

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

Ali Hortaçsu[†] Jakub Kastl[‡]

October 13, 2011

In many financial markets, dealers have the advantage of observing the orders of their customers. To quantify the economic benefit that dealers derive from this advantage, we study detailed data from Canadian treasury auctions, where dealers observe customer bids while preparing their own bids. In this setting, dealers can use information on customer bids to learn about (i) *competition*, i.e. the distribution of competing bids in the auction, and (ii) *fundamentals*, i.e. the ex-post value of the security being auctioned. We devise formal hypothesis tests for both sources of informational advantage. In our data, we do not find evidence that dealers are learning about fundamentals. In the language of auction theory, our data is therefore consistent with the private value paradigm. We find that the “information about competition” contained in customer bids accounts for 13–35% of dealers’ expected surplus.

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, Jon Levin, Rob Porter, Peter Reiss, Azeem Shaikh, Quang Vuong 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

Many financial markets are organized around dealer/specialists who have information on other traders' orders, and can potentially utilize this information to their advantage. 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."¹ Theoretical discussions regarding whether dealers should be able to trade on their own account as well their customers' (i.e. "dual trading") have also been quite lively in the context of futures markets (e.g. Grossman (1989), Fishman and Longstaff (1992)). Empirical studies of dealers' usage of customer order-flow information include analyses of specialist/dealer behavior in the NYSE (Madhavan and Smidt (1991), in the foreign exchange market (Lyons (1995) and Evans and Lyons (2002)), and in options markets (Easley, O'Hara and Srinivas (1998)).

This paper seeks to study and quantify the economic benefits that dealers derive from having access to customer order information in the setting of Canadian Treasury auctions, where the potential for informational advantage is particularly transparent. In these auctions, government securities dealers route the bids of non-dealer bidders, called "customers."² We should note that the institutional setup of Canadian Treasury auctions is echoed around the world: out of 39 countries surveyed by Arnone and Iden (2003), 29, including the U.S., employed a similar "primary dealership" system that limited participation a small number of bidders who also place bids on behalf of their customers.

In the (sealed bid, discriminatory price) auction setting we analyze, there are two sources of economic benefit to dealers from observing customer bids. In a general interdependent values setting, customer bids are informative about the ex-post or *fundamental* value of the securities being auctioned, and thus may induce the dealer to revise her willingness-to-pay for the security on sale. Such learning about fundamentals is not necessarily socially undesirable, as it leads to information aggregation across bidders (see e.g., Hayek (1945) or Grossman (1976)) . Along with

¹<http://www.nytimes.com/2009/03/05/business/05specialist.html>.

²Most "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 2.1 for summary statistics.

such fundamentals related information, however, customer bids are also useful to the dealer since they provide information regarding the *competition* they will face in the auction. As such, customer bids may still compel a dealer to modify her bid, even in a purely private values setting where learning about fundamentals does not play any role.

Our data from Canadian Treasury auctions allow us to study the above two mechanisms in 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. Observing bid modifications *per se* does not allow us to infer the nature of the information that dealers extract from customer bids, however. Indeed, consider a situation in which bidder i is about to submit her bid y_i , but before submitting y_i she observes a bid actually submitted by bidder j . Bidder j 's "competitive" information allows bidder i to improve her estimate of the location and shape of residual supply³ she will be facing in the auction. Using this additional information, she revises her initial bid y_i , and submits an alternative bid y'_i . In a discriminatory auction, this additional information allows the bidder to submit a bid that is "closer" to the expected market clearing price, thereby reducing payments on the inframarginal units. If "fundamentals" related information is also relevant, which would be the case in an interdependent values environment, bidder i will also update her prior on the ex-post value of the securities by inverting bidder j 's bid, and will submit a new bid, y''_i , taking into account both of these new pieces of information.

In our empirical analysis, we start with the null hypothesis of a purely private values setting, where learning about fundamentals does not play any role. We then test whether the observed modifications to dealer bids in response to customer bids can be rationalized in this setting. Specifically, we build on our earlier work (Hortaçsu (2002a) and Kastl (2011)) 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 marginal valuation that we estimate for a dealer's bid before observing a customer's bid vs. after should be the same.

³Residual supply is defined as the total quantity for the auction minus the bids of other bidders.

Our empirical tests, reported in Section 5, strongly suggest that learning about fundamentals is not needed to rationalize observed bid modifications in our data on auctions of 3- and 12-month Treasury bills. This result also enables us to calculate the economic benefit to a dealer from observing customer bids. The estimated marginal valuations allow us to calculate the ex-post surplus of the dealer. Thus, we can 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. Consistent with views of practitioners we find, in Section 6, that observing customer order flow contributes significantly to dealers' overall profits from participating in Canadian Treasury auctions.

Outside of quantifying the economic benefits that dealers derive from observing customer bids, our empirical strategy also allows us to test whether the informational environment of Canadian treasury auctions is one of private values. Empirically distinguishing between private vs. 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.⁴ To our knowledge, this is the first attempt to test the null hypothesis of private values in a multi-unit divisible good auction 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. 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.

The rest of the paper proceeds as follows: in Section 2, we present the data, and some descriptive

⁴As pointed out by Ausubel and Cramton (2002), neither the revenue equivalence theorem of Vickrey (1961) nor Milgrom and Weber (1982) 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 (Fevrier, Pregel and Visser (2002), Armantier and Sbaï (2006), Hortaçsu (2002a), Kastl (2011), Kang and Puller (2006), Chapman, McAdams and Paarsch (2006, 2007)) 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.

evidence suggesting that dealers modify their bids in a way that appears to reflect the information in customer bids. This evidence alone, however, does not allow us to test between the “learning about fundamentals” vs. “learning about competition” explanations. In Section 3, we construct a model of bidding which allows us to make quantitative predictions about how dealers should modify their bids in response to the information in customer bids. The model also leads to a statistical test between the “learning about fundamentals” vs. “learning about competition” hypotheses, which we describe in Section 4, with the results of the tests reported in Section 5. In Section 6, we use our estimated model to calculate the economic benefit that dealers derive from observing customer bids.

2 Description of Data and Institutional Background

The Government of Canada sells Treasury bills and other securities 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 large banks that for some reason choose not to be registered as dealers. 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.

2.1 Data

Our data 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. 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. This allows us 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.

Table 1 offers some summary statistics of our data set.

