric independent private values (S-IPV) case, in which s_i are distributed identically and independently across bidders, and bidder valuations are ex-ante symmetric: $v_i(q, s_i, s_{-i}) = v(q, s_i)$. I will later relax this assumption to consider what happens when a bidder's valuation may depend on other bidders' information, i.e. the case where valuations are of the form $v(q, s_i, s_{-i})$. This is the case where a "winner's curse" might arise.

In the Wilson (1979) setup, bidders' strategies are constrained to be strictly decreasing, differentiable demand functions, $y_i(p)$.⁷ Let Q be the amount of Treasury bills for sale, which, when it is not announced before the auction by the Treasury, will be treated as an exogenous random variable whose distribution is common knowledge among the bidders. I will assume that the number of bidders in the auction, N , is common knowledge among the bidders; although this assumption will be relaxed in section 6.1.

With N competing bidders submitting strictly decreasing demand functions, the market clearing price, p^c , will be at the point where:

$$y_i(p^c) = Q - \sum_{j \neq i}^N y_j(p^c)$$

i.e. where bidder i 's demand curve intersects her "residual supply curve." Now, following Wilson (1979), let us define the most central theoretical object in this paper:

$$\begin{aligned} H(p, y_i(p)) &= \Pr\{y_i(p) \leq Q - \sum_{j \neq i}^N y_j(p)\} \\ &= \Pr\{p^c \leq p | y_i(p)\} \end{aligned}$$

This is the probability distribution function of the market clearing price, conditional on submitting the demand function $y_i(p)$. Observe that for every p and $y(p)$, $H(\cdot)$ is a probability distribution generated by the random variables $\{y_j(p), j \neq i\}$, and, potentially, Q and N . The requirement that the demand functions are strictly decreasing ensures that H does not have atoms over the support of prices. I will further assume that H is differentiable in both of its arguments in the following discussion.

Let us first look at the case of symmetric independent private values. In this case, the surplus of a given bidder i conditional on a particular realization of the market clearing price, p^c , will be given

⁷In section 2.4, I will also consider the more realistic case in which bidders are constrained to submit bids on a discrete price grid

by

$$\int_0^{y_i(p^c)} v(q, s_i) - \int_0^{y_i(p^c)} y_i^{-1}(q) dq \quad (1)$$

where the first term is the expected surplus a bidder gets from winning $y_i(p^c)$ units of T-bills and the second term is the payment of the bidder, which is equal to the area under his bid function to the point it intersects the market clearing price.

Since $H(\cdot)$ defines the probability distribution over the set of market clearing prices, the bidder's expected profit maximization problem becomes:

$$\max_{y_i(\cdot)} \int_0^\infty \left\{ \int_0^{y_i(p^c)} [v(q, s_i) - y_i^{-1}(q)] dq \right\} dH(p^c, y_i(p^c))$$

where $dH(p, y(p))$ (denoting the total derivative of H) is the probability density function of the market clearing price. The Euler-Lagrange necessary condition for the functional optimization problem above is, as derived in Appendix 8.1:

$$v(y_i(p), s_i) = p + \frac{H(p, y_i(p))}{\frac{\partial}{\partial p} H(p, y_i(p))} \quad (2)$$

That is, the optimum price bid, p , for $y_i(p)$ units of T-bills should equal the bidder's marginal valuation for that many units, minus the "bid-shading factor," $\frac{H(p, y_i(p))}{\frac{\partial}{\partial p} H(p, y_i(p))}$. The denominator of the bid-shading factor can be interpreted as the "probability of winning at least $y_i(p)$ units," and the numerator can be interpreted as "the probability of winning exactly $y_i(p)$ units."⁸

The necessary condition derived above implicitly defines a type- or signal-dependent bid function, $y(p, s_i)$, that maps the signal space into the space of decreasing, differentiable functions of price, which can constitute a Bayesian Nash equilibrium.⁹ In the case where there are a finite number of

⁸Those familiar with auction theory will recognize that this relation is very similar to the inverse hazard ratio relationship between the bid and the value of the bidder in a private value first price auction. The intuition behind this relationship can also be cast (loosely) in terms of the inverse elasticity rule in monopoly pricing. Recall that the monopoly markup is $-\frac{D(p)}{D'(p)}$, the demand at the monopoly price divided by the derivative of demand at the monopoly price. In the auction setting, each bidder can be thought of a monopolist restricting its demand in the face of an uncertain residual supply curve. So, if a bidder is asking for q units of T-bills at price p , the "average supply of T-bills" at the bid price, p , will be $H(p, q)q$. If the bidder increases his bid by dp , the change in "average supply" will be $\frac{d(H(p, q)q)}{dp} = H_p(p, q)q$. The monopoly markup condition in this case, the bid price of p for q units is optimal if $v(q) = p + \frac{H(p, q)}{H_p(p, q)}$.

⁹An analytically solved example of the equilibrium of the above game when there are two bidders and the signals are scalars with negative exponential distribution is in Hortaçsu (2001). Wilson (1979) derives the first-order condition for the uniform price auction, and calculates equilibrium bid functions for several functional form assumptions.

units for sale, and consequently bids take the form of finite dimensional price vectors, Reny (1999) shows the existence of a Bayesian Nash equilibrium of the (asymmetric) independent private value discriminatory auction game. Although Reny’s results in the discrete version of this model need not necessarily carry through to the continuous version analyzed here, in section 2.4, I show that a similar first-order characterization of equilibrium strategies can be obtained in a discretized version of the model discussed in this section.

Another question not addressed by the above discussion is whether a given set of model primitives (the vector of marginal valuation functions and joint distribution of bidder signals), leads to a unique set of equilibrium bid functions. Unlike the independent private value first-price auction, there are no results regarding the uniqueness of the Bayesian Nash equilibrium of a discriminatory auction. However, regardless of which equilibrium is selected, the equilibrium bid functions must still obey the necessary condition derived in equation (2).

2.2 Identification of marginal valuations from bid data

The identification strategy pursued below builds on the insight of Elyakime et al. (1994) and Guerre et al. (2000), who utilized the necessary condition for optimality characterizing the optimal bid function in a first-price auction to identify and estimate the valuation distribution of the bidders.

The necessary condition for optimality, equation (2), allows us to nonparametrically identify the marginal valuations of the bidders using observed bids – if we make the assumption that the bid functions we observe in the data are indeed generated by a Bayesian Nash equilibrium of the discriminatory auction game, satisfying necessary condition (2).

To see this, define the following quantity:

$$G_i(p, x) = \Pr\left\{x \leq Q - \sum_{j \neq i}^N y(p, s_j)\right\} \quad (3)$$

which is the probability that, at price p , a demand of quantity x will be less than the (stochastic) residual supply faced by bidder i . This probability can be estimated for all (p, x) pairs if the joint distribution of $\{y(p, s_j), j \neq i\}$ can be estimated from the data – given this distribution, one can calculate the probability in (3).

Then, since:

$$H(p, y(p, s_i)) = G_i(p, x)|_{x=y(p, s_i)} \quad (4)$$

and

$$\frac{\partial}{\partial p} H(p, y(p, s_i)) = \frac{\partial}{\partial p} G_i(p, x)|_{x=y(p, s_i)} \quad (5)$$

all components of the necessary condition (2) are identified from the data. Hence, the marginal valuations, $v(y(p, s_i), s_i)$, corresponding to each point on the bid function, $y(p, s_i)$, are identified as:¹⁰

$$v(y(p, s_i), s_i) = p + \frac{G_i(p, x)|_{x=y(p, s_i)}}{\frac{\partial}{\partial p} G_i(p, x)|_{x=y(p, s_i)}} \quad (6)$$

I should note that the above identification argument did not assume anything about the symmetry of the bidders. Similarly, the above discussion also does not make require the independence of bidder's signal distribution. Necessary condition (2) would also hold true if bidders' signals are correlated, subject to the caveat that a Bayesian Nash equilibrium may not exist in this case.

2.3 Discrete strategy space

In the Turkish Treasury auction market, and plausibly in many other real-world Treasury auction markets, there is a smallest price increment (0.001 TL), and a minimum bid quantity restriction. We also observe bidders submitting a finite number of price-quantity pairs as their bids, making up “step-function” bids, rather than continuous, strictly decreasing bid functions as analyzed above – which is probably the case in almost all multi-unit auction markets.

I will maintain the divisibility of the quantities, but restrict the prices to lie on a discrete grid. In particular, following the model of Nautz (1995), let $p_0 < p_1 < \dots < p_{K+1}$ denote the set of possible prices on the grid. Hence, the “bid vector” submitted by each bidder takes the form of a series of quantities specified for each of these prices: $\vec{y}_i : \{y_{i0} \geq y_{i1} \geq \dots \geq y_{iK+1}\}$. A Bayesian Nash equilibrium in this setting would be a vector valued mapping of the form $\vec{y}_i(s_i)$, where the quantity requested at each point of the price grid is a function of a bidder's private information.

After all bids are submitted, the auctioneer determines the market clearing price by aggregating the quantity bids for each point on the price grid and finding the price at which total demand just

¹⁰Observe that if we observe bidder i bidding for y units at price p , then, for $y' \neq y$, $v(y', s_i) \neq p + \frac{G_i(p, y')}{\frac{\partial}{\partial p} G_i(p, y')}$. The equality holds if we replace p with $p' = y^{-1}(y', s_i)$.

falls short of the total supply:¹¹

$$p_{k^*} : k^* = \min\{k : \sum_{i=1}^N y_{ik} \leq Q\} \quad (7)$$

Analogous to the continuous case, define the probability that the market clearing price is below p_k , conditional on the bid vector of bidder i , \vec{y}_i , to be: $H(p_k, \vec{y}_i)$.

Given this setting, a reasonable conjecture of the relevant first-order condition of optimality would be:

$$v_i(y_{ik}, s_i) = p_k + \frac{H(p_{k-1}, \vec{y}_i)[p_k - p_{k-1}]}{H(p_k, \vec{y}_i) - H(p_{k-1}, \vec{y}_i)} \quad (8)$$

which is just the discretized version of the first-order conditions given by equation (2). Although this intuition is likely to be correct for most practical purposes, in Appendix 8.2, I investigate this conjecture more carefully, and conclude that subtle modifications are needed to account for the cases in which the monotonicity constraint on the bid function is binding, which might have observable consequences on both the observed bids and the marginal valuation function that rationalizes these bids.

3 Two Estimation Methods

Bidding data for multi-unit auctions typically take the form of a sequence of price-quantity pairs for each bidder in a series of auctions. That is, for each auction t , we observe the series of bid functions: $\{y_{it}(p, s_i), i = 1, \dots, N_t; t = 1, \dots, T\}$. Utilizing equation (2) requires us to estimate the distribution of the market clearing price distribution that each bidder expects to face before the auction, given the data we have.

In this section, I describe two different methods to estimate the components of equation (2). The first method is geared towards minimizing the impact of a priori distributional assumptions on the data, and has the additional benefit of being computationally straightforward. The second method illustrates how to use more standard estimation methods to relax the symmetry and exchangeability assumptions needed for the first estimation method to work.

¹¹Observe that defining the market clearing price in this way enables me to avoid technical problems associated with having to ration quantities when demand exceeds supply at the clearing price. This might be regarded as an unnatural assumption in an auction environment where rationing occurs frequently. My conversations with auction insiders, however, indicate that rationing has not been a big concern in the Turkish market

3.1 A resampling strategy

In this section, I will use the notation of section 2.4, since the method discussed will be applicable in the (realistic) case where the “bid functions” in the data set are actually discrete price-quantity pairs. Recall that in this notation, the price grid is set to be $p_0 < p_1 < \dots < p_{K+1}$, and the bids take the form of a quantity vector specified for each of these prices: $\vec{y}_i : \{y_{i0} \geq y_{i1} \geq \dots \geq y_{iK+1}\}$.

One way to approach the problem of estimating the unobserved marginal valuations through the optimality condition (2) is to ask how bidder i would calculate $H(p_k, \vec{y}_i)$, the probability that the market clearing price is less than p_k , conditional on bidder i 's bid vector, \vec{y}_i . In a Bayesian Nash equilibrium of the game, bidder i knows other bidders' bid vectors up to their signals, $s_j, j \neq i$. Hence, in equilibrium, the residual supply curve that bidder i faces is a function of $N - 1$ random variables, whose distribution is common knowledge among bidders. But bidder i could simulate the gamut of residual supply curves he is likely to encounter by repeating the following procedure many times: First, generate a random draw of the $N - 1$ element signal vector of the competing bidders, $j \neq i$. Then, evaluate the $N - 1$ equilibrium “opponent” bids $y_k(s_j)$ corresponding to these signals. Using these, form the residual supply function, $RS_k = Q - \sum_{j \neq i} y_k(s_j)$. Then calculate the market clearing price by finding the intersection of the residual supply function with own bid vector.

After several thousand iterations, the above procedure will get the bidder a very good approximation to $H(p_k, \vec{y}_i)$. Unfortunately, the distribution of bidders' signals is not observed by the empirical researcher. More importantly, the exact mapping from the signals, s_i to the equilibrium bid vectors \vec{y}_i is generally not available in closed form. Therefore, a “direct estimation strategy” that explicitly inverts the equilibrium mapping from the signals to the bid vectors to form a likelihood function requires significant computational effort.¹²

Fortunately, the following “resampling” procedure, illustrated in figure 3 and figure 4, goes a great distance towards simplifying the estimation problem:

1. Fix bidder i among the total N_t bidders in auction t . (In figure 3, this is bidder 2 in my data set).
2. From the sample of N_t bid vectors in the data set, draw a random sample of $N_t - 1$ bid vectors

¹²This is not to say that such a strategy is infeasible or undesirable. There are several examples of such a direct, likelihood-based approach to estimate structural models of *single-unit* auctions, such as Paarsch (1992), Laffont, Ossard and Vuong (1995), Bajari and Hortacısu (2000).

with replacement, giving equal probability of $\frac{1}{N_t}$ to each bid vector in the original sample.

