SIEPR Discussion Paper No. 09-031

Informational Advantage and Information Structure: An Analysis of Canadian Treasury Auctions

By
Ali Hortaçsu and Jakub Kastl
Stanford University

August 2010

Stanford Institute for Economic Policy Research
Stanford University
Stanford, CA 94305
(650) 725-1874

Informational Advantage and Information Structure: An Analysis of Canadian Treasury Auctions *

Ali Hortaçsu[†] Jakub Kastl[‡]

August 10, 2010

Several important auction settings, including treasury auctions in Canada and the U.S., have the feature that some bidders (dealers) observe the bids of a subset of other bidders (customers). Quantifying the economic advantage that informationally advantaged bidders derive from this institutional feature requires that we empirically distinguish between private vs. interdependent values paradigms. Bidders with private values who obtain information about rivals' bids use this information to update their beliefs about the distribution of residual supply. With interdependent values, bidders also update their beliefs about the value of the good being auctioned. We use these differential updating effects to construct formal hypothesis tests of the presence of private vs. interdependent values. Using data from Canadian treasury auctions, we cannot reject the null hypothesis of private values in auctions of 3- and 12-month treasury bills. We also do not find evidence supporting the alternative hypothesis of interdependent values. We use the estimated model to quantify the value of observing customer bids to a dealer. We find that the extra information contained in customers' bids leads on average to an increase in payoff equal to $13-35\%$ of the expected surplus of dealers.

Keywords: multiunit auctions, treasury auctions, structural estimation, nonparametric identification and estimation, test for common values

JEL Classification: D44

*We would like to thank Liran Einav, Ken Hendricks, Han Hong, Rob Porter, Peter Reiss, Azeem Shaikh for many helpful comments and suggestions and especially Philip Haile for detailed comments on an earlier draft. Participants at The University of Minnesota Applied Microeconomics Workshop, SED 2008 meetings in Prague, seminars at the Bank of Canada, Chicago, CERGE-EI, NYU Stern, Queens U., Stanford, Yale and conferences on auctions at University College London and Penn State provided helpful feedback. Kastl is grateful for the hospitality and financial support from the Cowles Foundation at Yale University and for the financial support of the NSF (SES-0752860). Hortaçsu is also grateful for the hospitality of the Cowles Foundation and the financial support of the NSF (SES-0449625) and the Sloan Foundation. All remaining errors are ours.

[†]Department of Economics, University of Chicago, and NBER

[‡]Department of Economics, Stanford University, and NBER

1 Introduction

In certain important auction settings, a subset of bidders have the informational advantage of observing a (subset of) their rivals' bids before submitting their bids. In Canadian Treasury auctions, the empirical setting we study, "primary dealers" and "government securities dealers" ("dealers" for short) route the bids of non-dealer bidders, called "customers."[1] The particular organization of Canadian Treasury auctions is not unique: many countries, including the U.S., use a primary dealership system in the sale of government securities, which limits participation in auctions to a small set of financial institutions, who often bid on behalf of their "customers."[2] In the context of U.S. Treasury auctions, Nyborg and Sundaresan (1996) have pointed out that, despite the presence of an active "when-issued market" that aggregates information prior to the auction, customer orders are an important source of private information for primary dealers, as they allow dealers to "obtain private information about the aggressiveness of bidding." (p.68)

Many other financial markets are organized around dealer/specialists who observe customer "order flow," and can potentially utilize this information to their advantage. The financial microstructure literature has studied the behavior of dealers with superior information in a number of contexts; including specialists in the NYSE (Madhavan and Smidt (1991), dealers in the foreign exchange market (Lyons (1995) and Evans and Lyons (2002)), and in options markets (Easley, O'Hara and Srinivas (1998)). Perhaps not surprisingly, allegations of "front-running," i.e. utilizing customer order information to make profitable trades on securities markets, are commonplace in financial news. For example, on March 4, 2009, 14 trading firms paid \$69 million to settle charges by the SEC that they engaged in various types of "front-running."[3]

Quantifying the economic value of observing customer orders, however, requires one to take a stance as to whether the informational environment of the auctions is one of *private* vs. *interdependent* values. In an interdependent value environment, customer bids are informative about

[1] Most "customers" in our data set are Canadian or international banks (such as BNP Paribas or Bank of America), or large institutional investors (such as pension funds) who demand a substantial portion of the marketed securities. See Section 2.2 for summary statistics.

[2] Out of 39 respondent countries to the IMF survey conducted by Arnone and Iden (2003), 29 employed a primary dealership system that limited participation a small number of bidders who committed to minimum participation and/or market-making responsibilities. The median number of primary dealers in this sample was 13. The U.S. currently has 18 primary dealers who can bid on behalf of their customers.

[3] http://www.nytimes.com/2009/03/05/business/05specialist.html.

the ex-post value of the securities being auctioned, and thus may induce the dealer to revise her willingness-to-pay for the security on sale. In a pure private value environment, however, the observation of customer bids does not change the willingness-to-pay of the dealers. However, customer bids may still contain strategic information that can compel a dealer to modify her bid.

Our data from Canadian Treasury auctions allow us to study the above two mechanisms in some detail. In particular, we observe dealers' bids before and after they route customer bids; thus we can track modifications in dealer bids made in response to the observation of customer bids. In this context, consider a situation in which bidder i is about to submit her bid (demand) function y_i, but before submitting y_i she observes a bid actually submitted by bidder j. In a purely private values environment, bidder i obtains better information only about the location and shape of residual supply[4] she will be facing in the upcoming auction. Using this additional information, she revises her initial bid y_i, and submits an alternative bid y_i'. In an auction with interdependent values or a common value component, on top of the additional information about the location and shape of the residual supply curve, she also obtains new important information about a rival's private information, and therefore will update her prior on the ex-post value of the securities being auctioned. Therefore, she submits a new bid y_i'' taking into account both of these new pieces of information. In general, the way she will revise her bid y_i will differ under the two scenarios and hence $y_i'' \neq y_i'$.

The above distinction motivates our formal testing strategy to determine whether a private vs. interdependent value framework is appropriate for our analysis. In Section 5, we start with the null hypothesis of conditionally independent private values,[5] and test whether the observed modifications to dealer bids in response to customer bids can be rationalized by this specification. To conduct this test, we build on our earlier work (Hortaçsu (2002a) and Kastl (2008)) to characterize the necessary conditions for equilibrium bidding under private values, and estimate the marginal valuations that rationalize a dealer's bid under equilibrium beliefs about her competitors' bids. As above, under private values, information about a customer's bid only changes the dealer's beliefs about the distribution of competitors' bids, but not her marginal valuation. Thus, the rationalizing

[4]Residual supply is defined as the total quantity for the auction minus the bids of other bidders.

[5]We test the conditional independence assumption in Section 4.4.

marginal valuation that we estimate for a dealer's bid before observing a customer's bid vs. after should be the same.

Since our null hypothesis is a composite one of private values and model specification, we also provide, in Section 4.2, a test for the alternative specification of interdependent values. This test relies on a novel characterization of necessary conditions for optimal bidding in the interdependent value multi-unit discriminatory auction. What we can recover using these necessary conditions, however, is no longer a model primitive as the marginal valuations in the private values case. Nevertheless, we argue that the presence of interdependent values induces testable sign restrictions on the observables.

Our main empirical result, as described in Section 5, is that the null hypothesis of (conditionally independent) private values can not be rejected in our data on auctions of 3- and 12-month Treasury bills. However, we find, across a number of specifications, that rejection of the null hypothesis is slightly more likely in our sample of 12-month Treasury bills vs. 3-month Treasury bills. We therefore implement our test for the alternative hypothesis in the 12-month sample, and fail to find evidence in favor of interdependent values.

A private values specification is convenient for calculating the economic benefit to a dealer from observing a particular customer's bid, as the estimated marginal valuations allow us to calculate the ex-post surplus of the dealer. Thus, we can simply calculate the profit the dealer would have made if she had submitted her bid before observing the customer bid, vs. the profit she made with her updated bid.[6] Consistently with views of the practioners we find, in Section 5.4, that the customers' order flow contributes significantly to dealers' overall profits from participating in primary auctions.

Outside of the dealer-customer setting, determining whether the informational environment of an auction is one of private or interdependent values has important implications for the choice of optimal auction mechanism, a question that has been addressed frequently in the auctioning of securities (especially Treasuries) context. Whereas the earliest theoretical analyses of securities auctions by Vickrey (1961) and Smith (1966) have pursued the private value model, later models

[6]Unfortunately, a similar "value of information" calculation is much more difficult in the interdependent values case. See Section 5.4.

by Wilson (1979), Kyle (1985), or Bikhchandani and Huang (1989) emphasized the common value component. Bikhchandani and Huang (1989), for example, build on the seminal revenue-ranking results of Milgrom and Weber (1982) to argue for a shift away from the discriminatory auction format. However, as pointed out by Ausubel and Cramton (2002), neither the revenue equivalence theorem nor Milgrom and Weber's revenue ranking results apply to the multi-unit auction setting (with multi-unit demands). In the absence of general theoretical results on revenue ranking in either the private or interdependent value settings, empirical answers have been sought to answer the question on a case-by-case basis. In particular, a number of recent papers have utilized a structural econometric modelling approach to answer the revenue ranking question. However, these papers impose interdependent vs. private values as an *a priori* assumption that is not tested empirically.[7]

To our knowledge, this is the first attempt to formally test between private vs. interdependent values in a multi-unit divisible good auction setting.[8] In particular, we provide nonparametric tests of the null hypothesis of private values and the alternative hypothesis of interdependent values in this setting. Our approach is most similar to the contribution in the single-unit auction context by Haile, Hong and Shum (2003) (henceforth HHS). HHS pose a nonparametric test for common value in first-price auctions making use of variation in the number of bidders across auctions. They use nonparametric techniques developed in recent empirical auctions literature (e.g., Laffont and Vuong (1995), and Guerre, Perrigne and Vuong (2000)) to estimate the distribution of values given the observed bids, under the null hypothesis of private values. Theory predicts a certain ordering

[7]While Fevrier, Preguet and Visser (2002) and Armantier and Sbaï (2006) base their econometric models on a pure common value specification, other studies, including Hortaçsu (2002a), Kastl (2008) and Kang and Puller (2006), utilize independent private value models. Chapman, McAdams and Paarsch (2006, 2007) focus on bidding in Canadian cash auctions also assuming a private values framework. Yet another important part of the literature on estimation of multiunit auctions are papers analyzing the electricity auctions which are usually modelled in a private value framework (e.g., Wolak (2003, 2005)).

[8]For single object auctions, Gilley and Karels (1981) were first to propose a reduced form testing approach based on examining how bids vary with the number of participants. For second-price sealed-bid and English auctions, Paarsch (1991) and Bajari and Hortaçsu (2003) employed tests for CV using standard regression techniques: under the maintained null hypothesis of independent private values and (weakly dominant strategy) truthful bidding, bids should not respond to information about the number of participants. Pinkse and Tan (2005) establish, however, that such a reduced form test cannot distinguish unambiguously a CV from PV model in first price auctions. Paarsch's (1992) seminal paper showed that a more detailed structural model can achieve the goal of distinguishing CV and PV in first-price auctions. Paarsch's method, however, relies on parametric assumptions about the distribution of bidder's private information, and hence it is hard to disentangle the influence of the parametric assumptions on the actual outcomes of the testing procedure.

between the distribution of pseudovalues (the expected value of the object conditional on winning) under the common value paradigm as the number of bidders varies, while this should not vary with the number of participants under PV. We exploit a different source of variation in our data to base our test on. Specifically, our data set from Canadian treasury bill auctions allows us to observe the modifications that a subset of bidders (*dealers*) make to their submitted bids upon observing the bids of some of their competitors (*customers*). Thus, we are able to observe how bids change *within* an auction in response to new information about competition. In contrast, the testing strategy utilized in the literature so far focuses on *across-auction* responses to changes in the number of competing bidders. Our *within-auction* testing strategy makes our results less susceptible to the presence of unobserved heterogeneity across auctions. We should note that Hill and Shneyerov (2009) also propose an alternative nonparametric test for common values which allows for unobserved auction heterogeneity. However, their test is based on the behavior of the CDF of bids near the reserve price, and thus crucially relies on the one-dimensionality of bids, and cannot easily be extended to the multi-unit setting.

2 Institutional Background and Data Description

2.1 Institutional Background

Treasury bills and other Bank of Canada securities are issued in the primary market through sealed-bid discriminatory auctions. Bids consist of price-quantity schedules and define step functions, with minimum price increment of 0.1 basis points and minimum quantity increment of C$1 million. Bids are submitted electronically and can be revised at any point before the submission deadline. There are two major groups of potential bidders: *dealers* (primary dealers and government securities distributors) and *customers*. The customers are typically not individual investors. Many of them are actually large banks that for some reason choose not to be registered as dealers, but whose demands are sufficiently important for the Bank of Canada to require their separate identification. The major distinction between customers and dealers, however, is that customers cannot bid on their own account in the auction, but have to route their bids through one of the dealers. The dealers are required to identify bids submitted by customers in the electronic bidding system.

Our two samples consist of all submitted bids in 116 auctions of 3-months and 12-months treasury bills of the Canadian government issued between 10/29/1998 and 3/27/2003. The auctions were conducted as sealed-bid discriminatory price auctions. Along with the set of bids taken into consideration when making the final allocation, we also have the entire record of electronic bid submissions by dealers (under their own bidder ID and their customers' IDs) during the bid submission period. Thus, we are able to observe any modifications made by the dealers to their own bids up until the bidding deadline. Each electronic submission has a time stamp, thus we are able to observe whether a dealer's bid modification was preceded by the entry of a customer bid.

Dealers are required to submit customer bids as soon as they receive them.[9] Of course, it might be difficult for customers to verify whether their dealer abides by this requirement. However, many customer bids are observed to be submitted very close to the bidding deadline. This suggests that a customer can ensure the timely submission of her bid by sending her bid very close to the deadline and demand that the dealer routes it as soon as possible (and complain if their bids end up not being submitted in time for the deadline). Appendix A discusses the timing of bids in further detail.

This "last minute" bidding behavior may also explain why many dealers submit multiple bid modifications before the bidding deadline, and do not simply wait to submit their bid until they have seen the bids of their customers. If entering bid modifications is not very costly, the uncertainty in the arrival time of customer bids may render it optimal for the dealers to follow the strategy of entering their best response conditional on their current information set.[10] Moreover, such "last minute" bids may also explain why some customer bids were not followed by updates to the dealer's bid. [11]

In Appendix A, we also conduct a short descriptive analysis of the relationship between customer

[9]Investment Industry Regulatory Organization of Canada (IIROC) Rule 2800 , the "Code of Conduct for IIROC-Regulated Firms Trading in Wholesale Domestic Debt Market" prohibits the following practice (sec. 4.3(b-2)): "(. . .) no Dealer Member (. . .) shall engage in any trading which takes unfair advantage of customers, counterparties or material non-public information, such as (. . .) executing proprietary trades ahead of client orders on the same side of the market without first disclosing to the client the intention to do so and obtaining the client's approval."

[10]In previous work studying the same data set, Hortaçsu and Sareen (2006) report that although some customers appear to be in long-term relationships with their dealers, some of them appear to change dealers frequently; which may render some customer bids as a surprise.

[11]Hortaçsu and Sareen (2006) find that some dealers' modifications to their own bids in response to these late customer bids narrowly missed the bid submission deadline, and that such missed bid modification opportunities had a negative impact on dealers' ex-post profits.

and dealer bids, and find evidence that customer bids are informative of the modification that dealers make to their own bids. Of course, this is a purely descriptive finding, and our model of bidding to follow will make very precise predictions about the type of modification that a dealer should make under different informational environments.[12]

2.2 Summary Statistics

Table 1 offers some summary statistics of our data set, which encompasses all submitted bids in 116 auctions of 3-month treasury bills of the Canadian government issued between 10/29/1998 and 3/27/2003.

Table 1: Data Summary

Summary Statistics for 3-month T-bill auctions				
# of Auctions	116			
	Mean	St.Dev.	Min	Max
# of Dealers	12.34	1.64	9	15
# of Customers	4.66	2.3	0	12
# of Participants	17	2.83	11	23
# of Submitted steps	2.88	1.69	1	7
Price bid	989,353	3,244	984,515	994,969
Quantity bid[a]	0.092	0.07	0.00023	0.2512
Issued Amount (billions C$)	3.881	0.552	2.8	5

[a] As a percentage of total supply.

On average, 12 dealers and less than 5 customers participate in every auction. An average bid function consists of less than 3 steps and the average quantity demand is for about 9% of the supply.

As usual in most government securities auctions, bids can be submitted both as competitive tenders and as noncompetitive tenders. Each participant is allowed to submit a single noncompetitive tender. A noncompetitive tender specifies a quantity that the bidder wishes to purchase

[12]Hortaçsu and Sareen (2006) also report various descriptive measures suggesting that obtaining customer information has a causal impact on dealers' bidding patterns. They find that the direction of changes in a dealer's (quantity-weighted price) bid typically follows the direction of discrepancy between the dealer's pre-customer information bid, and the customer's bid. Hortaçsu and Sareen point out that both common value and private value models are consistent with their descriptive patterns, however, and do not conduct tests to distinguish between these informational environments.

at the price at which the auction clears. In our data, there are on average 3.6 noncompetitive tenders in an auction of the preannounced amount for sale – however, the biggest noncompetitive tenders were placed by the central bank itself. In our estimation approach we thus treat separately non-competitive bids by the central bank and non-competitive bids by regular participants. All these non-competitive bids operate by shifting the available supply to the left.[13]

Table 2 presents the summary statistics for the 12-month T-bill auctions. Relative to auctions of 3-month treasury bills (which are sold in parallel auctions), there is less participation both by dealers and especially by customers. Price bids exhibit larger variation. The amount offered for sale in each auction is also significantly lower.

Table 2: Data Summary

Summary Statistics for 12 month T-bill auctions				
# of Auctions	116			
	Mean	St.Dev.	Min	Max
# of Dealers	5.69	0.76	4	8
# of Customers	2.84	1.78	0	8
# of Participants	8.53	1.93	5	15
# of Submitted steps	2.91	1.69	1	7
Price bid	958,997	10,972	940,004	979,928
Quantity bid[a]	0.096	0.074	0.00028	0.25
Issued Amount (billions C$)	1.67	0.194	1.3	2

[a] As a percentage of total supply.

3 The Model and Test Description

Our analysis is based on the share auction model of Wilson (1979) with private information, in which both quantity and price are assumed to be continuous. We modify Wilson's model to take into account the discretness of bidding (i.e., finitely many steps in bid functions) as in Kastl (2008). We further adapt this model into the context of our application where some bidders (the "dealers") observe the bids of others ("customers"). Formally, suppose there are two classes of bidders:

[13]The non-competitive bids are not as important in Canadian treasury bill auctions as in other settings. In particular, while the average non-competitive bid is for about 4.4% of the supply, most of this is driven by non-competitive bids that were placed by the central bank itself. The average non-competitive bid conditional on being placed by a dealer or a customer is for less than 0.06% of the supply and hence quite negligible.

N_d potential *dealers* (in index set $\mathcal{D}$) and N_c potential *customers* (in index set $\mathcal{C}$) who are each bidding for a share of a perfectly divisible good. Customers and dealers observe private (possibly multidimensional) signals, $S_1^c, ..., S_{N_c}^c, S_1^d, ..., S_{N_d}^d$.

