

Order Flow and the Formation of Dealer Bids: Information Flows and Strategic Bidding in the Government of Canada Securities Auctions *

Ali Hortaçsu, University of Chicago and NBER

Samita Sareen, Duke University

October 10, 2005

Abstract

Is privately observed order flow an important component of the strategic price formation process in government securities markets? Utilizing a detailed data set on Government of Canada securities auctions, we argue that the answer is yes. Government securities dealers who participate in these auctions have access to private order flow information through two channels: non-dealer customer bids that are, by the rules of the auction, routed through dealers, and when-issued market transactions undertaken by the dealers. Dealer bids respond strongly to privately observed customer bids and when-issued market positions, and dealers who observe a larger number of customer bids can predict the auction cutoff price better. We also document patterns of customer bidding behavior consistent with a strategic response to dealers' use of the information contained in customer bids.

*We thank the editor, an anonymous referee, Bruno Biais, Kjell Nyborg, Robert Porter, Rafael Romeu, Leonardo Rezende, Christian Upper, S. Viswanathan and participants at the HIOC 2004, AMD 2004 and MTS 2004 conferences for their very valuable comments and suggestions. For correspondence, e-mail: hortacsu@uchicago.edu or ssareen@econ.duke.edu.

A large theoretical literature in financial economics studies how trading mechanisms aggregate individual traders' private information regarding asset values into asset prices.¹ However, it is not always easy to identify the precise sources of private information in some important asset markets. For example, in the vast markets for government securities, most, if not all, traders have access to the same financial and political data and news needed to value the fixed cash flow streams, making it difficult to argue for private information playing a role.

Recent empirical research, however, yields evidence consistent with the existence of private information in government securities markets. Green (2004) finds that sensitivity of prices to order flow in the U.S. treasury market increases following economic announcements. This suggests that such announcements increase the amount of information asymmetry in this market. Brandt and Kavajecz (2004) point out that there are significant price changes in the U.S. treasury market in days without news announcements. Following the methodology of Evans and Lyons (2002), they show that orderflow imbalances (as measured by signed trades) can account for 26% of the day-to-day variation of yields in days without news announcements.

The existence of private information in government securities markets is also a key assumption underlying theoretical and empirical analyses of Treasury auctions.² However, the existing literature on government securities auctions, while ascribing a central role to private information, is not able to identify its source directly. Isolating what drives private information is important in these markets, since theoretical and policy analyses regarding the design of the auction and surrounding market rules rely very sensitively on the exact specification of the informational and strategic environment.³

This paper shows that an important source of private information for government securities dealers is (privately observed) customer order flow. This mechanism was first noted, but not empirically documented, by Nyborg and Sundaresan (1996, page 68). We provide direct evidence for it using very detailed data from Government of Canada securities auctions, where we have the unique ability to relate the bidding behavior of government securities dealers to two sources of private information.

The first source of dealers' private information are non-dealer customers who place their own bids in the auctions, but who, by regulation, have to route their bids through dealers. Our data set allows us to observe these customer bids at the dealer level. Moreover, since the computerized

¹O'Hara (1995) and Brunnermeier (2001) provide comprehensive surveys.

²See, for example, Gordy (1994), Nyborg and Sundaresan (1996), Nyborg, Sundaresan and Rydqvist(2002), Hortacsu (2002), Fevrier, Preguet and Visser (2002), Goldreich (2004), Keloharju, Rydqvist, Nyborg (2004).

³See Binmore and Swierbinski (2003), Sareen (2003) for recent surveys of the theoretical and empirical literature on the design of Treasury auctions.

auction system of the Bank of Canada allows bidders to modify their bid entries until the bid submission deadline,⁴ we can track changes made to customer and dealer bids before they become official.

The second source of private information is the “when-issued” market, where dealers and customers sign forward contracts. Our data set allows us to track the evolution of dealers’ net long or short positions in this market during the period preceding each auction, thus allowing us to reconstruct this dimension of dealers’ private information as well.

Our analysis uncovers a strong response of dealers’ bidding behavior to customer order flow, and a rich array of strategic behavior that arises from the manner in which private information is gathered and utilized in this market. In Section 1, we note that “when-issued” trading is done exclusively through the informationally opaque channel of bilateral negotiations, and not through the more transparent medium of electronic interdealer brokers.⁵ This suggests that valuable information is gathered on the when-issued market. Section 2 provides more details of the “information-gathering” process that precedes each auction. We show that bidding activity is very much concentrated in the last few minutes preceding the submission deadline, that dealers who route customer bids are the latest to submit their official bids (and that their bids lag their customers’ bids), and that bids submitted closer to the deadline are much less dispersed and more informative of the auction outcome than earlier bids.

Section 3 presents our key results, where we use bid-level data to examine the response of dealers’ bids to privately observed customer order flow. We find that dealer’s bids, and especially the modifications they make to their bids reflect the information contained in customer bids. Net positions obtained in the when-issued market also affect dealers’ bidding in a manner consistent with the model of Nyborg and Strebulaev (2004). Moreover, we demonstrate that access to customer order flow increases dealers’ bidding effectiveness: dealers with access to order flow information from a larger number of customers win a higher proportion of the securities they bid for, but do not pay more (per unit) than other bidders.

If access to customer order flow provides an informational advantage to dealers, we may expect customers to devise strategies to lessen any adverse consequences of this advantage on their own trading outcomes. Indeed, in section 4, we document evidence suggestive of several customer “strategies.” One strategy is to divide bids across multiple dealers; which appears to be used

⁴At the bid submission deadline, bids are deemed “official,” and securities are allocated using a discriminatory (pay-as-bid) pricing mechanism. Until the auction results are announced, bids are kept secret from competing bidders. We explain the institutional features in more detail in Section 1.

⁵However, dealers have to report large changes to their net positions to the Bank of Canada; our data captures these reports.

more often when the customer has a larger order. Another strategy is to establish a long-term relationship with a dealer. Yet another customer strategy that we have suggestive evidence for is to delay one's bid as much as possible, so that the dealer does not have time to revise her own bid based on the customer's bid.