3. Construct the residual supply function generated by these $N_t - 1$ “resampled” bid vectors. (Figure 3 shows about 30 “resampled” residual supply functions.)
4. Intersect with bidder i ’s bid to find the market clearing price.
5. Repeat B (a large number) times for each bidder and for all bidders in the data set.

This generates B market clearing prices *conditional* on \vec{y}_i (these are the points of intersection on figure 3). One can then estimate $H(p_k, \vec{y}_i)$ by counting the frequency with which a given p_k remained above the market clearing prices generated above: this is what I do in figure 4, in which I superimpose the histogram of market clearing prices, which resulted from 5000 resampled draws of the residual supply curve, horizontally on top of bidder #2’s “staircase” bid vector (the histogram is a density estimate, but the distribution can be easily recovered from it).

Clearly, the above resampling procedure relies heavily on the following assumption:

Assumption 1. *Bids in the data are generated by a symmetric pure strategy equilibrium of the independent private value discriminatory price auction. That is, $\vec{y}_i(s_i) = \vec{y}(s_i)$, with s_i , and consequently, $\vec{y}(s_i)$, i.i.d.*

I will discuss how this assumption can be relaxed in section 6. However, taking this assumption as given for now, let us now investigate statistical properties of this resampling method. Call the above estimator of $H(p_k, \vec{y}_i)$, $\hat{\mathbf{H}}^{\mathbf{R}}(p_k, \vec{y}_i)$, the *resampling estimator* of the probability distribution of the market clearing price. The following result can be proved about the properties of the resampling estimator:

Proposition 1. *Suppose our data consists of T auctions and for the t -th auction, there are N_t bid vectors $\{\vec{y}_t(s_1), \dots, \vec{y}_t(s_{N_t})\}$ generated by N_t i.i.d. draws from the distribution of private bidder signals, where $\vec{y}_t(s)$ denotes the symmetric pure strategy equilibrium of the t -th discriminatory price auction with N_t bidders. Then:*

1. *If the bids in different auctions are generated from the same distribution (as is the case when the auctions are repetitions of the same experiment), then, as $L = NT \rightarrow \infty$, $\hat{\mathbf{H}}^{\mathbf{R}}(p_k, \vec{y}_i)$ converges to $H(p_k, \vec{y}_i)$ almost surely.*

2. If we use data from a single auction, that is, set $T = 1$, but let the number of bidders, N , go to infinity, then $\hat{\mathbf{H}}^{\mathbf{R}}(p_k, \vec{y}_i)$ is a consistent estimator of $H(p_k, \vec{y}_i)$ if $E\vec{y}(s)$ can be estimated consistently at a rate faster than \sqrt{N} . This can be done by a two step procedure in which we first compute $\hat{\mathbf{H}}^{\mathbf{R}}(p_k, \vec{y}_i)$ using data from a single auction, and then correct this estimate using an estimator of $E\vec{y}(s)$ that uses data across auctions (hence convergence is at rate \sqrt{NT} where T can be taken to infinity independent of N).

Appendix 8.3 offers a detailed argument for the first claim. The basic intuition for this claim comes from viewing resampling as a way to make random draws from the empirical distribution of bid vectors. By the Glivenko-Cantelli Theorem, the empirical distribution converges (almost surely) to the true distribution as the sample size grows large. The assumption is that auctions take place in a static environment where the information commonly observed by the bidders stays the same, but bidders draw different signals (conditional on commonly observed information) in each repetition of the auction. This assumption can be realistic in certain divisible good auctions settings like electricity markets, in which the auction is run every day, and day to day variation in the commonly observed information available to bidders can be ignored (for example, data from days with observably similar weather can be pooled).

However, in the Treasury auction setting, it is reasonable to expect that the economic environment inhabited by the bidders changes drastically from auction to auction. For example, the summary statistics in Table 2 show that the number of bidders in the auction and the supply of securities vary greatly across auctions in my data set. Therefore, as in the proposed algorithm above, one might want to use data from only a single auction.

The second claim in the proposition suggests that this might not be a very bad idea. As mentioned in section 2.3, the market clearing price is a statistic that aggregates bidders' signals. Hence the resampling algorithm can be interpreted as "bootstrapping" the market clearing price about its observed value (\hat{p}^c). Letting p^{c*} denote the "bootstrapped" market clearing price generated in each iteration of resampling algorithm above, and letting $E(p^c)$ denote the population mean of the market clearing price statistic, standard bootstrap consistency results as in van der Waart (1998) and Hall (1992) suggest that $p^{c*} - \hat{p}^c \stackrel{D}{=} \hat{p}^c - E(p^c)$ as the number of bidders in the auction grows large. That is, the distribution of the bootstrapped market clearing price around the realized value of the market clearing price is a consistent estimate of the distribution of the realized market clearing price about

its population mean.

Appendix 8.4 makes the above intuition much more explicit to establish the “bootstrap consistency” of the market clearing price statistic. Establishing this result requires a few additional steps on top of a standard demonstration of bootstrap consistency, since the market clearing price is an “equilibrium” statistic whose distribution changes as the number of bidders in an auction changes.

Appendix 8.4 also clarifies the need for the “second-step correction” in the case where we resample data from a single auction. Intuitively, we need this correction to “recenter” the bootstrapped distribution of the market clearing price about the true mean of the market clearing price, $E(p^c)$, rather than about the observed value of the market clearing price, \hat{p}^c . However, as Appendix 8.4 shows, this recentering can be done consistently only if \hat{p}^c converges to $E(p^c)$ at a rate faster than \sqrt{N} . This can be achieved by using a separate estimate of the mean market clearing price that relies on the full sample of auctions; this allows T to be taken to infinity independent of N .

Although the centering problem can be fixed by using data from the full set of auctions, one might wonder how large an estimation bias we might expect when we forego the second step correction. In Appendix 8.5, I conduct a Monte-Carlo experiment to assess the performance of the resampling scheme when data from a single auction is used. In this experiment, I use data from a simulated auction for which I know (or can compute) the true marginal valuations, and see how estimated marginal valuations using data from only a *single* auction matches up with the known true marginal valuations of the bidders. I report that even without the recentering correction, the resampling algorithm that uses data from a single auction performs quite well.

In Appendix 8.6, I discuss a “jackknife-after-bootstrap” method to assess the standard errors of $\hat{\mathbf{H}}^{\mathbf{R}}(p_k, \vec{y}_i)$, and $\hat{v}(y_{ik}, s_i)$.¹³

3.2 A parametric strategy

The resampling strategy described in the previous section has definite advantages in terms of its minimal use of distributional assumptions; however, it also relies quite heavily on the assumption that bid functions are i.i.d. I relax this assumption somewhat in section 6.1, where asymmetries will be handled by assuming the sample can be broken down to a small number of i.i.d. sub-samples.

¹³I have also developed a more standard “delta-method” approach to compute standard errors, a description of which is available to interested readers upon request. In practice, these two methods yielded very similar answers.

However, it is difficult to extend the resampling strategy in a way, for example, to control for bidder-level fixed-effects.

Another practical concern with the resampling strategy outlined in the previous section is the lack of precise estimates in the case where the number of bidders, N_t , is small. Clearly, one would not want to perform this algorithm if, for example, $N_t = 3$. Therefore, there is need for an estimation strategy that takes advantage of the panel nature of the data set.

The following is an alternative estimation scheme that addresses some of these concerns at the cost of imposing somewhat stronger functional form restrictions on the data. First we reduce the dimensionality of the data by, for example, fitting a linear function to each observed bid function, $y(p, s_{it})$.¹⁴ This linear fit yields a series of estimated slope and constant coefficients of the form: $\{(\hat{\alpha}_{it}, \hat{\beta}_{it}, i = 1, \dots, N_t; t = 1, \dots, T)\}$; where $\hat{\alpha}_{it}$ is the estimated constant and $\hat{\beta}_{it}$ the estimated slope of bidder i 's linearized bid function in auction t .

Observe that assuming $y(p, s_{it}) = \alpha_{it} + \beta_{it}p$ yields the following expression for $G_i(p, x)$ in equation (3):

$$G_{it}(p, x) = \Pr\left\{x \leq Q_t - \sum_{j \neq i}^{N_t} y(p, s_{jt})\right\} \quad (9)$$

$$= \Pr\left\{x \leq Q - \sum_{j \neq i}^{N_t} \alpha_{it} + p \sum_{j \neq i}^{N_t} \beta_{it}\right\} \quad (10)$$

Given this expression the following procedure can be used to estimate the market clearing price distribution that bidder i faces conditional on demanding x units of the security:

1. Estimate the joint distribution of β_{it} and α_{it} .
2. Make simulation draws from this estimated distribution, and use these simulation draws to calculate the probability in equation (10). Or, equivalently, given an x , simulate the distribution of the (expected) market clearing price distribution.

It is evident that asymmetries could be handled by allowing for bidder-level covariates or fixed effects to enter into the specification of β_{it} and α_{it} . However, in order to use the panel structure of the

¹⁴In principle, any finite order polynomial or basis-function representation could be used. The attention to the linear approximation is motivated by its relative ease of exposition, and its success in capturing an average 92% of the variation in the actual data.

data, one has to make further, and potentially quite strong, assumptions regarding the way auction-level covariates enter the specification for β_{it} and α_{it} . Furthermore, it is quite conceivable that not all relevant covariates affecting the bidding decisions of the bidders are observed by the econometrician, leading to potential omitted-variable bias. Also, by linearizing the bid functions, or approximating them in some other way, one introduces additional error into the estimation. One should also account for this approximation error in the resulting estimates for $G_i(p, x)$.

My limited set of experiments with both of these estimation strategies has yielded very similar results in the empirical application about to be discussed. Therefore, to highlight what is novel, from this point on, I will focus almost entirely on results obtained using the resampling strategy. However, more extensive Monte Carlo experimentation juxtaposing the tradeoffs involved in employing the two different estimation strategies would be very beneficial from the standpoint of the practitioner.

4 The Turkish Treasury auction market and the data

The empirical application uses data from 13-week Treasury bill auctions conducted by the Turkish Treasury between October 1991 and October 1993. The data consist of price and quantity pairs for each bidder (coded by an ID number), the quantity of Treasury bills supplied, and the market clearing prices for the auctions.

The Turkish Treasury has auctioned government debt since May 1985. Short term (13 week to 52 week) securities offered by the Treasury are pure discount bills and bonds and do not bear coupons. The auctions are held on Wednesdays, rotating through 52, 39, 26 and 13 week bonds.¹⁵ Bidding is open to the general public; however the majority of bids come from banks, brokerages, and other financial institutions.¹⁶

The Turkish Treasury uses the discriminatory auction format. Each bidder is asked to specify a price and a quantity demanded at that price. Prices are quoted as the amount a bidder is willing to pay for an imaginary T-bill with face value of 100 Turkish Liras (TL). Prices can be specified up

¹⁵Although I do not have data on auctions of all maturities, the 13 week bills I focus on comprise about three-eighths of the bond issues of the period; therefore, the data set provides a good representation of the workings of the domestic debt market at the time.

¹⁶Financial institutions have to put up a collateral equaling 1% of the face value of their bids; the general public is subject to a 100% collateral requirement, causing most private parties to bid through their broker or banker.

to 3 significant digits. Quantities are specified in terms of the face value of T-bills the bidder wants to buy. The minimum quantity a bidder can request to buy is 50 million TL (about \$6000). There is no limit to the number of price-quantity pairs submitted by a bidder. In fact, the average number of price-quantity pairs submitted in my sample is 6.9, with one bidder submitting over 60 price-quantity pairs per auction. Bids are submitted by noon on Wednesdays. Auction results are announced at 5:00 PM and winning bids are settled 1 week after the auction.¹⁷

In the period I study, October 1991 to October 1993, the Treasury followed two different procedures in the allocation of the bills. Until February 1993, the total quantity of bills sold in the auction was not announced until after the bids were submitted. However, bank managers that I have interviewed claimed they could estimate the supply of Treasury bills quite accurately by tracking debt service requirements of the Treasury, and through contacts in the Treasury (Table 5 shows that even a simple AR(1) regression can explain 81% of the variance in supply). Beginning February 1993, the Treasury began to pre-announce the quantity in its 3 and 6 month Treasury bill auctions in an attempt to commit to restricting the supply of short term securities and increase the average maturity of outstanding government debt.¹⁸

The main venue for trading in government securities, the Istanbul Stock Exchange Bonds and Bills Market (ISEBBM), was established in June 1991 as a computerized double-auction market. Using daily transaction data from the ISEBBM, I have calculated that, on average, 4% of the total volume of auctioned Treasury bills are traded on the ISE in a 3-day post-auction window. Based on this figure one could draw the conclusion that liquidity in the resale market is not very large compared to the primary auction market, especially before 1993. This would imply that the case for a common resale value of the auctioned securities is not as strong as it is believed in previous theoretical and empirical literature. I should note, however, that trading volumes on the ISEBBM showed a strong growth trend during the period. Further, ISEBBM was not the only resale market available to the bidders,

¹⁷There are three main differences between the institutional setup of Turkish Treasury auctions and U.S. Treasury auctions, the subject of most research in this area: First, there is no when-issued (forward) trading in the Turkish T-bill market. Second, unlike the U.S. Treasury auctions, non-competitive bids (bids that do not specify a price, and are automatically filled at the prevailing market-clearing price) are not allowed in the Turkish Treasury auctions. Third, bidders in the Turkish T-bill auctions are not subject to maximum quantity constraints. In the U.S. bidders are constrained to at most 35% of the auction issue.