Assumption 1 *Customers' and dealers' private signals, $S_1^c, ..., S_{N_c}^c, S_1^d, ..., S_{N_d}^d$, are drawn from a common support $[0,1]^M$ according to an atomless joint d.f. $F\left(S_1^c, ..., S_{N_c}^c, S_1^d, ..., S_{N_d}^d\right)$ with strictly positive density f. $F(\cdot)$ is exchangeable with respect to its first N_c arguments (i.e. over customer signals), and also with respect to its N_c+1-st to N_c+N_d-th arguments (i.e. over dealer signals).*

Dealers also observe an additional piece of "order flow" information, Z. Z may equal the customer bid observed by that dealer, or is null when the dealer does not observe any additional information. We assume that Z is not observed by the dealer's competitors, but is observed by the econometrician. The joint distribution of dealers' private information conditional on (the vector of) customers' information and customers' equilibrium strategies is $F^d\left(\left(S_1^d, Z_1\right), ..., \left(S_{N_d}^d, Z_{N_d}\right) | \mathbf{S}^c, \{y_i^c(p|s_i^c)\}_{i=1}^{N^c}\right)$ where $y_i^c(p|s_i^c)$ is the equilibrium strategy of a customer observing signal s_i^c.

Winning q units of the security is valued according to a marginal valuation function $v_i(q, S_i, S_{-i})$. We assume that the marginal valuation function is symmetric within each class of bidders. We will impose the following assumptions on the marginal valuation function $v^g(\cdot,\cdot,\cdot)$ for $g \in \{c, d\}$:

Assumption 2 *$v^g(q, S_i, S_{-i})$ is non-negative, measurable, bounded, strictly increasing in (each component of) S_i $\forall (q, s_{-i})$ and weakly decreasing in q $\forall (s_i, s_{-i})$ for $g \in \{c, d\}$.*

We will denote by $V(q, S_i, S_{-i})$ the gross utility: $V(q, S_i, S_{-i}) = \int_0^q v(u, S_i, S_{-i})\, du$. Throughout the paper we will distinguish between private values and other valuation structures, where bidders' values could be interdependent (for example could have a common value component). The following definition states what we mean by these terms.

Definition 1 *(i) Bidders have private values when $\forall i$ and $g \in \{c, d\}$: $v^g(q, S_i, S_{-i}) = v^g(q, S_i)$.*
(ii) Bidders have interdependent values if $\forall i, j$ and $g \in \{c, d\}$ and a.e. s_i :
$\exists \Omega_j', \Omega_j'' : \Omega_j' \cap \Omega_j'' = \emptyset$ such that $\Pr\left(S_j \in \Omega_j'\right) > 0, \Pr\left(S_j \in \Omega_j''\right) > 0$ and
$\mathbb{E}_{S_{-i}}\left(v^g(q, S_i, S_{-i}) | S_j \in \Omega_j', S_i = s_i\right) \neq \mathbb{E}_{S_{-i}}\left(v^g(q, S_i, S_{-i}) | S_j \in \Omega_j'', S_i = s_i\right)$.

Our definition of interdependent values simply states that each bidder possesses with positive probability some private information that is relevant for determining the value of each of his rivals. In particular, in the context of our empirical application it implies that at least some information that customers possess is useful for primary dealers' estimates of the value of the underlying security.[14]

Finally, we will assume that under interdependent values, the expected utility of dealers is increasing in customers' signals. This assumption is satisfied, for example, when signals are affiliated and v^d is strictly increasing in each S_{-i}.

Assumption 3 *When values are interdependent,* $E_{S_{-i}\setminus S_j^c}\left[v^d\left(q, S_i^d, S_{-i}\right)|S_i^d = s_i^d, S_j^c = s_j^c\right]$ *is strictly increasing in (each component of)* S_j^c, $\forall\left(q, s_i^d\right)$.

To ease notation, let θ_i denote private information of bidder i, i.e., for a customer $\theta_i \equiv S_i$, whereas for a dealer $\theta_i \equiv (S_i, Z_i)$. Bidders' pure strategies are mappings from private information to bid functions $\sigma_i : \Theta_i \to \mathcal{Y}$, where the set $\mathcal{Y}$ includes all admissible bid functions. Since in most divisible good auctions in practice, including the Canadian treasury bill auctions, the bidders' choice of bidding strategies is restricted to non-increasing step functions with an upper bound on the number of steps, $\overline{K}$, we impose the following assumption:

Assumption 4 *Each player* $i = 1, ..., N$ *has an action set:*

$$A_i = \left\{ \begin{array}{c} \left(\vec{b}, \vec{q}, K\right) : \dim\left(\vec{b}\right) = \dim\left(\vec{q}\right) = K \in \left\{0, ..., \overline{K}\right\}, \\ b_{ik} \in B = \left[0, \overline{b}\right],\ q_{ik} \in Q = [0, 1],\ b_{ik} \geq b_{ik+1}, q_{ik} \leq q_{ik+1} \end{array} \right\}$$

where l denotes a losing bid or non-participation and $\overline{b}$ the largest possible bid, for example the face value of a treasury bill. Therefore the set $\mathcal{Y}$ includes all non-decreasing step functions with at most $\overline{K}$ steps, $q : \mathbb{R}_+ \to [0, 1]$, where $q_i(p) = \sum_{k=1}^{K} q_{ik} I\left(p \in (b_{ik+1}, b_{ik}]\right)$ where I is an indicator function. A pure strategy (bid function) for a bidder from group $g \in \{c, d\}$ with private information $\tilde{\theta}_i$ will

[14] Strictly speaking, we will not be able to test whether values are interdependent across dealers; we can only test whether there is interdependency between customers and dealers. It is important to reiterate that the banks which we label as customers are not necessarily smaller than primary dealers or a-priori very different than dealers in other dimensions. They just choose not to participate as primary dealers. In fact one of the customers in our sample was a primary dealer for a short time.

be denoted by $y_i^g\left(p|\tilde{\theta}_i\right)$ and it specifies for each price p, how big a share of the securities offered in the auction (type $\tilde{\theta}_i$ of) bidder i demands.

Q will denote the amount of T-bills for sale, i.e., the good to be divided between the bidders. Q might itself be a random variable if it is not announced by the auctioneer ex ante. In the auctions we study, the Government of Canada has the right to cancel the auction or restrict the announced supply. We assume that the distribution of Q is common knowledge among the bidders.

Furthermore, the number of potential bidders of each type participating in an auction, which we denoted by N_c, N_d for potential customers and dealers respectively, is also commonly known. This assumption is reasonable in the context of our empirical application as all participants have to register with the auctioneer before the auction and the list of registered participants is publicly available.

Since bidders use step functions, a situation may arise in which multiple prices would clear the market. If that is the case, we assume that the auctioneer selects the most favorable price from his perspective, i.e., the highest price. Moreover, because bidders' strategies are step functions, the residual supply will be a step function and hence but for knife-edge cases any equilibrium will involve rationing with probability one. We only consider the rationing rule pro-rata on-the-margin, under which the auctioneer proportionally adjusts the marginal bids so as to equate supply and demand.[15] This particular rationing rule is important for our equilibrium characterization as Kastl (2008) shows that it ensures that no bidder would prefer to tie with another bidder when gaining strictly positive marginal surplus at the quantity allocated after rationing.

The natural solution concept to apply in this setting is Bayesian Nash Equilibrium. The expected utility of a bidder i with private information $\tilde{\theta}_i$ who employs a strategy $y_i\left(\cdot|\tilde{\theta}_i\right)$ in a

[15]The rationing rule employed is pro-rata on-the-margin, under which the rationing coefficient satisfies

$$R\left(p^c\right)=\frac{Q-TD_+\left(p^c\right)}{TD\left(p^c\right)-TD_+\left(p^c\right)}$$

where $TD\left(p^c\right)$ denotes total demand at price p^c, and $TD_+\left(p^c\right)=\lim_{p\downarrow p^c}TD\left(p\right)$. Only the bids exactly at the market clearing price are adjusted.

discriminatory auction given that other bidders are using $\{y_j(\cdot|\cdot)\}_{j\neq i}$ can be written as:

$$EU_i\left(\tilde{\theta}_i\right) = E_{Q,\theta_{-i}|\theta_i=\tilde{\theta}_i}\left[\begin{array}{c} \int_0^{Q_i^c(Q,\theta,\mathbf{y}(\cdot|\theta))} v_i(u,s_i)\,du \\ -\sum_{k=1}^{K}\mathbf{1}\left(Q_i^c(Q,\theta,\mathbf{y}(\cdot|\theta)) > q_k\right)(q_k - q_{k-1})\,b_k \\ -\sum_{k=1}^{K}\mathbf{1}\left(q_k \geq Q_i^c(Q,\theta,\mathbf{y}(\cdot|\theta)) > q_{k-1}\right)\left(Q_i^c(Q,\theta,\mathbf{y}(\cdot|\theta)) - q_{k-1}\right)b_k \end{array}\right] \tag{1}$$

where $Q_i^c(Q,\theta,\mathbf{y}(\cdot|\theta))$ is the (market clearing) quantity bidder i obtains. It is random from perspective of bidder i since it depends on the state (bidders' private information and the supply quantity), (Q,θ) and on bidders bidding according to strategies specified in the vector $\mathbf{y}(\cdot|\theta) = [y_1(\cdot|\theta_1),...,y_{N_d+N_c}(\cdot|\theta_{N_d+N_c})]$ with $y_i\left(\cdot|\theta_i=\tilde{\theta}_i\right)$. The first term in (1) is the gross utility the type θ_i enjoys from his allocation, the second term is the total payment for all units allocated at steps at which the type $\tilde{\theta}_i$ was not rationed and the final term is the payment for units allocated during rationing. A Bayesian Nash Equilibrium in this setting is thus a collection of functions such that (almost) every type $\tilde{\theta}_i$ of bidder i is choosing his bid function so as to maximize his expected utility: $y_i\left(\cdot|\tilde{\theta}_i\right) \in \arg\max EU_i\left(\tilde{\theta}_i\right)$ for a.e. $\tilde{\theta}_i$ and all bidders i.

Given the within group symmetry assumption, we will assume that the bidding data is generated by a Bayesian Nash equilibrium of the game in which customers submit bid functions that are symmetric up to their private signals, i.e. $y_i^c(p|s_i) = y^c(p|s_i) = y^c\left(p|\tilde{\theta}_i\right)$, $i \in \mathcal{C}$. Dealers' bid functions are also symmetric, but up to their private signal *and* order flow information, i.e. $y_i^d(p|s_i,z_i) = y^d(p|s_i,z_i) = y^d\left(p|\tilde{\theta}_i\right)$, $i \in \mathcal{D}$.

With this equilibrium assumption, we can define the distribution of the market clearing price. Market clearing price is a random variable, P^c, mapping the state of the world, $\left(Q,\mathbf{s^c},\mathbf{s^d},\mathbf{z}\right)$, or simply $\left(Q,\overrightarrow{\theta}\right)$, into prices through equilibrium strategies. This random variable is thus summarized by a function $P^c\left(Q,\mathbf{s^c},\mathbf{s^d},\mathbf{z},\mathbf{y^c}(\cdot|s),\mathbf{y^d}(\cdot|s,z)\right)$ (or simply $P^c(Q,\theta)$ which we will sometimes abbreviate as P^c). The distribution of P^c from the perspective of dealer i, for whom $z_i=\emptyset$ is given by:

$$\Pr(p \geq P^c|s_i,z_i) = E_{\left\{S_{j\in\mathcal{C}\cup\mathcal{D}\setminus i},Z_{k\in\mathcal{D}\setminus i}\right\}} I\left(Q - \sum_{j\in\mathcal{C}} y^c(p|S_j) - \sum_{k\in\mathcal{D}\setminus i} y^d(p|S_k,Z_k) \geq y^d(p|s_i,\emptyset)\right) \tag{2}$$

where $E_{\{\cdot\}}$ is an expectation over other dealer and customers' private information, and $I(\cdot)$ is the indicator function.

If dealer i instead observes customer m's bid function, i.e. $z_i = \{y^c(p|s_m)\}$

$$\Pr(p \geq P^c|s_i, z_i) =$$
$$E_{\{S_{j\in\mathcal{C}\setminus m\cup\mathcal{D}\setminus i}, Z_{k\in\mathcal{D}\setminus i}\}} I\left(Q - \sum_{j\in\mathcal{C}\setminus m} y^c(p|S_j) - \sum_{k\in\mathcal{D}\setminus i} y^d(p|S_k, Z_k) \geq y^d(p|s_i, z_i) + y^c(p|s_m)\right) \quad (3)$$

These probabilities will play a crucial role in equilibrium characterization. Note that the main difference of equation (2) compared to equation (3) is that the dealer conditions on her customer's bid, instead of taking an expectation over the customer's private information.

3.1 Equilibrium strategy of a bidder in a private value auction

In this subsection we describe equilibrium behavior of a bidder in a private value setting, which is our null hypothesis. The discriminatory auction version of Wilson's model with private values has been previously studied in Hortaçsu (2002a). Kastl (2008) extends this model to the empirically relevant setting in which bidders are restricted to use step functions with limited number of steps as their bidding strategies. He shows that an equilibrium in this game exists and provides an implicit characterization of equilibrium bidding strategies:

Proposition 1 *Suppose values are private and assumptions 1-4 hold, then in any Bayesian Nash Equilibrium of a Discriminatory Auction, for almost all θ_i, with a bidder of type $\tilde{\theta}_i$ submitting $\hat{K}\left(\tilde{\theta}_i\right) \leq K$ steps, every step k in the equilibrium bid function $y\left(\cdot|\tilde{\theta}_i\right)$ has to satisfy:*

$$v\left(q_k, \tilde{\theta}_i\right) = b_k + \frac{\Pr\left(b_{k+1} \geq P^c|\tilde{\theta}_i\right)}{\Pr\left(b_k > P^c > b_{k+1}|\tilde{\theta}_i\right)}(b_k - b_{k+1}) \quad (4)$$

$\forall k \leq \hat{K}\left(\tilde{\theta}_i\right)$ such that $v\left(q, \tilde{\theta}_i\right)$ is continuous in a neighborhood of q_k.

Notice that if signals were independent, the probabilities in (4) would not be conditional on $\tilde{\theta}_i$, but would still be of course a function of the submitted bid as can be seen in equations (2) and (3).

3.2 Best-response condition for a bidder in an auction with interdependent or common values

Let $\theta_i = (S_i, Z_i)$ and $\theta_{-i} = (S_{-i}, Z_{-i})$. In Appendix B, we show that the necessary condition for optimality at k^{th} step in an interdependent value environment with step function bidding is:

$$\begin{aligned} E_{\theta_{-i}|\tilde{\theta}_i, P^c \in (b_{k+1}, b_k), q_k} \left[v\left(q_k, \theta_i, \theta_{-i}\right) \right] + \frac{\Upsilon(\tilde{\theta}_i, b_k, b_{k+1}, q_k)}{\Pr\left(b_k > P^c > b_{k+1} | \tilde{\theta}_i\right)} \\ = b_k + \frac{\Pr\left(b_{k+1} \geq P^c | \tilde{\theta}_i\right)}{\Pr\left(b_k > P^c > b_{k+1} | \tilde{\theta}_i\right)} \left(b_k - b_{k+1}\right) \end{aligned} \tag{5}$$

for $v\left(q, \tilde{\theta}_i, \theta_{-i}\right)$ continuous in a neighborhood of q_k for a.e. $\left(\tilde{\theta}_i, \theta_{-i}\right)$. In this expression,

$$\begin{aligned} \Upsilon(\tilde{\theta}_i, b_k, b_{k+1}, q_k) = \frac{\partial \Pr\left(b_k = P^c | \tilde{\theta}_i\right)}{\partial q_k} E_{\theta_{-i}|\tilde{\theta}_i, P^c = b_k, q_k} \left[V\left(Q_i^c, \theta_{-i}, \theta_i\right) \right] \\ + \frac{\partial \Pr\left(b_{k+1} = P^c | \tilde{\theta}_i\right)}{\partial q_k} E_{\theta_{-i}|\tilde{\theta}_i, P^c = b_{k+1}, q_k} \left[V\left(Q_i^c, \theta_{-i}, \theta_i\right) \right] \\ + \frac{\partial \Pr\left(b_k > P^c > b_{k+1} | \tilde{\theta}_i\right)}{\partial q_k} E_{\theta_{-i}|\tilde{\theta}_i, P^c \in (b_{k+1}, b_k), q_k} \left[V\left(q_k, \theta_{-i}, \theta_i\right) \right] \end{aligned}$$

where Q_i^c is the quantity won by bidder i if his bid q_k is not fully-allotted at b_k or b_{k+1} (since the amount of allotment depends on other bids, Q_i^c is a random variable).

The term at the right-hand side of (5) is identical to the term in (4) for the private values case, and reflects bid-shading in response to the expected distribution of residual supply. However, the left-hand side of (5) can no longer be interpreted as a model primitive, as in the private values case. In an interdependent value setting, the realization of market clearing price, P^c, conveys information to the bidder regarding the ex-post value of the security. Some of this effect is captured in the first-term, $E_{S_{-i}, Z_{-i}|s_i, z_i, P^c \in (b_{k+1}, b_k), q_k} \left[v\left(q_k, s_i, z_i, S_{-i}, Z_{-i}\right) \right]$, which denotes the expected marginal value, conditional on winning quantity q_k at a market-clearing price P^c within $[b_{k+1}, b_k]$. This term is a direct analog of the *pseudo-value* term that emerges in the single unit first-price auction context (Haile, Hong, and Shum (2003)).[16] However, there is an additional term which captures how the

[16]Note, however, that P^c is potentially a complicated function of bidder signals that depends on equilibrium bidding

inference about the expected value from winning changes when the bidder marginally changes his quantity demand q_k (but keeps his price bids b_k and b_{k+1} the same), which may lead to a shift in the distribution of market clearing prices. The shift in the distribution of market clearing prices leads the bidder to change her inference regarding the value of the securities she is winning; thus she has to adjust her expected valuation for the inframarginal units as well as the marginal unit.

In a recent paper, De Castro and Riascos (2009) have also characterized necessary conditions for optimality in the multi-unit discriminatory auction with interdependent values (see their Example 11, p. 566). Their formulation of the game differs in that they restrict the quantity bids to be discrete, but the price bids to be continuous and strictly decreasing for every quantity increment. Under this specification, they show that the necessary condition (5) does not have the extra term, $\Upsilon(\cdot)$. Unfortunately, the assumption that price bids have to be strictly decreasing does not hold in our data, as the bid functions are characterized by horizontal sections where the price is constant for a wide range of quantities. However, as we will show in Section 5.2, the extra $\Upsilon(\cdot)$ term in our derivation will not be quantitatively important.

As usual in interdependent value environments, when the value is increasing in rivals' signals, a "winner's curse" effect arises as the event of winning is informative about information held by the rivals. We will make use of this observation in Section 4.2 to construct a test of a restriction that the theory model under interdependent values predicts. The idea is similar to Haile, Hong and Shum (2003), and utilizes the stochastic ordering of the distributions of pseudo-values when z_i changes.

3.3 Estimation of marginal values

Under the null hypothesis of private values, using the necessary conditions (4), we can obtain point estimates of marginal valuations at submitted quantity-steps nonparametrically as described in Hortaçsu (2002a) and Kastl (2006). The "resampling" method that we employ in these papers is to draw from the empirical distribution of bids to simulate different realizations of the residual supply function that can be faced by a bidder, thus obtaining an estimator of the distribution of

strategies, and thus conditioning on it may not be straightforward. In the single unit context, the conditioning can be done directly on order-statistics of bidder signals.

the market clearing prices.[17]

We now discuss how we adapt this method into the present context with dealers and customers. Recall that we have two classes of bidders: N_d potential *dealers* (in index set $\mathcal{D}$) and N_c potential *customers* (in index set $\mathcal{C}$). If we further assume (conditionally) independent private values, customers have iid signals with marginal distribution $F^C(S_i^c)$. Each dealer also observes a private signal, S_i^d, which is also iid across dealers. We also assume that (S_i^d, Z_i) are iid across dealers $i \in \mathcal{D}$, but we allow (S_i^d, Z_i) to be correlated within dealer. In this context, the resampling algorithm should be modified in the following manner: to estimate the probability in equation (2), we draw N_c customer bids from the empirical distribution of customer bids (we augment the data with zero bids for non-participating customers). Now, to account for the asymmetry induced across dealer bids due to the observation of customer signals, we do the following: conditional on each customer bid, $y^C(p, S_j)$, drawn, draw a corresponding dealer's bid as follows: (i) If a zero customer bid is drawn, draw from the pool of dealers' bids, which have been submitted without observing any customer bid, or (ii) If a non-zero customer bid is drawn, draw from the pool of dealers' bids, which have been submitted having observed a "similar" customer bid.[18] After drawing N_c customer bids, continue drawing from the pool of bids submitted by uninformed dealers until $N_d - 1$ dealer bids are drawn. Obtain the market clearing price, and repeat.