Table 1: Data Summary

Summary Statistics for 3-month T-bill auctions				
	Mean	St.Dev.	Min	Max
<hr/> # of Auctions: 116 <hr/>				
# 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
# of Non-Competitive Bids	4.88	1.25	3	9
Issued Amount (billions C\$)	3.88	0.55	2.8	5
<hr/> # of Bids: 3,511 <hr/>				
Quantity-Weighted Price Bid	989,234	3,224	984,515	994,964
Maximum Quantity Share Demanded ^a	0.12	0.09	0.00023	0.2512
<hr/> # of Non-Competitive Bids: 566 <hr/>				
Size of Non-Competitive Bids ^a	0.03	0.07	0.000002	0.27

^a As a percentage of total supply.

On average, 12 dealers and about 5 customers participate in every auction. An average bid function consists of less than 3 steps and the average maximum quantity demanded is about 13% 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. Like a market order, a noncompetitive tender specifies a quantity that the bidder wishes to purchase at the price at which the auction clears. In our data, there are on average 4.88 noncompetitive tenders in an auction – 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. The non-competitive bids essentially reduce the available supply to competitive bidders.⁵

⁵The 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 3% 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 0.07% of the supply.

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), customers participate slightly more. 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				
	Mean	St.Dev.	Min	Max
# of Auctions: 116				
# of Dealers	12.28	1.59	8	16
# of Customers	6.03	3.29	0	18
# of Participants	18.31	4.29	10	32
# of Submitted Steps	2.91	1.69	1	7
# of Non-Competitive Bids	4.60	1.14	2	7
Issued Amount (billions C\$)	1.67	0.194	1.3	2
# of Bids: 4, 126				
Quantity-Weighted Price Bid	958,261	10,788	940,181	979,851
Maximum Quantity Share Demanded ^a	0.13	0.09	0.0005	0.25
# of Non-Competitive Bids: 534				
Size of Non-Competitive Bids ^a	0.03	0.05	0.00003	0.20

^a As a percentage of total supply.

2.2 Preliminary evidence on dealer information advantage

A preliminary indicator of whether dealers are utilizing the information in their customers' bids is whether dealers are modifying their bids in response to observing a customer bid. In our 3 month T-bill sample, out of 660 dealer bids (which were also accompanied by a customer bid), 216 previously placed dealer bids were "updated" after seeing a customer bid. In our 12 month sample, out of 659 dealer bids, 275 were updated after routing a customer bid. It is also possible that dealers wait to see all of their customer bids before submitting their own bids: indeed, in the 3-month T-bill sample, 154 dealer bids were submitted for the first time after receiving the customer bid, and in the 12 month sample 250 bids were submitted for the first time after routing a customer bid.

Another suggestive indicator, reported in Table 3, is that customer bids are a statistically

significant (at the 1% level for the 3M sample and 10% level for the 12M sample) correlate of the (quantity-weighted average) price level of the updated dealer bid, controlling for the dealer’s bid before observing the customer bid. The updated dealer bid appears to load more heavily on the customer bid in the 3-month sample as opposed to the 12 month sample.

Table 3: Correlation between dealer bid updates and customer bids

	3 Month Updated bid	12 Month Updated bid
Customer bid	0.146*** (0.0289)	0.0166* (0.00863)
Dealer’s orig. bid	0.853*** (0.0289)	0.982*** (0.00860)
Constant	1,399** (609.5)	946.6* (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 *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

What about dealers who did not update/modify their bids in response to observing a customer bid (290 instances in the 3-month sample, and 134 instances in the 12-month sample)? One possibility is that dealers might not have time to update their bids in response to customer bids: 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, with the distributions (of customer bids that are followed by a dealer update and those that are not) following a clear first order stochastic dominance relation, as plotted in Figure 2.2 and Figure 2.2 below. Note that these figures include the negative domain – these are customer bids that arrive after the deadline passes and consequently can not participate in the auction.

To check whether 5 minutes is enough for a dealer to change his bid, we calculated the average




Figure 1: Comparison of customer bid submission times with and without dealer updates

(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 – which suggests that it may take some time for dealers to incorporate any information in customer bids. Thus, it is possible that a customer bid submitted 5 minutes before the deadline might be “too late” for a dealer to update her own bid.⁶

Indeed, “last minute” bidding behavior by customers can be a best response by strategically sophisticated customers who do not want dealers to utilize the information in their bids. In previous work studying the same market, 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. This may also explain why dealers bid early, and do not simply wait to submit their bid until they have seen the bids of their customers.

We also investigated how customers choose which dealer to route their bids through – in particular whether customers are in exclusive relationships with their dealers, or whether they switch providers. To do this, we focused on customers who were in the top 50% of customers in terms of participation (i.e participating in 14 auctions or more in 3 month auctions, 10 auctions or more in

⁶We 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 customer bids or not; some dealers may be following relational agreements with their customers, and purposely avoiding to update their bids.

12 month auctions). We defined a “exclusive” relationship if they used a given dealer more than 75% the time. We found that 21% (25%) of these customers in 3(12)-month auctions are in an exclusive relationship by this definition. Of the total number of bids sent to dealers by customers, 15%(21%) are routed by customers in an exclusive relationship in 3(12)-month auctions.⁷

As we will discuss in more detail in Section 3.1, these routing patterns are important in formulating how dealers calculate their bids before the arrival of customer information. If most customer bids come from non-exclusive customers, then dealers can infer that the absence of a customer bid at a given time means that the customer has likely routed her bid through other dealers. The absence of an exclusive customer’s bid by the time of the dealer’s bidding, however, suggests that this customer may not be participating in the current auction.

What we report above are purely descriptive findings, and the model of bidding in the next section will make precise and quantitative predictions about the type of bid modification that a dealer should make upon seeing her customer’s bids.⁸ Most importantly, the descriptive analyses above can not be used to distinguish between the “learning about the fundamentals” vs. “learning about competition” explanations of dealer behavior. The model below will be used as the basis of a formal hypothesis test between the two mechanisms.

3 Model of Bidding

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 discreteness of bidding (i.e., finitely many steps in bid functions) as in Kastl (2011). 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: 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

⁷The fraction of bids by customers in exclusive relationships is 3%(5%) if we increase the exclusivity threshold to 90% of bids and 35%(29%) if we set the threshold at 60% of bids.

⁸Hortaçsu and Sareen (2006) report further descriptive measures suggesting that obtaining customer information impacts dealers’ bidding patterns. For example, 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.

multidimensional) signals, $S_1^c, \dots, S_{N_c}^c, S_1^d, \dots, S_{N_d}^d$.

