Valuing Dealers' Informational Advantage: A Study of Canadian Treasury Auctions*

Ali Hortaçsu[†] Jakub Kastl[‡]

May 12, 2012

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. We find that the "information about competition" contained in customer bids accounts for $13 - 27\%$ of dealers' expected profits.

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

JEL Classification: D44

*We would like to thank Steve Berry, Liran Einav, Ken Hendricks, Han Hong, Jon Levin, Rob Porter, Peter Reiss, Azeem Shaikh, Quang Vuong and four anonymous referees 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, CERGE-EI, Chicago, Mannheim, NYU Stern, Queens, 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 & Lonstaff 1990). Empirical studies of dealers' usage of customer order-flow information include analyses of specialist/dealer behavior in the NYSE (Madhavan & Smidt 1990), in the foreign exchange market (Lyons 1995, Evans & Lyons 2002), and in options markets (Easley, O'Hara & 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." Most "customers" in this setting 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. This institutional setup is not unique to Canada: out of 39 countries surveyed by Arnone & 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) auctions we analyze, there are two sources of economic benefit to dealers from observing customer bids. Customer bids may be 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, Grossman 1976). Customer bids may also provide information regarding the *competition* a

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

dealer will face in the auction. As such, customer bids may compel a dealer to modify her bid, even in a 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 submitted by rival bidder j. Bidder j's "competitive" information allows bidder i to improve her estimate of the location and shape of residual supply (i.e. total quantity being auctioned minus the bids of other bidders). 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, 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 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 under this null hypothesis. To do this, we build on our earlier work (Hortaçsu & McAdams 2010, Kastl 2011*a*) to characterize necessary conditions for equilibrium bidding, and estimate the marginal valuations that rationalize a dealer's bid under equilibrium beliefs about her competitors' bids. Under our null hypothesis, 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.

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 bids contributes significantly to dealers' overall profits from participating in Canadian Treasury auctions.

Our analysis is very closely related to the literature on empirical tests of private vs. interdependent values in auctions. Indeed, our analysis can be regarded as testing whether there are value interdependencies between dealers and customers. To the extent one is willing to consider customers, which are often large financial institutions themselves, as representative of other bidders, our results can be seen as a test of the private vs. interdependent value hypothesis more broadly. To our knowledge, this is the first attempt to test the null hypothesis of private values in a multi-unit divisible good auction setting. 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.[2] Our testing strategy is most similar to the contribution in the single-unit auction context by Haile, Hong & 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 evidence suggesting that dealers modify their bids in a way that appears to reflect the information

[2] As pointed out by Ausubel & Cramton (2002), neither the revenue equivalence theorem of Vickrey (1961) nor Milgrom & 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 (Février, Préget & Visser 2004, Armantier & Sbaï 2006, Hortaçsu & McAdams 2010, Kastl 2011*a*, Kang & Puller 2008, Chapman, McAdams & Paarsch 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.

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. For example, Desjardin Securities, which is the securities division of one of the largest Canadian banks with overall assets of over C$173 billion, is a primary dealer in the bond market, but only a customer in the treasury bill market. Similarly, Casgrain & Company or JPMorgan are not registered as primary dealers and yet are very important players in the Canadian government securities markets. The major distinction between customers and dealers 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. Similarly to primary dealers, customers are required to report their net positions in the government securities before each auction.

2.1 Data

Our data consist of all submitted bids in 116 auctions of 3-month and 12-month 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
# of Auctions: 116				
# 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
# of Bids: $3,511$				
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
# of Non-Competitive Bids: 566				
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. 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

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

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

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 explanation 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 1. Note that these graphs include the negative domain – these are customer bids that arrive after the deadline passes and consequently cannot participate in the auction.

To check whether 5 minutes is enough for a dealer to change his bid, we calculated the average (median) number of minutes that a dealer takes to "follow" a customer bid with his own bid. In situations where there is updating it takes, on average, 5.49 (med = 4.93) minutes for the updated dealer bid to be entered – 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.

Indeed, a closer look at when dealers update their bids as a function of when the customer bid arrives reveals some stark patterns. As reported in Table 4, as customer bids arrive closer and closer to the deadline, the frequency of times that the dealer submits a bid update following the customer bid declines. Indeed, while dealers update their bids 79% (81%) of the time when the customer bid comes earlier than 15 minutes to the deadline in the 3(12)-month auctions, they only update their bids 25% (37%) of the time when the customer bid comes with less than 5 minutes to the deadline.

Table 4: Timing of customer bids and frequency of dealer bid updates

Customer Bid Submission Time	Percent of times dealer updates (3M sample)	Percent of times dealer updates (12M sample)
> 15 min before deadline	78.87%	81.01%
$(10, 15]$ min before deadline	52.58%	67.65%
$(5, 10]$ min before deadline	32.82%	47.03%
< 5 min before deadline	25.17%	37.93%

Note: All entries are statistically significantly different at the 1% level except for the > 15 min and $(10, 15]$ min difference in 12M sample, which is significant at the 7% level.