To estimate the probability in (3), we need to take into account the full information set of the dealer. This is achieved by a slight modification of the above procedure: fixing a dealer, who has seen a customer bid, we draw $N_c - 1$, rather than N_c, customer bids, and take the observed

[17] Specifically, in the case where *all N bidders are ex-ante symmetric, private information is independent across bidders* and the data is generated by a symmetric Bayesian Nash equilibrium, the resampling method operates as follows: Fix a bidder. From all the observed data (all auctions and all bids), draw randomly (with replacement) $N-1$ actual bid functions submitted by bidders. This simulates one possible state of the world from the perspective of the fixed bidder, a possible vector of private information, and thus results in one potential realization of the residual supply. Intersecting this residual supply with the fixed bidder's bid we obtain a market clearing price. Repeating this procedure a large number of times we obtain an estimate of the full distribution of the market clearing price conditional on the fixed bid. Using this estimated distribution of market clearing price, we can obtain our estimates of the marginal value at each step submitted by the bidder whose bid we fixed using (4).

[18] Ideally, we would draw from dealers' bids that have been submitted after observing *exactly* the same customer bid. However, our data has the practical limitation that customer bids are typically unique within or across auctions. Thus, the "conditional" draws often consist of repeatedly drawing the same customer and "informed" dealer pair. Of course, asymptotically, we expect the number of dealer bids corresponding to a given customer bid to increase; however, in small samples, this is rarely true. In practice, we thus employ a "smoothing" strategy: instead of drawing informed dealer bids that exactly correspond to a given customer bid, draw bids from dealers who saw customer bids that are "close" to the given customer bid (similarly to kernel estimation). We describe this procedure and the asymptotic properties of our estimator in section 3.4.

customer bid along with the dealer's own bid as given when calculating the market clearing price, i.e., we subtract the *actual observed* customer bid from the supply before starting the resampling procedure.

Although handling a general affiliated private values setting is difficult, our resampling approach carries over to the case when dealers' and customers' signals are *conditionally* independent within class; i.e. conditional on auction-level covariates observed by all bidders, their private signals are independent.[19] Of course, an important concern is whether the econometrician can condition on the same set of covariates that bidders observe; we will discuss this concern in Section 4.3 below. We can, however, address this concern by resampling using bids from a single auction at a time. Moreover, in Section 4.4, we conduct several tests for independence, and find considerable support for the conditional independence assumption.

Note that the necessary conditions described in Proposition 1 apply only to dealers, whose bids are not revealed to anyone. We base our test exclusively on the behavior of dealers. Of course, customers may have a (potentially complex) strategic response to the fact their bids are observed by dealers. This may generate an alternative testing strategy: the equilibrium effect of how a bidder adjusts her bid when she knows that her bid will be observed by her rival is likely to differ when values are private and when values are interdependent. We do not pursue this approach, however, and utilize the (empirical) distribution of customer bids only to characterize dealers' best-responses.

3.4 Inference

As suggested above our test will be based on comparing two sets of marginal valuation estimates. Therefore, we have to be able to account for the sampling error in these estimates. It is easy to see from equation (4) that these estimates are a non-linear function of the distribution of the market

[19] The resampling procedure can, in principle, be modified for a more general affiliated private values setting, where we specify the vector of customers' signals, $\mathbf{S}^c$ to have a joint distribution $F^C(\mathbf{S}^c)$, and let $F^D\left((S_1, Z_1), ..., (S_{N_d}, Z_{N_d})\,|\mathbf{S}^c\right)$ be the joint distribution of dealer signals and orderflow information conditional on the vector of customers' signals, $\mathbf{S}^c$ (and implicitly also conditional on customers' equilibrium strategies). To simulate possible states of the world, and thus the distribution of the market clearing price from the perspective of a dealer of type who submitted a bid $y^d\left(p|s_i, z_i\right)$, we draw with replacement whole vectors of $N_d + N_c - 1$ bids of bidders other than i, where we draw only from those auctions in which the *exact* same bid, $y^d\left(p|s_i, z_i\right)$, and orderflow information, z_i, was submitted by i. Unfortunately, in our data set, we do not observe more than 1 auction with bidder i having observed the same z_i and submitting the same bid $y^d\left(p|s_i, z_i\right)$.

clearing price, which is estimated by the resampling method described above. Let us rewrite (4) as

$$v\left(q_k,\tilde{\theta}_i\right)=b_k+\frac{H\left(b_{k+1}\right)}{G\left(b_k\right)-H\left(b_{k+1}\right)}\left(b_k-b_{k+1}\right)$$

where $H\left(X\right)$ (resp. $G\left(X\right)$) is the probability that market clearing price is weakly (resp. strictly) lower than X.

Define an indicator of excess supply at price X given bid functions $y_1,...,y_{N_c+N_d-1}$ and i's own bid $y_i\left(X|\tilde{\theta}_i\right)$ as follows:

$$\Phi\left(y_1,...,y_{N_c},...,y_{N_c+N_d-1};X\right)=I\left(Q-\sum_{j=1}^{N_c+N_d-1}y_j\left(X|\theta_j\right)\geq y_i\left(X|\tilde{\theta}_i\right)\right)$$

Consider the following statistic (on which we will be base our estimator of $H(X)$) based on all subsamples (with replacement) of size $(N_c+(N_d-1))$ consisting of N_c customers' bids and N_d-1 dealers' bids from the full sample of $\left(N_c+N_d\right)T$ datapoints:

$$\begin{aligned}\xi\left(\hat{F};X,h_T\right)=&\frac{1}{\left(N_cT\right)^{N_c}}\frac{1}{\left(N_dT\right)^{(N_d-1)}}\sum_{\alpha_1=(1,1)}^{(T,N_c)}...\sum_{\alpha_{N_c}=(1,1)}^{(T,N_c)}\sum_{\alpha_{N_c+1}=(1,N_c+1)}^{(T,N_c+N_d)}...\sum_{\alpha_{N_c+N_d-1}=(1,N_c+1)}^{(T,N_c+N_d)}\\&\Phi\left(y_{\alpha_1},...,y_{\alpha_{N_c}},y_{\alpha_{N_c+1}},...,y_{\alpha_{N_c+N_d-1}},X\right)W\left(\alpha_1,...,\alpha_{N_c},\alpha_{N_c+1},...,\alpha_{N_c+N_d-1};h_T\right)\end{aligned}\tag{6}$$

where $\alpha_i\in\{(1,1),(1,2),...\left(1,N_c+N_d\right),...,\left(T,N_c+N_d\right)\}$ is the index of the bid in the subsample and $\hat{F}$ is the empirical distribution of bids.[20] To understand the summations and indexes, observe that the data can be viewed as a table with T auctions and N_c+N_d bidders (hence index (t,i) corresponds to auction t and bidder i) and we are drawing subsamples of size $(N_c+(N_d-1))$ since we are constructing residual supplies from perspective of a dealer. The first N_c sums are over the indices of customers' bids and last N_d-1 sums are over the indices of dealers' bids. We have N_dT dealer bids, and there are $\left(N_dT\right)^{(N_d-1)}$ subsamples of size (N_d-1). Finally, $W\left(\cdot\right)$ denotes

[20] Note that since in our case a bid is a point in at most $2K$-dimensional space, $\hat{F}$ is simply the empirical probability distribution over such points.

the kernel weights defined by

$$W\left(\alpha_1,...,\alpha_{N_c},\alpha_{N_c+1},...,\alpha_{N_c+N_d-1};h_T\right) =$$
$$= \frac{\prod_{j=\alpha_{N_c+1}}^{\alpha_{N_c+N_c}}\left(K\left(\frac{\left\|z_j-y_{j-N_c}\right\|}{h_T}\right)\right)\prod_{j=\alpha_{2N_c+1}}^{\alpha_{N_c+N_d-1}}\left(K\left(\frac{\left\|z_j-0\right\|}{h_T}\right)\right)}{\sum_{\alpha'_1=(1,1)}^{(T,N_c)}\cdots\sum_{\alpha'_{N_c+N_d-1}=(1,N_c+1)}^{(T,N_c+N_d)}\prod_{j=\alpha'_{N_c+1}}^{\alpha'_{N_c+N_c}}\left(K\left(\frac{\left\|z_j-y_{j-N_c}\right\|}{h_T}\right)\right)\prod_{j=\alpha'_{2N_c+1}}^{\alpha'_{N_c+N_d-1}}\left(K\left(\frac{\left\|z_j-0\right\|}{h_T}\right)\right)} \tag{7}$$

where $K(\cdot)$ is a bounded kernel with compact support[21] and h_T is the bandwidth satisfying $h_T \to 0$, $Th_T^5 \to 0$ and $Th_T \to \infty$ as $T \to \infty$. Given that the bids are multidimensional, the kernel should be multidimensional with the dimension equal to that of the price grid. In practice, we use the difference in quantity-weighted average bids as the norm $\|\cdot\|$ where we let $\|\emptyset\| = 0$ and $\|x-\emptyset\| = x$. The subsample with the highest kernel weight would have the actual observed customer bids associated with the first N_c drawn dealer's bids exactly correspond to the N_c drawn customer bids and the last $N_d - 1 - N_c$ dealer bids would have $z = \emptyset$.

Notice that the statistic ξ defined as above is for an uninformed dealer, i.e., one with $z = \emptyset$. For an informed dealer, the above test statistic must be slightly modified by drawing one less customer bid and by fixing the observed customer bid when evaluating the indicator $\Phi(\cdot)$ in each subsample.

Observe also that it is not feasible to compute ξ by summing over all permutations of dealer and customer bids. However, our resampling estimator, $\hat{H}^R(X)$, is a simulator of $\xi\left(\hat{F};X\right)$, for which M subsamples are randomly drawn rather than all $(N_cT)^{N_c}(N_dT)^{(N_d-1)}$. We choose $M = 5,000$ in our application to make sure the simulation error is not an important factor.

We study the asymptotics of our estimator in detail in Appendix C. Here we only state our main result, which follows from Theorem 1 in Stute (1991).

Proposition 2 *Assume that (i) $h_T \to 0$, $Th_T \to \infty$, $Th_T^5 \to 0$, (ii) $\frac{M}{T} \to \infty$ and (iii) K is bounded and has compact support, then*

$$(Th_T)^{\frac{1}{2}}\left(\hat{H}^R(X) - H(X)\right) \to N\left(0,\sigma^2\right)$$

where $\sigma^2 = \sum_{j=1}^{N_c+N_d-1}\sum_{l=1}^{N_c+N_d-1} I\left[y_j = y_l\right]\left[\Phi_{jl}(\mathbf{y}) - \Phi^2(\mathbf{y})\right]\int K^2(u)\frac{du}{f(y_j)}$,

[21] Examples of possible kernel functions are Epanechnikov, triangular, rectangular etc. In our application, we use the uniform kernel: $K(u) = \frac{1}{2}I_{(|u|\leq 1)}$.

and where when $y_j = y_l$,

$\Phi_{jl}(\mathbf{y}) = E\left[\Phi\left(y_1, ..., y_{j-1}, Y, y_{j+1}, ..., y_{N_c+N_d-1}\right) \times \Phi\left(y_{N_c+N_d}, ..., y_{N_c+N_d+l-2}, Y, y_{N_c+N_d+l}, ..., y_{2(N_c+N_d-1)}\right)\right]$

and when $y_j \neq y_l$, $\Phi_{jl}(\mathbf{y}) = 0$.

In our empirical application, we will use bootstrap confidence intervals, which are readily generated by iterating the resampling scheme used for point estimates on bootstrap samples of the bid data. The validity of the bootstrap in this case follows from the results for V- and U-statistics of Bickel and Freedman (1981, Theorem 3.1) using the fact that the variance and any covariances of our kernel $\Phi(.)$ in the V-statistic (6) are bounded.

4 Test Specification

The main idea behind our test for private values, as described in the introduction, is to find instances where a dealer observes customer information, and to test whether the estimated marginal valuation rationalizing that dealer's bid remains constant before and after accounting for the residual supply information provided by that customer bid. A first practical challenge in implementing such a test arises from the fact that bids in multiunit auctions are submitted as discrete price-quantity pairs. Unfortunately, we can only obtain point-identification for marginal values at the discrete price-quantity points (McAdams (2008), Kastl (2008)). Since bidders may change the discrete bid steps they submit after they receive extra information, we face the challenge of testing the equality of non-point identified parameters. Therefore our test is based on comparing the two sets of estimates of marginal values, $k = BI, AI$, in situation where a bid has been submitted for the same quantity. In such cases, under private values the two estimates of marginal values should coincide asymptotically and thus could differ in a finite sample only due to a sampling error.

Consider the test statistic:

$$T_i(q) = \left|\hat{v}_i^{BI}(q, s_i) - \hat{v}_i^{AI}(q, s_i)\right| \tag{8}$$

where $\hat{v}_i^{BI}(q, s_i)$ is the estimated marginal value for share q before the information was revealed and similarly $\hat{v}_i^{AI}(q, s_i)$ is the estimated marginal value for share q after the additional information arrived. We obtain the critical value for this test statistic using bootstrap. Using B bootstrap draws, the critical values are computed as follows:

$$\tilde{c}_{1-\alpha} = \inf \left\{ x : \frac{1}{B} \sum_{b=1}^{B} 1 \left\{ \tilde{T}_b \leq x \right\} \geq 1 - \alpha \right\} \tag{9}$$

where $\tilde{T}_b$ is the re-centered value of the test statistic for bootstrap draw b, where the re-centering is by the value of the test statistic under $\hat{F}$, i.e., value of the test statistic evaluated on the sample, $\tilde{T}_0$. For each bootstrap draw of the test statistic, the marginal value is re-estimated by the resampling method described earlier, where a new sample of bid functions from which this resampling is performed is drawn.[22]

4.1 Joint hypothesis test

So far we discussed how we can test for the null of private values based on a set of estimated values for a given bidder, whom we observe submitting a bid before and after information arrives. In our application, however, we observe multiple such bidders. Therefore, the proper way to test whether the null hypothesis of private values can be rejected is to test this null jointly for all bidders. The most direct way to test a joint hypothesis is to adjust the desired p-value to satisfy the Bonferroni inequality, i.e., instead of comparing each individual hypothesis to a p-value of 0.05, compare it to a p-value of $\frac{0.05}{\#\ of\ hypotheses}$. This method, while controlling the familywise error rate (type-I errors), is usually regarded as being too conservative.[23] An alternative way is to construct a test statistic which takes into account all of the tested hypotheses at once. Motivated by a χ^2 test, we construct

[22]To construct a bootstrap sample of bid functions, we have to follow a procedure similar to the conditional resampling. In constructing these bootstrap samples we need to include also the 'zero' bids for those potential bidders that do not end up actually submitting a bid. We start by drawing N_c customer bids with replacement giving $\frac{1}{TN_c}$ probability to each (where $T \geq 1$ is the number of auctions which we pooled together for resampling). Conditional on having drawn a non-zero customer's bid, we draw from the observed sample N_d^c dealer bids submitted following the same customer's bid with replacement giving $\frac{1}{N_d^c}$ probability to each such dealer bid. Conditional on drawing a zero customer bid, we draw from dealers' bids submitted without knowledge of any customer's bid putting equal probability on each.

[23]See Romano, Shaikh and Wolf (2008) for a detailed discussion of multiple testing.

the following "sum of squares" of the test statistics T_i:

$$SSQ_T = \sum_i \left(\frac{T_i}{\sigma_{T_i}}\right)^2 \tag{10}$$

where we studentize these test statistics by scaling T_i by their (bootstrap) standard deviations.

Another test statistic that takes into account all of the hypotheses at once is the maximum (first-order statistic), among bidders i, of T_i:

$$FOS_T = \max_{i \in \mathcal{D}} \frac{T_i}{\sigma_{T_i}} \tag{11}$$

We approximate the asymptotic distribution of these test statistics again by bootstrap. For example, in the case of SSQ, the null hypothesis of private values is rejected on level α, when $S\hat{S}Q > SSQ_{1-\alpha}$, where $S\hat{S}Q$ is the sample value of the statistic, and $SSQ_{1-\alpha}$ is the $(1-\alpha)^{th}$ quantile of the boostrap distribution of the test statistic.

4.2 Restriction under the alternative hypothesis

Since our null hypothesis is a joint one of private values and of other assumptions of the model, we now propose an additional testable restriction which the theoretical model predicts under the alternative hypothesis of common or interdependent values. Consider two (ex ante symmetric) dealers i and j with the same (ex post) private information signal, $s_i = s_j = s$, but different observed customer bids that can be totally ordered and hence the econometrician can infer also the ordering of customers' signals.[24] Notice that since the dealers are ex ante symmetric and have the same private signal s, they would submit the same bid (in a symmetric pure strategy BNE), $y(p|s, \emptyset)$, in the absence of any order flow information. The ordering of customer bids implies that customer submitting the bid z_i has a higher private signal (is more optimistic) than the one submitting z_j. Therefore, if values were interdependent, we would be able to rank the pseudovalues (the first term on the left-hand side of equation (5)) of the dealers i and j with respective information

[24]This implicitly assumes that customers use non-decreasing pure strategies, so that $b(q, z_i) > b(q, z_j)\ \forall q \Rightarrow z_i > z_j$. Although we provide an example of non-decreasing, thus fully revealing, strategy in the private value case in Appendix D, proving existence of such an equilibrium more generally and especially in the interdependent values case is outside of the scope of this paper.

sets $\tilde{\theta}_i = (s, z_i)$ and $\tilde{\theta}_j = (s, z_j)$:

$$E_{\theta_{-i}|\tilde{\theta}_i, P^c \in (b_{k+1}, b_k), q_k}\left[v\left(q_k, \theta_i, \theta_{-i}\right)\right] > E_{\theta_{-j}|\tilde{\theta}_j, P^c \in (b_{k+1}, b_k), q_k}\left[v\left(q_k, \theta_j, \theta_{-j}\right)\right] \tag{12}$$

The implementation of this testing approach, however, is complicated by the fact that we do not observe the term $\Upsilon(\cdot)$ in the second term of equation (5), and that this term cannot be ranked unambiguously across dealers with the information sets, $\tilde{\theta}_i$ and $\tilde{\theta}_j$, as defined above. Our empirical strategy in Section 5.3 will be to place bounds on this quantity based on observables, and assess whether this term is quantitatively important.

In our application of the test, we have to first find uninformed dealer bids (within an auction) that are close to each other. Since bids are multi-dimensional vectors, we define two bids as being "close" if their quantity-weighted average prices are within 1 basis point. Among these "close" dealer bids, we then look to see if the customer bids received by the dealers can be totally ordered and we construct the difference in quantity-weighted pseudo-values (corrected for the extra Υ term): $\frac{\sum_{k=1}^{K_i} q_k \hat{v}_i(q_k, s|z_i)}{\sum_{k=1}^{K_i} q_k} - \frac{\sum_{k=1}^{K_j} q_k \hat{v}_j(q_k, s|z_j)}{\sum_{k=1}^{K_j} q_k}$. We then test whether this difference is significantly greater than 0.