Should the Government of Canada eliminate the informational advantage of dealers by allowing customers to participate directly in the auctions? Section 5 of the paper and the Appendix discuss the revenue and efficiency consequences of such a policy change within the context of a simple, stylized model.

Along with establishing an economic mechanism through which order flow affects asset prices in government securities markets, our study contributes to the empirical microstructure literature in two ways.⁶ First, the cross-sectional nature of our data set allows us to conclude that access to order flow is an important factor explaining the variation in bidding behavior and performance *across* participants in the Government of Canada securities market. Such a cross-sectional conclusion could not have been obtained from previous dealer-level studies of microstructure hypotheses, such as Lyons (1995) and Madhavan and Smidt (1991), as these studies utilized transactions data from a single dealer.⁷

Second, the panel nature of our data allows us to account for one of the chief empirical confounds of the empirical microstructure literature in a convincing manner. Any study attempting to link (dealer) trading behavior to proxies for order flow has to account for the possibility that both order flow and dealer trading may be correlated solely through their response to public information flows that are observed by all market participants, but unobservable to the econometrician.

One way to address this problem is to study the response of trading to privately observed order-flow *across* multiple dealers. Differences in dealer behavior, controlling for other factors leading to dealer heterogeneity, should be attributable to privately observed order flow; not to public information that all dealers observe. This is exactly our identification strategy: since we observe multiple dealers bidding in the same auction (and revising their bids within the same pre-auction period), we look at the differences across dealers' bids and bid modifications, and see how these covary with privately observed order flow.⁸

⁶Examples from this literature include work on foreign exchange markets (Lyons (2001) and the references within, Evans and Lyons (2002), Ito, Lyons and Melvin (1998)), equity markets (Hansch, Naik and Viswanathan (1998)) and option markets (Easley, O'Hara and Srinivas (1998)).

⁷Lyons (1995) uses data obtained from a large currency dealer to test whether the prices quoted by this dealer responds to changes in the inventory holdings and customer orders observed by this dealer. Madhavan and Smidt (1991) analyze NYSE specialists' trades.

⁸An alternative identification strategy is pursued by the single dealer analyses of Madhavan and Smidt (1991)

I. Description of Data and Surrounding Institutions

The Bank of Canada, on behalf of the Government of Canada (GoC), issues bonds and treasury bills. Bonds are long-term securities and treasury bills are short-term securities.¹⁰ The average issue size for bonds and treasury bills is 5.5 and 2.4 billion Canadian dollars, respectively. Besides the large absolute magnitude of the issuance size, government securities are the backbone of the fixed-income securities market in Canada. Government securities account for 72% of the fixed income market in Canada.¹¹

The process for issuing a “typical” GoC debt instrument links three markets: the *when-issued market*, the *primary market* or the *auction*, and the *secondary market*. The bidding process begins with the Bank of Canada soliciting bids in the *primary market*. The *when-issued market* for the auctioned security also begins with this solicitation. Government securities distributors¹² typically take short positions in the when-issued market through forward contracts for the yet-to-be auctioned security. Subsequently, dealers attempt to cover these short positions by buying the security from other dealers in the when-issued market, from the issuer in the primary market, and finally from another dealer in the secondary market after the auction. Trading for the auctioned security begins in the *secondary market* 15-20 minutes after the bid submission deadline.

It is significant to note the overlap of trading in the when-issued market and the bid preparation process in the primary market in the period preceding the bid submission deadline. The overlap of these two processes makes it possible for a dealer to gather orderflow information from customers in *both* the when-issued market and the primary market, and respond strategically in one market to orderflow information gathered in the other. We now detail some further aspects of these markets.

and Lyons (1995), who utilized an explicit structural specification of the evolution of public information and the dealer’s optimal response to it, which leads to a reduced-form specification in which the contribution of unobserved public information to *changes* in price quotes is an orthogonal error term.⁹ Note that while we use the first-differencing strategy of these papers when analyzing the response of *bid modifications* to order flow in Section 3.A, our ability to observe multiple dealers allows us to eschew the orthogonality assumption.

¹⁰Bonds have a maturity of 30, 10, 5 and 2 years, and treasury bills with a maturity of 1 year, 3 months and 6 months. Treasury bills are zero-coupon bonds.

¹¹The figure reported refers to the amount of securities outstanding that are issued by the Government of Canada to that issued by all issuers in Canada. Source: pp. 92, Table 16A, BIS Quarterly Review, International Banking and Financial Market Developments, September 2004, Statistical Annex.

¹²The Bank of Canada designates certain institutions as distributors of government securities. These institutions are obligated to buy and sell securities to individual investors.

A. Primary Market

The Bank of Canada issues GoC securities in the primary market through sealed-bid discriminatory (pay-as-bid) auctions. The auctions are conducted on an electronic system: a bidder communicates his bid in real-time on a secure electronic system that can be viewed only by the Bank of Canada.

On average, 26.2 bidders participate in an auction. Potential bidders in the government securities auctions can be classified into three groups: primary dealers (PDs), other government securities distributors (GSDs) and customers. “Other government securities distributors” refers to government securities distributors excluding primary dealers. On an average, 12.7 PDs, 4.9 GSDs, and 8.6 customers participate in an auction. Of these, PDs account for as large as 90% of the primary market in terms of the volume of securities issued. Individually, PDs are the largest bidders, though GSD and customer bids are similar in size. The average PD bid is for 13.4% of the total issue amount, the average GSD bid is for 2.7%, and the average customer bid is for 2.4%.

The distinction between government securities distributors (PDs and GSDs) and customers is that the latter cannot bid on their own account in the auction; rather government securities distributors submit bids on behalf of customers.¹³ Thus, Government securities distributors submit bids on their account and on behalf of the customers, being “bidders” in the former and “submitters” in the latter case. We see that an average dealer services 0.8 customers in an auction. Customers can choose to route their bids through more than one dealer in the auction. On an average, a customer routes bids through 1.5 dealers in an auction.