Assumption 1 *Customers' and dealers' private signals, $S_1^c, \dots, S_{N_c}^c, S_1^d, \dots, S_{N_d}^d$, are drawn from a common support $[0, 1]^M$ according to an atomless joint d.f. $F(S_1^c, \dots, S_{N_c}^c, S_1^d, \dots, S_{N_d}^d)$ with strictly positive density f . The distribution of private signals also satisfies the within group symmetry assumption, i.e. $F(\cdot)$ 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 $F^d\left((S_1^d, Z_1), \dots, (S_{N_d}^d, Z_{N_d}) \mid \mathbf{S}^c, \{y_i^c(p|s_i^c)\}_{i=1}^{N_c}\right)$ 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 $g \in \{c, d\}$.*

An important special case we will consider with regards to valuation structures is the case where bidders have *private values*, i.e. when $\forall i$ and $g \in \{c, d\} : v^g(q, S_i, S_{-i}) = v^g(q, S_i)$. This is the case where learning other bidders’ signals does not affect one’s own valuation – thus “learning about fundamentals” does not play a role.

In the more general case where bidders’ marginal valuation functions are allowed to depend on other bidders’ signals, and thus “learning about fundamentals” does matter, we will make the assumption that 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 *$E_{S_{-i} \setminus S_j^c} \left[v^d(q, S_i^d, S_{-i}) \mid S_i^d = s_i^d, S_j^c = s_j^c \right]$ is strictly increasing in (each component of) $S_j^c, \forall (q, s_i^d)$.*

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 $\sigma_i : \Theta_i \rightarrow \mathcal{Y}$, 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, \bar{K} , we impose the following assumption:

Assumption 4 *Each customer and dealer, $i = 1, \dots, N_c, N_c + 1, \dots, N_c + N_d$ has an action set:*

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

where a bid of 0 denotes non-participation and a bid of ∞ denotes a non-competitive bid. The set \mathcal{Y} includes all non-decreasing step functions with at most \bar{K} steps, $q : \mathbb{R}_+ \rightarrow [0, 1]$, where $q_i(p) = \sum_{k=1}^K q_{ik} I(p \in (b_{ik+1}, b_{ik}])$ 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 be denoted by $y_i^g(p|\tilde{\theta}_i)$ 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. Since bidders use step functions as their bids, which leads to rationing except in very rare cases, we will assume pro-rata on-the-margin rationing, which is the rule used by the Bank of Canada to proportionally adjust the marginal bids so as to equate supply and demand. Also, in situations where multiple prices clear the market, we assume that the auctioneer selects the highest market clearing price.

We also assume that 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 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.

3.1 Characterization of equilibrium dealer bids

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. $y_i^c(p|s_i) = y^c(p|s_i) = y^c(p|\theta_i)$, $i \in \mathcal{C}$. Dealers' bid functions are also symmetric, but up to their private signal *and* order flow information, i.e. $y_i^d(p|s_i, z_i) = y^d(p|s_i, z_i) = y^d(p|\theta_i)$, $i \in \mathcal{D}$.

In this setting, market clearing price is a random variable, P^c , mapping the state of the world, $(Q, \mathbf{s}^c, \mathbf{s}^d, \mathbf{z})$, or simply $(Q, \vec{\theta})$, into prices through equilibrium strategies. This random variable is thus summarized by a function $P^c(Q, \mathbf{s}^c, \mathbf{s}^d, \mathbf{z}, \mathbf{y}^c(\cdot|s), \mathbf{y}^d(\cdot|s, z))$ (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:

$$\Pr(p \geq P^c | s_i, z_i) = E_{\{S_j \in \mathcal{C} \cup \mathcal{D} \setminus i, Z_k \in \mathcal{D} \setminus i | z_i = \emptyset\}} I \left(Q - \sum_{j \in \mathcal{C}} y^c(p|S_j) - \sum_{k \in \mathcal{D} \setminus i} y^d(p|S_k, Z_k) \geq y^d(p|s_i, \emptyset) \right) \quad (1)$$

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. $z_i = \{y^c(p|s_m)\}$

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

These probabilities will play a crucial role in equilibrium characterization. Note that the main difference of equation (1) compared to equation (2) is that the dealer conditions on her customer's bid, instead of taking an expectation over the customer's private information conditional on the fact that she has observed no customer bid. Indeed, this is exactly where "learning about competition" occurs – the dealers' expectations about the distribution of the market clearing price are altered once she observes her customer's bid.

Given these probabilities, a necessary condition for bidding in the case where dealers do not “learn about fundamentals” is given by the following:

Proposition 1 (*Kastl (2008)*) *Suppose values are private, i.e. there is no “learning about fundamentals.” Under assumptions 1-4 in any Bayesian Nash Equilibrium of a Discriminatory Auction, for a bidder of type θ_i submitting $\hat{K}(\theta_i) \leq K$ steps, every step k in the equilibrium bid function $y(\cdot|\theta_i)$ has to satisfy:*

$$v(q_k, \theta_i) = b_k + \frac{\Pr(b_{k+1} \geq P^c | \theta_i)}{\Pr(b_k > P^c > b_{k+1} | \theta_i)} (b_k - b_{k+1}) \quad (3)$$

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

We should emphasize here that the necessary conditions above apply to the optimization problem of a dealer. The solution to customers’ optimization problem may be different as they should rationally anticipate that their bids reveal information to dealers and thus adjust their bids. Since we are interested in evaluating the impact of information contained in customers’ order flow on primary dealers’ rents, we only need to obtain estimates of dealers’ marginal values and not those of customers.

We now specify in more detail how to calculate the conditional market clearing price distribution in equation (1), which is the key component of equation (3). The expectation conditional on not observing a customer bid depends on dealers’ beliefs about how customers route their bids. To simplify the analysis, we assume that a dealer can route at most one customer bid in an auction. Since we observed in Section 2.2 that most customers switch dealers regularly while others appear to have exclusive relationships with one dealer, we entertain two scenarios. In the first case, we assume that each customer has an exclusive relationship with exactly one dealer and hence the beliefs about that customer’s signal conditional on observing no customer bid put all weight on customer signals as being low enough to induce non-participation. In the second case, we assume that customers get randomly and independently matched to dealers.