Another important implementation detail is how we apply the resampling algorithm in the interdependent values case. Although, as explained in Section 3.3, we can allow for a general affiliated value specification using our estimation approach, we only have limited panel of auctions. We will therefore restrict our attention to a specification where dealers' and customers' signals are iid within class of bidders *conditional* on variables that are observed by all bidders (but not necessarily by the econometrician). To preserve the conditional iid assumption in our estimation, our resampling algorithm will only utilize data from a single auction at a time.

4.3 Unobservable heterogeneity across auctions

A practical challenge in implementing the testing procedure is the presence of auction-level covariates that are observed by the bidders, but not by the econometrician. A big challenge in the test for common values in a single unit setting proposed in Haile, Hong and Shum (2003) lies in the ability of the econometrician to isolate exogenous variation in the number of competing bidders across auctions. HHS thus propose a method that deals with the endogeneity of the number of participants. In contrast, our testing strategy is based on looking at modification of bids by a given bidder, *within* the same auction. Thus, at least in principle, we do not have to rely on across auction information to construct our test statistic. However, our estimates of marginal values (under the null of independent private values) will be more precise if we can pool bid data across auctions. Pooling data across auctions, on the other hand, may lead to biases in our estimation of bid shading if auction-level unobservables are present. We will therefore experiment with different levels of data pooling.

A more subtle concern regarding omitted variable bias may arise in our application if the changes in dealer bids are caused by innovations to the bidders' information set that are not observed by the econometrician. The presence of such omitted pieces of information biases our estimates of bidders' optimal bid shading. We can test for the presence of such omitted information flows if the information is *public*, i.e. observed by multiple bidders. In our data, as we observe exact time of each bid submission, we can distinguish a change in bid due to more information coming from a customer from a change in bid due to some new public information. In the latter case, conditional on some small time window, all adjustments by dealers should be positively correlated, whereas in the former case they should be independent. Therefore if we subject to our test only those changing bids that are not accompanied by similar changes in rival's bids, we can be more confident that no commonly observed (but unobserved by us) piece of information is biasing our test.[25]

4.4 Testing independence of private information

Since we specify the null hypothesis as independent private values to facilitate estimation using our data, we first test the independence part directly using the bid data. In particular, we now offer several alternative tests for independence. To test whether dealer bids within an auction are independent, we first compiled all the dealer bids that were submitted before the dealer saw any customer bid. We then randomly split the (quantity-weighted) bids into two halves. (When the number of bids was odd, we dropped one bid.) We then computed four test statistics for the bivariate samples (one for each auction in our data set) constructed using the random split. The

[25]See Appendix A for more explicit treatment of this issue.

first three of these are correlation measures: the Pearson correlation, Spearman rank correlation, and Kendall's tau. The fourth test statistic is the Blum-Kiefer-Rosenblatt (1961) nonparametric test for independence.[26] We use Mudholkar and Wilding's (2003) tabulation of critical values for the BKR test statistic.

Before we report our results, let us emphasize that the test statistics are computed separately for *each* auction in the data set. Since we are running the test separately for each auction, the null hypothesis is independence at the auction level; i.e. conditional independence. When we report the results, however, we will count the number of auctions for which we reject the null hypothesis, rather than reporting the result of the test separately for each auction. An important issue with our testing strategy, however, is that the way we split the data in each auction into two is arbitrary. In order to make sure we did not –by chance– split the data in a way that favors independence, we repeat the (auction-by-auction) test 100 times.

We first report the results from the 3 month sample. Since the Blum-Kiefer-Rosenblatt (BKR) test is only applicable when $N > 4$, we only considered auctions with at least 10 dealer bids, which reduced our sample to 64 auctions. We considered 100 random splits of the sample when constructing the test statistics, and recorded the number of times the bids from an auction rejected the null hypothesis of independence. Over 100 iterations, the average number of auctions (among 64 auctions) for which the Pearson coefficient was significantly (at the 5% level) different from zero was 3.68. The Spearman rank correlation was different from zero on average for 1.9 auctions, Kendall's tau test led to rejection for 1.33 auctions, and the BKR test led to rejection in 3.67 auctions. We also looked at the maximum number of auctions for which the test statistic was rejected for any given random split of the data. The Pearson test was rejected for a maximum of 8 auctions, Spearman 5 auctions, Kendall's tau 5 auctions, and BKR test was rejected for a maximum of 9 auctions (out of 64 auctions in the data).

In the 12 month sample, only 33 auctions had at least 10 dealer bids. We again consider 100 random bivariate splits of the dealer bids to construct the independence test statistics. Over 100 iterations, the average number of auctions (among 33 auctions) for which the Pearson coefficient

[26]Mudholkar and Wilding (2003) conduct an extensive Monte Carlo analysis of these 4 different test statistics for testing independence and find that none of them strictly dominates the others in terms of power.

was significantly (at the 5% level) different from zero was 1.63. The Spearman rank correlation was different from zero on average for 0.87 auctions, Kendall's tau test led to rejection for 0.56 auctions, and the BKR test led to rejection in 1.65 auctions. Within 100 random splits across the 33 auctions in the data, the Pearson test was rejected for a maximum of 5 auctions, Spearman 4 auctions, Kendall's tau 3 auctions, and BKR test was rejected for a maximum of 5 auctions.

Although the above does not constitute a formal joint hypothesis test (that the independence hypothesis is correct for all 64 or 33 auctions), the fact that very few auctions in our data violate the null hypothesis of independence individually suggests that the independence assumption is reasonable.

Finally, we conduct a Wilcoxon Rank Sum test of the null hypothesis that two random independently drawn samples are identically distributed, where we test equality of two conditional distributions of signals. We report results of our tests only for dealers, but even stronger results obtain for customers. Our main findings are: (i) the data are consistent with private information being independently distributed *within* auctions across bidders; (ii) unobserved heterogeneity *across* auctions is an important factor which leads to violation of private information being identically distributed *across* auctions; (iii) data are not inconsistent with the independence assumption across auctions when the unobserved heterogeneity is indirectly accounted for. We first perform the Wilcoxon test *within* each auction t, so as to avoid any concern for unobservable heterogeneity across auctions. To do this, we split the sample of bids *within* each auction into two halves, leave out the first dealer in each half (i.e., condition on his bid), and we test whether the two samples are identically distributed: $H_0 : F\left(b|b_{1t}\right) = F\left(b|b_{\frac{N^*_{dt}}{2}+1}\right)$, where the first sample consists of $b_{2t}, ..., b_{\frac{N^*_{dt}}{2}}$ and the second sample of $b_{\frac{N^*_{dt}}{2}+2}, ..., b_{N^*_{dt}}$, where N^*_{dt} is the number of non-zero dealer bids in auction t.[27] This test rejects the null hypothesis in 9 out of 116 auctions of 3-months T-bills.

A concern about the *within* auction test is that N is not very large. Thus, we run the Wilcoxon test *across* consecutive auctions by specifying the null hypothesis as: $H_0 : F\left(b|b_{1t}\right) = F\left(b|b_{1(t+1)}\right)$. Here, the test rejects in over 100 pairs of consecutive auctions. However, unobserved heterogeneity

[27] By leaving out the first dealer in each half, we wanted to make it explicit that we are testing for the equality of two conditional distributions: $F(B|B_{it} = b_{it})$ and $F(B|B_{jt} = b_{jt})$ where $b_{it} \neq b_{jt}$ and B is a rival's bid. Of course since B_{it} is continuously distributed, $b_{it} \neq b_{jt}$ will be satisfied with probability one and hence we could have in principle performed an unconditional test.

may confound the interpretation of this *across* auction Wilcoxon test: if the true distribution of bids, conditional on the unobservable U, is $F(B|U)$ and $u_t \neq u_{t+1}$, this test might rejected because of the inequality of the distribution rather than because of the lack of independence. Therefore, in our next test, we combine one half of bids from auction t and one half from $t+1$ as one sample, and combine the other halves to create another sample. By combining these halves of data samples, we have effectively created a mixture of $F(B|U = u_t)$ and $F(B|U = u_{t+1})$, which creates two homogeneous samples to apply the Wilcoxon test to.[28] The test rejects in only 4 of the consecutive auction pairs (out of 115). We interpret the results from the performed Wilcoxon ranksum tests as providing evidence that unobserved heterogeneity might indeed be an important factor and that independence of bids cannot be rejected.

5 Results

5.1 Results from 3 month T-bill auctions

In the 116 auctions of 3 month T-bills in our sample, we observed 216 dealer bids that were updated after a customer bid arrived.[29] Figure 1 depicts updating of a bid by one dealer. After observing a relatively low bid by one customer, the dealer submits a new bid which is uniformly weakly below his original bid.[30]

Before updating, these 216 bids consisted of 802 bidsteps (price-quantity pairs) and after updating they consisted of 859 bidsteps. We focus on these updated bids to conduct our tests. We construct a bootstrap sample of 400 replications of the test statistics (using always 5000 resampling draws for estimating each bidder's marginal value) for each of these bidders as defined by (8) and construct the corresponding critical values given by (9).

To illustrate the marginal value estimation procedure, Figure 2 depicts the marginal value

[28] The advantage of this test over the within auction test (i) is the increased sample size: we have $N-1$ realizations from each distribution rather than $N/2-1$.

[29] See Appendix A for further summary statistics regarding bid updates. We did not include instances where the dealer did not revise his bid. If updating decisions are such that conditional on deciding to update, the dealer updates in a manner that is optimal with respect to the private values model, our test is not biased. This assumption holds true if any "costs" to updating are fixed costs that only affect the decision of whether to update or not.

[30] This "parallel shift" of the updated bid is not a general feature of the data, however. Some updated bids cross the original bids.

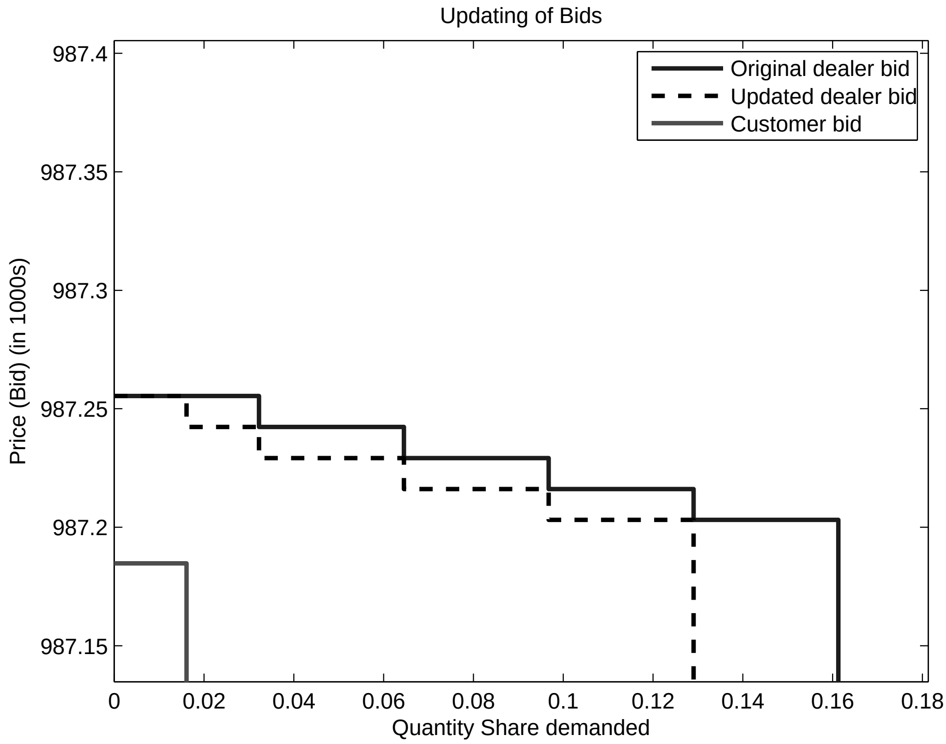

Figure 1: Updating of a dealer's bid

estimation results for the dealer in Figure 1. We run our test on the three steps where the bidder's quantity stayed the same.[31] As can be seen, the confidence intervals on the three steps appear to overlap. However, these confidence intervals do not take into account the covariance between the two marginal valuation estimates. Our test statistic, which is computed using the bootstrap, can account for the covariance, and yields that the null hypothesis that the marginal values rationalizing the original and updated bids are the same cannot be rejected.

To construct the various test statistics, we first estimated bidder's marginal values under different data pooling schemes and bandwidth selections. We first used resamples of bids from the same auction only (the "1 auction" case). This does not pool bid data across auctions, and hence minimizes a potential unobserved heterogeneity problem. However, since the data set used to estimate marginal values is small, the estimation error is potentially large, and, as in our Monte Carlo

[31] The fourth step is not included as the bid equals marginal valuation for the update bid curve as per proposition 1. We estimated the model and constructed the tests using both assumptions on $v(\cdot)$ and the results were qualitatively robust. The subsequently presented results utilize equation (4) for identification.

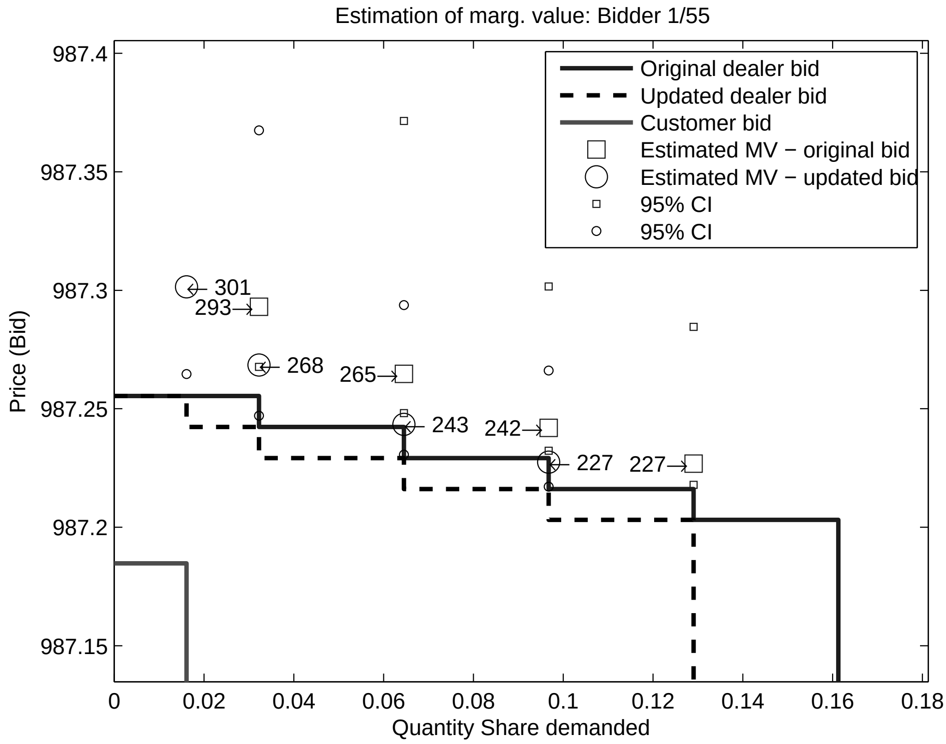

Figure 2: Estimated values - dealer 1

exercises in Appendix E, the power of our tests might be lower than desired. Resampling from a pooled set of auctions that are similar may decrease estimation error, but unobserved heterogeneity across auctions may result in rejections of the null hypothesis for other reasons. To explore this tradeoff, we report the same estimates that are obtained if we pool data across 2 and 4 consecutive auctions respectively (this assumes that the economic environment is stable across 2 and 4 week periods, respectively) and for three different bandwidths.

Our Monte Carlo exercise described in Appendix E suggests that the various testing approaches we utilize have complementary size-power tradeoffs in small samples, thus we will display results across an array of test statistics. It is not our goal to derive a "hard" treshold which will allow us to reject the null hypothesis, instead we will use the conventional level $\alpha = 0.05$ to obtain our critical values and we will focus on contrasting patterns that emerge after applying our testing procedure to data from 3-month and 12-month T-bill auctions respectively. In Table 3, we report results from hypothesis tests on individual bidders' updating behavior. First, we report results based on the equality test statistic (equation (8)) computed separately for each updated bid (the critical values were obtained using bootstrap using $\alpha = 0.05$). We find that we are able to reject the null (at the 5% level) only in about 10% of the individual hypotheses when we estimate marginal values using data from a single auction. When we increase the number of auctions used to estimate marginal values, we are able to reject more of the individual hypotheses: 13% of the individual hypotheses are rejected when we resample from 2 neighboring auctions (i.e., auctions in 2 adjacent weeks), and 16% are rejected when we resample from 4 neighboring auctions.

The individual hypothesis tests suggest that the null of private values is not easily rejected in our data. It appears that as we increase the size of the sample used to estimate marginal values, the rejection rates increase. However, as we noted earlier, increasing sample size may lead to overrejection of the null due to the introduction of unobserved heterogeneity as well.

In Table 4, we report results from the tests based on the joint test statistics defined in Section 4.1.[32] Once again, we estimate marginal values using 1, 2, and 4 neighboring auctions. Our Monte Carlo exercises in Appendix E suggested that studentization increases power against the alternatives. We thus decided to studentize the test statistics.

For each case, we use the bootstrap to calculate the critical values of the sum-of-squared studentized differences statistic (SSQ) and the studentized first-order test statistic (FOS).[33] Since the FOS may be overly demanding, we also report the 95^{th} percentile of the studentized test statistic distribution. Based on the studentized joint hypothesis tests and using the treasury bill with 3-months maturity, Table 4 shows that we fail to reject the null hypothesis across all resampling specifications.

5.2 Results from 12-months T-bill auctions

Tables 3 and 4 also include the results of our tests using updated dealer bids from 12 month auctions. There were 275 updated dealer bids in this sample, comprising of 937 bidsteps (price-quantity pairs) and after updating they consisted of 996 bidsteps.

[32] We use a bandwidth of approximately 4 basis points. Table 10 in the appendix reports the results for various other bandwidths.

[33] For example, the joint hypothesis test based on the first order statistic is constructed by first dividing each individual test statistic evaluated on the sample by its standard deviation estimated by bootstrap and then taking the largest of these.

Table 3 shows that based on individual hypotheses, we get larger rejection rates in the 12 month sample. In particular, we reject around 20% of the individual tests. The joint hypothesis tests based on studentized test statistics, results of which are reported in Table 4, show a consistent pattern (relative to the 3-months treasury bills): almost all tests result in similar critical values and larger values of the test statistic in case of T-bills with 12-months maturity than for T-bills with 3-months maturity. Nevertheless, we fail to reject the null hypothesis in all joint tests for the 12-months treasury bills.

Overall, we view these patterns as evidence that the null hypothesis of private values is consistent with observed bidder behavior in Canadian T-bill auctions of both 3-months and 12-months maturities.

Table 3: Individual Hypothesis Test Results

	Percent rejected					
	3-months			12-months		
Bandwidth	Auctions for resampling					
	1	2	4	1	2	4
100	10.9	13.4	16.6	20.1	18.9	20.2
500	11.6	12.6	15.5	19.8	19.3	19.9
5000	11.2	13.7	16.4	19.6	19.2	19.3

Notes: The entries are percent of individual hypotheses that are rejected based on the equality test. We report results using different numbers of consecutive auctions in our resampling. The bandwidth parameter determines the width of the kernel in equation (7), and is denoted in price points over a face value of CA$ 1 million. Thus 100 price points corresponds to approximately 1 basis points in annual interest for 12-M bills, and 4 basis points in annual interest for 3M-bills.