In general, the maximum amount that a dealer can bid either for himself or his customers, is based on his past primary market winning share and secondary market trading share, net of his current holdings of the auctioned security. But the Bank of Canada stipulates that no bidder in an auction can bid for more than 25% of the issue amount, and no dealer as a “submitter” can bid for more than 40% of the issue amount whether on his own account or that of his customers.¹⁴

Bids can be submitted as *competitive* tenders and *noncompetitive* tenders.¹⁵ Typically, a

¹³Government securities distributors do not charge a fee to route customer bids.

¹⁴The distinction between primary dealers and other government securities dealers is based on the differences in their shares of the primary and secondary markets. Thus, while primary dealers can bid the maximum amount allowed by the Bank of Canada to a “bidder” or a “submitter”, government securities distributors cannot. As an example in Treasury bill auctions, government securities distributors can bid only a maximum of 10% of the issue amount of an auction net of their current holdings of the auctioned security.

¹⁵A *noncompetitive* tender comprises a quantity subject to an upper bound of \$3 million, with a participant being allowed to submit a single *noncompetitive* tender.

participant's competitive tender will comprise of price-quantity pairs and the participant's *net position* of the yet-to-be auctioned security at the point of time the tender is submitted.¹⁶ The reported net position is the participant's net holdings (whether long or short) of the security being auctioned at that point of time, and capture the forward contracts of a dealer with other dealers or customers, prior to the auction. On average, dealers' contracted positions comprise 54.1% of the issue amount.

All tenders have to be submitted before the expiry of the bid submission deadline. However, participants can revise or cancel previously submitted tenders prior to the auction deadline; there are no limits on the number of revisions that an auction participant can make. Tenders submitted after the bid submission deadline are cancelled unless they are on account of transmission failure. After the expiry of the bid submission deadline, submitted tenders are allotted through a discriminatory price auction. The awards are announced 15 minutes after the bid submission deadline, with the announcement including the cutoff price, the amount issued, the quantity weighted average price, and the low and high yields.

Our data captures several aspects of the primary market. For the primary market we have data over the period October 1998 to March 2003. In particular, we have 347 treasury-bill auction and 66 bond auctions in our sample. For each auction in the sample, we have auction-level information reflecting the total issue and bid amounts and the cutoff yield (or market clearing price). In addition, we have detailed information regarding the entire sequence of tenders submitted by the bidders. For each tender, we observe a time-stamp, indicating the exact time at which the tender was submitted; the participant type of the bidder (primary dealer, other government securities distributor, customer, Bank of Canada); if the bid belongs to a customer, the submitter who entered the bid into the system (who has to be a primary dealer or other government securities distributor); net position amount indicating a participant's net holdings of the yet-to-be auctioned security at that time the allotted tenders of each bidder. Finally, for each tender submitted in an auction by a bidder, we have the bid amount and yield pairs; the maximum amount a participant can bid as a "submitter", and as a "bidder"; and the amount allotted to each participant.

Note that a bidder can revise a submitted tender before the auction deadline, and our data set captures the entire *history* of these revisions. For each auction participant, the unique tender and the constituent bids in this tender that are used to determine the cutoff yield and the allotment of the auctioned security is referred to as *official* data.¹⁷ Thus, the official data is a subset of the

¹⁶A customer's tender may comprise only price-quantity pairs as a customer has the option of submitting his net positions directly to the Bank of Canada instead of communicating it through the tender(s) he submits through dealers.

¹⁷An official tender is the last submitted or cancelled tender of a participant.

history data. While several empirical studies of treasury auctions have used data on official bids, this is the first paper that makes use of the history aspect of treasury auction data.

B. “When-Issued” and Secondary Markets

The *when-issued* market officially begins a week before the bid submission deadline in the primary market. In this period both the primary and when-issued market are open. Participants engage in forward trading for the yet-to-be auctioned security in the *when-issued* market. Also, the *primary* market is followed by trading in an active resale market for the “new issue” called the *secondary* market.

Both the when-issued and the secondary markets operate through two channels. The first is a bilateral, telephone-based over-the-counter market, and the second is an interdealer market operated by electronic interdealer-brokers.

Non-dealer customers’ when-issued and secondary market transactions are conducted on the over-the-counter market. Institutional investors (for example, pension funds, mutual funds) trade with dealers on a bilateral basis over the telephone, with the result of these transactions known only to the two counterparties participating in the transaction. In 2002, customer-dealer over-the-counter transactions accounted for 54% of the total volume traded of Government of Canada bonds.

Dealers, on the other hand, can use both the over-the-counter channel or the electronic interdealer-broker channel. Electronic interdealer brokers post on an electronic screen bid, offers, and trade outcomes communicated to him by the dealers.¹⁸ In 2002, 86% of the total interdealer transactions volume was brokered through these electronic interdealer brokers. As said, only dealers can post quotes or trade through the interdealer brokers. However, both customers and dealers have viewing access to the electronic screens of an interdealer broker via CanPX, a data service that consolidates and disseminates the trade and quote information provided by the electronic interdealer brokers.¹⁹

We do not have real-time quote or trade data for over-the-counter transactions. However, we do observe this market through the net positions that bidders report when they submit tenders in the primary market. This allows us to construct a bidder-level proxy of the transactions flow leading up to an auction.

In contrast, we do have access to transactions level data from CanPX, allowing us to observe

¹⁸Typically, the identity of the dealer is not revealed.

¹⁹Prior to the establishment of CanPX, only dealers could observe the trade and quote activity on the electronic interdealer-broker screens.

trade and quote activity in real-time in the interdealer market over the period July 4, 2001 to September 10, 2001, and February 25, 2002 to February 27, 2003. This data, at least in principle, allows us to track activity on the when-issued and secondary markets.

However, to our initial surprise, when we looked at CanPX data to investigate interdealer trading patterns in the when-issued market for a one-week window prior to the auction, we observed *no* trade or quote activity for the yet-to-be-auctioned security during the *entire* set of security issues covered by our sample. We should note that Nyborg and Sundaresan (1996) found a very similar lack of liquidity in the U.S. when-issued interdealer market in the period prior to the auction.