Such "last minute" bidding behavior by customers can be rationalized as a strategic response by customers who do not want dealers to utilize the information in their bids. In previous work studying the same market, Hortaçsu & 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

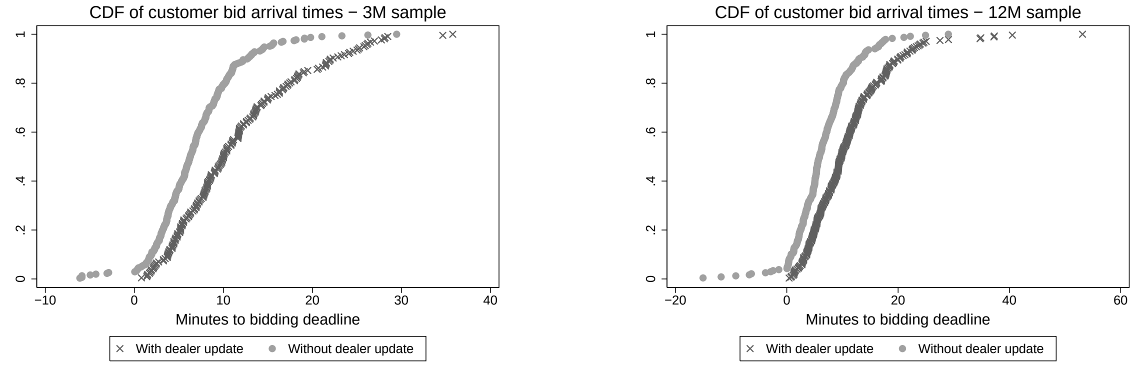

Figure 1: Comparison of customer bid submission times with and without dealer updates

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.

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.[3] 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*a*). We further adapt this model into the context of our application where some bidders (the "dealers") observe the bids of others ("customers") and may submit their own bids both before

[3]Hortaçsu & 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.

and after the customers' bids arrive.

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 perfectly divisible good of (random) Q units. We assume that the number of potential bidders of each type participating in an auction, N_c, N_d, is commonly known. Indeed, in our empirical application all participants have to register with the Bank of Canada before the auction as dealers and customers and the list is publicly available.

Before the bidding commences, both dealers and customers observe private (possibly multidimensional) signals, $S_1^c, ..., S_{N_c}^c, S_1^d, ..., S_{N_d}^d$. The bidding then proceeds in three stages. In stage 1, dealers may submit "early" bids, $y^{1d}\left(p|S^d\right)$, which specify for each price p, how big a share of the securities offered in the auction (type S_i^d of) dealer i demands. In stage 2, each customer who wants to bid gets matched to a dealer and this dealer submits her customer's bid, $y^c\left(p|S^c\right)$. In stage 3, dealers may submit a "late" bid to supersede the "early" bid from stage 1, $y^{3d}\left(p|S^d, Z^d\right)$, where Z^d contains all additional information observed by this dealer after stage 1.

To rationalize "early" bidding in stage 1, we assume that one dimension of dealer i's signal S_i^d is a random variable $\Psi_i \in [0,1]$ corresponding to the mean of another Bernoulli random variable, Ω_i, determining whether that dealer's "late" bid in stage 3 will make it in time to be accepted. In stage 3, Z_i^d includes the actual customer's bid or the fact that one did not arrive and the realization $\omega_i \in \{0,1\}$ of Ω_i, where $\omega_i = 1$ means the "late" bid will make it in time. We assume that Z_i^d is not observed by the dealer's competitors.

The above setup can deliver most of the "bid timing" patterns in the data we discussed in Section 2.2: if $\Psi = 1$ (corresponding to an early arrival of customer's bid), a dealer skips stage 1, and places only a "late" bid in stage 3. However, if Ψ is sufficiently low (i.e., as the auction deadline comes closer), the dealer submits an "early" bid in stage 1 to make sure she participates in the auction. If the realization of Ω_i in stage 3 is 1, the dealer can update her bid based on the information in customer bids, but if $\omega_i = 0$, the dealer does not have time to react to the customer bid.[4]

[4]To rationalize the submission of "late" dealer bids that nevertheless arrive after the bidding deadline, we could add further uncertainty over the realization of Ω_i, but since such bids are rare, we do not add this detail.

We will impose the following additional assumptions:

Assumption 1 *Customers' and dealers' private signals, $S_1^c, ..., S_{N_c}^c, S_1^d, ..., S_{N_d}^d$, are independent and drawn from a common support $[0,1]^M$ according to an atomless distribution functions $F^d\left(S^d\right)$ and $F^c\left(S^c\right)$, with strictly positive densities f^d and f^c.*

Strictly speaking, independence is not necessary for our characterization of equilibrium behavior in this auction, but we impose it in our empirical application. We will provide a test for independence in Appendix B.

Assumption 2 *The supply Q is a random variable and its distribution is common knowledge among the bidders. The (per-bidder) supply, $\tilde{Q} \equiv \frac{Q}{N_c+N_d}$, is independent of S_i^c, S_j^d $\forall i,j$ and is distributed according to $G\left(\underline{\tilde{Q}}, \overline{\tilde{Q}}\right)$ with strictly positive density g.*

The randomness in supply in our setting is driven by the substantial non-competitive bid of the Bank of Canada and by other non-competitive bids, which are not revealed to the bidders. The joint distribution of dealers' private information in the 3rd stage conditional on (the vector of) customers' information and customers' equilibrium strategies is $F^{3d}\left(\left(S_1^d, Z_1\right), ..., \left(S_{N_d}^d, Z_{N_d}\right) |\mathbf{S}^c, \{y_i^c\left(p|S_i^c\right)\}_{i=1}^{N^c}\right)$ where $y_i^c\left(p|S_i^c\right)$ 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 $v_i\left(q, S_i\right)$. 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 $v^g\left(\cdot,\cdot,\cdot\right)$ for $g \in \{c,d\}$:

Assumption 3 *$v^g\left(q, S_i^g\right)$ is non-negative, measurable, bounded, strictly increasing in (each component of) S_i^g $\forall q$ and weakly decreasing in q $\forall s_i^g$, for $g \in \{c,d\}$.*

Note that this assumption implies that learning other bidders' signals does not affect one's own valuation – thus "learning about fundamentals" does not play a role. For the more general case where dealers' marginal valuation functions are allowed to depend on other bidders' signals, and thus "learning about fundamentals" does matter, we will assume that the expected utility of dealers is non-decreasing in customers' signals. This assumption is satisfied, for example, when signals are affiliated and v^d is non-decreasing in each S^c_{-i}. Formally, $E_{S_{-i}\backslash S^c_j}\left[v^d\left(q, S^d_i, S_{-i}\right)|S^d_i = s^d_i, S^c_j\right]$ is non-decreasing in (each component of) S^c_j, $\forall\left(j, q, s^d_i\right)$.

To ease notation, let θ^k_i denote private information of bidder i in stage k, i.e., for a customer $\theta^2_i \equiv S^c_i$, and for a dealer $\theta^1_i \equiv S^d_i$ and $\theta^3_i \equiv \left(S^d_i, Z^d_i\right)$. Bidders' pure strategies are mappings from private information in each stage to bid functions $\sigma_i : \Theta_i \rightarrow \mathcal{Y}$, where the set $\mathcal{Y}$ includes all admissible bid functions. 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. $y^c_i\left(p|S^c_i\right) = y^c\left(p|S^c_i\right), i \in \mathcal{C}$. Dealers also bid in an ex-ante symmetric way in stage 1, with $y^d_i\left(p|S^d_i\right) = y^d\left(p|S^d_i\right)$, $i \in \mathcal{D}$. Dealers' bid functions in stage 3 are also symmetric, but up to their private signal *and* customer information, i.e. $y^{3d}_i\left(p|S^d_i, Z^d_i\right) = y^{3d}\left(p|S_i, Z_i\right)$, $i \in \mathcal{D}$.

Since in most divisible good auctions in practice, including the Canadian treasury bill auctions, the bidders' choice of bidding strategies is restricted to non-increasing step functions with an upper bound on the number of steps, $\overline{K}$, we impose the following assumption:

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

$$A = \left\{ \begin{array}{c} \left(\vec{b}, \vec{q}, K\right) : \dim\left(\vec{b}\right) = \dim\left(\vec{q}\right) = K \in \left\{1, ..., \overline{K}\right\}, \\ b_k \in B = [0, \infty)\,,\ q_k \in Q = [0, 1] \wedge \forall k < K :\ b_k > b_{k+1}, q_k > q_{k+1} \end{array} \right\}$$

where a bid of 0 denotes non-participation and a bid of ∞ denotes a non-competitive bid. The range of strategies, the set $\mathcal{Y}$, thus includes all non-decreasing step functions with at most $\overline{K}$ steps, $y : \mathbb{R}_+ \rightarrow [0, 1]$, where $y_i\left(p\right) = \sum_{k=1}^{K} q_{ik} I\left(p \in \left(b_{ik+1}, b_{ik}\right]\right)$ where I is an indicator function. When bidders use step functions as their bids, rationing occurs except in very rare cases; thus we will assume, consistently with the application, pro-rata on-the-margin rationing, which proportionally adjusts 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.

3.1 Characterization of equilibrium dealer bids

The key source of uncertainty faced by the bidders in the auction is the market clearing price, P^c, which maps the state of the world, $\left(Q, \mathbf{s^c}, \mathbf{s^d}, \mathbf{z}\right)$, or simply $\left(Q, \overrightarrow{\theta}\right)$, into prices through equilibrium strategies. This random variable is thus summarized by a function $P^c\left(Q, \mathbf{s^c}, \mathbf{s^d}, \mathbf{z}, \mathbf{y^c}\left(\cdot|s\right), \mathbf{y^{1d}}\left(\cdot|s\right), \mathbf{y^{3d}}\left(\cdot|s,z\right)\right)$ (or simply $P^c\left(Q,\theta\right)$ which we will sometimes abbreviate as P^c).

Let us now define the probability distribution of the market clearing price from the perspective of a dealer in stage 1, who is preparing to make the "early" bid $y^{1d}\left(p|s_i\right)$. This bidder does not observe any customer information at this stage, but can rationally anticipate that the other dealers will observe customer information, and may submit updated stage 3 bids. Thus, the probability distribution of the market clearing price from the perspective of such a dealer will be:

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

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

In stage 3, we have two possibilities. Either the dealer has observed a customer bid, or not. The distribution of P^c from the perspective of dealer i, who observes customer m's bid function, i.e. $y^c\left(p|s_m\right) \in z_i$ is:

$$\Pr\left(p \geq P^c|s_i, z_i\right) =$$
$$E_{\left\{Q, S_{j\in\mathcal{C}\setminus m}, S_{k\in\mathcal{D}\setminus i}, Z_{k\in\mathcal{D}\setminus i}|z_i\right\}} I\left(Q - \sum_{j\in\mathcal{C}\setminus m} y^c\left(p|S_j\right) - \sum_{k\in\mathcal{D}\setminus i} y^d\left(p|S_k, Z_k\right) \geq y^{3d}\left(p|s_i, z_i\right) + y^c\left(p|s_m\right)\right) \tag{2}$$

Note that the main difference in equation (2) compared to equation (1) is that the dealer conditions on customer m's bid, instead of taking an expectation over that customer's private information. 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 a customer's bid.