5.3 Test under the alternative hypothesis of interdependent values

In Section 4.2 we described a test that should help us distinguish the reason for rejecting the null: whether the rejection is due to common (or interdependent) values or due to some other violation of the model. Even though our joint test failed to reject the hypothesis of private values for both the 3-month and 12-month samples, looking at Table 4, 12-month auctions appear closer to rejecting private values. Therefore, we decided to subject our 12-month sample to our test of

Table 4: Joint Hypothesis Studentized Test Results

	3-months			12-months		
Bandwidth[a]	100			500		
	Auctions for resampling			Auctions for resampling		
	1	2	4	1	2	4
SSQ[b]	49.37	199.06	188.81	135.81	397.45	225.08
Critical Value	1265.74	1589.46	1555.34	1481.45	1882.04	1598.21
Std Dev	424.16	492.07	583.21	454.79	542.84	555.79
p-value	1	1	0.95	0.95	0.85	0.8
FOS[c]	3.86	9.87	6.04	5.89	16.78	5.46
Critical Value	19.86	19.97	19.98	19.97	19.98	19.97
Std Dev	5.25	3.89	5.17	5.15	1.43	5.89
p-value	0.96	0.51	0.74	0.86	0.39	0.78
95th percentile[d]	0.23	0.94	1.37	0.16	0.54	1.33
Critical Value	1.72	2.3	2.63	1.63	1.48	2.36
Std Dev	0.37	0.44	0.45	0.44	0.31	0.41
p-value	1	1	1	1	0.95	0.67
Fraction trimmed[e]	0.05	0.06	0.05	0.09	0.1	0.09

[a] The bandwidth parameter determines the width of the kernel in equation (7), and is denoted in price points over a face value of CA$ 1 million. Thus 100 price points corresponds to approximately 1 basis points in annual interest for 12-M bills, and 4 basis points in annual interest for 3M-bills.

[b] Test based on sum of squares.

[c] Test based on first-order statistic.

[d] Test based on 95th percentile of the test statistic distribution.

[e] Trimming was done by eliminating marginal value estimates exceeding the maximum bid + 100 basis points.

interdependent values. For two ex-ante identical dealers observing customer signals that can be ordered, the interdependent values predicts a concordant ordering of the estimated pseudovalues, as described in equation 12 in Section 4.2 .

To check whether this prediction holds, we identified 15 pairs of dealers, who submitted very similar uninformed bids in a given auction, and who then observed customers' bids that can be ordered (along both dimensions of bids). We then estimate the right-hand side of equation (5) for each dealer-bid submitted *after* receiving the customer bid in our matched sample. We used our resampling algorithm, utilizing data from a single auction at a time (to preserve the conditional independence assumption). We then calculate the difference, Δv, of the estimated terms across the dealers, preserving the ordering of the customer bids observed by the dealers. As we pointed out in Section 4.2, however, a potential caveat to this calculation is that we cannot isolate the pseudovalues using our first-order condition (5); we also need to account for the extra term $\Upsilon(\cdot)$. Fortunately, as we explain in Appendix B, our data allows us to estimate an upper bound for Υ, up to a parameter C, which measures the marginal impact of market clearing price realizations on the bidder's average value from winning.[34] To see the importance of the $\Upsilon(\cdot)$ term, we plot the distribution of Δv's corrected using estimated upper-bounds for Υs with $C \in \{0.1, 1, 10\}$ in Figure 3.[35] The figure shows that Υ makes a very small difference to the distribution of Δvs.

Figure 3 also provides preliminary evidence that there is not much support for the one-sided hypothesis that the dealer observing the larger customer signal should have the higher pseudo-value. To perform a joint test of the hypothesis that the Δvs are greater than zero for the multiple bid pairs under consideration, we calculate the mean/median of studentized Δvs across the 15 matched bid pairs, and bootstrap this mean/median to approximate its sampling distribution. Under the null hypothesis, this mean should be greater than zero. The results of the test performed on our sample from 12-month T-bills are displayed in Table 5. Based on the results of this test, we do not find evidence for interdependent values in our data. In particular, the mean and the

[34]Since we need to calculate the difference in pseudovalues, we use 2 times the maximum upper-bound on Υ as an upper-bound for the difference in the Υ terms.

[35]As we note in the Appendix B, in analytically solved examples in the multi-unit auctions literature, $C \leq 1$. We thus consider $C = 10$ to be a very conservative allowance, which means that a 1 basis point increase in market clearing price would compel a bidder to increase her average valuation for the T-bills by 10 basis points. While it is possible to specify information structures in which such dramatic swings in one's posterior are possible, we deem these to be unlikely in the present context.

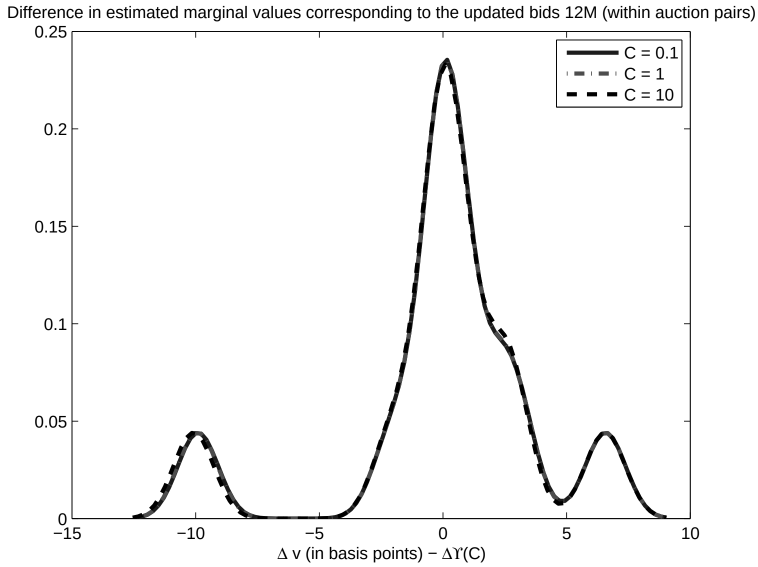

Figure 3: Distribution of Δv's in 12M T-bills for within auction bid-pairs.

median of the distribution of Δv's are almost equally likely to be greater or smaller than zero. We should also note that the tests in Table 5 when performed after accounting for the Υ term with $C \in \{0.1, 1, 10\}$ revealed virtually equivalent results, though this is perhaps not surprising given that these correction terms are negligible in magnitude.

Table 5: Test of the Restriction under Interdependent Values

Bandwidth (for resampling)	Mean	Median
100	0.47	0.52
500	0.52	0.43
5000	0.49	0.61
Number of tested pairs	15	

The entries are p-values of mean (median) of the distribution of $\Delta v = \hat{v}_i\left(q, s|z_i\right) - \hat{v}_j\left(q, s|z_j\right)$ being larger than 0.

5.4 Value of Information

Given that our tests failed to reject private values in our sample, in what follows we will use our estimates of marginal values generated by assuming the private values paradigm to estimate the value of information. In particular, we try to answer the question what is the effect of the additional information on a dealer's interim (expected) and ex post utility. Let $U_d^{EP}\left(B^I\right)$ denote the ex post utility of a dealer d, when submitting a bid B^I where the superscript I denotes information. In particular, as before $I = AI$ denotes additional information is incorporated, i.e., when the updated bid is used to compute the utility, and $I = BI$ when the original bid is used instead.

We can measure the value of information in terms of this notation as follows:

$$VI_d^{EP} = U_d^{EP}(B^{AI}) - U_d^{EP}(B^{BI}) \tag{13}$$

To compute the value of information given by equation (13), for each bidder we take the upper envelope of her estimated marginal values and the actual realized residual supply this bidder faced in the auction.[36] We then evaluate this bidder's ex-post utility using the original bid (submitted before the arrival of the customer's bid) to obtain $U_d^{EP}(B^{BI})$ and then evaluate her utility using the updated bid to obtain $U_d^{EP}(B^{AI})$. Performing this exercise for each bidder whom we observe updating her bid, we obtain a full distribution of the value of information, the mean of which we call "value of information."

Our value of information calculation assumes that customers do not alter their bidding behavior in response to making customer information unavailable to the dealers.[37] Studying the effects of a "ban" on dealers' utilization of customer information, or a full separation of dealer and customer bids requires that we re-compute the equilibrium strategies of customers and dealers under the new rules of the game. Unfortunately, computing equilibrium strategies in (asymmetric) discriminatory multi-unit auctions is still an open question, and we will leave this calculation to future research.

We should also note that the "value of information" calculation would be much more difficult to conduct in an interdependent values environment. To see why, observe that our data only allows us to identify the right-hand side of equation (5). Thus, even if the $\Upsilon(.)$ term is numerically negligible (as we argued above for our empirical application), what we can calculate are the changes in the "pseudo-value" of the dealer with respect to (presence vs. lack of) customer information.

[36]Notice that, in principle, VI_d^{EP} need not be positive for every observation.

[37]The illustrative example discussed in Appendix D is a case where this assumption is satisfied in equilibrium.

Recovering structural parameters of bidders' information/valuation structure (which is needed to assess the "value of information") remains, for now, an open problem in the interdependent values case.[38]

Using our estimates we find that that the ex post value of information, VI^{EP}, is on average about 0.45 of a basis point per T-bill for sale (0.64 when using 4 auctions for resampling, 0.46 when 2 and 0.26 when 1).[39] Since the average ex post utility amounts to about 1.65 basis points, the extra information contained in customers' order flow generates about 27% of the payoff of the dealers. In monetary terms, the order flow generates rents of around C$1.35 Million annually. Thus, access to customer bids is a significant component of dealer surplus from participating in Government of Canada securities auctions. Again, we do not attempt a detailed calculation of how this surplus would change if dealers are no longer allowed to route customer bids. This would involve a recalculation of equilibrium bids in the auction, which we leave for future research.

The value of information in 12-months treasury bills seems to be slightly lower: customers' information results in an increase in dealers' expected payoff of 13% (17% when using one auction for estimation; 12% when using 2 auctions and 10% when using 4) and the ex post payoff of dealers is 0.69 basis points per T-bill. In monetary terms, the order flow in 12-months Tbill auctions amounts to about C$0.4 Million. The discrepancy between the value of information in 3-months and 12-months Tbill auctions is perhaps caused by the fact that the 3-months Tbill auctions are about three times larger.

6 Conclusion

In Canadian Treasury auctions, dealers observe the bids of their customers. This institutional feature (which is also shared with a number of other countries) allows us to conduct tests for private

[38]It is known from the single-unit auction literature (Laffont and Vuong (1995), Athey and Haile (2002)) that recovering pseudo-values is not sufficient to identify the underlying structural parameters of the information/valuation structure; the well-known result is that there exists an affiliated private value model that is observationally equivalent to the interdependent value model. Here, we have the advantage of observing exogenous variation in dealers' information sets that allows us to rule out private values. This may, in principle, aid us in identifying a common/interdependent value model. However, the multi-unit setting introduces an additional complication that the conditional expectation defining the pseudo-value term conditions on equilibrium strategies as well (we can no longer rely on monotonicity of bid strategies to express the expectation as being conditional on the second order-statistic of rivals' signals); this makes inverting the pseudo-value term even more difficult than in the single-unit context.

[39]The standard deviation of ex post payoff is slightly over 2.5 basis points.

and interdependent values in multiunit auctions. The private values test is based on comparing two estimated distributions of marginal values corresponding to distributions of bids submitted by uninformed bidders and bidders who are informed about actions taken by their rivals. The interdependent values case is based on comparing the bids of two ex-ante similar dealers who observe different customer signals: theory predicts that the updated bids should preserve the ordering of the received customer information. We fail to reject the null hypothesis of private values for both 3-month and 12-month treasury bills auctioned by the Government of Canada, and do not find evidence supporting the alternative hypothesis of interdependent values. Under the empirically supported hypothesis of private values, we can also calculate the economic value of the informational advantage possessed by dealers in this market: we find that access to customer information provides dealers with 27% and 13% of their expected surplus, respectively, in 3-month and 12-month T-bill auctions.

References

[1] Armantier, O., and Sbaï, E., "Estimation and Comparison of Treasury Auction Formats when Bidders are Asymmetric," *Journal of Applied Econometrics*, Vol.21, Issue 6, pp.745-779, 2006.

[2] Arnone, M. and Iden, G. "Primary Dealers in Government Securities: Policy Issues and Selected Countries' Experience," IMF Working paper, WP/03/45, 2003.

[3] Ausubel, L., and Cramton, C., "Demand Reduction and Inefficiency in Multi-Unit Auctions," mimeo 2002

[4] Athey, S., and Haile,. P., "Nonparametric Approaches to Auctions," forthcoming in J. Heckman and E. Leamer, eds., *Handbook of Econometrics*, Vol. 6, Elsevier, 2005.

[5] Bajari, P., and Hortaçsu, A., "Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions," *RAND Journal of Economics*, Vol. 34, pp. 329-355, 2003.

[6] Bickel, P., and Freedman, D., "Some Asymptotic Theory for the Bootstrap," *The Annals of Statistics*, Vol. 9, pp.1196-1217, 1981.

[7] Bikhchandani, S., and Huang, C., "The economics of treasury security markets," *Journal of Economic Perspectives*, Vol. 7, pp. 117-134, 1989.

[8] Blum, J. R., Kiefer, J., and Rosenblatt, M., "Distribution free tests of independence based on the sample distribution function," *Annals of Mathematical Statistics*, Vol. 32, pp. 485-498, 1961.

[9] Chapman, James T.E., David McAdams, and Harry J. Paarsch, "Bounding Best-Response Violations in Discriminatory Auctions with Private Values." Typescript. Iowa City, Iowa: Department of Economics, University of Iowa, 2006.

[10] Chapman, James T.E., David McAdams, and Harry J. Paarsch, "Bounding Revenue Comparisons across Multi-Unit Auction Formats under ε-Best Response," *American Economic Review Papers and Proceedings*, Vol. 97, pp. 455-458, 2007.

[11] de Castro, Luciano, and Riascos, Alvaro, "Characterization of bidding behavior in multi-unit auctions," *Journal of Mathematical Economics*, Vol. 45, pp. 559-575, 2009

[12] Easley, David, Maureen OHara, and P. S. Srinivas. "Option Volume and Stock Prices: Evidence on Where Informed Traders Trade," *Journal of Finance*, Vol. 53, No. 2, pp. 431-465, 1998.

[13] Evans, Martin D.D. and Richard K. Lyons, "Order Flow and Exchange Rate Dynamics," *Journal of Political Economy*, pp. 170-180, 2002.

[14] Falk, M. and R.-D. Reiss, "Statistical Inference for Conditional Curves: Poisson Process Approach," *Annals of Statistics*, 20 (2), pp. 779-796, 1992.

[15] Fevrier, P., Preget, R., and Visser, M., "Econometrics of Share Auctions," mimeo 2004.

[16] Gilley, O., and Karels, G., "The Competitive Effect of Bonus Bidding: New Evidence," *The Bell Journal of Economics*, Vol. 12, pp. 637-648, 1981.

[17] Guerre, E., Perrigne, I., and Vuong, Q., "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica*, Vol. 68, No. 3, 2000.

[18] Haile, P., Hong, H., and Shum, M., "Nonparametric Tests for Common Values In First-Price Sealed-Bid Auctions," mimeo, 2003.

[19] Hill, J., and Shneyerov, A., "Are There Common Values in BC Timber Sales? A Tail-Index Nonparametric Test," mimeo, 2009.

[20] Hoeffding, W., "A Class of Statistics with Asymptotically Normal Distribution," *The Annals of Mathematical Statistics*, Vol. 19, pp. 293-325, 1948.

[21] Hortaçsu, A., "Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market," mimeo, 2002a.

[22] Hortaçsu, A., "Bidding Behavior in Divisible Good Auctions: Theory and Empirical Evidence from Turkish Treasury Auctions," mimeo, 2002b.

[23] Hortaçsu, A., and Sareen, S.,"Order Flow and the Formation of Dealer Bids in Treasury Auctions," mimeo, 2006.

[24] Kang, B., and Puller, S., "The Effect of Auction Format on Efficiency and Revenue in Divisible Goods Auctions: A Test Using Korean Treasury Auctions," *Journal of Industrial Economics*, forthcoming.

[25] Kastl, J., "Discrete Bids and Empirical Inference in Divisible Good Auctions", mimeo, 2006.

[26] Kastl, J., "On the Properties of Equilibria in Private Value Divisible Good Auctions with Constrained Bidding", mimeo, 2008.

[27] Koul, H. L., and Schick, A., "Testing for equality of two nonparametric regression curves," *Journal of Statistical Planning and Inference*, 65, pp. 293-314, 1997.

[28] Kyle, A., "Informed Speculation with Imperfect Competition," *Review of Economic Studies*, Vol. 56, pp. 317356, 1989.

[29] Laffont, J.J., and Vuong, Q., "Structural Analysis of Auction Data," *American Economic Review*, Papers and Proceedings, Vol. 86, pp. 414-420, 1996.

[30] Lehmann, E., *Elements of Large Sample Theory*, Springer, New York,1999.

[31] Lyons, Richard K., Tests of Microstructural Hypotheses in the Foreign Exchange Market, *Journal of Financial Economics*, 39, 321-351, 1995.

[32] Madhavan Anand and Smidt, Seymour, "A Bayesian Model of Intraday Specialist Pricing," *Journal of Financial Economics*, 30, pp. 99-134, 1990.

[33] Matthews, S., "Information Acquisition in Discriminatory Auctions," in *Bayesian Models in Economic Theory* ed. by M. Boyer and R. Kihlstrom, North Holland, 1984.

[34] McAdams, D., "Partial Identification and Testable Restrictions in Multi-Unit Auctions," *Journal of Econometrics*, pp. 74-85, 2008.

[35] Milgrom, P., and Weber, R., "A Theory of Auctions and competitive Bidding," *Econometrica*, Vol. 50, pp. 1089-1122, 1982.

[36] Mudholkar, G.S. and Wilding, G. E., "On the conventional wisdom regarding two consistent tests of bivariate dependence," *The Statistician*, 52(1), pp. 41-57, 2003.

[37] Nyborg, Kjell G. and Suresh M. Sundaresan. "Discriminatory versus Uniform Treasury Auctions: Evidence from When-Issued Transactions," *Journal of Financial Economics* 42, 63-105, 1996.

[38] Paarsch, H., "Empirical Models of Auctions and an Application to British Columbian Timber," University of Western Ontario, Department of Economics Technical Report 91-19, 1991.

[39] Paarsch, H., "Deciding Between the Common and Private Value Paradigms in Empirical Models of Auctions," *Journal of Econometrics*, Vol. 51, pp.191-215, 1992.

[40] Pinkse, J., and Tan, G., "The Affiliation Effect in First-Price Auctions," *Econometrica*, Vol. 73, pp. 263-277, 2005.

[41] Romano, J.P., A. Shaikh, and M. Wolf, "Formalized Data Snooping Based on Generalized Error Rates," *Econometric Theory*, Vol. 24, pp. 1-44, 2008.

[42] Smith, V., "Bidding Theory and the Treasury Bill Auction: Does Price Discrimination Increase Bill Prices?," *Review of Economics and Statistics*, Vol. 48, pp. 141-146, 1966.

[43] Stute, W., "Conditional U-Statistics," *The Annals of Probability*, Vol. 19, pp. 812-825, 1991.

[44] Vickrey, W., "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, Vol. 16, pp. 8-37, 1961.

[45] Wilson, R., "Auctions of Shares," *Quarterly Journal of Economics*, 1979.

[46] Wolak, F., "Identification and Estimation of Cost Functions Using Observed Bid Data: An Application to Electricity," *Advances in Econometrics: Theory and Applications*, Eighth World Congress, Volume II, M. Dewatripont, L. P. Hansen, and S. J. Turnovsky (editors), Cambridge University Press, pp.133-169, 2003.