Our finding that no when-issued trades are observable on CanPX was, at first, puzzling, given that the basis for a when-issued market is price discovery. However, as evidenced from our observation of the levels and changes in the declared net positions of the bidders, there is indeed an active market for when-issued trading. Yet, none of this trading shows up on CanPX, and appears to be conducted exclusively on the over-the-counter market. Since transactions on the over-the-counter market are bilateral, and are not, to our knowledge, publicly observable, this finding appears to be consistent with the intuition of Bikhchandani et al. (1994), Simon (1994), and Nyborg and Sundaresan (1996) that pre-auction when-issued transactions carry important informational content that dealers do not want to share with other market participants.

II. Dynamics of Bid Submission

Our finding in the previous section that the “when-issued” market operates exclusively through the opaque over-the-counter market suggests that dealer-customer interactions carry economically important information. A corollary of this observation is that dealers should wait to gather as much information from their customers before finalizing their bidding decisions.

We find strong evidence in support of this hypothesis. Bank of Canada securities auctions have a fixed bid submission deadline, and bidders are allowed to submit bids for a particular auction two weeks ahead of time. However, bidding activity is concentrated within the last few minutes before the deadline. In Figure 1, we use the time stamps of tenders to plot the cumulative distribution function of official ²⁰ bid submission times for different subsamples of our data set. Observe that ninety percent of all competitive bids, those that specify a price as well as a quantity, are submitted in the last 20 minutes of the submission deadline. In contrast, observe that non-competitive bids, i.e. those that do not specify a price, tend to come much earlier than competitive

²⁰Recall that “official” bids are the final bid submissions that are taken into account for the allocation.

bids. The median competitive bid comes in 7.9 minutes before the deadline, whereas the median non-competitive bids comes in 26.5 minutes before the deadline.

One explanation for the last-minute concentration of bids is that, especially for competitive bids, new information is very important in forming expectations about the appropriate value of the security being auctioned. Hence, bidders are reluctant to commit to a price bid until they are certain that no new information will be released.²¹ Note that bidding in bond auctions is much more concentrated at the very last minutes of the auction, with the median official bid for bonds coming 2.5 minutes before the auction deadline, as opposed to 9.3 minutes for T-bills. Pricing longer maturity securities depends quite sensitively on expectations about the future, and since many more factors enter into forming expectations about the long-term rather than the short-term, one may expect bidding decisions to be more responsive to arrival of new information.

Yet another pattern that supports the information gathering hypothesis is that bids that come within time intervals closer to the bidding deadline are less dispersed, and, controlling for bid volume within a given time interval, are closer to the market clearing price than bids submitted within earlier time intervals.²² To show this, we first categorized official bids by clustering their submission times into 5 minute periods leading up to the deadline. Since bids come in the form of multiple price-quantity pairs, i.e. demand schedules, we calculated the quantity weighted average price of each bid schedule (made up of K price-quantity pairs) by the formula:

$$p^{QW} = \frac{\sum_{i=1}^K p_i q_i}{\sum_{i=1}^K q_i} \quad (1)$$

We then calculated two dispersion measures of the bids submitted in different time intervals. The first measure is the standard deviation of (quantity-weighted) price bids submitted within a given interval. The second measure we use is the absolute dispersion of the (quantity-weighted) price bid around the cutoff/market-clearing price of the auction, i.e. lowest price at which the securities

²¹The explanation for last-minute can not be pure “procrastination” by dealers, since we observe some dealers submitting both competitive and non-competitive bid in a given auction, where the non-competitive bid is submitted earlier than the competitive bid. Note also that the phenomenon we are observing is similar to the “sniping”, or “last minute bidding” phenomenon that has been documented in the context of Internet auctions by Roth and Ockenfels (2002) and Bajari and Hortaçsu (2003). Internet auctions are open-ascending auctions where bidders can see and respond to each others bids; this creates several strategic reasons for “sniping.” Although Bank of Canada securities auctions are sealed bid auctions in which dealers do not observe and cannot respond to other dealers’ bidding activity, dealers can observe and respond to the bidding activity of customers who route bids through them, and have the incentive to wait for these bids to arrive.

²²These patterns echo the findings of Biais, Hillion and Spatt (1999) in the opening periods of the Paris Bourse.

were sold. That is, if M_{jt} bids $\{p_1^{QW}, \dots, p_{M_{jt}}^{QW}\}$ arrived in auction j in time period t , we compute

$$ABS_DISPERSION_{jt} = \frac{1}{M_{jt}} \sum_{m=1}^{M_{jt}} |p_m^{QW} - p_j^c| \quad (2)$$

where p_j^c is the cutoff price in auction j . Again, this is a measure of how close the bids submitted in a given time interval are to being the marginal bid in the auction.²³

We then regress our dispersion measures on time-interval indicators and bid volume in time period t , to see if bid-dispersion around the market clearing price declines over time. The results are reported in Table 1, where we break down the results into the two subsamples of T-bill and bond auctions. In each regression, we control for auction-level fixed effects to isolate within-auction temporal patterns. We see that in bond auctions, the dispersion of bids around the market clearing price declines monotonically as the bidding interval gets closer to the deadline. A similar trend also reveals itself when we look at the standard deviation of bids within a given interval. However, in treasury bill auctions, the trend almost reverses, with later bids displaying a higher degree of dispersion than earlier bids. Thus, our results indicate that in bond auctions, bids that are submitted closer to the bidding deadline are less-dispersed than earlier bids, and that they tend to be closer to the market-clearing price. This supports the presence of an information gathering process leading up to bond auctions. However, the same process does not appear to be important in T-bill auctions.

We next examine dealer-level data to test the information-gathering hypothesis. Table 2(a) looks at all instances in which a dealer submitted a bid on her own behalf as well as for customers. We then compare the submission time of the latest customer bid to the submission time of the dealer's own official bid using a pairwise t-test (i.e. the within dealer difference). The test, when conducted for the entire spectrum of maturities, reveals that dealer bids lag the latest customer bid by 0.43 minutes, the difference being statistically significant at the 1.7% level. Furthermore, we find that dealer bids lag customer bids 55% of the time. However, as the next column of Table 2(a) shows, the difference in the timing between dealer and customer bids do not appear to statistically different for Treasury bills. Dealer bids lag customer bids 52% of the time – not visibly (or, as it appears, statistically) different from an even split.