For our empirical exercise, we will only focus on dealers who observe a customer bid. The case of dealers who get to stage 3 without observing a customer bid is slightly more complicated as we would have to take a stand on what a dealer believes that a lack of customer bid signifies. It could mean that one customer has drawn a bad signal (and has not participated in the auction) for sure, but it could also mean that all customers just happened to be matched to other dealers. Modelling a dealer's beliefs in this case would depend crucially on how customers are matched to dealers. For example, if a customer is in an exclusive relationship with its dealer, her absence would signify a low signal for that customer. However, if customers randomize across dealers, then the absence of a particular customer bid will not cause the bidder to infer that that customer's signal is low. Indeed, in our data, we find that only a minority of customers are in exclusive relationships with their dealers.

Before we state the main proposition, which links the observed bids to the unobserved marginal values, we define "learning about fundamentals" formally for our case, where dealers observe customers' bids.

Definition 1 *Dealers do not learn about fundamentals from customers' bids if* $\forall q_k, s_i^d, \forall s_j^c : y^c\left(q, s_j^c\right) \in z_i$*:*

$$E_{S_{-i} \setminus S_j^c}\left[v^d\left(q_k, S_i^d, S_{-i}\right) | S_i^d = s_i^d, z_i\right] = E_{S_{-i}}\left[v^d\left(q_k, S_i^d, S_{-i}\right) | S_i^d = s_i^d\right]. \tag{3}$$

Note that private values (Assumption 3) rules learning about fundamentals out. In the next section we will devise a formal hypothesis test of equation (3).

Given the probabilities defined in (1) and (2), a necessary condition for bidding in the case where dealers do not "learn about fundamentals" from customers' bids is given by the following:

Proposition 1 *(Kastl 2011*b*) Suppose 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* $\hat{K}(\theta_i)$ *steps, every step* k *in the equilibrium bid function* $y(\cdot|\theta_i)$ *has to satisfy:*

$$v(q_k, S_i) = b_k + \frac{\Pr(b_{k+1} \geq P^c|\theta_i)}{\Pr(b_k > P^c > b_{k+1}|\theta_i)}(b_k - b_{k+1}) \tag{4}$$

$\forall k \leq \hat{K}(\theta_i)$ *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.

4 Test Specification

The main idea behind our test for whether "learning about fundamentals" (as stated in definition 1) 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 (4) above. It is important to note that given our setup we are only able to test for interdependency of values (or learning about fundamentals) between dealers and customers. We cannot rule out the possibility that there might be value-related information contained in other dealers' signals.

4.1 Estimating Marginal Valuations

To implement this test, we first have to estimate the rationalizing marginal valuations, for which we extend Hortaçsu & McAdams (2010) and Kastl (2011*a*). The asymptotic behavior of our estimator is described in detail in Appendix A. The "resampling" method that we employ 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 (4).

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 $\mathcal{C}$). 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.[5] 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

[5]Ideally, we would draw from dealers' bids that have been submitted after observing *exactly* the same customer bid. However, our data has the practical limitation that customer bids are typically unique within or across auctions. Thus, the "conditional" draws often consist of repeatedly drawing the same customer and "informed" dealer pair. 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 A.1.

conditionally independent within class; i.e. conditional on auction-level covariates observed by all bidders, their private signals are independent.[6] Of course, an important concern is whether the econometrician can condition on the same set of covariates that bidders observe; we will discuss this concern in section 4.3 below. We can, however, address this concern by resampling using bids from a single auction at a time.

It is worth stressing that the above-described resampling method rests heavily on the assumption of ex-ante within group symmetry of dealers and customers and independence of private information. In Appendix B, 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*a*). 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: $v(q_k, s_i)$ and $v(q_k, s_i, z_i)$ in the situation where a bid has been submitted for the same quantity q_k. 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. Appendix A analyzes the asymptotic behavior of the

[6]The resampling procedure can, in principle, be modified for a more general affiliated private values setting, where we specify the vector of customers' signals, $\mathbf{S}^c$ to have a joint distribution $F^C(\mathbf{S}^c)$, and let $F^D((S_1, Z_1), ..., (S_{N_d}, Z_{N_d})|\mathbf{S}^c)$ 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 $y^d(p|s_i, z_i)$.

marginal valuation estimates formally.

Consider the test statistic:

$$T_i(q) = \left| \hat{v}^{1d}(q, s_i) - \hat{v}^{3d}(q, s_i, z_i) \right|. \tag{5}$$

In our application, we observe multiple bidders and to take into account potential correlation in individual test statistics we proceed with three joint hypothesis tests:

$$SSQ_T = \sum_i \left(\frac{T_i}{\sigma_{T_i}} \right)^2 \tag{6}$$

$$FOS_T = \max_{i \in \mathcal{D}} \frac{T_i}{\sigma_{T_i}} \tag{7}$$

$$95^{th} PERC_T = \left\{ \max_{X \in \cup T_i} X : \Pr[T_i \leq X] \leq 0.95 \right\} \tag{8}$$

where the first is motivated by a χ^2 test, the second is the maximum (first-order statistic) and the third is based on the 95^{th} percentile of the individual hypothesis test statistics.