¹⁸From the February 3, 1993 issue of *Cumhuriyet*.

who also participated in over-the-counter deals. Unfortunately, I do not have access to data on these transactions.

4.1 The bidders

Banks are the main players in the market, as they capture 93% of the bills sold in the sample. Brokerages buy 6%, and other bidders (institutional investors, insurance firms, pension funds) share the remainder. The top 5 bidders capture 30% about evenly, and the top 15 capture about 70% of the issues. However, the division of market share among these banks does not differ by more than 3%. This is about the same level of concentration as in the overall banking sector, in which top 15 banks own 65% of banking assets about evenly.

Participation varies widely among the total 134 participants observed in the auctions in the sample. 48 bidders entered more than 20 of the 27 auctions and 50 entered less than 10 auctions in total.

According to a survey of 51 banks and other financial institutions conducted by Alkan (1991) in late 1989, the main reasons for participating in the Treasury auctions (of all maturities) are as follows:

1. 42% of total purchases in the auctions are to meet the *liquid asset reserve requirements* monitored by the Central Bank. In the period I study, at least 30% of bank portfolios had to be held as government bonds and bills, with an average maturity of 210 days. Failure to comply with this requirement resulted in monetary fines and possibly a suspension of bank operations.
2. 37% of total purchases are for resale in the secondary market.
3. 10% of total purchases are to fill customer orders.
4. 10% of purchases are to fulfill collateral requirements, for investment funds administered by the bank, and for buy-and-hold purposes.

According to Alkan (1991), the relative importance of fulfilling reserve requirement vs. resale seems to vary a lot among survey respondents. Resale incentives might be more pronounced in 13 week T-bill auctions than in auctions of higher maturity securities. However, “buy-and-hold” strategies might also have been an important factor in the period I study, as short-term government securities yielded outstanding returns to investors.

Alkan (1991) also reports that bidders find the following bits of information as being useful in their bidding decision (listed in order of importance): Treasury’s borrowing requirements and repayment schedule, liquidity in money markets, the bidder’s own liquid asset reserve requirement and the reserve positions of other bidders, conversations with other bidders, and results of previous auctions.

Information on the Treasury’s borrowing requirement and repayment schedule can be used to reduce supply uncertainty. Banks who enter the auction to meet their liquid asset reserve requirements need this information as a critical input for assessing the probability of winning their minimum required quantity. Those who bid for speculative purposes need the information to estimate the resale price of the security following the auction.

Liquidity in money markets is an important determinant of the value of securities being auctioned. In the money market, Turkish banks routinely enter into “repo” and “reverse-repo” agreements with other banks or individuals, where they engage in short-term borrowing (“repo”) or lending (“reverse repo”), using Treasury securities as collateral.

Conversations among bidders up to minutes before the auction is a shared characteristic of Treasury auction markets around the world. Although one might be inclined to suspect collusion in the presence of such pre-auction communication, market participants and Treasury officials I have interviewed have expressed that they did not think collusion was an issue in this market. Therefore, in my analysis I will assume bidding is competitive.

As for the strategic and analytical sophistication of the bidders: since the government securities market is the largest organized financial market in Turkey, bidders expend significant resources to strategize. Personal interviews with managers of two mid-scale private banks have revealed that these banks have developed proprietary analytic and decision support software to aid their bidding decisions.

4.2 Description of the data

Table 1 displays various summary statistics for the auctions in the data set. Macroeconomic data at different frequencies was obtained from the Central Bank of Turkey.¹⁹ The total number of bidders

¹⁹The macroeconomic conditions during the sample period, October 1991–October 1993, were quite volatile, but not atypical of the overall macroeconomic environment of the past decade. The average annual inflation rate in this period was 68%. The budget deficit grew from 8% of GDP to 25% of GDP. I should also note that a major financial crisis and an ensuing recession occurred in Turkey circa May 1994.

corresponds to the total number of unique bidders showing up in the data set. Revenues are converted to US dollars using daily exchange rate data. Cover ratio is the ratio of the number of T-bills sold in the auction to the number of T-bills demanded by bidders. The auction yield is the quantity weighted average yield of the T-bills that were sold in the auction. To get an indication as to the amount of dispersion of opinion among bidders regarding the auction interest rate, I calculate the quantity-weighted variance of the price bids.

As we can see from Table 1, the ex-post *real* interest rates realized in the auction are very high (about 23%).²⁰ The average variance in bid prices translates to about 2% in annual yield. This bid spread is also very high compared to U.S. standards, where the spread is several basis points (one-hundredth of a percent).

The second observation from Table 1 is that the percentage of demand fulfilled in an auction varies widely. As noted before, in response to the shortening maturity structure of the debt, the Treasury began a conscious effort to limit the 3-month bill supply beginning in February 1993.

Next, I provide an empirical description of individual bidding behavior. In order to reduce the dimensionality of the bid data, I follow the strategy suggested in section 3.2: I fit a line through each bidder's multiple price-quantity pairs.²¹ The resulting fit is quite good: the average R^2 across all bidders and all auctions is 0.92.

The institutional setup suggests two sources for private information: the first is a bidder's information about his own inventory/liquidity needs, and the second is his private forecast of the secondary market price of the security. Unfortunately, I do not observe any of this information directly. However, I can construct a variable that proxies for the individual liquidity needs of the

²⁰While interpreting the real interest rates implied by the above table, one should remember that they are calculated ex-post, and do not take into account the risk premium associated with a highly variable inflation rate and a government with a growing debt burden. The default risk premium is probably not high enough to account for the overall premium, considering that the Turkish government was borrowing at an average 9% in dollar and DM denominated debt in European and U.S. markets in this period. The primary explanation for the high real interest rates should be sought in a large inflation/depreciation risk premium, which was partly justified by a devaluation of the Turkish lira in May 1994. In fact, Berument and Malatyali (1999) find empirical support for a significant inflation/devaluation risk premium by estimating a model of the Treasury auction rate that controls for expected inflation risk.

²¹Except I regress price on quantity, not vice versa as suggested in section 3.2. Therefore the intercept of a bid function is a measure of the "level" of individual demand, and the slope information can either be used to calculate the elasticity of the linear demand function, or as an indicator of size.

bidders by utilizing the panel structure of my data: using the total quantity demanded by each bidder in a given auction, and the quantity he actually ended up winning, I construct a variable called $SHORTFALL_{i,t-1}$ to denote the percentage of demand shortfall that the i -th bidder experienced in auction $t - 1$.²² Although bidders can cover acute shortfalls in their inventory need through the secondary market, since the secondary market was not very deep in this period, it is conceivable that a bidder might want to cover some of this shortfall in the next auction.

In table 2, I regress the fitted intercept and slopes of individual bidders demand functions on $SHORTFALL_{i,t-1}$ and $SHORTFALL_{i,t-2}$. In regressions (1) and (3), I use auction dummies. In regressions (2) and (4), I use auction and bidder dummies.

Regressions (1) and (2) indicate that the shortfall variable enters the regression in a strongly significant manner: suffering a shortfall in the previous auction is correlated with submitting a bid function with a higher price intercept in the next auction. Regression (1) manages to capture 61% of the variation in the fitted price intercept of individual bid functions. The R^2 from regression (2) indicates that allowing for bidder fixed effects does not add that much to the explanatory power of the regression. Setting the bidder coded as #1 as the base case, I found that of the 133 estimated bidder dummies, only 2 dummy variables were significantly different from zero. This indicates that bidder level asymmetries do not play a large role in determining the price level of demand.

Regressions (3) and (4) indicate that suffering a demand shortfall in the previous auction does not affect the slope of the bid functions in a statistically significant manner. Given that the R^2 of the regression with bidder dummies is significantly higher than the regression without, this makes us wonder whether bidder asymmetries are captured in the slope parameter. However, once again setting the bidder coded as #1 as the base case, I found that of the 133 estimated bidder dummies, only 8 dummy variables were significantly different from zero.

These results indicate that inventory/liquidity concerns are important determinants of bidding behavior in this market, and are an important component of private information that a bidder possesses before each auction. This is consistent with the institutional setup of the market. Therefore, I conclude that a private values model of the auction market can be a good benchmark specification.

Unfortunately, I do not have a good empirical test of the existence of a common value component in bidder's valuation functions. In single-unit auction models, this can be done using a test of the

²²Observe that this information is not available to competing bidders, and can be safely regarded as private information.

rational bid-shading due to the “winner’s curse” effect, which increases as the number of bidders in an auction increases.²³ Unfortunately, similar tests have not yet been developed for divisible good/multi-unit auction models.

5 Empirical Results

The previous section suggests that the market being analyzed is not far from satisfying the theoretical assumptions needed to take the empirical methods developed in section 3 to the data. Still, the benchmark symmetric independent private values specification might strike some readers as being overly restrictive. Therefore section 6 will discuss the robustness of the results under alternative specifications of the underlying model.

I will start by showing the results for the first auction in my data, the 3-month T-bill auction of October 1991. In this auction, there were 67 bidders for what amounted to 1959.5 billion Turkish liras worth of Treasury bills – about 400 million U.S. dollars. 58% of the bids that were submitted were successful. The market clearing price was 84.388 TL for a 100 TL face value 3 month T-bill.

I will focus on bidder #2, who submitted 14 price-quantity pairs totalling to a demand of 12% of the issue. As discussed in section 3.1, 5000 iterations of the resampling procedure yielded the market clearing price distribution that is plotted in figure 4. We see that all bids lie within the support of resampled market clearing price distribution. Looking at the figure, we also see that the high bid at 84.825 TL looks like an outlier among the next 13 bids, and that the probability of the market clearing price being above this bid is quite small. An explanation for this high bid is that the bidder has a really high value for the initial “step” of quantity. Such a high valuation makes sense when we consider that banks in Turkey have to satisfy a liquid asset reserve requirement which is monitored very closely by the Central Bank. If a bank can not win enough T-bills in the auction, then it has to buy securities in the resale market or in the following week’s auction to close its reserve shortfall. Since the resale market had much less volume than the primary market in this period, the bank runs the risk of paying a hefty fine to the government for falling short of the reserve requirement and possibly facing suspension of its operations.²⁴

²³See Paarsch (1992), Hendricks, Pinkse and Porter (1999), Haile, Hong and Shum (2000), Athey and Haile (2000) for discussions and various implementations of these tests.

²⁴In fact, in some banks, two different departments are in charge of preparing the bid vector: the department in charge

To calculate the marginal valuations of bidder #2 that correspond to each price-quantity pair submitted, I used the set of “corrected” first-order conditions derived in Appendix 8.2; rather than equation (2). However, in practice, I found that the correction terms derived in the Appendix for the discrete case to negligible in almost all instances. Given this, I conjecture that in most practical settings where the quantity of objects for sale are very large, equation (8) can be used as the estimating equation without much loss of accuracy. I used the unsmoothed empirical distribution of resampling draws of $H(p_k, \vec{y}_i)$ and computed standard errors using the “jackknife-after-bootstrap” method discussed in Appendix 8.6.

Figure 5 displays the estimation results for bidder #2, using 20,000 realizations of the “re-sampled” market clearing price. The horizontal axis is the percent of total supply that bidder #2 requested, with the prices on the vertical axis. Once again, we have the “staircase” representation of the bid vector. The points marked with a “v” correspond to the point estimates of the marginal valuation for the quantity bids, calculated using the equations above. I also plot the 5-95% confidence band around my point estimates.

Excluding the marginal value estimate for the first and last price-quantity pairs, marginal valuations are on average 0.16 TL (about 0.15% in annual yield terms) higher. The first (leftmost) bid, we see that the marginal valuation that rationalizes this is indeed quite high, in concordance with the intuition that banks with reserve price requirements may indeed value initial units very highly. The confidence intervals are wide for two intermediate quantities, where the point estimates of the marginal valuation seem to go against declining marginal value – given these wide confidence bands, the estimated marginal valuations can be safely seen as being declining in quantity.

5.1 Counterfactual comparisons

To perform a counterfactual revenue comparison between the uniform price auction and the discriminatory price auction, I first repeat my analysis for bidder #2 for every bidder in auction #1. This gives me a set of point estimates for the marginal valuations rationalizing each price quantity pair observed in this auction.

One way to do the counterfactual comparison is to use the point estimates of the marginal

of treasury operations puts in several bids to meet the reserve requirement, along with more “speculative” bids given by the trading desk.

valuations as the true marginal valuations of the bidders. This yields the perfectly competitive outcome of this auction, in which bidders reveal their true marginal valuations. The revenue of the auctioneer in the competitive outcome provides an upper bound to the revenue from the uniform price auction, since bidders typically exercise “demand reduction” in the uniform price auction. The competitive outcome also provides an upper bound to the revenue from the Vickrey auction.