²³Since Bank of Canada securities auctions follow the discriminatory (pay-your-bid) pricing format, bidders have an incentive to bid prices that are as close to the cutoff price as possible. If the cutoff price was known to the bidders, all of them would submit bids that are equal to the cutoff price. Of course, in reality, bidders do not know the cutoff price and have to balance the risk of paying too much vs. not winning the auction. However, one has to control for the bid volume submitted within a time interval to make sure that bid volume does not cause these bids to become marginal. We thank the referee for pointing this out.

In contrast, dealer bids appear to lag customer bids much more visibly in auctions for longer-term securities. Not only both customer and dealer appear to come much later for these auctions; dealer bids, on average, are submitted 2.5 minutes later than the latest customer bid. The lag is statistically significant at the 1% level. Moreover, dealer bids lag customer bids 74.9% of the time in these auctions.

As a second test, we compare the official bid submission times of dealers who route customer bids with dealers who do *not* route customer bids.²⁴ Although both types of dealers conduct transactions on the when-issued market, dealers with customers have additional information coming from customer bids. This might drive them to bid later than dealers without customers, especially if customers are also transacting in the when-issued market before submitting bids through their dealers.

To implement this test, we ran the following regression:

$$\begin{aligned} TimeToDeadline_{ij} = & a_1 * I(i \in \{Dealers\ without\ Customers\}) + a_2 * I(i \in \{Customers\}) \\ & + b * BidQuantity_{ij} + u_j + \varepsilon_{ij} \end{aligned} \quad (3)$$

where $TimeToDeadline_{ij}$ is the minutes before the bidding deadline that bidder i submitted her official bid in auction j ; $BidQuantity_{ij}$ is the quantity bid of i , and u_j is an auction fixed-effect. Dealers with customers are the excluded category of agents in this regression. With the auction fixed effects, this regression isolates within-auction differences in official bid submission times across customers, dealers without customers, and dealers with customers.

Table 2(b) reports the results of this regression, again broken down into the subsamples of T-bill and bond auctions. We see, consistent with previous results, that in T-bill auctions there are no statistically significant differences across dealer (with or without customer) and customer bid timings. However, in bond auctions, it appears that dealers without customers submit their official bids 1.32 minutes earlier than dealers with customers, and that customer bids are routed 2.18 minutes earlier. Furthermore, in both types of auctions, we find that larger bids are submitted later than bids for small quantities – which suggests that bidders with larger payoffs at stake wait longer to bid, possibly in order to benefit from any extra information they may gather before making their final bidding decision. Hence, the evidence once again suggests that dealers react to the information contained in their customers’ bids when and where it matters. The findings also confirm our intuition regarding the difference in the importance of “customer information” across bond vs. T-bill auctions.

²⁴We thank the referee for suggesting this test.

III. How Does Private Order Flow Information Affect Dealer Bidding Behavior?

We now investigate whether the observed differences in dealer bids *within* a given auction can be explained by the order flow information that is observed privately by each dealer. As explained in Section 1, a dealer may obtain private information about order-flow from two sources: through the transactions she conducts on the over-the-counter when-issued market, and through customer bids routed through her.

Specifically, the regression specification we use to test this is:

$$\begin{aligned} DealerP_{ij}^{QW} = & \beta_1 CustP_{ij}^{QW} + \beta_2 \frac{NetPos_{ij}}{BidLimit_{ij}} + \beta_3 \frac{DealerQ_{ij}}{BidLimit_{ij}} \\ & + \beta_4 \frac{CustQ_{ij}}{BidLimit_{ij}} + \beta_5 BidLimit_{ij} + u_j + \gamma_i + \varepsilon_{ij} \end{aligned} \quad (4)$$

In this specification, j indexes auctions in our data set, while i indexes dealers within an auction. The dependent variable, $DealerP_{ij}^{QW}$, is dealer i 's quantity-weighted price bid in auction j , minus the cutoff price in auction j . $CustP_{ij}^{QW}$ is dealer i 's customer's quantity-weighted price bid, minus the cutoff-price (this is averaged over different customers if dealer i submits bids for multiple customers). $NetPos_{ij}$ is dealer i 's declared net position in auction j . We normalize the net position amount of the dealer by the dealer's maximum submission limit, $BidLimit_{ij}$, which we also take as a proxy of the size of the dealer's demand. To account for any constraints that the submission limit may put on a dealer's price bid, we also control for the ratios $\frac{DealerQ_{ij}}{BidLimit_{ij}}$, and $\frac{CustQ_{ij}}{BidLimit_{ij}}$. These respectively account for the fraction of submission capacity that a dealer uses for her own bid vs. customers' bids.

As pointed out in the introduction, dealer bids may respond to public as well as private information flows; thus variation in dealer bids across auctions may reflect unobserved (to the econometrician) variation in public information flows as well as variation in private information flows. In the above specification, the auction fixed-effect, u_j , serves to capture the unobservable component of dealer bids common to all dealers' participating in the auction, presumably due to public information flows.²⁵ We also include a dealer fixed-effect term, γ_i , to control for systematic

²⁵One could also proxy for public information flows using real-time trading prices on the "when-issued" market, or prices of securities that might be close substitutes. As noted in Section 1.B, however, CanPX is conspicuously silent during the period preceding the auction. Although there is some trading activity in other securities, in most cases, they have very different maturities and/or coupon structures. Nyborg and Sundaresan (1996) note the problems attached to using yield-curve based valuation benchmarks. Even in the cases where there is a maturity and cash-flow structure match, the auctioned security and the security trading on CanPX are not direct substitutes due to the liquidity premium attached to the benchmark issue, mimicking the "on-the-run" premium in the U.S.

differences across the bidding patterns of dealers that do not vary from auction-to-auction. Note that, in contrast, “order flow” information possessed by a dealer varies from auction to auction. In Section 3.B, we control for unobserved dealer heterogeneity that is allowed to vary from auction to auction.