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. To construct a bootstrap sample of bid functions, we 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.

4.3 Unobservable heterogeneity

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.

4.3.1 Arrival of New Common Knowledge Information During an Auction

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

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. 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.[7]

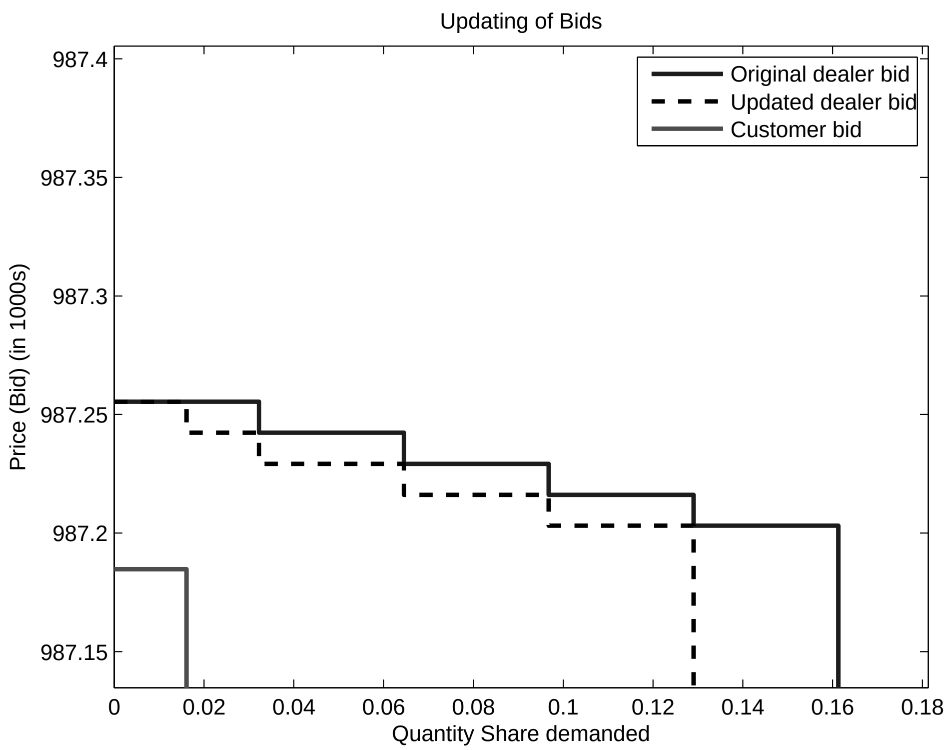

Figure 2: Updating of a dealer's bid

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

[7]This "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 (5) 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.

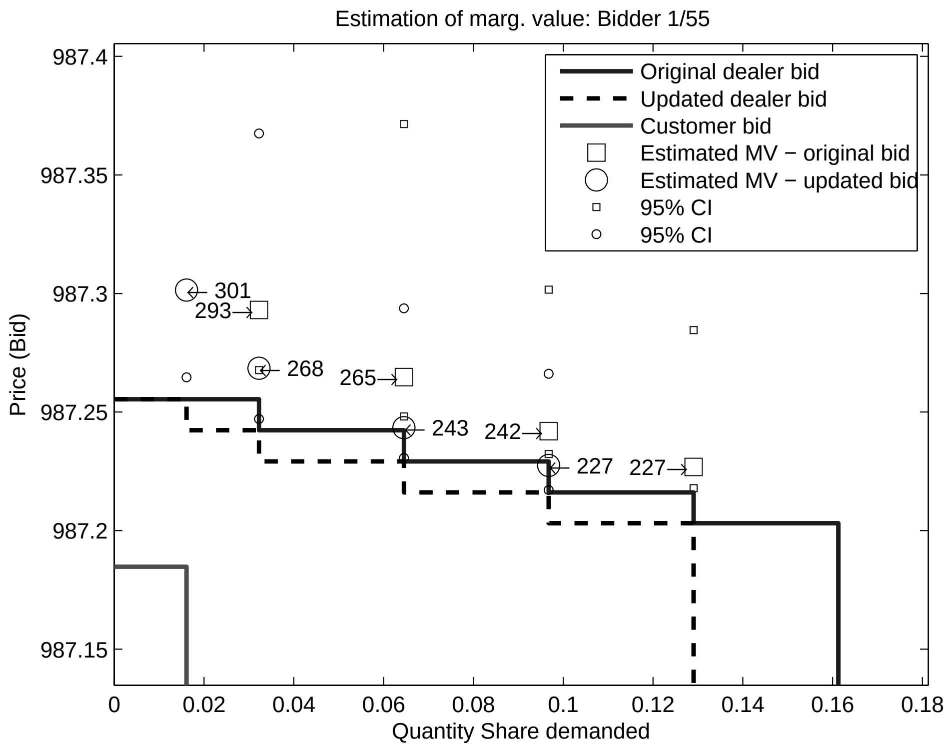

Figure 3: Estimated values - dealer 1

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 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.