Performing this counterfactual comparison correctly requires that certain practical issues be resolved. Since I estimate only certain points on the marginal value function, I have to account for the marginal valuations in between the points I estimate. I resolve this by restricting the marginal valuation to be the “upper envelope” of the pointwise estimates of the marginal valuations. In line with this proposition, I obtain the “upper envelope” of the step function interpolation of the marginal value estimates. As a robustness check, I also calculate the “lower envelope” of the point estimates of the marginal valuations. For the point estimates in Figure 5, the upper and lower envelopes are illustrated in Figure 6. With a general weakly decreasing marginal valuation function connecting these point estimates, the upper (lower) envelope provides an upper (lower) bound as to what the marginal value function should be. Comparing Figure 5 to Figure 6, we can see that except for the first bid, the envelope of the marginal valuation lies within the confidence intervals of the estimates.

Using the upper envelope for each bidder’s marginal valuation functions, I can calculate the revenue in the best-case uniform price auction by finding the market clearing price that would have resulted if the bidders had revealed their marginal valuations truthfully. To do this, I aggregate the upper envelope of the estimated marginal valuations across bidders, and find the intersection of this aggregate schedule with the total supply.

Figure 7 illustrates the result of this procedure for auction #1. In this figure, I overlay the aggregated bid schedule and the aggregated upper envelope of the marginal valuations. With supply normalized to 1 in this figure, the counterfactual uniform price auction revenue is given by the rectangle formed by the intersection of the aggregate marginal valuation schedule and total supply. Since the area under the aggregate bid schedule up to the total supply is greater than this rectangle, the actual discriminatory price auction revenue is greater than the counterfactual uniform price auction revenue. Specifically, I calculate that the market clearing price in the counterfactual “best-case” uniform price auction would have been 84.465 TL, as opposed to 84.388 TL in the discriminatory price auction. The revenue in the counterfactual case would be 1.3271 trillion TL, as opposed to the actual revenue of

1.3522 trillion TL generated by the discriminatory price auction, which is about a \$5.2 million revenue loss. In percentage terms, the counterfactual uniform price auction causes a revenue loss of 1.86%.

A more careful way to do the counterfactual revenue comparison is to remember that according to our model, the observed bids and the estimated marginal valuations are in fact random variables. Therefore, the auction revenue is also a random variable. Hence, a better comparison of the two mechanisms would be compare the expected revenues from each mechanism.

To compare the expected revenues under both mechanisms, I once again use a resampling procedure. For the discriminatory price auction, I draw 10,000 resamples from the original set of bid vectors and calculate the auctioneer’s revenue with the resampled set of bids. I do the same with the “best-case” uniform price auction, where the “bids” are the estimated marginal valuations.

Table 3 reports the results of the two counterfactual experiments for each auction in the data set. The first counterfactual experiment is the *ex-post* revenue comparison that treats the actual revenue in a given auction and the revenue that would have been generated in a uniform price auction in which the bids were the estimated marginal valuations. The second counterfactual experiment is the *ex-ante* revenue comparison that compares the expected revenues from both mechanisms.

We see that in all *ex-post* comparisons auctions, the “best case” uniform price auction would have yielded lower revenues for the Treasury. The average revenue loss would have been 3.8% of the total revenue, or about \$ 9 million per auction. According to these estimates, the cost of a switch to the uniform price auction would have been about 450 million dollars out of the 12 billion dollars of debt that was auctioned in this period!

The *ex-ante* comparison yields similar results for mean estimates of the revenue differential. However, the standard errors are high, therefore, on the basis of these results we can not reject a revenue equivalence result between the discriminatory price auction and the “best-case” uniform price auction.

The above revenue comparisons also say something about counterfactual revenues in Vickrey auctions. In a private values setting, it is well-known that the Vickrey auction achieves truthful revelation of marginal valuations.²⁵ Therefore one might be curious about the revenue performance of this mechanism. The answer to this question can be gleaned from a comparison of the payment rules in the uniform price and the Vickrey auctions: if bidders reveal their true marginal valuations

²⁵See Ausubel and Cramton (1998) for a discussion of the Vickrey auction.

in the uniform price setting, the revenue from the uniform price auction should be an *upper bound* to the revenue from a Vickrey auction, since the residual supply curve is increasing. Hence, the results in Table 3 suggest that the discriminatory price auction outperforms the Vickrey auction in terms of revenue.

6 Robustness of model assumptions and extensions

In this section I discuss the robustness of the empirical results if some of the model assumptions are relaxed. First, I describe a way to account for certain kinds of asymmetries and random participation decisions. I then investigate the effect of a common value specification – a familiar assumption used in modelling Treasury auctions and other financial markets. Finally, I extend the empirical procedure to account for uncertain supply.

6.1 Asymmetries among bidders and stochastic participation

The assumption that bidders are ex-ante symmetric plays a pivotal role in the resampling procedure. However, the data suggests that there can be potentially important asymmetries among bidders. Fortunately, certain kinds of asymmetries are easy to cope with in the resampling framework. If bidders can be grouped into $K \ll N$ classes, where bidders are ex-ante symmetric within a class, the resampling procedure can be modified to take separate resamples from each class.

Another closely related relaxation of model assumptions is to allow the number of bidders to be a random variable. If a stochastic process for N is specified, the resampling procedure can be modified to create resamples of differing sizes, where the resample size is governed by draws from the distribution of N .

To incorporate asymmetry and stochastic participation into my estimation method, I divide the bidders into three categories. Type 1 bidders are those who participated in more than 21 auctions. Type 2 bidders were those who participated in 11 to 20 auctions. Type 3 bidders were those who participated in less than 10 auctions. In terms of market shares, this categorization was also quite similar. Summary statistics for participation from different categories is given in Table 4.

When simulating the residual supply function using samplings from the bid functions, I first divided the bid functions into the categories they belonged to. Then I generated three random variables,

N_1, N_2, N_3 drawn from normal distributions with means and standard deviations given in Table 4 (and rounded to the next integer). To form my sample of bid functions for the residual supply function, I then drew N_1, N_2, N_3 bid functions from their respective categories.

Results under this set of modifications do not change the implication of the results in Table 2. I find that with the addition of asymmetries and stochastic participation, the estimated average ex-post revenue loss from a switch to the uniform price auction would be 4.5% of auction revenue. However, comparison of ex-ante expected revenues still does not reject revenue equivalence between the discriminatory auction and the “best-case” uniform price auction.

6.2 Specification bias in the presence of common values

An important extension to the above modelling framework is to allow for a common value component in the bidders’ marginal valuations – or, more generally, allowing marginal valuation functions to be of the form $v(q, s_i, s_{-i})$, where the signals may be correlated.

This extension comes with some cost to tractability, as the objective function (14) becomes:

$$\max_{y_i(\cdot)} \int_0^\infty \left\{ \int_0^{y_i(p^c)} [E_{s_{-i}|p^c, s_i} v(q, s_i, s_{-i}) - y_i^{-1}(q)] dq \right\} dH(p^c, y_i(p^c)) \quad (11)$$

This change is due to the fact that conditioning on a particular value of the market clearing price, p^c , enables the bidder to make an inference regarding the distribution of s_{-i} , since $p^c \equiv p^c(s_i, s_{-i})$ is an “equilibrium statistic” aggregating all bidder signals. The objective function suggests that what we can identify from the data is not $v(q, s_i, s_{-i})$, but the conditional expectation, $E_{s_{-i}|p^c, s_i} v(q, s_i, s_{-i})$.

Letting $v^E(q, s_i, p^c) = E_{s_{-i}|p^c, s_i} v(q, s_i, s_{-i})$ and evaluating the Euler-Lagrange condition for the above functional optimization problem results in:

$$v^E(y(p, s_i), s_i, p) = \underbrace{p + \frac{H}{H_p}}_{\text{Term A}} + \underbrace{\frac{H_y}{H_p} \int_0^{y(p, s_i)} \frac{\partial v^E(q', s_i, p)}{\partial p} dq'}_{\text{Term B}} \quad (12)$$

where Term A is the familiar necessary condition (2) in the private value setting, and H_p and H_y denote the partial derivatives of H with respect to its first and second arguments respectively.

The following lemma ascribes economic content to $v^E(q, s_i, p^c) = E_{s_{-i}|p^c, s_i} v(q, s_i, s_{-i})$:

Lemma. *In a Vickrey auction with the above common values setup, the bidding strategy, $y_i(p, s_i)$, is characterized by the following relation: $p = E_{s_{-i}|p, s_i} v(y_i(p, s_i), s_i, s_{-i})$. That is, for $y_i(p, s_i)$ units of*

the good, the bidder's price bid is his expected marginal value conditional on p being the market-clearing price of the auction.

Proof. In Appendix 8.7. □

Observe that the analogy to the single-unit Vickrey auction (which is the second price auction) with common values is clear: in the second price common value auction, an agent bids his expected value *conditional* on his bid being tied to win (Milgrom and Weber (1982a)). In the multi-unit case, the agent bids his expected value *conditional* on his bid being the market clearing bid. This “bid-shading” enables bidders to protect themselves from the “winner’s curse.”

Since H, H_p, H_y can all be estimated from the data, equation (12) can be used to estimate $v^E(y(p, s_i), s_i, p)$, what the bids would have been if this were a Vickrey auction. This is, of course, subject to the constraint that equation (12) is solvable. Instead of trying to solve this equation (which is a partial differential equation), I follow a different strategy: by imposing some additional structure on $v^E(y(p, s_i), s_i, p)$, I try to sign specification biases resulting from using a private values specification when the underlying the model is in fact characterized by common values.

A natural restriction on $v^E(y(p, s_i), s_i, p)$ in a common value environment (with affiliated s_i) is that it be increasing in p , since a higher ex-post value of the market clearing price conveys the information that other bidders have also received high signals. However, a formal demonstration of this intuitive statement requires that the conditional expectation is spelled out explicitly in terms of signal levels. Instead of carrying out this exercise, however, I will make the “reduced form” assumption that $\frac{\partial v(q, s_i, p)}{\partial p} \geq 0$ for all q .²⁶

Yet another intuitive argument leads to $H_y \leq 0$. The proof of this is not difficult, however: fixing a distribution of residual supply realizations, we see that if a bidder unilaterally increases her quantity demand at each price point, the probability distribution of the market clearing price distribution must shift up, reducing the probability that the market clearing price is below a given level p .

Given these conditions, the following claim follows from equation (12):

²⁶Justifying this assumption is not all that difficult in the discrete quantities case, since the market clearing price corresponds to a well-defined order statistic of the price bids submitted by the bidders, and affiliation of bidder signals can be used to prove that bidders’ conditional expected valuations are increasing in this order statistic, just as in Milgrom and Weber (1982b).

Proposition 2. *If the true underlying model is a common values model, the marginal value estimates obtained without accounting for the common value component yields an upper bound to what the bids would have been in a Vickrey auction.*

Observe that this result is also quite intuitive: equation (2) backs out the counterfactual bid in a Vickrey auction with private values (which is their true marginal valuation), whereas equation (12) backs out the bid in a Vickrey auction with common values. Since there is a winner’s curse effect in the common values case, we would expect bidders to shade their valuations.

Given this result, we would expect the ex-post revenue losses reported in Table 3 to be a *lower bound* to the revenue losses that could be expected from switching to a Vickrey auction. Hence the empirical result regarding the ex-post revenue comparison between the discriminatory auction and the Vickrey auction is robust. However, equation (12) does not say much about the counterfactual revenue in a common value *uniform price* auction – observe that the “best-case” bound constructed in the private values case can not be used in this case, as it is not immediately obvious that bidders will submit demand curves that are lower than the right hand side of equation (12) in a common value uniform price auction. However, intuition tells us that this is probably correct, since bidders’ incentive to shade their bids is more pronounced in the uniform price auction than in the Vickrey auction, since bidders have a strategic incentive to lower the market clearing price.²⁷

6.3 Effect of supply uncertainty

The results reported in the previous section do not take into account any supply uncertainty that bidders may be facing. However, the first 17 auctions in my data set are in a period in which the Treasury did not preannounce the total quantity of T-bills for sale. I will argue that the manner in which supply uncertainty is modelled has important effects to the estimation framework.

One way to model supply uncertainty is to view the Treasury as a strategic actor just like the bidders. In this case, the Treasury becomes a first-degree price-discriminating monopolist facing the “revealed” aggregate demand curve comprised of individual bids. Since the Treasury’s revenue is the area under the aggregate demand curve, the demand curve becomes the marginal revenue curve and the profit maximizing quantity is at the point where the aggregate demand equals marginal cost.

²⁷In fact Wilson (1979), Back and Zender (1993), and McAdams (1999) construct equilibria of common value uniform price auction games in which the auctioneer’s revenue can be zero!

Another way to incorporate supply uncertainty is to assume that the Treasury has limited discretion in setting the supply of bills to be allocated. Discretion is limited because the Treasury has to meet the requirements of its repayment calendar. Bidders have information about the targeted supply of debt in each auction through contacts with Treasury officials (my interviews indicate this is the case), but there is residual uncertainty about the exact amount of supply.

In this case, the total supply of T-bills in auction t , Q_t , becomes a random variable which shifts the residual supply faced by each bidder and enters into the estimation of marginal valuations. The estimation method now has to incorporate this additional random variable, which I model as having a time series structure. In Table 5, I fit an AR(1) process for $QSOLD$, the quantity of T-bills sold in auction. Alternative specifications for the supply process, including those with additional variables to proxy for the borrowing requirement of the Treasury do not yield more informative results.