The results of the above regression are reported in Table 3. We first report the results of the regressions without dealer fixed-effects. The first coefficient estimate in Column (1) of the table shows that one cent increase in customers’ bids is associated with a 0.59 cent increase in a dealer’s own bid. This estimate is highly statistically significant. Note also that the net position of the dealer enters very significantly into the regression. The coefficient estimates indicate that dealers with long (positive) positions bid less aggressively. One explanation for this might be that dealers taking on large short positions are wary of being squeezed in the resale market, and consequently bid more aggressively.

Our findings regarding the impact of net positions on bidding behavior is consistent with the theoretical results of Nyborg and Strebulaev (2004) who analyze a strategic model of short squeezes in discriminatory auctions. In the equilibria they analyze, dealers with short positions bid more aggressively than dealers with zero net positions to avoid a squeeze in the resale market. Interestingly, their results also indicate that dealers who have long positions also bid more aggressively. Intuitively, this arises from the fact that dealers who are long have the incentive to create and therefore profit from a squeeze; but they can only do this if they outbid some of the shorts.²⁶

Columns (2) and (3) of this table replicate the same regressions on the subsamples of T-bill and bond auctions, respectively. Note that the correlation between dealer and customer bids is 0.598 for bonds, as opposed to 0.158 for T-bills, which, consistent with the results in the previous section, suggests that customer bids are more influential drivers of dealer bids in bond auctions.

Binding bid limits would suggest that dealers who are bidding for larger amounts (compared to their bidding limits) would bid more aggressively. While this appears to be the case in T-bill auctions, we find the reverse pattern in bond auctions. This is consistent with the asymmetry we have noted between long and short-term government securities: since dealers are less uncertain about the market clearing price for T-bills relative to bonds, they appear to be willing to bet

Treasury market documented by Fleming (2001, 2002).

²⁶To test this prediction, we ran the regressions in Table 3 by introducing long and short positions separately into the specification, and obtained that both short and long positions lead to more aggressive bidding. However, our results indicated that dealers who are short bid much more aggressively than dealers with comparably-sized long positions – thus the overall negative coefficient reported in Table 2. We should note that Nyborg and Strebulaev’s results also suggest that shorts should be expected to bid much more aggressively than longs.

on bid amounts closer to their bidding limits in T-bill auctions compared with bond auctions. Submission limit constraints brought on by customer orders, as proxied by $\frac{CustQ_{ij}}{BidLimit_{ij}}$, do not appear to affect the price level of dealer bids significantly.

In Columns (4)-(6), we run the same regressions with dealer-level fixed effects added into the specification. Some dealers may systematically bid higher than others due to differences in their demand for the security; dealer fixed-effects attempt to capture such unobserved drivers of dealer demand. Indeed, the R^2 's of the regressions indicate that dealer fixed effects capture a lot of the variation in the bids. However, the correlations between customer bids and dealer bids still remain at very similar levels, indicating that the estimated dealer response to customer bids is not attributable to unobserved systematic differences across dealers.

A. Dealer Bid Revisions

We can conclude from the previous results that there is a robust correlation between dealer and customer orderflow (as measured by customer bids and net positions), and that this helps explain the variation in dealer bids within a given auction. However, the variation in dealer bids may in fact be driven by other, unobserved sources of dealer-level heterogeneity with which our orderflow variables merely happen to be correlated.²⁷ Thus, we now refine the causal interpretation of our analysis by looking at bid *revisions*. In particular, we test whether these bid revisions can be explained by innovations to a dealer's information set through changes in the dealer's net position and bid revisions by customers.

To do this, we utilize the bid history aspect of our data set, which comprises of *all* bids submitted by the dealers, not just the official bids. This allows us to track the modifications that dealers make in their bids on Bank of Canada's computerized bidding system.

To calculate modifications in dealers' bids, we fix a time interval, starting $T = 10$ or $T = 30$ minutes prior to the bid submission deadline, until the submission deadline. We then calculate the "standing bid" of the dealer at time T , which is the dealer's most recent bid as of time T . We then calculate the difference between the dealer's official bid, and her standing bid at T minutes prior to the deadline.

We then perform the same calculation for the customers, taking the difference between their official bid and their standing bids at T minutes prior to the bid submission deadline. Thus, for a dealer who routes customer bids, the change in her information set between time T and the auction deadline includes the modifications in customer bids.²⁸

²⁷Although we controlled for dealer fixed effects in Table 3, as we noted there, these fixed effects are constant across auctions.

²⁸The dealer may choose not to change her bid in response to a modification in her customer's bid. This choice

Dealers who do not route customer bids also make revisions to their bids. As the auction deadline approaches, dealers may also respond to developments in the when-issued market; thus dealers who do not observe customer bids may modify their bids in response to the when-issued market. Since, by law, dealers have to report changes in their net long or short positions to the Bank of Canada along with any changes to their bids, we observe net impact of this trading activity through modifications to the dealers positions. Again, we code these modifications as the difference between a dealer’s “standing” net position at T minutes prior to the deadline, and the net position reported along with the dealer’s official bid.

We run the following regression to test whether dealer bid revisions are correlated with customer bid revisions:

$$\begin{aligned} \Delta DealerP_{ij}^{QW} = & \beta_0 RoutesCustomerBids_{ij} + \beta_1 RoutesCustomerBids_{ij} \Delta CustP_{ij}^{QW} \\ & + \beta_2 \frac{\Delta NetPos_{ij}}{BidLimit_{ij}} + \beta_3 \frac{\Delta DealerQ_{ij}}{BidLimit_{ij}} + \beta_4 \frac{\Delta CustQ_{ij}}{BidLimit_{ij}} + u_j + \gamma_i + \varepsilon_{ij} \end{aligned} \quad (5)$$

This regression is essentially the first-differenced (between $T = 10$ or $T = 30$ standing bid and the official bid) version of equation (4); with the additional feature that we control for the bid revision activities of dealers who do not observe customer bids (the indicator variable *RoutesCustomerBids* takes the value of 1 if dealer i routed customer bids in auction j). *BidLimit_{ij}* does not enter into this regression since it does not vary across bid revisions.