[47] Wolak, F., "Quantifying the Supply-Side Benefits from Forward Contracting in Wholesale Electricity Markets," *Journal of Applied Econometrics*, Vol. 22, pp.1179-1209, 2007.

A Further Data details

A.1 Bid revisions

In our analysis, we focus on observations where dealers update their bids after routing a customer bids. In our 3 month T-bill sample, out of 660 dealer bids (which were also accompanied by a customer bid), 154 dealer bids were submitted for the first time after receiving the customer bid. In 216 instances, existing dealer bids were "updated" after seeing a customer bid. In 290 instances, the dealer did not update her bid after routing a customer bid. In the 12 month sample, out of 659 dealer bids, 250 were submitted for the first time after routing a customer bid. 275 were updated after routing a customer bid, and in 134 instances, the bids were not updated. Thus, bid updating in response to customer bids is a frequent event, though it also appears that dealers often do not revise their bids in response to customer bids.

However, this may reflect the fact that dealers might not have time to update their bids in response to customer: the average (median) customer bid in situations where the dealer does not update his bid in response comes 5.55 (med=4.85) minutes before the deadline. The average (median) customer bid in situations where the dealer does update his bid in response comes 17.8 (med=10.3) minutes before the deadline. Indeed, the distributions (of customer bids that are followed by a dealer update and those that are not) follow a clear first order stochastic dominance, as plotted in Figure 4 and Figure 5 below.To check whether 5 minutes is enough for a dealer to change his bid, we calculated the average (median) number of minutes that a dealer takes to "follow" a customer bid with his own bid. In situations where there is updating it takes, on average, 5.49 (med = 4.93) minutes for the updated dealer bid to be entered. Thus, it is possible that a customer bid submitted 5 minutes before the deadline is "too late" for a dealer to follow.

We also ran regressions to understand why some customer bids were followed by a dealer update, and others were not. The main findings from these regressions were: (a) the timing of customer bids was the most significant predictor of updating, i.e. late customer bids were less likely to be responded to, (b) controlling for timing, larger customer bids were more likely to be followed by a dealer update, even after controlling for customer fixed effects, and (c) controlling for customer size, there is some heterogeneity across dealers as to whether they update their bids in response to

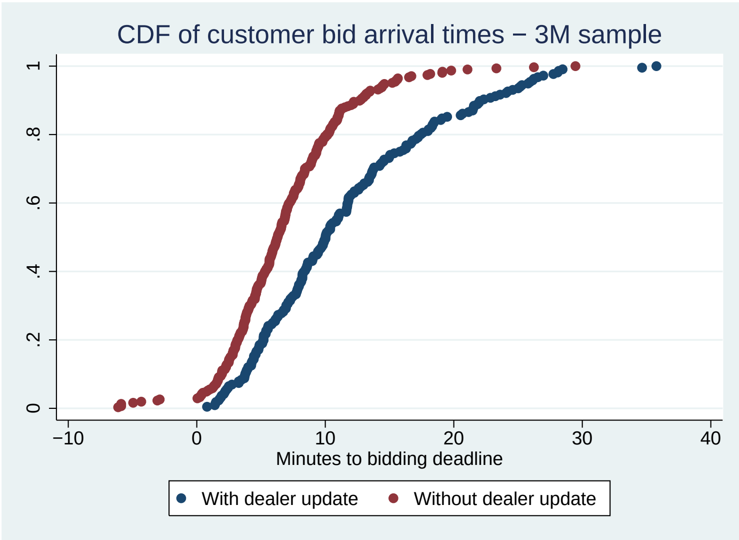

Figure 4: Comparison of customer bid submission times with and without dealer updates – 3M sample

customer bids or not; some dealers may be following relational agreements with their customers, and purposely avoiding to update their bids.

Finally, in instances where there was a bid update, we regressed the final (quantity-weighted) dealer bid on the dealer's original bid before seeing the customer bid, and the customer's bid. The results, reported in Table 6, show that customer bids are a statistically significant (at the 1% level for the 3M sample and 10% level for the 12M sample) determinant of the price level of the updated dealer bid, even though a huge portion of the variation is captured by the original dealer bid. Interestingly, the updated dealer bid appears to load more heavily on the customer bid in the 3M sample as opposed to the 12M sample.

A.2 Public information releases

An important potential caveat regarding our testing strategy is that privately observed customer bids *per se* are not the causal drivers of observed changes in dealer bids, and that customer bids are correlated with other unobservable information flows driving modifications to dealer bids. The presence of such unobservable information flows would confound our testing strategy, since these information flows may affect the dealer's marginal values, and/or allow them to observe an extra

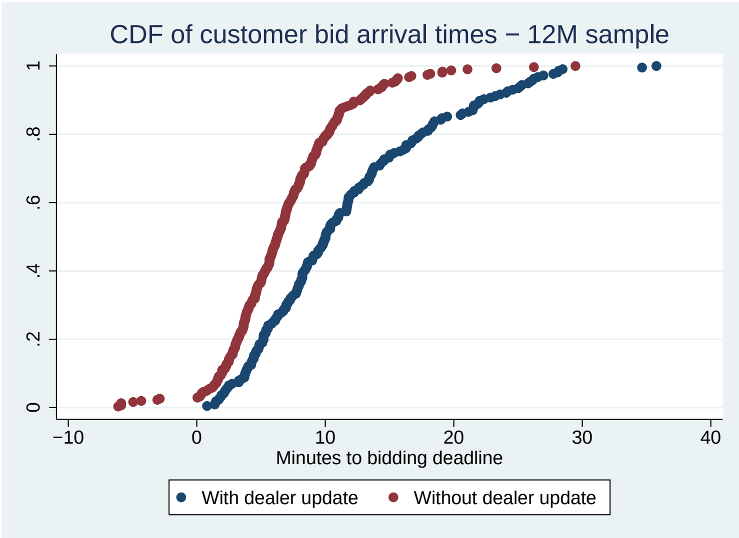

Figure 5: Comparison of customer bid submission times with and without dealer updates – 12M sample

piece of information regarding the auction environment that we are not able to account for in our marginal value estimation procedure. One source of unobservable information flows maybe in the form of news announcements or market movements during the bidding period that are observed by all dealers, but not the econometrician. To examine the plausibility of such unobserved public information flows, we examined the timing of changes in dealer bids in our data set. If information flows are publicly observed across dealers, we should observe some amount of clustering in the timing of bid modifications in our data set. We failed to find an important degree of clustering in this dimension – within any 5 minute window around a particular bid updating event, there was at most one other dealer changing his/her bid (and such a dealer was only found in 40 instances out of the total 216 updated bids in our sample). This suggests that it is unlikely that customer bids were driven by or accompanied with important public information releases that are unobservable to us. As a complement to this finding, Hortaçsu and Sareen (2006) report that unobservable public information releases by official sources are highly unlikely, as Bank of Canada and Treasury pay careful attention to avoid public disclosures during the bidding period.

Table 6: Correlation between dealer bid updates and customer bids

	3 Month Updated bid	12 Month Updated bid
Customer bid	0.146***	0.0166*
	(0.0289)	(0.00863)
Dealer's orig. bid	0.853***	0.982***
	(0.0289)	(0.00860)
Constant	1,399**	946.6*
	(609.5)	(490.1)
Observations	216	275
R-squared	0.998	0.998

This table contains regressions of the dealer's updated (quantity-weighted average price) bid on the customer bid and the dealer's original bid. The regressions are reported separately for the 3- and 12-M T-bill samples. Standard errors in parentheses *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$ 0.

A.3 Long-short T-bill positions

Along with customer bids, dealers may also adjust their net long-short positions during the bidding period in the when-issued market. The dealer's net long/short position may shift his/her bid and/or her marginal valuation curve independent of the information contained in customers bid; thus this is a potential confounding factor. Fortunately, the Bank of Canada requires dealers to report their net positions (short or long) in the when-issued market along with their bids. In our data, only 13 out of 216 updated bids in 3-months T-bill auctions were accompanied by a change in the net position of the dealer in the when-issued market. In 12-months T-bill auction, only 7 dealers (out of 275 updates) also changed their net position. Given that such a small fraction of bids are subject to this potential confound, we do not think that developments in the when-issued market affect our test results.

B Necessary Conditions under Interdependent Values

We start with the expected utility of a bidder who receives the signals $\tilde{\theta}_i = (s_i, z_i)$ and submits the price-quantity vector $\{(p_k, q_k), k = 1, \ldots, K\}$, defining a step function. Note first that the bidder's profit depends crucially on the (random) market clearing price and quantity, P^c and Q^c, which will be determined by the intersection of the residual supply function, and the bidder's step function. As illustrated in Figure 6, the residual supply function can intersect the bid function on either a horizontal segment (RS_1 and RS_3), which yield $P^c = b_k$ or $P^c = b_{k+1}$, or on a vertical segment (RS_2), in which case the realized P^c is between b_k and b_{k+1}. When the intersection is on a vertical segment of the bid curve, the bidder is awarded the quantity she asked for, i.e. $Q^c = q_k$ as in Figure 6. When the intersection is at a horizontal segment, however, the bidder gets only the amount that is up to the residual supply curve: for RS_1, the bidder will be awarded a quantity Q^c that is between 0 and q_k, and for RS_3, Q^c will be between q_k and q_{k+1}.

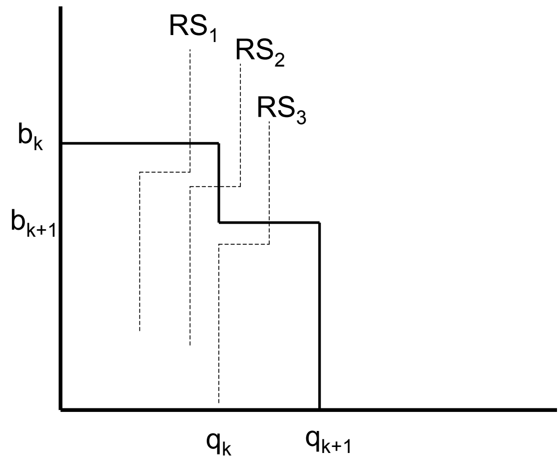

Figure 6: Illustration of necessary condition

With these preliminaries the bidder's expected utility, summed over different realizations of P^c, can be written as:

$$EU\left(\tilde{\theta}_i, \{b_k, q_k\}\right) = \sum_{k=1}^{K} \Pr\left(P^c = b_k|\tilde{\theta}_i\right) E_{\theta_{-i}|\tilde{\theta}_i, P^c=b_k, q_k}\left[V\left(Q^c\left(b_k, \theta_{-i}\right), \tilde{\theta}_i, \theta_{-i}\right) - b_k\left(Q^c\left(b_k, \theta_{-i}\right) - q_{k-1}\right)\right] +$$

$$+\sum_{k=1}^{K} \Pr\left(b_k > P^c > b_{k+1}|\tilde{\theta}_i\right) E_{\theta_{-i}|\tilde{\theta}_i, P^c\in\left(b_{k+1}, b_k\right), q_k}\left[V\left(q_k, \tilde{\theta}_i, \theta_{-i}\right)\right] - \sum_{k=1}^{K} \Pr\left(b_k > P^c|\tilde{\theta}_i\right) b_k\left(q_k - q_{k-1}\right)$$

The first summation is over the set of events when the residual supply curve intersects bidder i's bid on a horizontal portion of her bid function (as in RS_1 and RS_3 in Figure 6). The bidder gets information about θ_{-i} from the realization of P^c, which she incorporates into her expectation of the value of winning. Note also that when the residual supply curve is at the flat portion of the bid, the quantity awarded to the bidder, Q^c, depends on the bids submitted by competing bidders, i.e. the residual supply curve at b_k or b_{k+1} (bids at the same price level only occur with zero probability).

The second summation denotes the events where the residual supply curve intersects the vertical portion of bidder i's bid function (as in RS_2). In this event, the bidder is awarded the full quantity, q_k, that she requested. The third summation aggregates the payments that the bidder makes. However, this payment should be adjusted if the residual supply curve is at the flat portion of the bid, as Q^c does not equal q_k; this adjustment is reflected in the first summation term.

When we are taking the first-order conditions with respect to q_k, only the following terms need to be considered:

$$\Pr\left(b_k = P^c|\tilde{\theta}_i\right) E_{\theta_{-i}|\tilde{\theta}_i, P^c=b_k, q_k}\left[V\left(Q^c\left(b_k, \theta_{-i}\right), \tilde{\theta}_i, \theta_{-i}\right)\right] \tag{B-1}$$

$$+\Pr\left(b_k > P^c > b_{k+1}|\tilde{\theta}_i\right) E_{\theta_{-i}|\tilde{\theta}_i, b_k>P^c>b_{k+1}, q_k}\left[V\left(q_k, \tilde{\theta}_i, \theta_{-i}\right)\right] \tag{B-2}$$

$$+\Pr\left(b_{k+1} = P^c|\tilde{\theta}_i\right) E_{\theta_{-i}|\tilde{\theta}_i, b_{k+1}=P^c, q_k}\left[V\left(Q^c\left(b_{k+1}, \theta_{-i}\right), \tilde{\theta}_i, \theta_{-i}\right)\right] \tag{B-3}$$

$$-\Pr\left(b_k = P^c|\tilde{\theta}_i\right) b_k E_{\theta_{-i}|\tilde{\theta}_i, b_k=P^c, q_k}\left(Q^c\left(b_k, \theta_{-i}\right) - q_{k-1}\right) \tag{B-4}$$

$$-\Pr\left(b_{k+1} = P^c|\tilde{\theta}_i\right) b_{k+1} E_{\theta_{-i}|\tilde{\theta}_i, b_{k+1}=P^c, q_k}\left(Q^c\left(b_k, \theta_{-i}\right) - q_k\right) \tag{B-5}$$

$$-\Pr\left(b_k > P^c|\tilde{\theta}_i\right) b_k\left(q_k - q_{k-1}\right) \tag{B-6}$$

$$-\Pr\left(b_{k+1} > P^c|\tilde{\theta}_i\right) b_{k+1}\left(q_{k+1} - q_k\right) \tag{B-7}$$

While taking the derivatives of the above terms with respect to q_k, note that $Q^c(.)$ does not depend on q_k, as it is equal to the residual supply quantity at b_k or b_{k+1}.

The derivatives of the above terms with respect to q_k are thus as following:

$$\Pr\left(b_k > P^c > b_{k+1}|\tilde{\theta}_i\right) \frac{\partial E\left[V(q_k,\cdot)|b_k > P^c > b_{k+1}\right]}{\partial q_k} + \tag{B-8}$$

$$+\frac{\partial \Pr\left(b_k = P^c|\tilde{\theta}_i\right)}{\partial q_k} E_{\theta_{-i}|\tilde{\theta}_i, P^c=b_k, q_k}\left[V\left(Q^c(b_k,\theta_{-i}),\theta_{-i},\tilde{\theta}_i\right)\right] + \tag{B-9}$$

$$+\frac{\partial \Pr\left(b_{k+1} = P^c|\tilde{\theta}_i\right)}{\partial q_k} E_{\theta_{-i}|\tilde{\theta}_i, P^c=b_{k+1}, q_k}\left[V\left(Q^c(b_{k+1},\theta_{-i}),\theta_{-i},\tilde{\theta}_i\right)\right] + \tag{B-10}$$

$$+\frac{\partial \Pr\left(b_k > P^c > b_{k+1}|\tilde{\theta}_i\right)}{\partial q_k} E_{\theta_{-i}|\tilde{\theta}_i, P^c\in\left(b_{k+1},b_k\right), q_k}\left[V\left(q_k,\theta_{-i},\tilde{\theta}_i\right)\right] \tag{B-11}$$

$$+\Pr\left(b_{k+1} \geq P^c|\tilde{\theta}_i\right) b_{k+1} - \Pr\left(b_k > P^c|\tilde{\theta}_i\right) b_k \tag{B-12}$$

Noting that $\Pr\left(b_k > P^c|\tilde{\theta}_i\right) = \Pr\left(b_k > P^c > b_{k+1}|\tilde{\theta}_i\right) + \Pr\left(b_{k+1} \geq P^c|\tilde{\theta}_i\right)$, and defining

$$\Upsilon(\tilde{\theta}_i,\cdot) = \frac{\partial \Pr\left(b_k = P^c|\tilde{\theta}_i\right)}{\partial q_k} E_{\theta_{-i}|\tilde{\theta}_i, P^c=b_k, q_k}\left[V\left(Q^c(b_k,\theta_{-i}),\theta_{-i},\tilde{\theta}_i\right)\right] + \tag{B-13}$$

$$+\frac{\partial \Pr\left(b_{k+1} = P^c|\tilde{\theta}_i\right)}{\partial q_k} E_{\theta_{-i}|\tilde{\theta}_i, P^c=b_{k+1}, q_k}\left[V\left(Q^c(b_{k+1},\theta_{-i}),\theta_{-i},\tilde{\theta}_i\right)\right] + \tag{B-14}$$

$$+\frac{\partial \Pr\left(b_k > P^c > b_{k+1}|\tilde{\theta}_i\right)}{\partial q_k} E_{\theta_{-i}|\tilde{\theta}_i, P^c\in\left(b_{k+1},b_k\right), q_k}\left[V\left(q_k,\theta_{-i},\tilde{\theta}_i\right)\right] \tag{B-15}$$

we get the following first-order condition:

$$\begin{aligned} E_{\theta_{-i}|\tilde{\theta}_i, P^c\in\left(b_{k+1},b_k\right), q_k}\left[v\left(q_k,\tilde{\theta}_i,\theta_{-i}\right)\right] + \frac{\Upsilon(\tilde{\theta}_i,\cdot)}{\Pr\left(b_k > P^c > b_{k+1}|\tilde{\theta}_i\right)} \\ = b_k + (b_k - b_{k+1}) \frac{\Pr\left(b_{k+1} \geq P^c|\tilde{\theta}_i\right)}{\Pr\left(b_k > P^c > b_{k+1}|\tilde{\theta}_i\right)} \end{aligned} \tag{B-16}$$

Inspecting the first-order condition (B-16), we see that the first term on the left-hand side, $E_{\theta_{-i}|\tilde{\theta}_i, P^c\in(b_{k+1},b_k), q_k}\left[v\left(q_k,\tilde{\theta}_i,\theta_{-i}\right)\right]$, is similar to the "pseudo-value" term obtained in a first-price interdependent value auction by Haile, Hong, and Shum (2003). This is the expected marginal value of the q_k-th unit of the security, conditional on the bid b_k for this unit being the market-clearing price in the auction. Under our assumption 3, this term is increasing in one's own signal, and, in the case of a dealer, one's customer's signal.

The extra term on the left-hand side of equation (B-16) arises from the effect of an infinitesimal change in q_k on the distribution of the market clearing price; which, in turn, affects the bidder's value from *inframarginal* units. To further analyze this term, note first that:

$$\frac{\partial \Pr\left(b_k = P^c|\tilde{\theta}_i\right)}{\partial q_k} + \frac{\partial \Pr\left(b_{k+1} = P^c|\tilde{\theta}_i\right)}{\partial q_k} + \frac{\partial \Pr\left(b_k > P^c > b_{k+1}|\tilde{\theta}_i\right)}{\partial q_k} = 0$$

As can be seen in Figure 6, an infinitesimal change in q_k will not affect the total probability of the residual supply being on either of the horizontal segments of i's bid function, or on the vertical segment. However, the relative weights on the three different possibilities may change. Since the bidder's expected value from winning is different in these three different states of the world, $\Upsilon(\tilde{\theta}_i, .)$ captures the reweighting over these different states.