For the 17 auctions under the supply uncertainty regime, I used the estimated AR(1) specification of Q_t to generate random draws for the total supply, and used these draws in my resampling procedure. The results, which I do not report in detail here, are very close to those reported in Table 3: the average ex-post revenue loss of the uniform price auction is 2.22% of the total revenue for the 16 auctions that were conducted in this regime. Ex-ante revenue comparisons, however, still fail to find a statistically significant revenue advantage of the discriminatory price auction over the best-case of the uniform-price auction. Hence I conclude that the presence of supply uncertainty, when modelled as an exogenous as opposed to a strategic decision by the Treasury, does not affect the qualitative nature of my results.

7 Conclusion

This paper has developed a method to identify and estimate bidder valuations from observed bid data under the assumption that bidders are acting according to a strategic model of bidding. The arguments leading to the identification of bidder valuations from bid data build upon the idea propounded in Elyakime et al. (1994) and Guerre et al. (2000), but extend the applicability of these ideas to the strategically more complex setting of divisible good/multi-unit auctions. Furthermore, the two estimation methods proposed here are, to my knowledge, novel, and rely very little on specific parametric assumptions. Perhaps more importantly from an application point of view, they are also

quite straightforward to implement in different multi-unit or divisible good auction settings. Aside from other Treasury auction systems around the world, data from auctions conducted in recently deregulated power exchanges around the world can also be analyzed to answer similar questions about mechanism choice.²⁸ Another market that can be analyzed is the national exchange for sulfur-dioxide emission permits which uses a discriminatory auction format.²⁹

The main finding of this paper relevant for policy debates surrounding the choice of mechanism in Treasury auctions is that revenue equivalence between a discriminatory auction and the “best-case” uniform price auction can not be rejected, at least within the empirical application considered. This is not an obvious theoretical result hinging on the assumption of independent private values. As shown quite generally by Krishna and Perry (1999), revenue equivalence applies to mechanisms that are allocationally equivalent. However, there are no results that demonstrate the allocational equivalence of the uniform price, discriminatory price, and Vickrey auction mechanisms.³⁰ In fact, a calculation using my estimates of bidders’ marginal valuations reveals that the discriminatory auction caused an average efficiency loss that was 11% of the total surplus obtainable from the auction. Hence, we can say that the discriminatory auction is not allocationally equivalent to a Vickrey auction, and thus we should not expect revenue equivalence of these two auction formats.³¹ What is somewhat surprising is that my empirical results suggest that the allocational inefficiency introduced by the discriminatory mechanism leads to higher revenues than the revenue obtainable from the efficient (Vickrey) mechanism.³²

The paper has also discussed three potential sources of misspecification (bidder asymmetries, supply uncertainty, and the presence of a winner’s curse) and has investigated the robustness of the results under these alternative specifications of the underlying strategic environment. However, an important issue that I have completely side-stepped when constructing the policy counterfactuals is

²⁸See Wolak (1999), Wolfram (1998) for analyses of strategic bidding in electricity markets.

²⁹See Joskow, Schmalensee and Bailey (1998) for a descriptive analysis of bidding behavior in this market.

³⁰As pointed out in the introduction, Ausubel and Cramton (1997) and Engelbrecht-Wiggans and Kahn (1998) provide counterexamples to this claim.

³¹Another theoretical benchmark to consider is the result by Swinkels (1999) on the asymptotic efficiency of private value discriminatory auctions. Swinkels’ result implies that as the number of bidders in the auction goes to infinity, the amount of inefficiency caused by differential bid-shading in the discriminatory auction goes to zero. My results indicate that $N = 70$ might not be enough to achieve efficiency.

³²Recall that the counterfactual uniform price auction revenue I obtained is an upper bound to the Vickrey auction revenue.

the possibility that changing the auction mechanism may affect bidders' participation decisions. Many proponents of uniform price or Vickrey auctions, including Friedman (1960), have suggested that more bidders are likely to participate in a uniform price or Vickrey auction, since optimal bid-shading, and the bid-preparation costs this may induce, becomes less of an issue under these mechanisms.³³ My conjecture, however, is that analyzing the effect of endogenous bidder participation decisions on the mechanism comparisons will be difficult unless one abandons the "indirect" empirical approach utilized here in favor of a "direct" approach in which the strategic equilibria under the actual and counterfactual auction mechanisms have to be explicitly solved for.³⁴ However, this involves the development of analytical or computational methods to solve for equilibrium strategies in standard divisible good auction games. This, I believe, is a research endeavor that is extremely interesting in its own right.³⁵

8 Appendix

8.1 Derivation of the optimality condition

We would like to solve the following functional optimization problem using calculus of variations (where I have dropped the i subscript from the demand function):

$$\max_{y(\cdot)} \int_0^\infty \left(\int_0^{y(p)} v(q, s_i) - y^{-1}(q) dq \right) dH(p, y(p))$$

First we have to get rid of the inner integral. Let the profit of the bidder from submitting the bid function $y(p)$ be

$$\pi(y(p)) = \int_0^{y(p)} v(q, s_i) - y^{-1}(q) dq$$

Now, setting $y(\infty) = 0$:

$$\pi(y(\infty)) = \pi(0) = 0$$

and

$$\begin{aligned} \frac{d\pi}{dp} &= (v(y(p), s_i) - y^{-1}(y(p)))y'(p) \\ &= (v(y(p), s_i) - p)y'(p) \end{aligned}$$

³³This intuition is not very precise for the uniform price auction, since although truthful revelation of marginal valuations is a weakly dominant strategy in the Vickrey auction, the same can not be said about the uniform price auction.

³⁴See Bajari and Hortacısu (2000) for an attempt along these lines in the context of single-unit Internet auctions.

³⁵Although analytic characterizations of strategic equilibria in uniform price divisible good auctions are provided in the seminal works of Wilson (1979) and Kyle (1989), very little has been done on discriminatory auctions, except for Viswanathan, Wang and Witelski (2001) and Hortacısu (2001), who construct analytical examples in the two bidder case, but have to resort to analytic approximations to analyze more general cases.

Substituting in the expression for $\pi(y(p))$ in the objective function and integrating by parts, I get

$$\max_{y(\cdot)} - \int_0^\infty (v(y(p), s_i) - p)y'(p)H(p, y(p))dp$$

Observe that the integrand is a function of p, y and y' , denote it by $F(p, y, y')$.

The Euler equation (which is a necessary condition for optimality) is given by (Kamien and Schwartz (1993)):

$$F_y = \frac{d}{dp}F_{y'}$$

Evaluating the derivatives, I get:

$$H(p, y(p)) + (p - v(y(p), s_i))H_p(p, y(p)) = 0$$

Solving this equation results in the optimality condition.

The sufficiency condition requires the integrand of the objective function to be (quasi-)concave in y and y' . Since the objective function can be shown to be linear in the latter, the sufficiency condition becomes (after substituting the first-order condition into the last term):

$$v_{yy}H + 2v_yH_y + \frac{H}{H_p}H_{yy} \leq 0 \quad (13)$$

where the subscripts denote the corresponding partial derivatives. Observe that, similar to the monopoly pricing analogy, the second-order condition is far from obvious. In the equilibrium setting in which H is treated as an endogenous quantity, sufficiency requires that the primitives of the model (the functional form of the marginal valuation functions and the distribution of the signals) be specified so as to satisfy the above condition.³⁶ Since H is a function that can be estimated from observed data, the sufficiency condition can, at least in principle, be checked using the empirical analogs of the quantities in (13), hence offering a test of the theory.

8.2 The discrete model

As in the text, I will maintain the divisibility of the quantities, but restrict the prices to lie on a discrete grid. In particular, following the model of Nautz (1995), let $p_0 < p_1 < \dots < p_{K+1}$ denote the set of possible prices on the grid. Hence, the “bid vector” submitted by each bidder takes the form of a series of quantities specified for each of these prices: $\vec{y}_i : \{y_{i0} \geq y_{i1} \geq \dots \geq y_{iK+1}\}$. In reality, bidders do not submit a different quantity bid for each price increment, so to account for this, let $\{y_{ik^1} > y_{ik^2} > \dots > y_{ik^L}\}$ be the set of L “observed” quantity bids that the bidder submits on the price points $\{p_{k^1} < p_{k^2} < \dots < p_{k^L}\}$, where $\{p_{k^1} < p_{k^2} < \dots < p_{k^L}\}$ is a subset of the entire price grid, $\{p_0 < \dots < p_{K+1}\}$.

Figure 8 illustrates the notation: in this figure, there are 10 price points that the bidder can place a quantity bid on, p_0, \dots, p_9 . However, instead of submitting a different quantity bid for each price point, the bidder has given 3 price-quantity pairs, $\{(p_1, y_{i1}), (p_4, y_{i4}), (p_7, y_{i7})\}$. In the figure, I mark the “observed” price-quantity pairs as the solid circles, and use the notation above, i.e. $p_1 = p_{k^1}, y_{i1} = y_{ik^1}$, etc.

³⁶If we make the assumption that $H_y = H_{yy} = 0$, which amounts to assuming that bidders are price-takers (their quantity bids will not affect the probability distribution of the market clearing price) the sufficiency condition reduces to requiring $v_{yy} \leq 0$ – that is the demand functions of the bidders should be concave.

Given this setup, the expected payoff of a risk neutral bidder who submits the bid vector $\vec{y}_i = \{y_{i0}, y_{i1}, \dots, y_{iK+1}\}$ will be:

$$\begin{aligned} & \sum_{k=1}^K [\Pr\{p_{k-1} \leq \text{mkt. clr. price} \leq p_k, \text{ given } \vec{y}_i\}] \times \{\text{Profit on bids above } p_{k-1}\} \\ &= \sum_{k=1}^K [H(p_k, \vec{y}_i) - H(p_{k-1}, \vec{y}_i)] \times \sum_{j=k}^K \left(\int_{y_{ij+1}}^{y_{ij}} v_i(q, s_i) dq - p_j (y_{ij} - y_{ij+1}) \right) \end{aligned} \quad (14)$$

Now, observe that after an interchange of summations, and the addition of the monotonicity constraints, the Lagrangian of the objective function of the bidder given in equation (14) can be written as:

$$\mathcal{L} = \sum_{k=1}^K H(p_k, \vec{y}_i) \left(\int_{y_{ik+1}}^{y_{ik}} v_i(q, s_i) dq - p_k (y_{ik} - y_{ik+1}) \right) + \lambda_k (y_{ik} - y_{ik+1})$$

The first-order conditions for a maximum are, for each k :

$$\begin{aligned} H(p_k) [v_i(y_{ik}, s_i) - p_k] &+ \frac{\partial H(p_k)}{\partial y_{ik}} \left(\int_{y_{ik+1}}^{y_{ik}} v_i(q, s_i) dq - p_k (y_{ik} - y_{ik+1}) \right) + \lambda_k = \\ H(p_{k-1}) [v_i(y_{ik}, s_i) - p_{k-1}] &+ \frac{\partial H(p_{k-1})}{\partial y_{ik}} \left(\int_{y_{ik}}^{y_{ik-1}} v_i(q, s_i) dq - p_{k-1} (y_{ik-1} - y_{ik}) \right) + \lambda_{k-1} \end{aligned}$$

In reality, I see distinct quantity bids at only a subset of the possible price points: $\{y_{ik^1} > y_{ik^2} > \dots > y_{ik^L}\}$. I will interpret these observed quantity bids as having been submitted at price points at which the monotonicity constraint $y_{ik} \geq y_{ik+1}$ is not binding, i.e. where $\lambda_k = 0$. For all other k , i.e. at price points where the quantity request remains the same, I will assume that the monotonicity restriction is binding, so $\lambda_k > 0$.

To ease notation for the rest of the analysis, define:

$$A_{k^{m+1}} = \frac{\partial H(p_{k^{m+1}}, \vec{y}_i)}{\partial y_{ik^{m+1}}} \left(\int_{y_{ik^{m+2}}}^{y_{ik^{m+1}}} v_i(q, s_i) dq - p_{k^{m+1}} (y_{ik^{m+1}} - y_{ik^{m+2}}) \right)$$

and

$$B_{k^{m+1}} = \frac{\partial H(p_{k^m}, \vec{y}_i)}{\partial y_{ik^{m+1}}} \left(\int_{y_{ik^{m+1}}}^{y_{ik^m}} v_i(q, s_i) dq - p_{k^m} (y_{ik^m} - y_{ik^{m+1}}) \right)$$

Now, write the first-order conditions for all price points j such that $k^m + 1 \leq j \leq k^{m+1}$, and adding them from $k^m + 1 \leq j \leq k^{m+1}$, we see that the Lagrange multipliers conveniently cancel out from successive equations and the overall sum telescopes. Since for $k^m + 1 \leq j < k^{m+1}$, $y_{ij} = y_{ij+1}$, the integral terms vanish in A_{k^j} and B_{k^j} for such j . Also, what is left of A_{k^j} and B_{k^j} cancel out, since $\frac{\partial H}{\partial y_{ik^j}}$ is the same across consecutive A_{k^j} and B_{k^j} terms. Hence, we are left with:

$$H(p_{k^{m+1}}, \vec{y}_i) [v_i(y_{ik^{m+1}}, s_i) - p_{k^{m+1}}] + A_{k^{m+1}} = H(p_{k^m}, \vec{y}_i) [v_i(y_{ik^{m+1}}, s_i) - p_{k^m}] + B_{k^{m+1}}$$

where I have used the fact that $y_{ik^{m+1}} = y_{ik^{m+1}}$, $y_{ik^{m+1}+1} = y_{ik^{m+2}}$.