Let $M(i, j)$ denote the probability that customer j gets matched with dealer i , then (1) be-

comes:

$$\Pr(p \geq P^c | s_i, z_i = \emptyset) = E_{\{S_j \in \mathcal{C} \cup \mathcal{D} \setminus i, Z_k \in \mathcal{D} \setminus i\}} \left(Q - \sum_{j \in \mathcal{C}} (1 - M(i, j)) y^c(p | S_j) - \sum_{k \in \mathcal{D} \setminus i} y^d(p | S_k, Z_k) \right) \geq y^d(p | s_i, z_i) + y^c(p | s_m) \quad (4)$$

With exclusive relationships, $M(i, j)$ is only equal to 1 if dealer i and customer j are in an exclusive relationship, and zero if they are not. In this case, even if dealer i does not observe j 's bid, there is “information about competition:” the dealer will infer that j will not participate, which is equivalent to knowing that only $N_c - 1$ potential customers participate. If customers are iid randomly matched to dealers, the dealer will get less information about the competition in the auction: $M(i, j) = \frac{1}{N_d}$, and hence dealer i attaches only $\frac{1}{N_d}$ probability to j not participating, but with complementary probability $1 - \frac{1}{N_d}$ integrates over all N_c potential customer bids. Note that since $N_d = 12$ in practice, this latter case is approximately equivalent to the situation where the dealer does not get any information about the competition from not observing a customer bid, i.e. $M(i, j) = 0$ and the dealer just integrates over all N_c potential customer bids.

3.1.1 Learning about fundamentals

In the more general case where dealers' marginal valuations can depend on other dealer or customers' signals, i.e. dealers can “learn about fundamentals” by observing customer bids, Appendix A shows that the analogous necessary condition for optimality is:

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

where $\theta_i = (S_i, Z_i)$ and $\theta_{-i} = (S_{-i}, Z_{-i})$, and:

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

In the above expression, $V(q, \theta_i, \theta_{-i}) = \int_0^q v(u, \theta_i, \theta_{-i}) du$, and Q_i^c is the quantity won by bidder i if his bid q_k is rationed at b_k or b_{k+1} .

The term at the right-hand side of (5) is identical to the term in (3) for the private values case (when bidders only learn about competition), 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, $E_{S_{-i}, Z_{-i} | s_i, z_i, P^c \in (b_{k+1}, b_k), q_k} [v(q_k, s_i, z_i, S_{-i}, Z_{-i})]$, 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)).⁹ However, there is an additional term which captures how the 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.¹⁰

⁹Note, however, that P^c is potentially a complicated function of bidder signals that depends on equilibrium bidding 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.

¹⁰In a recent paper, De Castro and Riasco (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.

4 Test Specification

The main idea behind our test for whether “learning about fundamentals” is an important factor is to find instances where a dealer observes customer information, and to test whether the marginal valuation that rationalizes that dealer’s bid remains constant before and after accounting for the “information about competition” provided by that customer bid. If it does, we are in the case characterized by equation (3) above.

4.1 Estimating Marginal Valuations

To implement this test, we first have to estimate the rationalizing marginal valuations, for which we follow Hortacısu (2002a) and Kastl (2011). 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 the market clearing prices. 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 (3).

We now turn to the present context where 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 (1), 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.¹¹ 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 (2), 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 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.

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.¹² Of course, an important concern is whether the

¹¹Ideally, 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. Asymptotically, we expect the number of dealer bids corresponding to a given customer bid to increase; however, in small samples, this is rarely true. Similar to kernel estimation, instead of drawing informed dealer bids that exactly correspond to a given customer bid, we draw bids from dealers who saw customer bids that are “close” to the given customer bid. We describe this procedure and the asymptotic properties of our estimator in section B.1.

¹²The 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 $F^D((S_1, Z_1), \dots, (S_{N_d}, Z_{N_d}) | \mathbf{S}^c)$ 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 $y^d(p|s_i, z_i)$, 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, $y^d(p|s_i, z_i)$, 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 $y^d(p|s_i, z_i)$.

econometrician can condition on the same set of covariates that bidders observe; we will discuss this concern in Appendix 4.3 below. We can, however, address this concern by resampling using bids from a single auction at a time. Moreover, in Appendix E, we conduct several tests for independence, and find considerable support for the conditional independence assumption.

4.2 Test Statistic

The main idea behind our test for distinguishing between the two types of learning, 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 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 (2011)). 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 the null hypothesis that there is no “learning about fundamentals,” the two estimates of marginal values should coincide except due to sampling error.¹³

Consider the test statistic:

$$T_i(q) = |\hat{v}_i^{BI}(q, s_i) - \hat{v}_i^{AI}(q, s_i)| \quad (6)$$

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.

In our application, we observe multiple bidders, for whom we need to test the null hypothesis

¹³Appendix B discusses the asymptotic distribution of the marginal valuation estimates.

jointly. We utilize two joint test statistics

$$SSQ_T = \sum_i \left(\frac{T_i}{\sigma_{T_i}} \right)^2 \tag{7}$$

$$FOS_T = \max_{i \in \mathcal{D}} \frac{T_i}{\sigma_{T_i}} \tag{8}$$

where the first is motivated by a χ^2 test and second is the maximum (first-order statistic).

We obtain the critical value for these test statistics using bootstrap. 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.¹⁴ Monte Carlo evidence on the performance of the testing approach is provided in Appendix F.

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. Fortunately, our testing strategy is based on looking at the 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.¹⁵

¹⁴To 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.

¹⁵The more subtle concern that changes in dealer bids may be caused by innovations to dealers' information sets that are not observed by the econometrician is discussed in Appendix C.

5 Test 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.¹⁶ Figure 2 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.¹⁷