We will display results across an array of test statistics. It is not our goal to derive a "hard" treshold which will allow us to reject the null hypothesis, instead we will use the conventional level $\alpha = 0.05$ to obtain our critical values and we will focus on contrasting patterns that emerge after applying our testing procedure to data from 3-month and 12-month T-bill auctions respectively. In Table 5, we report results from hypothesis tests on individual bidders' updating behavior. First, we report results based on the equality test statistic (equation (5)) 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. Some overrejection should be expected due to potential correlations in test statistics. 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 6, we report results from the tests based on the joint test statistics defined in Section 4.2, which address the issue of potential correlation between individual hypothesis tests. For these tests, we use a bandwidth of approximately 4 basis points.[8] Once again, we estimate marginal values

[8]Table 7 available in the online appendix reports the results for various other bandwidths.

using 1, 2, and 4 neighboring auctions, and 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).[9] Since the FOS may be overly demanding, we also report the 95^{th} percentile of the studentized test statistic distribution. Based on the studentized joint hypothesis tests and using the treasury bill with 3-months maturity, Table 6 shows that we fail to reject the null hypothesis across all resampling specifications.

5.2 Results from 12-months T-bill auctions

Tables 5 and 6 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 5 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 6, 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.

6 The Value of Customer Information

Given that our tests failed to reject the null hypothesis of there being "no learning about fundamentals" in our sample, in what follows, we will use our estimates of marginal values to quantify

[9] 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 5: Individual Hypothesis Test Results

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

Notes: The entries are percent of individual hypotheses that are rejected based on the equality test. We report results using different numbers of consecutive auctions in our resampling. The bandwidth parameter determines the width of the kernel in equation (A-2), 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 6: Joint Hypothesis Studentized Test Results

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

[a] The bandwidth parameter determines the width of the kernel in equation (A-2), 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.

the benefit accruing to the dealers from being able to observe the customer bids. In particular, we attempt to calculate the interim (expected) profit gain that a dealer makes when updating her bid in response to a customer bid.

Let $\Pi^{3d}(s_i, z_i)$ denote the expected profit of a dealer d, when using the bidding strategy $y^{3d}(p, s_i, z_i)$, i.e., after incorporating the information from customers' orders. Similarly, let $\Pi^{1d}(s_i,)$ denote the expected profit corresponding to the bidding strategy $y^{1d}(p, s_i)$, i.e., before customers' orders arrive.

The value of information in terms of this notation is as follows:

$$VI^d = \int_0^\infty \Pi^{3d}(s_i, z_i)\, dH\left(P^c, y^{3d}(s_i, z_i)\right) - \int_0^\infty \Pi^{1d}(s_i)\, dH\left(P^c, y^{1d}(s_i,)\right) \tag{9}$$

where $H\left(P^c, y^{xd}(\cdot)\right)$ is the distribution of the market clearing price, P^c, given other bidders using equilibrium strategies and dealer d following the strategy $y^{xd}(\cdot)$. Recall that this distribution is defined in (1) and (2) for $x = 1$ and $x = 3$, respectively.

Equation (9) defines the value of information as the difference between the expected profit when the dealer uses different strategies and has different beliefs about the distribution of the market clearing price. In stage 1, the dealer does not observe a customer bid, and thus has to integrate over customers' and other dealers' (some who might see customer bids in stage 3) bids. In stage 3, 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.

Although the calculation above allows us to measure how much extra (expected) profit dealers make when updating their bids, we do not attempt to answer the following, more ambitious question: what is the value to dealers from the institutional structure that allows them to observe and react to customer bids? Answering this question would require us to calculate dealer profits under the counterfactual scenario where dealers do not observe customer bids, for which we would have to recompute the equilibrium bidding strategies of customers and dealers. 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 point out that the "value of information" calculation would be much more difficult to conduct in an environment where there is learning about fundamentals. As we noted in the introduction, the case where there is learning about fundamentals corresponds to an auction environment with interdependent values. 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 (Laffont & Vuong (1996) and Athey & 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). (The standard deviation of ex post payoff is slightly over 2.5 basis points.) Since the average expected profit 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 for each dealer 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 profit of 13% (17% when using one auction for estimation; 12% when using 2 auctions and 10% when using 4) and the expected profit 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 per dealer and year.

7 Conclusion

In Canadian Treasury auctions, like in many financial markets around the world, dealers observe the bids of their customers. Detailed data on dealer bid updates allowed us to test whether dealers use customer bids to learn about competition and/or whether they learn about the fundamental value of the securities being auctioned. Our tests indicate that the main source of learning in this market is about competition, and that the economic value of this information to the dealers is substantial. Of course, our results pertain to the specific context we have studied in this paper. However, provided similar data is available, our testing methodology can be applied in other markets where the asymmetric treatment of dealer and non-dealer participants is a concern.

References

Armantier, O. & Sbaï, E. (2006), 'Estimation and comparison of treasury auction formats when bidders are asymmetric', *Journal of Applied Econometrics* **21**(6), pp. 745–779.

Arnone, M. & Iden, G. (2003), Primary dealers in government securities: Policy issues and selected countries' experience. IMF Working paper WP/03/45.

Athey, S. & Haile, P. A. (2002), 'Identification of standard auction models', *Econometrica* **70**(6), pp. 2107–2140.

Ausubel, L. & Cramton, P. (2002), Demand reduction and inefficiency in multi-unit auctions. working paper.

Bickel, P. J. & Freedman, D. A. (1981), 'Some asymptotic theory for the bootstrap', *The Annals of Statistics* **9**(6), pp. 1196–1217.

Blum, J. R., Kiefer, J. & Rosenblatt, M. (1961), 'Distribution free tests of independence based on the sample distribution function', *The Annals of Mathematical Statistics* **32**(2), pp. 485–498.