Table 4 reports the results of the regression. The dependent variable in this table is the change in a dealer’s bid during the last 10 or 30 minutes of an auction. Note that in both time intervals, modifications in customer bids are positively correlated with modifications in dealer bids. For the entire sample of auctions, the estimates indicate that a 1 cent increase in a customer’s bid translates into a 0.24 cent increase in the dealer’s bid, if the change comes within the last 10 minutes. The response is somewhat less if the change comes within the last 30 minutes, and indicates a 0.06 cent upward revision of the dealer bid in response to a 1 cent increase in the customer’s bid. Note that from columns (2),(3),(5) and (6), we see that the correlation is much more pronounced for bond auctions as opposed to T-bill auctions.

Another interesting result of this regression is that dealer bids respond asymmetrically to changes in net position that come early vs. late. This we see from the reversal of the sign on the dealer’s net position variable across $T = 10$ and $T = 30$. We do not have a very good explanation for this finding, but one possibility is that the types of when-issued orders that a dealer receives early in the bidding process are from a different set of clientele than those who put in their when-

also has information about how dealers respond to customer bids. Hence, our data specification also includes those instances where the dealer does not change her bid in response to a customer bid modification.

issued orders late in the bidding period. If we reconcile these results with the bid level regressions in Table 2, however, we get the implication that most of the net position information is obtained by the dealer earlier in the bidding period, and that changes within the last 10 minutes are rare events, which appears to be the pattern in our data.

Controlling for the information contained in customer bids, are the bid revisions of dealers who route customer bids any different from bid revisions of dealers who do not route customer bids? The coefficient on *RoutesCustomerBids* in Table 4 suggests that the answer is “No” in bond auctions, but “Yes” in T-bill auctions. Thus, it appears that in bond auctions, the “extra” information source of dealers with customers is their customer bids. In T-bill auctions, however, these dealers may be responding to another information source that we are not controlling for.

B. Dealer Profitability and Access to Order Flow

Why is privately observed order flow information important for dealers? One intuitive answer to this question comes from a common value auction framework, in which the post-auction resale/liquidation value of a given security, v , is common to all bidders, but a priori unknown. The dealer has a prior belief about v , where the prior has mean μ and variance σ_v^2 . Suppose the dealer also observes a private signal about the security value, x , which is distributed about the true value of the security with mean zero error, i.e. $x = v + \varepsilon_x$, where ε is distributed with zero mean and variance $\sigma_{\varepsilon_x}^2$. If all of the aforementioned random variables are distributed normally, conditional on her private information, the dealer expects the value of the security to be:

$$E(v|x) = \frac{\tau_v}{\tau_v + \tau_{\varepsilon_x}} \mu + \frac{\tau_{\varepsilon_x}}{\tau_v + \tau_{\varepsilon_x}} x \quad (6)$$

where $\tau_v = \frac{1}{\sigma_v^2}$ and $\tau_{\varepsilon_x} = \frac{1}{\sigma_{\varepsilon_x}^2}$, are the precisions of the prior and private signals respectively.

If the dealer receives another signal about the security’s true value, in the form of a customer bid that she routes, for example, she will incorporate this information into her conditional expectation.²⁹ If we call this piece of information y , where $y = v + \varepsilon_y$, also distributed with mean zero noise about the true security value, the conditional expectation becomes:

$$E(v|x, y) = \frac{\tau_v}{\tau_v + \tau_{\varepsilon_x} + \tau_{\varepsilon_y}} \mu + \frac{\tau_{\varepsilon_x}}{\tau_v + \tau_{\varepsilon_x} + \tau_{\varepsilon_y}} x + \frac{\tau_{\varepsilon_y}}{\tau_v + \tau_{\varepsilon_x} + \tau_{\varepsilon_y}} y \quad (7)$$

i.e. a simple linear combination of the different pieces of information that the bidder received, weighted by their relative precisions. Furthermore, the variance of this conditional expectation as an estimate of true value of the security, v , is lower when the dealer observes the additional piece

²⁹We will say more about customers’ incentives to reduce the informational content of their bids in Section 4.

of information y , as opposed to when she does not. In particular, the variance of the estimator declines with the number of independent additional signals that the dealer observes.

Seeing customer bids also gives the dealer a “strategic” advantage. To see this, consider the very simple case where there is a single unit of the security being sold, and the customer is the dealer’s only competitor in the auction. To abstract from the Bayesian updating mechanism discussed above, assume also that both the customer and the dealer have independent private valuations for the security (thus seeing the customer’s bid will not affect the dealer’s valuation of the security). Here, seeing the customer’s bid will allow the dealer to bid a very small increment above the customer to win the auction in instances where winning is desirable for the dealer. More generally, a dealer who sees customer bids has better information on the distribution of competing bids (since customers are also competitors). Indeed, in the Appendix, we provide a more general model based on a private value single-unit first-price auction framework that has as its main prediction that dealers who observe customer bids are, on average, more profitable than customers, and, in particular, dealers who do not observe customer bids.³⁰

Despite the different economic mechanisms at work, both explanations have the intuitive prediction that dealers who observe more/better order-flow information, whether in the form of customer bids, or information flows from the when-issued market, should be more profitable bidders. We now turn to empirical tests of this prediction.³¹

Figure 2 compares the bidding patterns of three classes of bidders: customers, dealers who route customer bids, and dealers who do not route customer bids. Here, we plot the distribution of the difference of quantity-weighted price bids from the cutoff-price. Since we plot the distribution for the entire set of bids in our data, which encompasses securities of a wide variety of maturities, we normalized all of the deviations from the cutoff prices using the formula $10000(\frac{360}{\text{Maturity}} (\log(p^{bid}) - \log(p^c)))$, which converts the price differences reported in the data into the (negative) basis point difference of the annual yields implied by the prices.³²

³⁰Engelbrecht-Wiggans, Milgrom and Weber (1983) and Hendricks and Porter (1988) analyze common value first-price auctions with asymmetrically informed bidders. In their model, the informational advantage does not stem from some bidders’ directly observing others’ information, but their model also has the intuitive result that bidders who are better informed are on average more profitable.