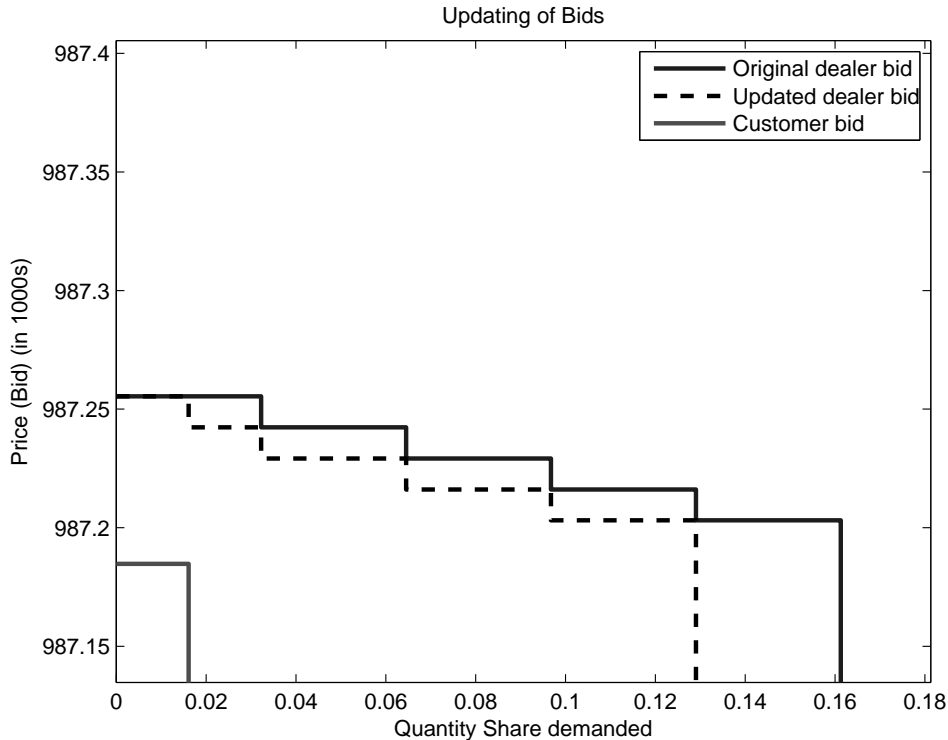


Figure 2: Updating of a dealer's bid

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

¹⁶We did not include instances where the dealer did not revise his bid. As described in section 2.2, we cannot be sure if an update did not arrive because the optimal bid did not change or because there was not enough time. 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.

¹⁷This “parallel shift” of the updated bid is not a general feature of the data, however. Some updated bids cross the original bids.

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 (6) and construct the corresponding critical values.

To illustrate the marginal value estimation procedure, Figure 3 depicts the marginal value estimation results for the dealer in Figure 2. We run our test on the three steps where the bidder’s quantity stayed the same. 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.




Figure 3: Estimated values - dealer 1

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 exercises in Appendix F, 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 F 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” threshold 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 4, we report results from hypothesis tests on individual bidders’ updating behavior. First, we report results based on the equality test statistic (equation (6)) 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 5, we report results from the tests based on the joint test statistics defined in Section 4.2.¹⁸ Once again, we estimate marginal values using 1, 2, and 4 neighboring auctions. Our

¹⁸We use a bandwidth of approximately 4 basis points. Table 10 in Appendix G reports the results for various

Monte Carlo exercises in Appendix F 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).¹⁹ Since the FOS may be overly demanding, we also report the 95th percentile of the studentized test statistic distribution. Based on the studentized joint hypothesis tests and using the treasury bill with 3-months maturity, Table 5 shows that we fail to reject the null hypothesis across all resampling specifications.

5.2 Results from 12-months T-bill auctions

Tables 4 and 5 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.

Table 4 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 5, 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.

other bandwidths.

¹⁹For 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 4: Individual Hypothesis Test Results

Bandwidth	Percent rejected					
	3-months			12-months		
	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 (B-24), 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.

Table 5: Joint Hypothesis Studentized Test Results

Bandwidth ^a	3-months			12-months		
	100			500		
	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 (B-24), 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.

6 The Value of Customer 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 quantify the benefit accruing to the dealers from being able to observe the customers' order flow, the value of information. In particular, we attempt to quantify the effect of the additional information on a dealer's interim (expected) utility. Let $U_d^{EP}(\sigma^I)$ denote the ex post utility of a dealer d , when using a bidding strategy σ^I where the superscript I denotes information. As before, $I = AI$ denotes the case when 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 approximate the value of information in terms of this notation as follows:

$$VI_d = \int_0^\infty U_d^{EP}(\sigma^{AI})dH(P^c, \sigma^{AI}) - \int_0^\infty U_d^{EP}(\sigma^{BI})dH(P^c, \sigma^{BI}) \quad (9)$$

where $H(P^c, \sigma^I)$ is the distribution of the market clearing price, P^c , given other bidders using equilibrium strategies σ and dealer d using the strategy σ^I . Recall that this distribution is defined in (1) and (2) for BI and AI, respectively.

Equation (9) defines the value of information as the difference between the expected utility when the dealer uses different strategies and has different beliefs about the distribution of the market clearing price. These beliefs are unambiguous for the case "AI" - the dealer observes a customer's bid, $y^c(p)$, and thus reacts to it by assuming that available supply at price p is $Q - y^c(p)$ and by integrating out over the remaining $N_c - 1$ customers' and $N_d - 1$ dealers' bids.

For the case "BI", we consider the following two polar cases we discussed in Section 3.1: the case where not observing a customer bid is (i) informative about competition in the sense that the dealer infers that one customer will not bid and integrates out over the remaining $N_c - 1$ customers' (and competing dealers') bids, (ii) non-informative about competition, and that the dealer integrates over all N_c customer and $N_d - 1$ dealer bids.

To calculate the proper value of information, we would like σ^{BI} to be the best response to the beliefs described in case 2. In case 1, σ^{BI} would still contain information: in particular, that one

customer submits a bid of zero. In our application, the difference in the distributions of market clearing price implied by these two alternative scenarios is very small. This is likely due to two reasons. First, the ex-ante probability of routing a (non-zero) customer bid by a dealer is only 38% in our sample. Second, a customer’s bid can be both “bad” and “good” news and this may thus mean that from an ex ante perspective, believing that an extra customer might participate does not have a significant impact on the expected market clearing price.

Figure 4 shows this in detail. The solid (blue) curve is the histogram of the difference between the mean of the estimated distribution of the market clearing price assuming a dealer integrates over all N_c customers (denoted as “Ex-ante” in the figure) and when he believes one customer bid is zero (denoted as “Bad Info”). An observation in this histogram is the difference in these means for a dealer who submitted a bid both before and after information arrived. These differences are clustered very close to zero. The two dashed lines plot these differences for the case when the dealer conditions on a particular observed (non-zero) customer bid and integrates over all N_c customer bids (red curve) or integrates over $N_c - 1$ customer bids (blue curve). It shows that while the actual knowledge of a customer’s bid has an effect on the beliefs about the market clearing price (dashed lines), assuming one less customer participates and averaging over all possible customer bids for one customer has very similar effect on those beliefs (solid line).