Cassola, N., Hortaçsu, A. & Kastl, J. (2011), The 2007 subprime market crisis in the euro area through the lens of ecb repo auctions. working paper.

Chapman, J. T. E., McAdams, D. & Paarsch, H. J. (2007), 'Bounding revenue comparisons across multi-unit auction formats under -best response', *The American Economic Review* **97**(2), pp. 455–458.

Easley, D., O'Hara, M. & Srinivas, P. S. (1998), 'Option volume and stock prices: Evidence on where informed traders trade', *The Journal of Finance* **53**(2), pp. 431–465.

Evans, M. D. & Lyons, R. K. L. (2002), ‘Order flow and exchange rate dynamics’, *Journal of Political Economy* **110**(1), pp. 170–180.

Février, P., Préget, R. & Visser, M. (2004), Econometrics of share auctions. working paper.

Fishman, M. J. & Lonstaff, F. (1990), ‘Dual trading in futures markets’, *Journal of Finance* **47**, pp. 643–671.

Grossman, S. (1976), ‘On the efficiency of competitive stock markets where trades have diverse information’, *The Journal of Finance* **31**(2), pp. 573–585.

Grossman, S. J. (1989), An economic analysis of dual trading, Technical report, Wharton School Rodney L. White Center for Financial Research.

Haile, P. A., Hong, H. & Shum, M. (2003), Nonparametric tests for common values in first-price sealed-bid auctions. working paper.

Hayek, F. A. (1945), ‘The use of knowledge in society’, *American Economic Review* **35**(4), pp. 519–530.

Hoeffding, W. (1948), ‘A class of statistics with asymptotically normal distribution’, *The Annals of Mathematical Statistics* **19**, pp. 293–325.

Hortaçsu, A. & McAdams, D. (2010), ‘Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the turkish treasury auction market’, *Journal of Political Economy* **118**(5), pp. 833–865.

Hortaçsu, A. & Sareen, S. (2006), Order flow and the formation of dealer bids in treasury auctions. working paper.

Kang, B.-S. & Puller, S. L. (2008), ‘The effect of auction format on efficiency and revenue in divisible goods auctions: A test using korean treasury auctions’, *Journal of Industrial Economics* **56**(2), pp. 290–332.

Kastl, J. (2011*a*), ‘Discrete bids and empirical inference in divisible good auctions’, *Review of Economic Studies* **78**, pp. 978–1014.

Kastl, J. (2011*b*), On the properties of equilibria in private value divisible good auctions with constrained bidding. working paper.

Laffont, J.-J. & Vuong, Q. (1996), 'Structural analysis of auction data', *The American Economic Review* **86**(2), pp. 414–420.

Lehmann, E. L. (1999), *Elements of Large Sample Theory*, Springer, New York.

Lyons, R. K. (1995), 'Tests of microstructural hypotheses in the foreign exchange market', *Journal of Financial Economics* **39**, pp. 321–351.

Madhavan, A. & Smidt, S. (1990), 'A bayesian model of intraday specialist pricing', *Journal of Financial Economics* **30**, pp. 99–134.

McAdams, D. (2008), 'Partial identification and testable restrictions in multi-unit auctions', *Journal of Econometrics* **146**, pp. 74–85.

Milgrom, P. R. & Weber, R. J. (1982), 'A theory of auctions and competitive bidding', *Econometrica* **50**(5), pp. 1089–1122.

Mudholkar, G. S. & Wilding, G. E. (2003), 'On the conventional wisdom regarding two consistent tests of bivariate independence', *Journal of the Royal Statistical Society. Series D (The Statistician)* **52**(1), pp. 41–57.

Stute, W. (1991), 'Conditional u-statistics', *The Annals of Probability* **19**, pp. 812–825.

Vickrey, W. (1961), 'Counterspeculation, auctions, and competitive sealed tenders', *Journal of Finance* **16**, pp. 8–37.

Wilson, R. (1979), 'Auctions of shares', *The Quarterly Journal of Economics* **93**(4), pp. 675–689.

A 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 (4) 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 (4) as

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

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

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

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

$$\xi\left(\hat{F}; X, h_T\right) = \frac{1}{(N_cT)^{N_c}}\frac{1}{(N_dT)^{(N_d-1)}} \sum_{\alpha_1=(1,1)}^{(T,N_c)} ... \sum_{\alpha_{N_c}=(1,1)}^{(T,N_c)} \sum_{\alpha_{N_c+1}=(1,N_c+1)}^{(T,N_c+N_d)} ... \sum_{\alpha_{N_c+N_d-1}=(1,N_c+1)}^{(T,N_c+N_d)}$$
$$\Phi\left(y_{\alpha_1}, ..., y_{\alpha_{N_c}}, y_{\alpha_{N_c+1}}, ..., y_{\alpha_{N_c+N_d-1}}, X\right) W\left(\alpha_1, ..., \alpha_{N_c}, \alpha_{N_c+1}, ..., \alpha_{N_c+N_d-1}; h_T\right) \quad \text{(A-1)}$$

where $\alpha_i \in \{(1,1), (1,2), ... (1, N_c + N_d), ..., (T, N_c + N_d)\}$ is the index of the bid in the subsample and $\hat{F}$ is the empirical distribution of bids.[10] 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

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

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

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

where $K(\cdot)$ is a bounded kernel with compact support[11] and h_T is the bandwidth satisfying $h_T \rightarrow 0$, $Th_T^5 \rightarrow 0$ 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 $\xi\left(\hat{F}; X\right)$, 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.