Observe that this is an expression entirely in terms of the observed bids $\{y_{ik^1} > y_{ik^2} > \dots > y_{ik^L}\}$.

Rewriting this to solve for $v_i(y_{ik^{m+1}}, s_i)$, we get:

$$v_i(y_{ik^{m+1}}, s_i) = p_{k^{m+1}} + \frac{H(p_{k^m}, \vec{y}_i) [p_{k^{m+1}} - p_{k^m}]}{H(p_{k^{m+1}}, \vec{y}_i) - H(p_{k^m}, \vec{y}_i)} + \frac{B_{k^{m+1}} - A_{k^{m+1}}}{H(p_{k^{m+1}}, \vec{y}_i) - H(p_{k^m}, \vec{y}_i)}$$

Now, assuming that the marginal valuations are given by a step function which assume constant values on $(y_{ik^m}, y_{ik^{m-1}})$, the integrals in $A_{k^{m+1}}$ and $B_{k^{m+1}}$ can be evaluated and these terms become:

$$A_{k^{m+1}} = \frac{\partial H(p_{k^{m+1}}, \vec{y}_i)}{\partial y_{ik^{m+1}}} \{[v_i(y_{ik^{m+2}}, s_i) - p_{k^{m+1}}](y_{ik^{m+1}} - y_{ik^{m+2}})\}$$

$$B_{k^{m+1}} = \frac{\partial H(p_{k^m}, \vec{y}_i)}{\partial y_{ik^{m+1}}} \{[v_i(y_{ik^{m+1}}, s_i) - p_{k^m}](y_{ik^m} - y_{ik^{m+1}})\}$$

Now let's analyze the "boundary condition" for the bid with the lowest price, y_{ik^1} : this bid implies that the bidder is also willing to accept y_{ik^1} units of the security at the minimum price on the price grid, p_0 . I will assume that this implied bid is also an "observed bid," where $y_{i0} = y_{ik^1}$. Therefore, the first-order relation in equation (8.2) applies. But, since $y_{i0} = 0$, $B_{k^1} = p_0 y_{k^1}$. Since bidders can submit any positive price, I can safely set $p_0 = 0$, so this makes $B_1 = 0$. As for A_{k^1} , we can calculate this like the other A_{k^j} 's, so $A_{k^1} = \frac{\partial H(p_{k^1}, \vec{y}_i)}{\partial y_{ik^1}} \{[v_i(y_{ik^2}, s_i) - p_{k^1}](y_{ik^1} - y_{ik^2})\}$.

Also, at this price, $H(p_0, \vec{y}_i) = 0$ by the assumption in the model setup. So:

$$v_i(y_{ik^1}, s_i) = p_{k^1} - \frac{\frac{\partial H(p_{k^1}, \vec{y}_i)}{\partial y_{ik^1}} \{[v_i(y_{ik^2}, s_i) - p_{k^1}](y_{ik^1} - y_{ik^2})\}}{H(p_{k^1}, \vec{y}_i)}$$

Observe that $\frac{\partial H(p_{k^1}, \vec{y}_i)}{\partial y_{ik^1}} \leq 0$ by the following intuitive argument: for every upward sloping residual supply curve that the bidder expects to face, increasing his quantity demand at a particular price point either increases the market clearing price or leaves it the same. Hence, the probability that the market clearing price is below a given price becomes less or equal than what it was.

As for the other "boundary condition" for the bid with the highest price, y_{ik^L} : this bid implies that the bidder is not willing to accept any units of the security for prices above p_{k^L} . Hence, the "implied bid" for price $p_{k^{L+1}}$ is $y_{ik^{L+1}} = 0$.

$$v_i(y_{ik^L}, s_i) = p_{k^L} + \frac{H(p_{k^{L-1}}, \vec{y}_i)[p_{k^L} - p_{k^{L-1}}]}{H(p_{k^L}, \vec{y}_i) - H(p_{k^{L-1}}, \vec{y}_i)} + \frac{B_{k^L} - A_{k^L}}{H(p_{k^L}, \vec{y}_i) - H(p_{k^{L-1}}, \vec{y}_i)}$$

where

$$A_{k^L} = -\frac{\partial H(p_{k^L}, \vec{y}_i)}{\partial y_{ik^L}} p_{k^L} y_{ik^L}$$

since $v_i(y_{ik^{L+1}}, s_i) = v_i(0, s_i) = 0$. Also:

$$B_{k^L} = \frac{\partial H(p_{k^{L-1}}, \vec{y}_i)}{\partial y_{ik^L}} \{[v_i(y_{ik^L}, s_i) - p_{k^{L-1}}](y_{ik^{L-1}} - y_{ik^L})\}$$

Hence, given estimates of $H(p_{k^m}, \vec{y}_i)$ and $\frac{\partial H(p_{k^m}, \vec{y}_i)}{\partial y_{ik^m}}$, we have a recursive set of L linear equations which we can solve for the "steps" of the marginal valuation function.³⁷

³⁷As for the second-order conditions: if $\frac{\partial H}{\partial y_{ik}} = 0$, then restricting $v_i(y, s_i)$ to be a decreasing function of y guarantees sufficiency. This assumption says that changing one's quantity bid does not shift the distribution of the market clearing price – which is the assumption in Nautz (1995). However, if a bidder commands a large fraction of the total demand, he will take into account the fact that the probability distribution of the market clearing price will shift with his bid. In this case, we also need $\frac{\partial H}{\partial y_{ik}} \leq 0$ and $\frac{\partial^2 H}{\partial y_{ik}^2} \leq 0$ to hold. We argued above that the first condition must be true. The second condition does not have a natural justification, therefore I will treat it as an additional assumption that could potentially be empirically tested.

What has been gained from this lengthy analysis? The expressions in obtained indicate that the monotonicity constraint on the bid functions can lead to the “coupling” of shading decisions on consecutive quantity bids on the price grid (which would not be the case in a straightforward discretization of equation (2)). That is, a bidder’s decision whether to submit a different quantity bid at a given price point also depends on his decision for the next price point. This is consistent with the analysis of Engelbrecht-Wiggans and Kahn (1998), who argue that in discriminatory auctions, the monotonicity constraint can cause price bids for consecutive quantities to be the same, even though the bidder might have unequal marginal valuations for these quantities.

A careful look at the above argument also reveals that we made perhaps too strong an assumption when assuming that the marginal valuation function is a step function. This leaves open the possibility that the marginal valuation function that “rationalizes” a given sequence of bids is not unique, and may cause concern regarding the empirical application. Section 5.1 outlines a method to address this problem by looking at the upper bound of allowable marginal valuation functions under the assumption that the marginal valuations are decreasing in quantity.

8.3 Proof of Proposition 1, claim 1

Proposition. *Suppose we observe L bid vectors $\{\vec{y}(s_1), \dots, \vec{y}(s_L)\}$ generated by L i.i.d. draws from the distribution of private bidder signals, where $y(s)$ denotes the symmetric pure strategy equilibrium of the discriminatory price auction with N bidders. Then, as $L \rightarrow \infty$, $\hat{H}^R(p_k, \vec{y}_i)$ converges to $H(p_k, \vec{y}_i)$ almost surely.*

Proof.

We first need the following result:

Claim 1. *If bidders are playing the symmetric Bayesian Nash equilibrium of the independent private value auction game, $H(p_k, \vec{y}_i) = \Pr\{p_{k^*} \leq p_k | \vec{y}_i\} = \Pr\{\sum_{j \neq i}^N y_k(s_j) \leq Q - y_k(s_i) | s_i\}$. That is, the probability distribution of the market clearing price can be represented as the probability distribution of the sum of $N - 1$ i.i.d. random variables.*

Proof. Since $k^* = \min\{k : \sum_{i=1}^N y_k(s_i) \leq Q\}$, $I\{p_k < p_{k^*}\} = I\{\sum_{i=1}^N y_k(s_i) > Q\}$, i.e. for any price strictly below the market clearing price, aggregate demand will strictly exceed supply. Taking the complement of this event, we get $I\{p_{k^*} \leq p_k\} = I\{\sum_{i=1}^N y_k(s_i) \leq Q\}$. Note that this argument assumes that y_k ’s are declining in k . \square

Henceforth, I will use the representation of $H(p_k, \vec{y}_i)$ in the above result. To simplify notation, I shall drop the dependence of the quantity bids on the signals, drop the component subscript, and refer to my data as the sample: $\{Y_1, \dots, Y_{N-1}\}$. Let $F(y)$ denote the probability distribution function that generated $\{Y_1, \dots, Y_{N-1}\}$, and in equilibrium, is *known* to the bidder. I will also replace the quantity $\frac{Q - y_k(s_i)}{N-1}$ by the constant c . Also, let $\bar{Y}_{N-1} = \frac{1}{N-1} \sum_{j=1}^{N-1} Y_j$, i.e. the $(N - 1)$ fold sample mean. Then

$$H(p_k, \vec{y}_i) = \Pr\{\bar{Y}_{N-1} \leq c\} \tag{15}$$

i.e. the distribution function of the sample mean of $\{Y_1, \dots, Y_{N-1}\}$. Now develop the “resampling” estimator for $H(p_k, \vec{y}_i)$ in the following way: Suppose we see L independent draws from the distribution $F(y)$, giving us the random variables $\{Y_1, \dots, Y_L\}$. This occurs, for example, if we see l repetitions of the same auction, with the signals of the bidders drawn

differently each time, giving us $L = lN$ data points in total. Draw a sample of $(N - 1)$ variables with replacement from $\{Y_1, \dots, Y_L\}$ with equal probability $\frac{1}{L}$ given to each sample point. Call this sample $\{Y_1^*, \dots, Y_{N-1}^*\}$. Then:

$$\Pr(Y_1^* \leq y_1, \dots, Y_{N-1}^* \leq y_{N-1}) = F_L(y_1) \dots F_L(y_{N-1}) \quad (16)$$

where

$$F_L(y_1) = \frac{1}{L} \sum_{i=1}^L I(Y_i \leq y_1) \quad (17)$$

i.e. the empirical distribution function generated by the sample $\{Y_1, \dots, Y_L\}$.

Let $\bar{Y}_{N-1}^* = \frac{1}{N-1} \sum_{j=1}^{N-1} Y_j^*$. The resampling estimator of $H(p_k, \bar{y}_i) = \Pr\{\bar{Y}_{N-1} \leq c\}$ is given by:

$$\hat{H}^R(p_k, \bar{y}_i) = \Pr\{\bar{Y}_{N-1}^* \leq c\} \quad (18)$$

To establish consistency, we would like to show the following:

Lemma 1. $\lim_{L \rightarrow \infty} \Pr\{\bar{Y}_{N-1}^* \leq c\} = \Pr\{\bar{Y}_{N-1} \leq c\}$ for all c if $F(y)$ is continuous.

Proof. Start with the conditional probability distribution function $\Pr\{\bar{Y}_{N-1}^* \leq c | Y_1, \dots, Y_L\}$. Define the conditional characteristic function³⁸ corresponding to this distribution function as:

$$E \left[\exp(it\bar{Y}_{N-1}^*) | Y_1, \dots, Y_L \right] = (E_{F_L}[\exp(\frac{itY_1^*}{N-1})])^{N-1} \quad (19)$$

$$= \left(\int \exp(\frac{ity}{N-1}) dF_L(y) \right)^{N-1} \quad (20)$$

By the Glivenko-Cantelli Theorem, $\lim_{L \rightarrow \infty} F_L(y) \stackrel{a.s.}{=} F(y)$, therefore

$\lim_{L \rightarrow \infty} \int \exp(\frac{ity}{N-1}) dF_L(y) \stackrel{a.s.}{=} \int \exp(\frac{ity}{N-1}) dF(y)$ at the points of continuity of F . Hence:

$$\lim_{L \rightarrow \infty} E \left[\exp(it\bar{Y}_{N-1}^*) | Y_1, \dots, Y_L \right] \stackrel{a.s.}{=} E \left[\exp(it\bar{Y}_{N-1}) \right] \quad (21)$$

By Levy's convergence theorem, this implies that $\lim_{L \rightarrow \infty} \Pr\{\bar{Y}_{N-1}^* \leq c | Y_1, \dots, Y_L\} \stackrel{a.s.}{=} \Pr\{\bar{Y}_{N-1} \leq c\}$ if $F(y)$ is continuous.

Now $\Pr\{\bar{Y}_{N-1}^* \leq c | Y_1, \dots, Y_L\}$ is bounded by 1, so by the bounded convergence theorem:

$$\lim_{L \rightarrow \infty} \Pr\{\bar{Y}_{N-1}^* \leq c\} \rightarrow \int_{\Omega} \lim_{L \rightarrow \infty} \Pr\{\bar{Y}_{N-1}^* \leq c | Y_1, \dots, Y_L\} d\omega = \Pr\{\bar{Y}_{N-1} \leq c\} \quad (22)$$

The lemma establishes the proposition.

8.4 Proof of Proposition 1, claim 2

The distribution of the market clearing price is equivalent to the distribution of the sum of $N - 1$ i.i.d. quantity bids.

Now, denoting the "resampled" quantity bids in step 2 of the procedure described above as $y_k^*(s_j)$, write:

$$Z_N = \frac{1}{N-1} \sum_{j=1}^{N-1} (y_k(s_j) - Ey_k)$$

$$Z_N^* = \frac{1}{N-1} \sum_{j=1}^{N-1} (y_k^*(s_j) - \bar{y}_k)$$

³⁸I thank Evarist Giné for suggesting the use of characteristic functions.

where Ey_k is the population mean of the quantity bids at price point k and \bar{y}_k is the sample mean.