To compute the value of information given by equation (9), for each bidder we take the upper envelope of her estimated marginal values and we use the distribution of residual supplies they face to evaluate the expectation with respect to the market clearing price this bidder faced in the auction. For each market clearing price in the support, 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 $U_d^{EP}(B^{AI})$, and finally weight the ex-post utilities by the density. 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. The illustrative example




Figure 4: Distribution of the Differences in Means of the Distribution of the Market Clearing Price for Various Assumptions about Beliefs

discussed in Appendix D is a case where this assumption is satisfied in equilibrium. 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 $\Upsilon(\cdot)$ term is not empirically important (which Appendix A shows is the case in our application), what we can calculate are the changes in the “pseudo-value” of the dealer with respect to (presence vs. lack of) customer information. 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.²⁰

²⁰See e.g. Laffont and Vuong (1995), and Athey and Haile (2002).

Using our estimates we find that that the value of information, VI_d , 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).²¹ Since the average expected 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 rents from order flow in 12-months T-bill auctions amount to about C\$0.4 Million.

7 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 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

²¹The standard deviation of ex post payoff is slightly over 2.5 basis points.

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 Hortacısu, 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] Grossman, S., "On the Efficiency of Competitive Stock Markets Where Traders Have Diverse Information," *Journal of Finance*, Vol. 31, No. 2, pp. 573-585, 1976.
- [18] Guerre, E., Perrigne, I., and Vuong, Q., "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica*, Vol. 68, No. 3, 2000.
- [19] Haile, P., Hong, H., and Shum, M., "Nonparametric Tests for Common Values In First-Price Sealed-Bid Auctions," mimeo, 2003.
- [20] Hayek, F., "The Use of Knowledge in Society," *American Economic Review*, Vol. 35, No. 4, pp. 519-530, 1945.
- [21] Hill, J., and Shneyerov, A., "Are There Common Values in BC Timber Sales? A Tail-Index Nonparametric Test," mimeo, 2009.

- [22] Hoeffding, W., "A Class of Statistics with Asymptotically Normal Distribution," *The Annals of Mathematical Statistics*, Vol. 19, pp. 293-325, 1948.
- [23] Hortaçsu, A., "Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market," mimeo, 2002a.
- [24] Hortaçsu, A., "Bidding Behavior in Divisible Good Auctions: Theory and Empirical Evidence from Turkish Treasury Auctions," mimeo, 2002b.
- [25] Hortaçsu, A., and Sareen, S., "Order Flow and the Formation of Dealer Bids in Treasury Auctions," mimeo, 2006.
- [26] 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.
- [27] Kastl, J., "Discrete Bids and Empirical Inference in Divisible Good Auctions," *Review of Economics Studies*, Vol. 78, pp. 978-1014, 2011.
- [28] Kastl, J., "On the Properties of Equilibria in Private Value Divisible Good Auctions with Constrained Bidding", mimeo, 2008.
- [29] Koul, H. L., and Schick, A., "Testing for equality of two nonparametric regression curves," *Journal of Statistical Planning and Inference*, Vol. 65, pp. 293-314, 1997.
- [30] Kyle, A., "Informed Speculation with Imperfect Competition," *Review of Economic Studies*, Vol. 56, pp. 317-356, 1989.
- [31] Laffont, J.J., and Vuong, Q., "Structural Analysis of Auction Data," *American Economic Review*, Papers and Proceedings, Vol. 86, pp. 414-420, 1996.
- [32] Lehmann, E., *Elements of Large Sample Theory*, Springer, New York, 1999.
- [33] Lyons, Richard K., Tests of Microstructural Hypotheses in the Foreign Exchange Market, *Journal of Financial Economics*, 39, 321-351, 1995.

- [34] Madhavan Anand and Smidt, Seymour, "A Bayesian Model of Intraday Specialist Pricing," *Journal of Financial Economics*, 30, pp. 99-134, 1990.
- [35] Matthews, S., "Information Acquisition in Discriminatory Auctions," in *Bayesian Models in Economic Theory* ed. by M. Boyer and R. Kihlstrom, North Holland, 1984.
- [36] McAdams, D., "Partial Identification and Testable Restrictions in Multi-Unit Auctions," *Journal of Econometrics*, pp. 74-85, 2008.
- [37] Milgrom, P., and Weber, R., "A Theory of Auctions and competitive Bidding," *Econometrica*, Vol. 50, pp. 1089-1122, 1982.
- [38] 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.
- [39] 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.
- [40] 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.
- [41] Paarsch, H., "Deciding Between the Common and Private Value Paradigms in Empirical Models of Auctions," *Journal of Econometrics*, Vol. 51, pp.191-215, 1992.
- [42] Pinkse, J., and Tan, G., "The Affiliation Effect in First-Price Auctions," *Econometrica*, Vol. 73, pp. 263-277, 2005.
- [43] Romano, J.P., A. Shaikh, and M. Wolf, "Formalized Data Snooping Based on Generalized Error Rates," *Econometric Theory*, Vol. 24, pp. 1-44, 2008.
- [44] 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.
- [45] Stute, W., "Conditional U-Statistics," *The Annals of Probability*, Vol. 19, pp. 812-825, 1991.

- [46] Vickrey, W., "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, Vol. 16, pp. 8-37, 1961.
- [47] Wilson, R., "Auctions of Shares," *Quarterly Journal of Economics*, 1979.
- [48] 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.
- [49] 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 The Case of Interdependent Values

In this appendix, we derive the necessary conditions for bid optimality under interdependent values, and propose an empirical test for this specification. We then describe the results of the test as implemented in our data set.

A.1 Necessary Conditions for Bid Optimality 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 $\{(p_k, q_k), k = 1, \dots, K\}$, 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 5, 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 5. 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} .

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

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

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 5). 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




Figure 5: Illustration of necessary condition

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:

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

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

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

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

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

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

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

While taking the derivatives of the above terms with respect to q_k , note that $Q^c(\cdot)$ 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:

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

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

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

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

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

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

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

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

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

we get the following first-order condition:

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

Inspecting the first-order condition (A-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 (A-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 5, 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, $\Upsilon(\tilde{\theta}_i, \cdot)$ captures the reweighting over these different states.

Note also that $\frac{\partial \Pr(b_k = P^c | \tilde{\theta}_i)}{\partial q_k} \geq 0$ and $\frac{\partial \Pr(b_{k+1} = P^c | \tilde{\theta}_i)}{\partial q_k} \leq 0$, as can be seen from Figure 5 (increasing quantity makes it more likely that residual supply intersects horizontally at b_k , and less likely at b_{k+1} .) Also, define:

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

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

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:

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

$\frac{\partial \Pr(b_k = P^c | \tilde{\theta}_i)}{\partial q_k} \geq 0$, $\frac{\partial \Pr(b_{k+1} = P^c | \tilde{\theta}_i)}{\partial q_k} \leq 0$ 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 (A-16) provides an upper bound on the "pseudo-value" term, $E_{\theta_{-i} | \tilde{\theta}_i, P^c \in (b_{k+1}, b_k), q_k} [v(q_k, \tilde{\theta}_i, \theta_{-i})]$. Moreover, we can use this expression to get an upper bound for $\Upsilon(\tilde{\theta}_i, \cdot)$, as

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

The first term in the product in equation (A-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 $V(\cdot)$ 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:

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

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 A.3.