Note also that $\frac{\partial \Pr\left(b_k = P^c|\tilde{\theta}_i\right)}{\partial q_k} \geq 0$ and $\frac{\partial \Pr\left(b_{k+1} = P^c|\tilde{\theta}_i\right)}{\partial q_k} \leq 0$, as can be seen from Figure 6 (increasing quantity makes it more likely that residual supply intersects horizontally at b_k, and less likely at b_{k+1}.) Also, define:

$$a = E_{\theta_{-i}|\tilde{\theta}_i, P^c = b_k, q_k}\left[V\left(Q^c(b_k, \theta_{-i}), \theta_{-i}, \tilde{\theta}_i\right)\right] - E_{\theta_{-i}|\tilde{\theta}_i, P^c \in \left(b_{k+1}, b_k\right), q_k}\left[V\left(q_k, \theta_{-i}, \tilde{\theta}_i\right)\right] \quad \text{(B-17)}$$

$$b = E_{\theta_{-i}|\tilde{\theta}_i, P^c \in \left(b_{k+1}, b_k\right), q_k}\left[V\left(q_k, \theta_{-i}, \tilde{\theta}_i\right)\right] - E_{\theta_{-i}|\tilde{\theta}_i, P^c = b_{k+1}, q_k}\left[V\left(Q^c(b_{k+1}, \theta_{-i}), \theta_{-i}, \tilde{\theta}_i\right)\right] \quad \text{(B-18)}$$

where a is the incremental change in expected value conditional on winning at $P^c = b_k$ vs. $P^c \in (b_{k+1}, b_k)$ and b is the incremental change in expected value conditional on winning at $P^c \in (b_{k+1}, b_k)$ vs. $P^c = b_{k+1}$. With these definitions, we can rewrite Υ as:

$$\Upsilon(\tilde{\theta}_i, .) = \frac{\partial \Pr\left(b_k = P^c|\tilde{\theta}_i\right)}{\partial q_k} a - \frac{\partial \Pr\left(b_{k+1} = P^c|\tilde{\theta}_i\right)}{\partial q_k} b \quad \text{(B-19)}$$

$\frac{\partial \Pr\left(b_k = P^c|\tilde{\theta}_i\right)}{\partial q_k} \geq 0$, $\frac{\partial \Pr\left(b_{k+1} = P^c|\tilde{\theta}_i\right)}{\partial q_k} \leq 0$ by above. Under Assumption 3, which states bidders' expected valuations are increasing in competitors' signals, and under the further assumption that a weakly increasing equilibrium exists (which leads to a higher P^c being better news about one's rivals' signals), we can get the intuitive condition that a and b are greater than zero. Under this further condition, we get that $\Upsilon \geq 0$ and that the right hand side of equation (B-16) provides an upper bound on the "pseudo-value" term, $E_{\theta_{-i}|\tilde{\theta}_i, P^c \in (b_{k+1}, b_k), q_k}\left[v\left(q_k, \tilde{\theta}_i, \theta_{-i}\right)\right]$. Moreover, we can use this expression to get an upper bound for $\Upsilon(\tilde{\theta}_i, .)$, as

$$\begin{aligned}\Upsilon(\tilde{\theta}_i, .) &\leq \left(\frac{\partial \Pr\left(b_k = P^c|\tilde{\theta}_i\right)}{\partial q_k} - \frac{\partial \Pr\left(b_{k+1} = P^c|\tilde{\theta}_i\right)}{\partial q_k}\right)(a-b) \qquad \text{(B-20)} \\ &= \left(\frac{\partial \Pr\left(b_k = P^c|\tilde{\theta}_i\right)}{\partial q_k} - \frac{\partial \Pr\left(b_{k+1} = P^c|\tilde{\theta}_i\right)}{\partial q_k}\right)\left(E_{\theta_{-i}|\tilde{\theta}_i, P^c=b_k, q_k}\left[V(\cdot)\right] - E_{\theta_{-i}|\tilde{\theta}_i, P^c=b_{k+1}, q_k}\left[V(\cdot)\right]\right)\end{aligned}$$

The first term in the product in equation (B-20) can be estimated from our data. The second term is the difference in the bidder's expected total valuation of the (up to q_k) units won depending on whether the market clearing price, P^c is b_k or b_{k+1}. In a private value model, of course, learning the market-clearing price does not impact the bidder's valuation, thus this term is zero. The exact magnitude of this term in the interdependent value model is not known, as this requires knowledge of $V(.)$ and the joint distribution of private information, but we will parametrize our bound on this term using an expression for the average valuation of units won:

$$\frac{E_{\theta_{-i}|\tilde{\theta}_i, P^c=b_k, q_k}\left[V(\cdot)\right] - E_{\theta_{-i}|\tilde{\theta}_i, P^c=b_{k+1}, q_k}\left[V(\cdot)\right]}{q_k} \leq C\left(b_k - b_{k+1}\right)) \qquad \text{(B-21)}$$

This means that learning that the market clearing price is b_k vs. b_{k+1} can increase the bidder's average valuation of units won by not more than C times $b_k - b_{k+1}$; i.e. C parametrizes the marginal impact of news about the market clearing price on the bidder's average valuation. The exact characterization of C requires solving for the equilibrium of the model, though analytically solvable (linear strategy) examples in the literature (Kyle (1989), Wang and Zender (2002)) where bidders are assumed to have normally distributed signals about a normally distributed common value derive $C \leq 1$, where C approaches 1 as the number of bidders grows large. The intuition for this result is that in a linear strategy equilibrium, the market clearing price can be expressed as the average of other bidders' signals, and under normal updating, the bidder's updated expected valuation is a convex combination of his signal and the market clearing price. Unfortunately, these derivations are for the uniform price auction, and rely rather heavily on the normality assumption, thus it is difficult to generalize them to more flexible functional forms. We will thus use a range values of for C in our analysis in Section 5.3.

C Asymptotics

To study the asymptotic properties of ξ, we start with the case where $W = 1$, i.e. all customer bids are weighted equally in resampling. We begin with a useful Lemma for this special case:

Lemma 3 *Suppose $W = 1$, data are iid across the T auctions and bidders, all bidders are symmetric and N is fixed. Then as $T \to \infty$, $\frac{M}{T} \to \infty$:*

$$\sqrt{T}\left(\hat{H}^R(X) - H(X)\right) \to N\left(0, \frac{(N-1)^2}{N}\zeta\right) \tag{C-22}$$

where $\zeta = E_{\theta_{-i}}\left[(\Phi(y_1, ..., y_{N-1}; X))^2\right] - \left(\binom{NT}{N-1}^{-1} \sum_{(1,1)\le\alpha_1<\alpha_2<...<\alpha_{N-1}\le(T,N)} \Phi\left(y_{\alpha_1}, ..., y_{\alpha_{N-1}}; X\right)\right)^2$ and where that last summation is taken over all combinations of $N-1$ indices $\alpha_i \in \{(1,1), (1,2), ..., (1, N-1), ..., (T, N)\}$ such that $\alpha_1 < \alpha_2 < ... < \alpha_{N-1}$.[40]

Proof. Consider the following statistic based on all subsamples of size $(N-1)$ from the full sample of NT datapoints:

$$\beta\left(\hat{F}; c\right) = \binom{NT}{N-1}^{-1} \sum_{1\le\alpha_1<\alpha_2<...<\alpha_{N-1}\le NT} \Phi\left(y_{\alpha_1}, ..., y_{\alpha_{N-1}}, c\right)$$

where $\hat{F}$ is the empirical distribution of bids (recall that in our case a bid is at most a $2K$-dimensional vector of K price-quantity pairs). β is a *U-statistic* and the result thus follows from applying Theorem 7.1 of Hoeffding (1948) which provides a useful version of a central limit theorem for this class. A sufficient condition for asymptotic normality is the existence of the second moment of the kernel of the functional Φ, which in this case is equivalent to finiteness of $E\left[\Phi(\cdot)^2\right]$, which is satisfied since $\Phi(\cdot)$ is an indicator function. As discussed earlier, the resampling estimator $\hat{H}^R(X)$ (in the case where $W = 1$ and all bidders are symmetric with iid bids) is a simulator of the closely related *V-statistic*:

$$\xi\left(\hat{F}; c\right) = \frac{1}{(NT)^{N-1}} \sum_{\alpha_1=1}^{NT} ... \sum_{\alpha_{N-1}=1}^{NT} \Phi\left(y_{\alpha_1}, ..., y_{\alpha_{N-1}}, c\right)$$

[40] The asymptotic distribution of the resampling estimator $\hat{G}^R$ can be established analogously, by replacing the weak inequality in the definition of $\Phi(\cdot)$ by a strict one.

where the averaging is over every permutation of the NT observations.[41] Lehmann (1999, Theorem 6.2.2, p.388) shows that the asymptotic distribution of this V-statistic is identical to that of the U-statistic. Finally, since $\frac{M}{T} \to \infty$, $\hat{H}^R \to \xi$. ■

To account for the asymmetry in dealer and customer bid distributions, we can appeal to Hoeffding (1948), Theorem 8.1, which extends the asymptotic normality result from Lemma 3 to the case where all y_{it} are allowed to have different distributions. This extension requires a slightly stronger condition on the third moment of $\Phi(\cdot)$ to use the Liapunoff Central Limit Theorem, but this condition is still satisfied since $\Phi(\cdot)$ is an indicator function which is uniformly bounded and our estimator is asymptotically normally distributed.

This last result would apply in our setting if all possible dealer's bids were independent from customer bids. Yet in our setting some dealer bids are of course submitted only after observing a particular customer bid and therefore it is necessary to use proper weighting of each subsample in the V-statistic. We propose to achieve this through our estimator (6) which uses the kernel weights $W(.)$ as defined in (7). The asymptotic properties of the estimator thus will depend on the properties of the kernel and assumptions on the bandwidth parameter h_T. Fortunately, the asymptotic properties of *conditional U-statistics* have been derived in Stute (1991)[42], whose Theorem 1 we adapted to our application and stated in the text as Proposition 2.

Using the asymptotic variance of the distribution of the market clearing prices, $H(X)$ (and $G(X)$), we can use the delta-method to derive the asymptotic variance of the estimates of the marginal values, i.e., $Var(\hat{v}_k) = J_v' \Sigma J_v$, where J_v is the matrix of partial derivatives with respect to $H(b_{k+1})$, $H(b_k)$ and $G(b_{k+1})$ and Σ is the asymptotic variance/covariance matrix for those estimates. To obtain the asymptotic covariance matrix of $\{H(c_1), H(c_2), G(c_3)\}$ at particular three values of c we can appeal to Hoeffding's and Stute's asymptotic theorems since they apply also to vector-valued random variables and the off-diagonal elements of the asymptotic variance/covariance matrix are the asymptotic covariances between two corresponding (conditional) *U-statistics*.

[41] Note that the sample size is NT (bidders per auction × auctions) and we are constructing subsamples of size $N-1$, hence the denominator $(NT)^{N-1}$.

[42] As noted earlier, *V-statistics* and corresponding *U-statistics* are well known to have the same asymptotic distribution (e.g. Lehmann 1999, Theorem 6.2.2). To see that the conditional U- and V- statistics are also asymptotically equivalent, observe, as in Stute (1991, page 813) that both can be written as ratios of unconditional U- and V-statistics all of which are asymptotically equivalent.

D Example of a fully-revealing equilibrium

We present a fully-revealing equilibrium in the case of an IPV first-price auction with 2 bidders, a customer, and a dealer who observes the customer's signal:

Proposition 4 *Let dealer's signal, x_d, and customer's signal, x_c, be uniform on $[0,1]$ and suppose any ties are broken in favor of the dealer. Then the following strategies constitute a BNE:*

$$b_d(x_d, b_c) = b_c \text{ } if \text{ } x_d \geq b_c \text{ } and \text{ } b_d(x_d, b_c) = 0 \text{ } otherwise$$

$$b_c(x_c) = \frac{x_c}{2}$$

Proof. Clearly, dealer's strategy is ex-post optimal and is thus a best-response. Thus we need to show that customer's strategy is a best response as well. It follows from the dealer's strategy that $b_d(x_d, b_c) < b_c(x_c)$ iff $x_d < b_c(x_c)$ and hence $\Pr(Customer\ x_c\ wins) = \Pr(x_d < b_c(x_c))$:

$$EU_c(x_c) = \Pr(x_d \leq b_c)(x_c - b_c) = b_c(x_c - b_c)$$

Hence maximization of the expected utility delivers the required result. ■

Notice that the customer fully reveals his information. Moreover, the customer's bidding strategy is the same regardless of whether the dealer observes her bid vs. when the dealer can not observe her bid. However, she receives only half of her expected payoff, since $\Pr(win|x_c) = \frac{x_c}{2}$ in the current auction (which is also the slope of the information rent) and it is $\Pr(win|x_c) = x_c$ in the symmetric equilibrium of a FPA. The auction is ex ante inefficient - we may have $x_c > x_d > \frac{x_c}{2}$ in which case the dealer wins even though he has lower value. The dealer, on the other hand, improves his payoff relative to the regular first price auction, since he can reduce his bid in states he wins and wins more often relative to a symmetric equilibrium of a FPA.

E Monte Carlo Study

Our ability to test the performance of the above described testing procedure in multiunit auctions is limited by the fact that in most general cases we do not have closed form solutions for equilibrium strategies, either in the private or in the affiliated values settings. We circumvent this problem by conducting a set of Monte Carlo exercises in a first price auction with independent private values, with interdependent values and pure common values, where we endow some bidders with the ability of observing a rival's bid. As mentioned earlier, the fact that some bidders' bids might be observed by their rivals is likely to have an equilibrium effect on the formation of these bids to begin with. Since we want to focus on the updating of bids associated with gaining information contained in a rival's bid, we instead consider Monte Carlo exercises where we shut down this equilibrium adjustment effect caused by the bid being revealed to a rival. In particular, in all examples, we generate the data from an equilibrium model of bidding with 3 bidders, where bidder 1 is a strategic player while bidders 2 and 3 are automatons playing as in a sealed bid first price auction with 3 bidders. We non-parametrically estimate the marginal values (of bidder 1) implied by the bids using Guerre, Perrigne and Vuong (2002) (henceforth GPV). In line with the spirit of the test we propose we assume that bidder 1 observes bidder 2's bid and submits an updated bid which supersedes her original bid. We generate the data from an equilibrium bidding function, which of course differs from the regular FPSB auction with 3 bidders. We again estimate the implied values of (informed) bidder 1 as if he were facing one less rival using GPV which assumes private values. We construct our test statistics described in the previous section and bootstrap the critical values.

E.1 First Price Auction with Informed Bidders

E.1.1 Independent private values (IPV)

The first exercise we consider is a first price auction with 3 bidders, values $v(s_i) = s_i$ and signals distributed uniformly on [0,1]. The unique equilibrium in strictly increasing differentiable strategies when all bidders are uninformed is $b^U(s_i) = \frac{2}{3}s_i$. Now consider the case that bidder 1 would be able to observe bidder 2's bid and bidders 2 and 3 are automatons that continue to bid as in a regular FPSB auction. In that case the optimal bid by the informed bidder would be:

$$b^I(s_1,s_2)=\begin{cases}\frac{s_1}{2}\,if\,\frac{s_1}{2}>\frac{2s_2}{3}\\ \frac{2s_2}{3}\,if\,s_1>\frac{2s_2}{3}>\frac{s_1}{2}\end{cases}$$

Figure 7 compares the estimated values of a bidder before and after she is informed. The Figure suggests that except at the boundaries of the support of the bids/values distribution, the values estimated with and without conditioning on observed information are likely to be very close. To correct the behavior at the boundaries, some adjustment would be necessary due to the bias in the kernel estimation, which we do by trimming $\frac{1}{9}$ of the estimated values on the bottom and $\frac{2}{9}$ of the estimated values on the top.[43]

Of course, this Figure depicts what happens only in one randomly chosen data set on bids. We then implement a joint test of the null hypothesis of the private values using randomly drawn data sets. We begin with testing for the equality of the median of estimated value distributions before and after information is received. In particular, in the first column of Table 7, we calculate the difference in the median of the distributions of values, $Med\left(\hat{v}_{informed}\right)-Med\left(\hat{v}_{uninformed}\right)$, in each Monte Carlo sample and construct the 2.5th and 97.5th bootstrap quantiles (using 400 resamples of the Monte Carlo data set) of these differences (re-centering the boostrap distribution by the test statistic computed on the original sample). The null hypothesis is rejected when the sample test statistic is not within this confidence interval. To make use of the fact that we observe two bids by the same bidder type, we also construct the "sum of squared differences" test, $SQ=\sum_i\left(\hat{v}_{i,informed}-\hat{v}_{i,uninformed}\right)^2$, and the "first order statistic" test, $FOS=\max_i\left(\hat{v}_{i,informed}-\hat{v}_{i,uninformed}\right)$, both of which we described earlier. The rejection probabilities are reported in the second and third columns of table 7. In order to understand how sampling error affects the rejection performance, we replicated the exercise for data sets of size 50, 100, and 200 – which are data sets of similar size to the empirical exercise. The (true) null of IPV is rejected less than 10% for all the cases for the median and SQ test statistics, with the SQ test statistic displaying particularly good performance. The FOS appears to overreject the null, and

[43]As the Figure suggests, the estimation bias arises mostly in the upper tail of the distribution and therefore we focus the trimming there. Trimming at the lower tail does not affect the results of our Monte Carlo experiments.

gives particularly bad results when the data is not trimmed.

Finally, to evaluate an joint hypothesis test , we also report the results of the pointwise test, in which we compare the absolute value of the difference between each $\hat{v}_{informed}$ and $\hat{v}_{uninformed}$ pair, to the difference at a given quantile of the (re-centered) bootstrap distribution of the difference between *that pair*. In particular, we report the percentage of points where we obtain rejection where each pointwise difference is compared to the 95^{th} percentile, and to the quantile correponding to the Bonferroni's method, $100 - \frac{5}{\#\ of\ hypotheses}$.

The null rejection frequencies of this alternative testing procedure is displayed in the first two columns of Table 8. While the Bonferroni method is generally regarded as too conservative, in our setting it still would lead to a wrongful rejection of the null hypothesis, even though it (unsurprisingly) leads to a large reduction in rejection rates.