Given this notation, the true distribution of the market clearing price, $H(p_k, \vec{y}_i)$ can be written as $\Pr\{Z_N \leq \frac{Q - y_k(s_i)}{N-1} - Ey_k\}$. The “resampling” estimate of the market clearing price distribution, $\hat{\mathbf{H}}^{\mathbf{R}}(p_k, \vec{y}_i)$ can be written as $\Pr\{Z_N^* \leq \frac{Q - y_k(s_i)}{N-1} - \bar{y}_k\}$.

Now, we have to prove:

Lemma 2. *Let F_{Z_N} be the d.f. (distribution function) of $(\sqrt{N-1})Z_N$ and $F_{Z_N^*}$ the d.f. of $(\sqrt{N-1})Z_N^*$. Then both distributions have the same normal limit as $N \rightarrow \infty$, given appropriate assumptions on the limiting behavior of equilibrium strategies in the N person game (which are made explicit below).*

Proof. To clear up the notation, drop the price indices, i.e. $Y = y_k$, and let $Y_j = y_k(s_j)$. Since Y_j are equilibrium bid vectors of an N -person game, the distribution of Y will depend explicitly on N . Denote this dependence by Y_{jN} , and let $\mu_N = EY_{jN}$ and $\sigma_N^2 = \text{Var}(Y_{jN})$. I impose the following conditions on the limiting behavior of these quantities:

1. Assumption 1: $\lim_{N \rightarrow \infty} \mu_N = \mu$.
2. Assumption 2: $\lim_{N \rightarrow \infty} \sigma_N^2 = \sigma^2 > 0, \sigma^2 < \infty$

Assumption 2 has more economic meaning than Assumption 1. It is satisfied if, even with infinitely many bidders, the bidders do not forego their private information when formulating their strategies, i.e. bidders do not submit identical bid functions. This might be problematic when Q is fixed and the support of the marginal valuation distribution is finite: as N goes to infinity, one might expect the distribution of the price bids to converge to the upper limit of the marginal valuation distribution. One way to get around this is to assume that the supply, Q , also grows fast enough with N . This is a reasonable assumption, since, with positive entry costs, one would expect the number of bidders to hit an upper bound with fixed supply.

Given these assumptions, we can start by showing that $F_{Z_N} \rightarrow \Phi_\sigma$. Since the distribution of Z_N changes with N , we can not use the usual CLT, but we can use the central limit theorem for triangular arrays (van der Waart (1998)). To use this theorem, we have to check the Lindeberg condition:

$$\sum_{j=1}^{N-1} E\left(\frac{Y_{jN}}{\sqrt{N-1}}\right)^2 I\left\{\left|\frac{Y_{jN}}{\sqrt{N-1}}\right| > \varepsilon\right\} = EY_{jN}^2 I\{|Y_{jN}| > \varepsilon\sqrt{N-1}\} \stackrel{a.s.}{\rightarrow} 0$$

for all ε , as $N \rightarrow \infty$. Therefore, by the triangular array CLT, $F_{Z_N} \rightarrow \Phi_\sigma$.

Now we have to show that $F_{Z_N^*}$ has the same normal limit. Let \mathcal{P}_N be the empirical distribution of Y_{jN} 's; i.e., the distribution that Y_{jN}^* 's are resampled from. Before we check the Lindeberg condition, let's show that:

$$\begin{aligned} E[Y_{jN}^* | \mathcal{P}_N] &= \sum_{j=1}^N \frac{1}{N} Y_{jN} \\ &= \bar{Y}_N \\ E[(Y_{jN}^* - \bar{Y}_N)^2 | \mathcal{P}_N] &= \sum_{j=1}^N \frac{1}{N} (Y_{jN} - \bar{Y}_N)^2 \\ &= \frac{1}{N} \sum_{j=1}^N Y_{jN}^2 - \left(\frac{1}{N} \sum_{j=1}^N Y_{jN}\right)^2 \end{aligned}$$

which converge to μ and σ by the law of large numbers. Now let's check the Lindeberg condition to be able to use the triangular array theorem to establish asymptotic normality:

$$\begin{aligned} \sum_{j=1}^{N-1} E \left[\left(\frac{Y_{jN}^*}{\sqrt{N-1}} \right)^2 | \mathcal{P}_N \right] I \left\{ \left| \frac{Y_{jN}^*}{\sqrt{N-1}} \right| > \varepsilon | \mathcal{P}_N \right\} &= E \left[Y_{jN}^{*2} | \mathcal{P}_N \right] I \{ |Y_{jN}^*| > \varepsilon \sqrt{N-1} | \mathcal{P}_N \} \\ &= \sum_{j=1}^{N-1} \frac{1}{N-1} Y_{jN}^2 I \{ |Y_{jN}| > \varepsilon \sqrt{N-1} \} \\ &< \frac{1}{N-1} \sum_{j=1}^{N-1} Y_{jN}^2 I \{ |Y_{jN}| > M \} \end{aligned}$$

for some N . By the law of large numbers, the last term converges to $\lim_{N \rightarrow \infty} E \left[Y_{jN}^2 I \{ |Y_{jN}| > M \} \right]$, which can be made arbitrarily small by the choice of M . Hence $\sum_{j=1}^{N-1} E \left[\left(\frac{Y_{jN}^*}{\sqrt{N-1}} \right)^2 | \mathcal{P}_N \right] I \left\{ \left| \frac{Y_{jN}^*}{\sqrt{N-1}} \right| > \varepsilon | \mathcal{P}_N \right\}$ can be made arbitrarily small for any ε by an appropriate choice of N . \square

A practical problem with the above is that an additional error will be introduced by the fact that we have to rely on the *sample* mean of y_k rather than the *population* mean Ey_k when centering the distribution estimate. That is, we should be calculating $\Pr \{ Z_N^* \leq \frac{Q - y_k(s_i)}{N-1} - Ey_k \}$, not $\Pr \{ Z_N^* \leq \frac{Q - y_k(s_i)}{N-1} - \bar{y}_k \}$. Observe that taking N to infinity does not help here: \bar{y}_k converges to Ey_k at rate \sqrt{N} , so scaling Z_N^* by \sqrt{N} outdoes this \sqrt{N} convergence, since we have to scale $Ey_k - \bar{y}_k$ by the same amount.

One way to get around this centering problem for the purpose of establishing asymptotic validity is to estimate Ey_k at a rate faster than \sqrt{N} . To do this, we could use the panel nature of the data, and estimate Ey_k using bids across auctions.

8.5 A simulation study

To investigate the performance of the resampling scheme proposed above, I conduct a simulation study in which I compare estimated marginal valuations with the known “true” marginal valuations of the bidders.

Unfortunately, as mentioned in the previous section, explicit computation of equilibrium strategies for the discriminatory price auction game is very difficult. Therefore, I took the approach of assuming a particular functional form for the equilibrium strategies, and constructed the “true” valuations that would rationalize these strategies by the simulation procedure outlined for the bidder in section 3.

Inspection of the data reveals that bid vectors in the price quantity plane can be interpolated quite accurately with linear functions. Therefore, for my experiment, I took a market with $N = 70$ bidders who I assumed to be following symmetric linear bid strategies:

$$y(p_k, s_i) = \alpha + \beta s_i - \gamma p_k \tag{23}$$

I allowed the bidders to submit quantity bids on 20 price points regularly spaced between a minimum price of 0.815 and a maximum of 0.8654.³⁹ I set $\alpha = 0.71, \beta = 1, \gamma = -0.83$ and $Q = 1$, and I specified the private signals to be i.i.d. normally distributed with zero mean and standard deviation 0.015. With the given constants, the resulting bid functions are roughly calibrated to resemble linear approximations to the bid functions in the data set.

³⁹In the data set, the range of prices is between 0.83 and 0.88.

For demonstration purposes, I take one realization of this auction – for which I took $N = 70$ draws from the distribution of s_i . This gave me 70 bid functions, which I treat as my “data.”

To calculate the “true” marginal valuations corresponding to the “data,” I followed the “bidder’s procedure” outlined in the previous section: I simulated 10,000 residual supply functions for each bidder by making draws from the known distribution of the signals (observe the crucial distinction between drawing from the known “true” distribution and the “empirical” distribution). I then plugged this distribution into the first-order condition in Proposition 1 to calculate the true marginal valuation corresponding to each bid.

In the empirical exercise, we do not know the distribution of s_i ’s. All we see is data from one auction – in this case, the 70 bid functions in the generated data set. The object is to see if estimating the marginal valuations by bootstrapping these 70 bid functions will be successful in recovering the “true” marginal valuations.

I proceed exactly in the way proposed: Holding one bidder’s bids fixed at a time, I draw 10000 samples of size 69 from the empirical distribution of bid vectors. I then calculate the residual supply curve that these resampled bids generate, and calculate the probability distribution of the market clearing price that these residual supply curves and the fixed bidder’s bid vector yield. I then plug point estimates of this probability distribution into the first-order condition to calculate the marginal valuation of the bidder.

Figure 9 displays the result of this experiment. The thick line plotted with x’s is the aggregate bid schedule in this auction, the solid line is the “true” marginal valuation calculated using draws from the known signal distribution, and the broken line is the estimated marginal valuation calculated using the bootstrap method. We see that the estimated marginal valuations are quite close to the true marginal valuations for a wide range of prices about the actual market clearing prices.

Repeating this experiment 200 times with different “data sets” did not yield very different results from this demonstration. Hence, I conclude that the performance of the estimation method is likely to be quite good.

8.6 Standard errors for the resampling method

Small sample properties of bootstrap and other resampling schemes are discussed in depth by the books by Hall (1992) and Efron and Tibshirani (1998). In particular, Hall (1992) shows that the bootstrap estimate of the distribution of the sample mean is, in most cases, as good or better than a standard normal approximation using the sample mean and variance. In the following section, I discuss two methods that can be used to assess the standard errors of my estimates of bidders’ marginal valuations.

The choice of the appropriate number of resampling draws is discussed widely in the applied statistical literature. Efron and Tibshirani (1998) suggest that for bootstrap distribution estimators in most applied settings, B should be at least 1000. I use $B = 10000$ in my application.

Recall that the estimated marginal valuation for each bid y_{ik} is given by the formula:

$$\hat{v}(y_{ik}) = p_k + \frac{\hat{H}^R(p_k, \vec{y}_i)[p_k - p_{k-1}]}{\hat{H}^R(p_k, \vec{y}_i) - \hat{H}^R(p_{k-1}, \vec{y}_i)} \quad (24)$$

where $\hat{H}^R(p_k, \vec{y}_i)$ has non-zero variance about the true probability distribution of the market clearing price, $H(p_k, \vec{y}_i)$. Unfortunately, the consistency proof does not lead the way to a *asymptotic* formula for this variance. Calculating the

standard errors introduced into the marginal valuations is doubly difficult since it is a non-linear function in $H(\cdot)$.

Given the complexity of the problem, I use the “jackknife-after-bootstrap” method suggested by Efron (1992) to compute the standard errors of my estimates of the marginal valuation. Specifically, $\hat{v}_B(y_{ik})$ be the estimate for the marginal valuation of bidder i for y_{ik} units of Treasury bills, calculated using B bootstrap simulations of the market clearing price. Then, the “jackknife-after-bootstrap” estimator of the variance of the estimated marginal valuation is defined to be:⁴⁰

$$v\hat{a}r_{jack}(\hat{v}_B(y_{ik})) = \frac{N}{N-1} \sum_{j=1}^N (\hat{v}_{B(j)}(y_{ik}) - (\frac{1}{N} \sum_{i=1}^N \hat{v}_{B(j)}(y_{ik})))^2 \quad (25)$$

where $\hat{v}_{B(j)}(y_{ik})$ is the bootstrap estimate of $\hat{v}_{B(j)}(y_{ik})$ calculated over a set of resamples that do not contain the bid vector \vec{y}_j . The validity of using this procedure is discussed at length in Efron (1992).

8.7 Proof of lemma 1

Recall that the payment rule in the Vickrey auction is that each bidder pays the area under his residual supply function up to the market clearing price. With this in mind, the objective function of the bidder in the Vickrey auction is:

$$\max_{y_i(\cdot)} \int_0^\infty \left\{ \int_0^{y_i(p^c)} E_{s_{-i}|p^c, s_i} [v(q, s_i, s_{-i}) - RS^{-1}(q, s_{-i})] dq \right\} dH(p^c, y_i(p^c)) \quad (26)$$

where $RS_i(p, s_{-i}) = Q - \sum_{j \neq i}^N y_j(p, s_j)$ is the “residual supply function” that bidder i faces.

Now let us look at the optimal quantity y^* a bidder would demand if she knew that the market-clearing price equalled p :

$$\max_y \int_0^y E_{s_{-i}|p, s_i} [v(q, s_i, s_{-i}) - RS^{-1}(q, s_i, s_{-i})] dq \quad (27)$$

Taking the first-order condition with respect to y , we get:

$$E_{s_{-i}|p, s_i} v(y^*, s_i, s_{-i}) = E_{s_{-i}|p, s_i} RS^{-1}(y, s_{-i}) \quad (28)$$

$$= p \quad (29)$$

which is the mapping stated in the claim. Since this mapping between price and quantity holds for every ex-post value of the market clearing price, it should also characterize the optimal bidding strategy that maximizes the expected utility of the bidder. Observe that the argument is similar to that used to demonstrate truth-telling in private value Vickrey auction: for every ex-post realization of the residual supply function, the best-response of the bidder is to reveal his marginal valuation. The modification here is that the bidder conditions his expected marginal valuation on the realization of the market clearing price.