A.2 Testable restriction under 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.²² 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 sets $\tilde{\theta}_i = (s, z_i)$ and $\tilde{\theta}_j = (s, z_j)$:

$$E_{\theta_{-i}|\tilde{\theta}_i, P^c \in (b_{k+1}, b_k), q_k} [v(q_k, \theta_i, \theta_{-i})] > E_{\theta_{-j}|\tilde{\theta}_j, P^c \in (b_{k+1}, b_k), q_k} [v(q_k, \theta_j, \theta_{-j})] \quad (\text{A-22})$$

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 A.3 will be to place bounds on this quantity based on observables, and assess

²²This 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.

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 4.1, 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.

A.3 Results of the test in our data

Even though our 3-month and 12-month samples failed to reject the test for the null hypothesis of private values, 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 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 A-22 in Appendix A.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 A.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 A, 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.²³ To see the importance of the $\Upsilon(\cdot)$ term, we plot the distribution of Δv ’s corrected using estimated upper-bounds for Υ s with $C \in \{0.1, 1, 10\}$ in Figure 6.²⁴ The figure shows that Υ makes a very small difference to the distribution of Δv s.

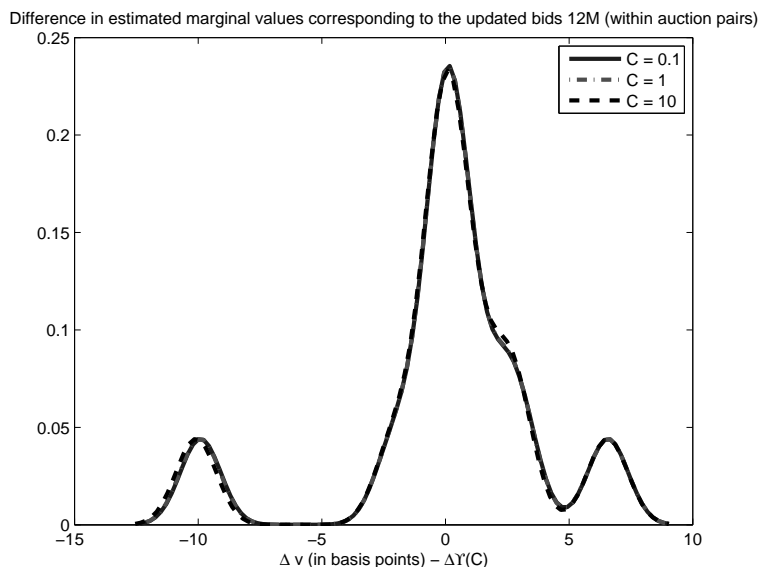


Figure 6: Distribution of Δv ’s in 12M T-bills for within auction bid-pairs.

Figure 6 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 Δv s are greater than zero for the multiple bid

²³Since 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.

²⁴As we note in the Appendix A, 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.

pairs under consideration, we calculate the mean/median of studentized Δv s 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 6. Based on the results of this test, we do not find evidence for interdependent values in our data. In particular, the mean and the 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 6 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 6: 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 $\Delta v = \hat{v}_i(q, s|z_i) - \hat{v}_j(q, s|z_j)$ being larger than 0.

B Properties of marginal valuation estimates

Our main empirical 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 (3) that these estimates are a non-linear function of the distribution of the market clearing price, which is estimated by the resampling method described above. Let us rewrite (3) as

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

where $H(X)$ (resp. $G(X)$) 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 $y_1, \dots, y_{N_c+N_d-1}$ and i 's own bid $y_i(X|\tilde{\theta}_i)$ as follows:

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

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 $(N_c + N_d)T$ datapoints:

$$\xi(\hat{F}; X, h_T) = \frac{1}{(N_c T)^{N_c}} \frac{1}{(N_d T)^{(N_d-1)}} \sum_{\alpha_1=(1,1)}^{(T, N_c)} \dots \sum_{\alpha_{N_c}=(1,1)}^{(T, N_c)} \sum_{\alpha_{N_c+1}=(1, N_c+1)}^{(T, N_c+N_d)} \dots \sum_{\alpha_{N_c+N_d-1}=(1, N_c+1)}^{(T, N_c+N_d)} \Phi(y_{\alpha_1}, \dots, y_{\alpha_{N_c}}, y_{\alpha_{N_c+1}}, \dots, y_{\alpha_{N_c+N_d-1}}, X) W(\alpha_1, \dots, \alpha_{N_c}, \alpha_{N_c+1}, \dots, \alpha_{N_c+N_d-1}; h_T) \quad (\text{B-23})$$

where $\alpha_i \in \{(1, 1), (1, 2), \dots, (1, N_c + N_d), \dots, (T, N_c + N_d)\}$ is the index of the bid in the subsample and \hat{F} is the empirical distribution of bids.²⁵ 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 $N_d T$ dealer bids, and there are $(N_d T)^{(N_d-1)}$ subsamples of size $(N_d - 1)$. Finally, $W(\cdot)$ denotes the kernel weights defined by

$$W(\alpha_1, \dots, \alpha_{N_c}, \alpha_{N_c+1}, \dots, \alpha_{N_c+N_d-1}; h_T) = \frac{\prod_{j=\alpha_{N_c+1}}^{\alpha_{N_c+N_c}} \left(K \left(\frac{\|z_j - y_{j-N_c}\|}{h_T} \right) \right) \prod_{j=\alpha_{2N_c+1}}^{\alpha_{N_c+N_d-1}} \left(K \left(\frac{\|z_j - 0\|}{h_T} \right) \right)}{\sum_{\alpha'_1=(1,1)}^{(T, N_c)} \dots \sum_{\alpha'_{N_c+N_d-1}=(1, N_c+1)}^{(T, N_c+N_d)} \prod_{j=\alpha'_{N_c+1}}^{\alpha'_{N_c+N_c}} \left(K \left(\frac{\|z_j - y_{j-N_c}\|}{h_T} \right) \right) \prod_{j=\alpha'_{2N_c+1}}^{\alpha'_{N_c+N_d-1}} \left(K \left(\frac{\|z_j - 0\|}{h_T} \right) \right)} \quad (\text{B-24})$$

where $K(\cdot)$ is a bounded kernel with compact support²⁶ and h_T is the bandwidth satisfying $h_T \rightarrow$

²⁵Note 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.