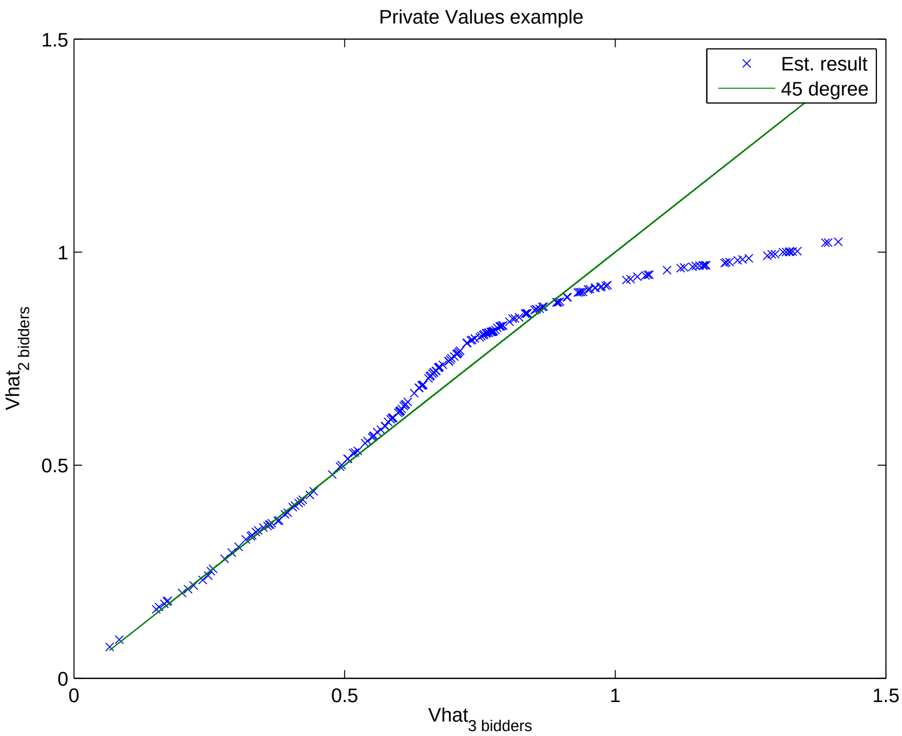

Figure 7: Estimated values for informed (i.e., facing only 1 rival) and uninformed bidders (i.e., facing 2 rivals) in a FPA with private values

E.1.2 First price auction with interdependent values and independent signals (IIV)

In the second exercise we look at a first price auction with interdependent values and independent signals (IIV). The valuation function is $v\left(s_i, s_{-i}\right) = \frac{s_i}{2} + \frac{\sum_{j \neq i} s_j}{2(n-1)}$ where $s_i \sim U[0,1]$. The unique symmetric equilibrium in strictly increasing differentiable strategies involved bidding according to $b^U\left(s_i\right) = \frac{7}{12} s_i$. In the appendix we show that the equilibrium strategy of an informed bidder who observes a bid of his rival (and thus for practical purposes another signal S_2) is bidding according to:

$$b_1^I(s_1, s_2) = \begin{cases} \frac{7}{11}\left(\frac{s_1}{2} + \frac{s_2}{4}\right) & if \quad s_1 > \frac{4}{3} s_2 \\ \frac{7}{12} s_2 & if \quad \frac{4}{3} s_2 \geq s_1 \geq \frac{25}{48} s_2 \end{cases}$$

Figure 8 depicts the results of estimating the implied values using GPV for a randomly selected data set. The null rejection frequencies of the testing procedures utilized in the IPV example are displayed in Tables 7 and 8. Observe that in this case, the FOS test appears to perform the best, in that for data sets of size exceeding 100, the test appears to work very well in that it rejects the null with close to 95% probability. The median and especially the SQ tests have lower power for smaller data sizes, though with $N = 200$, their performance increases dramatically. A similar pattern is observed in the pointwise tests.

In Table 9, we also report the studentized versions of the SQ and FOS test statistics, where the individual test statistics are scaled by their standard deviation. It appears that studentization increases the power of both FOS and SQ for all sample sizes – allowing FOS to reject the alternative in more than 92% of the time for all sample sizes considered.

E.1.3 First price auction with pure common values

In our third exercise we examine a first price auction with pure common values described in Matthews (1984). Let the utility be $u_i\left(s_i\right) = v$ where $v \sim \text{Pareto}\left(\alpha\right) : \, g(v) = \alpha v^{-(\alpha+1)}$ for $1 \leq v \leq \infty$ and $F(s|v) = \frac{s}{v}$.

In this case Matthews shows that there is a unique (symmetric) equilibrium in differentiable

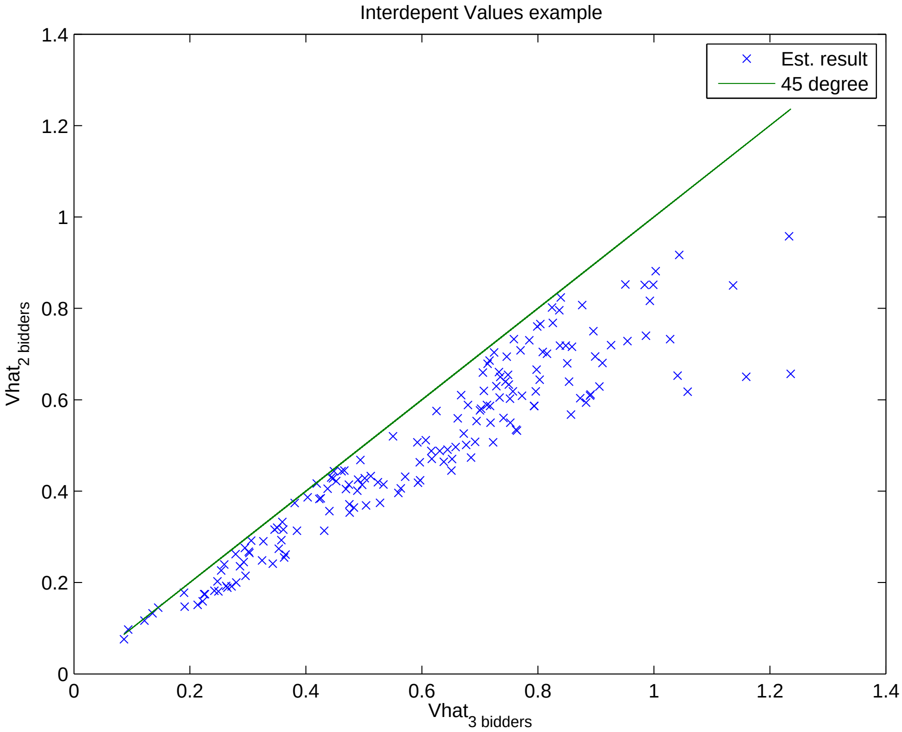

Figure 8: Estimated values for informed (i.e., facing only 1 rival) and uninformed (i.e., facing 2 rivals)bidders in a FPA with interdependent values

strictly increasing strategies of the form:

$$b(s) = \left(\frac{(N-1) + [\max\{1, s\}]^{-(N-1)-1}}{(N-1)+1} \right) \hat{v}(s, N)$$

where

$$\hat{v}(s, N) = \frac{N+\alpha}{N+\alpha-1} \max\{1, s\}$$

is the expected value conditional on winning. Notice that for $s \geq 1$ we have

$$b(s) = \frac{(N+\alpha)\, s\left[(N-1) + s^{-N}\right]}{N(N+\alpha-1)}$$

Now if bidder 1 were again to observe bidder 2's bid, two cases can occur: either he can infer s_2 or that $s_2 < 1$.

Suppose that $s_2 \leq s_1$. Then the optimal bid is as before since no such signal is informative about realized v conditional on winning ($s_{\max}$ is a sufficient statistic of the sample $(s_1, ..., s_N)$ for v). On the other hand, if $s_2 > s_1$, then the optimal bid becomes:

$$b(s_1, s_2) = \frac{(N+\alpha)\, s_2 \left[(N-1) + s_2^{-N}\right]}{N(N+\alpha-1)}$$

In other words, bidder uses just the highest signal he observes to base his bid upon and updates the prior on the distribution of v using the winning event.

We once again generated data for an informed and an uninformed bidder using the above described bidding strategies and used GPV to estimate the implied values under the null hypothesis of private values. The results, for a randomly chosen data set, are displayed in Figure 9.

In contrast to the IIV case, the median test appears to perform the best in this case. Figure 9 sheds some light into what might drive this result: it appears that boundary effects are particularly important in this example. Thus, trimming is particularly effective in increasing the power of SQ and FOS tests. Note that studentization also helps increase the power of SQ and FOS tests, as displayed in table 9.

Table 7: Monte Carlo Exercises: Joint test

	Rejection frequency								
	IPV			IIV			CV		
N	Median[a]	SSQ[b]	FOS[c]	Median	SSQ	FOS	Median	SSQ	FOS
50	0.30	0.06	0.20	0.52	0.52	0.52	0.76	0.14	0.32
100	0.16	0.28	0.62	0.68	0.78	0.90	0.99	0.18	0.15
200	0.30	0.26	0.80	0.82	0.96	0.98	0.99	0.12	0.18
50: trimmed[d]	0.20	0.00	0.00	0.46	0.56	0.74	0.70	0.60	0.46
100: trimmed	0.18	0.00	0.02	0.64	0.86	0.96	0.99	0.78	0.58
200: trimmed	0.14	0.00	0.06	0.76	0.94	0.98	0.99	0.99	0.82

[a] Test based on difference in medians of distributions.
[b] Test based on sum of squares of individual test statistics.
[c] Test based on the first-order statistic of individual test statistics.
[d] $\frac{1}{9}$ of estimated values discarded at the lower tail, $\frac{2}{9}$ in the upper tail.

Table 8: Monte Carlo Exercises: Pointwise test

	Rejection frequency					
	IPV		IIV		CV	
N	Ptwise^{a}_{95}	Ptwise^{b}_{Bonf}	Ptwise_{95}	Ptwise_{Bonf}	Ptwise_{95}	Ptwise_{Bonf}
50	0.09	0.01	0.34	0.07	0.69	0.09
100	0.11	0.01	0.47	0.15	0.79	0.39
200	0.10	0.01	0.61	0.35	0.84	0.68
50: trimmed[c]	0.03	0.01	0.36	0.10	0.84	0.20
100: trimmed	0.02	0.01	0.48	0.18	0.96	0.60
200: trimmed	0.02	0.01	0.61	0.33	0.99	0.83

[a] Comparing each hypothesis with 95^{th} percentile of the corresponding asymptotic distribution.
[b] Comparing each hypothesis with $\left(1-\frac{0.05}{N}\right)^{th}$ percentile of the corresponding asymptotic distribution.
[c] $\frac{1}{9}$ of estimated values discarded at the lower tail, $\frac{2}{9}$ in the upper tail.

Table 9: Monte Carlo Exercises: Studentized joint test

	Rejection frequency					
	IPV		IIV		CV	
N	SSQ^{a}_{Stud}	FOS^{b}_{Stud}	SSQ_{Stud}	FOS_{Stud}	SSQ_{Stud}	FOS_{Stud}
50	0.04	0.10	0.64	0.74	0.99	0.68
100	0.04	0.24	0.86	0.98	0.99	0.98
200	0.06	0.32	0.99	0.99	0.99	0.96
50: trimmed[c]	0.00	0.06	0.48	0.66	0.88	0.80
100: trimmed	0.00	0.02	0.82	0.98	0.99	0.99
200: trimmed	0.00	0.04	0.98	0.99	0.99	0.99

[a] Test based on sum of squares of studentized individual test statistics.
[b] Test based on the first-order statistic of studentized individual test statistics.
[c] $\frac{1}{9}$ of estimated values discarded at the lower tail, $\frac{2}{9}$ in the upper tail.

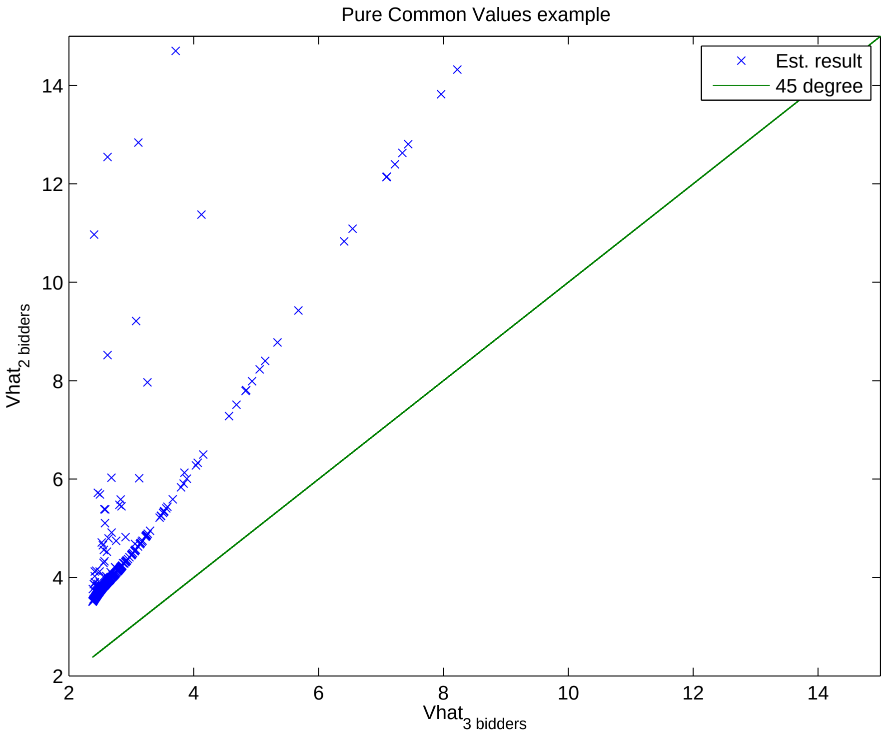

Figure 9: Estimated values for informed (i.e., facing only 1 rival) and uninformed (i.e., facing 2 rivals) bidders in a FPA with pure common values

E.1.4 Monte Carlo Summary

Our Monte Carlo exercises reveal several important observations regarding the joint test statistics FOS and SQ: due to the bias of the kernel estimates, without restricting attention to a subset of hypotheses which consider bidders with estimated values that are likely far enough from the boundary of the support, the tests can have much lower power in smaller samples than desired. While the joint hypothesis test based on the first-order statistic (FOS) seems to have higher power against the alternatives considered here, it also over-rejects the correct null unless the extent to which we trim the estimated marginal values increases substantially. The test based on sum of squares, SQ, on the other hand, does not over-reject the null, but has lower power against the alternatives, especially when no trimming is applied.[44] Both tests' power against the alternative

[44]One may speculate that the FOS may be more sensitive to outliers; thus in the IPV case, it may overreject. Indeed, the performance of FOS is markedly better with trimming (SSQ improves as well, but not as much). As for the IIV, one may reverse the argument, and say that because SSQ is the more conservative with respect to

appears to improve if they are studentized, though studentization leads to a slight amount of overrejection of the null. We also found that joint hypothesis testing based on multiple independent hypotheses and the Bonferroni correction does not generate significantly better performance than FOS or SQ.

Given that the various testing approaches appear to have different strengths and weaknesses, we will present results of all three main testing approaches in our application: (i) individual hypothesis tests, $\{S_i, T_i\}_{i=1}^N$, where N is the number of hypotheses, (ii) sum of squares, (SQ_T, SQ_S), and (iii) first-order statistic test, (FOS_S, FOS_T).

E.2 Closed form solutions for Monte Carlo exercises

Here we present the derivation of the closed form solution for bidding used to generate data in our Mont Carlo studies with 3 bidders.

E.2.1 First price auction with independent private values

Let the utility function be:

$$u_i = x_i$$

In this case bidder 1 maximizes $\Pr(b_1 > max\{b_2, b_3\})(x_1 - b_1)$ which implies that the symmetric equilibrium bidding function is:

$$b(x) = \frac{2}{3}x$$

If he observed 2's bid, he would bid in 2 cases (assuming any tie is broken in 1's favor and bidders 2 and 3 continue using the strategies given above) using the rule:

$$b(x_1, x_2) = \begin{cases} \frac{x_1}{2}\, if\, \frac{x_1}{2} > \frac{2x_2}{3} \\ \frac{2x_2}{3}\, if\, x_1 > \frac{2x_2}{3} > \frac{x_1}{2} \end{cases}$$

where the second case occurs whenever bid of bidder 1 using the rule for the first case would be

outliers, it is not rejected under the alternative hypothesis as frequently as FOS. Unfortunately, the difficulty of finding analytically tractable examples to work with (especially in the IIV and CV cases) make it very difficult for us to check the validity of the previous intuition with many more examples – thus we can not rule out the possibility that the relative performance of SSQ and FOS may be particular to this Monte Carlo exercise.

lower than 2's bid, but bidder 1 would prefer to win the object.

E.2.2 First price auction with interdependent values

Let the utility function be:

$$u_i = \frac{x_i}{2} + \frac{\sum_{j \neq i} x_j}{2(n-1)}$$

where

$$x_i \sim U[0,1]$$

With 3 bidders there exists a unique symmetric equilibrium in differentiable strictly increasing strategies:

$$b(x) = \frac{7}{12}x$$

Now suppose that bidders 2 and 3 follow these strategies. Suppose bidder 1 can observe bidder 2's bid. Since 2's strategy is strictly increasing, bidder 1 can recover the signal s_2. In this case, bidder 1's expected payoff when getting a signal s_1, observing s_2 and bidding b is:

$$\int_0^{\frac{12b}{7}} \left(\frac{s_1}{2} + \frac{s_2}{4} + \frac{\alpha}{4} - b\right) d\alpha$$

Maximizing this expression w.r.t. the bid b results in:

$$b^*\left(s_1, s_2\right) = \frac{7}{11}\left(\frac{s_1}{2} + \frac{s_2}{4}\right)$$

Finally, this bid is weakly higher than $b_2 = \frac{7}{12}s_2$ if and only if $s_1 \geq \frac{4}{3}s_2$. In the case that b_2 is higher than $b^*\left(s_1, s_2\right)$, bidder 1 will still prefer to win if $s_1 \geq \frac{25}{48}s_2$.

Table 10: Joint Hypothesis Studentized Test Results

		3-months			12-months		
		Auctions for resampling			Auctions for resampling		
	Bandwidth	1	2	4	1	2	4
SSQ[a]	100	49.37	199.06	188.81	48.99	376.5	277.58
Crit Value		1265.74	1589.46	1555.34	1438.91	1762.41	1834.39
Std Dev		424.16	492.07	583.21	452.25	550.4	605.73
pvalue		1	1	0.95	0.98	0.84	0.85
FOS[b]		3.86	9.87	6.04	5.54	15.31	8.96
Crit Value		19.86	19.97	19.98	19.97	19.97	19.97
Std Dev		5.25	3.89	5.17	5.44	1.99	4.62
pvalue		0.96	0.51	0.74	0.83	0.43	0.8
95th percentile[c]		0.23	0.94	1.37	0.2	0.64	0.97
Crit Value		1.72	2.3	2.63	1.48	1.68	2.01
Std Dev		0.37	0.44	0.45	0.35	0.38	0.42
pvalue		1	0.96	0.87	1	0.78	0.83
Fraction trimmed		0.05	0.06	0.05	0.09	0.1	0.09
SSQ	500	195.74	267.16	304.87	135.81	397.45	225.08
Crit Value		1464.42	1540.51	1565.46	1481.45	1882.04	1598.21
Std Dev		416.58	472.16	606.08	454.79	542.84	555.79
pvalue		0.99	0.97	0.58	0.95	0.85	0.8
FOS		11.33	12.93	10.04	5.89	16.78	5.46
Crit Value		19.92	19.97	19.97	19.97	19.98	19.97
Std Dev		3.32	3.21	5.39	5.15	1.43	5.89
pvalue		0.53	0.37	0.41	0.86	0.39	0.78
95th percentile		0.39	0.53	1.46	0.16	0.54	1.33
Crit Value		2.13	2.21	2.41	1.63	1.48	2.36
Std Dev		0.46	0.46	0.41	0.44	0.31	0.41
pvalue		1	1	0.67	1	0.95	0.67
Fraction trimmed		0.05	0.06	0.05	0.09	0.1	0.09
SSQ	5000	87.9	190.82	242.25	176.02	17463.34	545.96
Crit Value		1542.45	1473.43	1695.16	1517.09	18988.42	1715.39
Std Dev		467.9	482	584.54	450.88	567.29	401.92
pvalue		0.99	0.99	0.85	0.96	0.79	0.81
FOS		5.84	9.15	7.58	8.48	120.77	12.49
Crit Value		19.86	19.97	19.97	19.97	122.37	14.13
Std Dev		4.98	4.02	5.01	4.54	1	0.96
pvalue		0.84	0.63	0.58	0.71	0.52	0.75
95th percentile		0.33	0.86	1.45	0.23	1.07	1.56
Crit Value		1.93	2.69	2.49	1.51	2.53	3.12
Std Dev		0.44	0.62	0.44	0.37	0.52	0.51
pvalue		1	0.97	0.77	1	0.84	0.94
Fraction trimmed		0.05	0.06	0.05	0.09	0.1	0.09

[a] Test based on sum of squares.
[b] Test based on first-order statistic.
[c] Test based on 95th percentile of the test statistic distribution.