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

A.1 Asymptotic properties of ξ

A.1.1 Symmetric Case

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 $T \to \infty$, $\frac{M}{T} \to \infty$:*

$$\sqrt{T}\left(\hat{H}^R(X) - H(X)\right) \to N\left(0, \frac{(N-1)^2}{N}\zeta\right) \tag{A-3}$$

where $\zeta = E_{\theta_{-i}}\left[(\Phi(y_1,...,y_{N-1};X))^2\right] - \left(\binom{NT}{N-1}^{-1}\sum_{(1,1)\leq\alpha_1<\alpha_2<...<\alpha_{N-1}\leq(T,N)}\Phi\left(y_{\alpha_1},...,y_{\alpha_{N-1}};X\right)\right)^2$ *and where that last summation is taken over all combinations of $N-1$ indices $\alpha_i \in \{(1,1),(1,2),...,(1,N-1),...,(T,N)\}$ such that $\alpha_1 < \alpha_2 < ... < \alpha_{N-1}$.*[12]

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

$$\beta\left(\hat{F};c\right) = \binom{NT}{N-1}^{-1}\sum_{1\leq\alpha_1<\alpha_2<...<\alpha_{N-1}\leq NT}\Phi\left(y_{\alpha_1},...,y_{\alpha_{N-1}},c\right) \tag{A-4}$$

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

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

related *V-statistic*:

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

where the averaging is over every permutation of the NT observations.[13] Lehmann (1999), Theorem 6.2.2, p.388, shows that the asymptotic distribution of this V-statistic is identical to that of the U-statistic. Finally, since $\frac{M}{T}\to\infty$, $\hat{H}^R\to\xi$. ■

A.1.2 Consistency in the Number of Bidders

In subsection 4.3 we raised the issue of unobserved heterogeneity across auctions. Since our estimator is leveraging ex-ante symmetry of bidders (within groups of dealers and customers, respectively) and indendence of private information, there is a trade-off between pooling auctions together (i.e., performing the asymptotic exercise as the number of auctions grows to infinity) and thereby increasing the number of bids for resampling and potentially including confounding factors unobserved to the econometrician. An obvious alternative would be to use bids only from within a single auction. In Cassola, Hortaçsu & Kastl (2011) we prove that our estimator is consistent when the number of bidders within an auction grows without bounds provided that the available supply grows at an appropriate rate and that there is some non-degenerate supply uncertainty.

A.1.3 Accounting for Asymmetries

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

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

customer bids. Yet in our setting some dealer bids are of course submitted only after observing a particular customer bid and therefore it is necessary to use proper weighting of each subsample in the V-statistic. We propose to achieve this through our estimator (A-1) which uses the kernel weights $W(.)$ as defined in (A-2). 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)[14], whose Theorem 1 we adapted to our application and state as the following:

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

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

where $\sigma^2 = \sum_{j=1}^{N_c+N_d-1} \sum_{l=1}^{N_c+N_d-1} I [y_j = y_l] \left[\Phi_{jl} (\mathbf{y}) - \Phi^2 (\mathbf{y}) \right] \int K^2 (u) \frac{du}{f(y_j)}$,
and where when $y_j = y_l$,
$\Phi_{jl} (\mathbf{y}) = E \left[\Phi (y_1, ..., y_{j-1}, Y, y_{j+1}, ..., y_{N_c+N_d-1}) \times \Phi \left(y_{N_c+N_d}, ..., y_{N_c+N_d+l-2}, Y, y_{N_c+N_d+l}, ..., y_{2(N_c+N_d-1)} \right) \right]$
and when $y_j \neq y_l$, $\Phi_{jl} (\mathbf{y}) = 0$.

In our empirical application, we 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 & Freedman (1981), Theorem 3.1, using the fact that the variance and any covariances of our kernel $\Phi(.)$ in the V-statistic (A-1) are bounded.

B 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

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

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.[15] We use Mudholkar & Wilding (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 statistic tabulations are only available for $N > 4$, we only considered auctions with at least 10 dealer bids, which reduced our sample to 64 auctions. We considered 100 random splits of the sample when constructing the test statistics, and recorded the number of times the bids from an auction rejected the null hypothesis of independence. Over 100 iterations, the average number of auctions (among 64 auctions) for which the Pearson coefficient was significantly (at the 5% level) different from zero was 3.68. The Spearman rank correlation was different from zero on average for 1.9 auctions, Kendall's tau test led to rejection for 1.33 auctions, and the BKR test led to rejection in 3.67 auctions. We also looked at the maximum number of auctions for which the test statistic was rejected for any given random split of the data. The Pearson test was rejected for a maximum

[15] Mudholkar & 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.

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 $b_{2t}, ..., b_{\frac{N^*_{dt}}{2}}$ and the second sample of $b_{\frac{N^*_{dt}}{2}+2}, ..., b_{N^*_{dt}}$, where N^*_{dt} is the number of non-zero dealer bids in auction t.[16] This test rejects the null hypothesis in 9 out of 116 auctions of 3-months T-bills.

[16]By leaving out the first dealer in each half, we wanted to make it explicit that we are testing for the equality

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: $H_0 : F(b|b_{1t}) = F\left(b|b_{1(t+1)}\right)$. 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 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.[17] 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.

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.

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

Table 7: Joint Hypothesis Studentized Test Results

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

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