²⁶Examples of possible kernel functions are Epanechnikov, triangular, rectangular etc. In our application, we use

0, $Th_T^5 \rightarrow 0$ and $Th_T \rightarrow \infty$ as $T \rightarrow \infty$. 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 $\xi(\hat{F}; X)$, for which M subsamples are randomly drawn rather than all $(N_c T)^{N_c} (N_d T)^{(N_d - 1)}$. We choose $M = 5,000$ in our application to make sure the simulation error is not an important factor.

B.1 Asymptotic properties of ξ

For ease of exposition, 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 2 *Suppose $W = 1$, data are iid across the T auctions and bidders, all bidders are symmetric and N is fixed. Then as $T \rightarrow \infty$, $\frac{M}{T} \rightarrow \infty$:*

$$\sqrt{T} \left(\hat{H}^R(X) - H(X) \right) \rightarrow N \left(0, \frac{(N-1)^2}{N} \zeta \right) \quad (\text{B-25})$$

where $\zeta = E_{\theta_{-i}} \left[(\Phi(y_1, \dots, y_{N-1}; X))^2 \right] - \left(\binom{NT}{N-1}^{-1} \sum_{(1,1) \leq \alpha_1 < \alpha_2 < \dots < \alpha_{N-1} \leq (T,N)} \Phi(y_{\alpha_1}, \dots, y_{\alpha_{N-1}}; X) \right)^2$ and where that last summation is taken over all combinations of $N-1$ indices $\alpha_i \in \{(1, 1), (1, 2), \dots, (1, N-1), \dots, (T, N)\}$ such that $\alpha_1 < \alpha_2 < \dots < \alpha_{N-1}$.²⁷

the uniform kernel: $K(u) = \frac{1}{2} I_{(|u| \leq 1)}$.

²⁷The 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.

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

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

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 $E[\Phi(\cdot)^2]$, 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:

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

where the averaging is over every permutation of the NT observations.²⁸ 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 $\frac{M}{T} \rightarrow \infty$, $\hat{H}^R \rightarrow \xi$. ■

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 2 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

²⁸Note that the sample size is NT (bidders per auction \times auctions) and we are constructing subsamples of size $N - 1$, hence the denominator $(NT)^{N-1}$.

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 (B-23) which uses the kernel weights $W(\cdot)$ as defined in (B-24). 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)²⁹, whose Theorem 1 we adapted to our application and state as the following:

Proposition 3 *Assume that (i) $h_T \rightarrow 0$, $Th_T \rightarrow \infty$, $Th_T^5 \rightarrow 0$, (ii) $\frac{M}{T} \rightarrow \infty$ and (iii) K is bounded and has compact support, then*

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

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

and where when $y_j = y_l$,

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

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

In our empirical application, we 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 $\Phi(\cdot)$ in the V-statistic (B-23) are bounded.

C Unobserved heterogeneity within an auction

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

²⁹As 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.

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 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.

C.1 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.

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:*

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

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

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 $\Pr(\text{Customer } x_c \text{ wins}) = \Pr(x_d < b_c(x_c))$:

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

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 $\Pr(\text{win}|x_c) = \frac{x_c}{2}$ in the current auction (which is also the slope of the information rent) and it is $\Pr(\text{win}|x_c) = x_c$ 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 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 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.³⁰ 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,

³⁰Mudholkar 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.

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 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: $H_0 : F(b|b_{1t}) = F\left(b|b_{\frac{N_{dt}^*}{2}+1}\right)$, where the first sample consists of $b_{2t}, \dots, b_{\frac{N_{dt}^*}{2}}$ and the second sample of $b_{\frac{N_{dt}^*}{2}+2}, \dots, b_{N_{dt}^*}$, where N_{dt}^* is the number of non-zero dealer bids in auction t .³¹ 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: $H_0 : F(b|b_{1t}) = F(b|b_{1(t+1)})$. Here, the test rejects in over 100 pairs of consecutive auctions. However, unobserved heterogeneity 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 be 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.³² 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.

³¹By 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.

³²The 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 $N/2 - 1$.

F 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.

F.1 First Price Auction with Informed Bidders

F.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:

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

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.³³

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, $Med(\hat{v}_{informed}) - Med(\hat{v}_{uninformed})$, 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 bootstrap 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, $SQ = \sum_i (\hat{v}_{i,informed} - \hat{v}_{i,uninformed})^2$, and the “first order statistic” test, $FOS = \max_i (\hat{v}_{i,informed} - \hat{v}_{i,uninformed})$, 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

³³As 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 95th percentile, and to the quantile corresponding to the Bonferroni's method, $100 - \frac{5}{\# \text{ of hypotheses}}$.

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

F.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 $v(s_i, s_{-i}) = \frac{s_i}{2} + \frac{\sum_{j \neq i} s_j}{2(n-1)}$ where $s_i \sim U[0, 1]$. The unique symmetric equilibrium in strictly increasing differentiable strategies involved bidding according to $b^U(s_i) = \frac{7}{12}s_i$. 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:

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

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.

F.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 $u_i(s_i) = v$ where $v \sim \text{Pareto}(\alpha) : g(v) = \alpha v^{-(\alpha+1)}$ 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

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

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, \dots, 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

N	Rejection frequency								
	IPV			IIV			CV		
	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

N	Rejection frequency					
	IPV		IIV		CV	
	Ptwise ₉₅ ^a	Ptwise _{Bonf} ^b	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 95th percentile of the corresponding asymptotic distribution.

^b Comparing each hypothesis with $(1 - \frac{0.05}{N})^{th}$ 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

N	Rejection frequency					
	IPV		IIV		CV	
	SSQ _{Stud} ^a	FOS _{Stud} ^b	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

F.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.³⁴ Both tests' power against the alternative

³⁴One 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).

F.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.

F.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:

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

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.

F.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:

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

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 $b^*(s_1, s_2)$, bidder 1 will still prefer to win if $s_1 \geq \frac{25}{48}s_2$.

G More Detailed Joint Hypothesis Test Results

Table 10: Joint Hypothesis Studentized Test Results

	Bandwidth	3-months			12-months		
		Auctions for resampling			Auctions for resampling		
		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.