⁴⁰See Efron and Tibshirani (1998)

References

- Alkan, Ahmet**, “Treasury Domestic Debt Auctions (Hazine İç Borç İhaleleri),” Technical Report, Bogazici University 1991.
- Athey, Susan and Philip Haile**, “Identification in standard auction models,” August 2000. MIT working paper.
- Ausubel, Lawrence M. and Peter C. Cramton**, “Demand Reduction and Inefficiency in Multi-Unit Auctions,” March 1997. University of Maryland Working paper.
- and — , “The Optimality of Being Efficient,” May 1998. University of Maryland Working paper.
- Back, Kerry and Jaime F. Zender**, “Auctions of Divisible Goods: On The Rationale for the Treasury Experiment,” *Review of Financial Studies*, 1993, 6, 733–664.
- Bajari, Patrick and Ali Hortaçsu**, “Winner’s Curse, Reserve Prices, and Endogeneous Entry: Empirical Insights from eBay,” November 2000. Stanford Institute for Economic Policy Research Discussion Paper.
- Bartolini, Leonardo and Carlo Cottarelli**, “Treasury Bill Auctions: Issues and Uses,” in Mario I. Blejer and Teresa Ter-Minassian, eds., *Macroeconomic Dimensions of Public Finance: Essays in Honour of Vito Tanzi*, London: Routledge, 1997, pp. 267–336.
- Berument, Hakan and Kamuran Malatyali**, “Determinants of Interest Rates in Turkey,” Technical Report, The Central Bank of the Republic of Turkey February 1999.
- Bikchandani, Sushil and C. Huang**, “Auctions with Resale Markets: An Exploratory Model of Treasury Bill Markets,” *Review of Financial Studies*, 1989, 2, 311–340.
- Bikhchandani, Sushil and C. Huang**, “The economics of treasury security markets,” *Journal of Economic Perspectives*, 1993, 7, 117–134.
- Bresnahan, Timothy**, “Departures from Marginal-Cost Pricing in the American Automobile Industry,” *Journal of Econometrics*, 1981, 17, 201–227.
- Chari, V. V. and Robert Weber**, “How The U.S. Treasury Should Auction Its Debt,” *Quarterly Review of the Federal Reserve Bank of Minneapolis*, Fall 1992, pp. 3–12.
- Efron, B. and R. Tibshirani**, *An Introduction to the Bootstrap*, Chapman and Hall, 1998.
- Efron, Bradley**, “Jackknife-after-bootstrap standard errors and influence functions,” *Journal of the Royal Statistical Society, Series B*, 1992, 54, 83–127.
- Elyakime, Bernard, Jean Jacques Laffont, Patrice Loisel, and Quang Vuong**, “First-Price Sealed-bid Auctions With Secret Reservation Price,” *Annales D’Economie Et De Statistique*, 1994, 34, 115–141.
- Engelbrecht-Wiggans, Richard and Charles Kahn**, “Multi-Unit Pay-Your-Bid Auctions with Variable Awards,” *Games and Economic Behavior*, 1998, 23, 25–42.
- Friedman, Milton**, *A Program For Monetary Stability*, Fordham University Press, 1960.
- Gordy, Michael B.**, “Structural Analysis of Reservation Price Policy in a Treasury Bill Auction,” September 1994. Working paper.

- Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong**, “Optimal nonparametric estimation of first-price auctions,” *Econometrica*, 2000, 68, 525–574.
- Haile, Philip**, “Auctions with Private Uncertainty and Resale Opportunities,” *Journal of Economic Theory*, 2000. forthcoming.
- , **Han Hong, and Matthew Shum**, “Nonparametric Tests for Common Values in First-Price Auctions,” March 2000. Mimeo.
- Hall, Peter**, *The Bootstrap and Edgeworth Expansion*, Springer-Verlag, 1992.
- Heller, Daniel and Yvan Lengwiler**, “The Auctions of Swiss Government Bonds: Should the Treasury Price Discriminate or Not?,” January 1998. Working paper.
- Hendricks, Kenneth, Joris Pinkse, and Robert H. Porter**, “Empirical Implications of Bidding in First-Price, Symmetric, Common Value Auctions,” 1999. Northwestern University.
- Hortaçsu, Ali**, “Bidding Behavior in Divisible Good Auctions: Theory and Evidence from the Turkish Treasury Auction Market,” August 2001. Working paper.
- Joskow, P., R. Schmalensee, and E. Bailey**, “The market for sulfur-dioxide emissions,” *American Economic Review*, 1998, 88, 669–685.
- Kamien, M. and N. Schwartz**, *Dynamic Optimization*, North-Holland, 1993.
- Klemperer, Paul**, “Why every economist should learn some auction theory,” July 2000. invited paper at the World Congress of the Econometric Society.
- Krishna, Vijay and Motty Perry**, “Efficient Mechanism Design,” January 1999. Working paper.
- Kyle, Albert**, “Informed Speculation with Imperfect Competition,” *Review of Economic Studies*, 1989, 56, 317–356.
- Laffont, Jean Jacques, H. Ossard, and Q. Vuong**, “Econometrics of first price auctions,” *Econometrica*, 1995, 63, 953–980.
- Malvey, Paul F. and Christine M. Archibald**, “Uniform Price Auctions: Update of the Treasury Experience,” Technical Report, Department of the Treasury October 1998.
- McAdams, David**, “Infeasible Supply and “Collusive-Seeming Equilibria” in the Uniform Price Auction,” September 1999. Working paper.
- Milgrom, Paul R. and Robert J. Weber**, “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 1982, 50, 1089–1122.
- and —, “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 1982, 50, 1089–1122.
- Myerson, Roger B.**, “Optimal Auction Design,” *Mathematics of Operations Research*, 1981, 6, 58–73.
- Nandi, Saikat**, “Treasury Auctions: What Do the Recent Models and Results Tell Us?,” *Federal Reserve Bank of Atlanta Economic Review*, 1997, Fourth Quarter, 4–15.
- Nautz, D.**, “Optimal Bidding in Multi-Unit Auctions with Many Bidders,” *Economics Letters*, 1995, 48, 301–306.

- Noussair, Charles**, “Equilibria in a Multi-Object Uniform Price Sealed Bid Auction with Multi-Unit Demands,” *Economic Theory*, 1995, 5, 337–351.
- Nyborg, K. and S. Sundaresan**, “Discriminatory versus uniform treasury auctions: evidence from when-issued transactions,” *Journal of Financial Economics*, 1996, 42, 63–104.
- Nyborg, Kjell, Kristian Rydqvist, and Suresh Sundaresan**, “Bidder Behavior in Multiple Unit Auctions,” June 1997. mimeo.
- Paarsch, Harry J.**, “Deciding Between Common Values and Private Value Paradigms in Empirical Models of Auctions,” *Journal of Econometrics*, 1992, 51, 191–215.
- Reny, Philip**, “On The Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games,” *Econometrica*, 1999, 67 (5), 1029–1056.
- Rosse, James**, “Estimating Cost Function Parameters without Using Cost Data: Illustrated Methodology,” *Econometrica*, 1970, 38, 256–270.
- Simon, D.**, “The Treasury’s experiment with single-price auctions in the mid 1970’s: winner’s or taxpayer’s curse?,” *Review of Economics and Statistics*, 1994, 76, 754–760.
- Swinkels, Jeroen**, “Asymptotic Efficiency for Discriminatory Private Value Auctions,” *Review of Economic Studies*, 1999, 66, 509–528.
- Umlauf, Steven**, “An Empirical Study of the Mexican Treasury Bill Auctions,” *Journal of Financial Economics*, 1993, 33, 313–340.
- van der Waart, A. W.**, *Asymptotic Statistics*, Cambridge University Press, 1998.
- Vickrey, William**, “Counterspeculation, Auctions, and Competitive Sealed Tenders,” *Journal of Finance*, 1961, 16, 8–37.
- Viswanathan, S., James Wang, and Thomas Witelski**, “Optimal Bidding in Multi-Unit Discriminatory Auctions: Two Bidders,” January 2001. Mimeo.
- Wilson, Robert**, “Auctions of Shares,” *Quarterly Journal Of Economics*, 1979, pp. 675–689.
- Wolak, Frank A.**, “An Empirical Analysis of the Impact of Hedge Contracts on Bidding Behavior in a Competitive Electricity Market,” March 1999. Working paper.
- Wolfram, Catherine D.**, “Strategic Bidding in A Multi-Unit Auction: An Empirical Analysis of Bids to Supply Electricity in England and Wales,” 1998. Working paper.




Figure 1: Discriminatory and uniform price auctions




Figure 2: Bidder's problem




Figure 3: The resampling procedure

Figure 4: Bidder #2 and the simulated market clearing price distribution




Figure 5: Estimation results for bidder #2 in auction #1




Figure 6: Upper and lower envelopes of bidder #2's marginal valuations




Figure 7: Counterfactual revenue comparison for auction #1




Figure 8: Illustration of notation with discretized bids




Figure 9: Simulation study




Table 1: Summary statistics for 3-month T-bill auctions

Variable	Average	Std.dev.	Min	Max	Total
Inflation ^a	67.8%	5.6%	58.0%	78.7%	
No. of bidders	69	21.02	32	110	134 ^b
Revenue (million \$) ^c	456.3	443.1	31.4	1560.9	12369.1
Cover ratio ^d	28.9%	24.9%	0.2%	84.3%	
Auction yield ^e	90.8%	6.9%	79.0%	101.4%	
Variance of bids ^f	0.035	0.031	0.005	0.134	

^aMonthly CPI from the Central Bank of Turkey

^bTotal number of unique bidders who participated in the auctions.

^cConverted using daily exchange rate data

^dQuantity sold/Quantity demanded

^eQuantity weighted average yield

^fCalculated using the formula: $\frac{\sum_{i=1}^N (p_i - p_{avg})^2 q_i}{\sum_{i=1}^N q_i}$

Table 2: Regression analysis of intercepts and slopes of individual bid functions

	(1)	(2)	(3)	(4)
Dep. Variable	<i>INTCPT</i>	<i>INTCPT</i>	<i>SLOPE</i>	<i>SLOPE</i>
$\%SHORTFALL_{i,t-1}$.0014* (.0006)	.0019* (.0007)	-.00027 (0.00022)	-.00006 (.0002)
$\%SHORTFALL_{i,t-2}$	-.0008 (.0005)	.0008 (.0006)	-.00009 (.0002)	6.04e-06 (.0002)
	Auction dummies	Auction & bidder dummies	Auction dummies	Auction & bidder dummies
No. of obs.	1109	1109	1109	1109
Root MSE	.005	.005	.002	.002
R^2	0.62	0.68	0.036	0.35
$\overline{R^2}$	0.61	0.64	0.015	0.27

(Standard errors are reported in parentheses. Variables significant at the 99% level are marked with an asterisk.)

Table 3: Counterfactual Revenue Comparisons

Auction	Date	Revenue (million \$)	Ex-post ^a % Revenue loss	Ex-ante % Revenue loss	Std. dev. of Rev. loss
1	10/16/91	404	0.30%	2.68%	(3.40%)
2	11/13/91	388	3.10%	27.22%	(28.83%)
3	12/11/91	415	1.00%	0.12%	(5.07%)
4	01/08/92	368	2.49%	4.27%	(4.85%)
5	02/05/92	61	1.17%	5.81%	(6.22%)
6	03/04/92	45	14.77%	10.22%	(10.32%)
7	04/01/92	146	2.37%	1.49%	(4.48%)
8	04/29/92	32	1.47%	25.59%	(31.38%)
9	05/27/92	216	0.24%	1.48%	(2.40%)
10	06/24/92	322	7.13%	4.80%	(8.60%)
11	07/22/92	415	0.25%	3.99%	(5.19%)
12	08/19/92	663	0.18%	2.01%	(2.10%)
13	09/16/92	706	0.17%	2.34%	(2.39%)
14	10/14/92	808	0.17%	3.82%	(4.20%)
15	11/11/92	1144	1.84%	3.77%	(3.58%)
16	12/09/92	1255	0.13%	4.89%	(6.47%)
17	01/06/93	1550	0.06%	9.27%	(11.36%)
18	02/03/93	1405	3.77%	5.03%	(6.52%)
19	03/31/93 ^b	179	24.71%	14.13%	(16.04%)
20	04/28/93	734	0.69%	5.19%	(5.92%)
21	05/26/93	70	22.40%	24.79%	(25.09%)
22	06/23/93	196	0.52%	14.40%	(14.73%)
23	07/21/93	394	3.03%	15.29%	(15.21%)
24	08/18/93	73	22.05%	26.01%	(28.13%)
25	09/15/93	71	2.71	16.13	(21.15%)
Average		482	3.80%	14.23%	

^aFrom switching to uniform price auction^bThe 03/03/93 auction was effectively canceled since only 0.23% of bids appear to have been awarded.

Table 4: Participation statistics

Category	Mean no. from category	Std. dev.
Type 1	19.65	4.92
Type 2	33.55	11.16
Type 3	12.51	10.04

Table 5: AR(1) model of Treasury bill supply

Dependent variable: <i>QSOLD</i>			
Ind. variable	Coef.	Std. Err.	t-stat
<i>QSOLD</i> (-1)	1.088	.140	7.774
<i>CONSTANT</i>	195.88	416.34	0.470
No. of observations	16		
R^2	0.8119		