³¹One may also want to identify the relative importance of these two mechanisms separately, though we believe that this exercise will entail detailed modelling of the strategic environment faced by the dealers, which we leave for future research. Given that customer bids make up a small portion of bids, one may think that the value of obtaining more information about resale/liquidation values may be larger than the value of the “strategic advantage” gained by a dealer.

³²The formula is based on a conversion of the prices in our bid data into the implied prices of a 360 day zero coupon bond with the equivalent yield. To see this, suppose we have a zero coupon bond with face value P_M paid

As is apparent from the figure, there are stark differences across these three classes of bidders. Customers appear to bid the highest, followed by dealers who see customer bids. Dealers who do not see customer bids appear to bid the lowest prices. Indeed, the density of bids submitted by dealers who see customers appears to be sandwiched between the density of customer bids and the density of bids of dealers without customers – suggesting that dealers who see customer bids use a combination of some “prior” information possessed by all dealers, and the “private” information given by customer bids. Moreover, from this figure, it appears that dealers who observe customer bids are able to place bids that have a tighter spread around the cutoff price of the auction than dealers who do not observe customer bids. In particular, dealers who observe customer bids appear to “underbid” less, and consequently win the auction more often.³³

In Table 5, we regress this average prediction error measure on dealer characteristics, controlling for maturity class fixed effects. The first dealer characteristic we focus on is the number of customers served by the dealer over the auctions within the maturity class. The second dealer characteristic is the “size” of a dealer’s operations, calculated as the (log) average size of the quantity bids placed by the dealer *for its own account* across the auctions within a maturity class.

Table 5 reports a robust negative correlation between the number of customers served by a dealer and the dealer’s prediction error. This correlation appears robust across different subsamples in the data. The point estimate from the first column implies that each additional customer served by a dealer is correlated with price bids that are 8 cents closer to the cutoff price. Interestingly, however, we find a positive correlation between the dealer’s size and the absolute deviation from the cut-off price.

In the second column, we introduce two more proxies for the amount of “order flow information” that the dealer has access to: the (absolute value of the) size of the dealer’s net position, and the (total) size of customer bids that the dealer routes. One might expect that dealers who are more active in the when-issued market, or those who route larger customer orders may have learned more about the resale value of the security being auctioned, and hence will bid more out M days from now. Then, if we see a bid of P_{bid} for this bond, the continuous compounding discount rate of the bidder, r_{bid} , implied by this price can be solved from the formula: $P_{bid} = P_M e^{-r_{bid}M}$, or $r_{bid} = \frac{1}{M} \log(\frac{P_M}{P_{bid}})$. At this discount rate, the price that the bidder would pay for a 1 year zero coupon bond with face value P_{1yr} would have been $P_{bid}^{1yr} = P_{1yr} e^{-360r_{bid}}$. We can do the same to compute the daily discount rate implied by the cutoff price in the auction: $r_{cutoff} = \frac{1}{M} \log(\frac{P_M}{P_{cutoff}})$, and the price for a 1 yr zero coupon bond implied by this discount rate $P_{cutoff}^{1yr} = P_{1yr} e^{-360r_{cutoff}}$. Taking the log difference of these prices: $\log P_{bid}^{1yr} - \log P_{cutoff}^{1yr} = \frac{360}{M} \log(\frac{P_{bid}}{P_{cutoff}})$. Note that this is also equal to: $\log P_{bid}^{1yr} - \log P_{cutoff}^{1yr} = -360(r_{bid} - r_{cutoff})$, the (negative) difference in the implied annual zero-coupon bond yields. Since the magnitudes of this difference are on the order of one-one-hundredth of a percent, i.e. “basis points,” we scale up by 10000 to get a (negative) basis point equivalent.

³³Similar patterns were observable when we plotted the corresponding bid distributions for bond and bill auctions.

successfully. However, the regression shows only a weak negative correlation between net position size and bid prediction error, and almost no correlation with customer size (though the number of customers served is still negatively correlated with the bid prediction error). The result that (total) customer bid size is not as important as the number of unique customer bids that a dealer routes, seems consistent with the Bayesian updating hypothesis, where information from different sources is very important for improving forecasting accuracy. We should, however, be cautious about interpreting the weak result regarding the importance of net positions – consistent with the Bayesian updating hypothesis for customer bids size, the number of unique customer interactions in the when-issued market (for which we do not have data) may well be more important than the size of the net position.

In the third and fourth columns, we run the regression in the second column separately for T-bill and bond auctions. The main result from these regressions is that “number of customers served” is a more important driver of bidding performance in bond auctions than in T-bill auctions. This finding is consistent with our earlier results that order flow information appears more important in bond vs. T-bill auctions.

The next set of results in Table 5 report the regressions for two alternative measures of dealer bidding performance. The first measure we utilize is the difference between the dealer’s (quantity-weighted average price) bid and the cutoff price in auctions where the dealer bids *above* the cutoff price, i.e. this is a measure of the amount of “overbidding.”³⁴ We note that the correlation and explanatory power of the independent variables is low for this performance measure, except in bond auctions, where larger dealers appear to “overbid” less, though dealers with larger net positions appear to “overbid” more.³⁵

The second measure we utilize is the difference between the dealer’s bid and the cutoff price in auctions where the dealer bids *below* the cutoff price, i.e. a measure of the amount of “underbidding.” The results of the three regressions using this dependent variable are very similar to those when we used absolute deviations as the measure of bidding performance.

Consistent with the patterns in Figure 2, these last two sets of results indicate that customer information aids the dealer mostly through reducing the amount the dealer “underbids” in the auction. Since, by underbidding, the dealer wins a lower quantity of the securities she was demanding, this result suggests that customer information allows dealers to win more frequently. “Overbidding,” on the other hand, indicates how much a dealer pays over the market-clearing price for the quantity of securities he wins. Thus, our results show that dealers who serve a larger

³⁴We set this variable equal to zero when the dealer bids below the cutoff price.

³⁵The reported R^2 ’s reflect the explanatory power of maturity fixed effects.

number of customers consistently win a large fraction of the securities they were demanding, and do not consistently overpay for the units they win. If dealers have similar willingness-to-pay for the securities, this would indicate that dealers who see a larger number of customers are more profitable ex-post.

IV. Evidence for Strategic Response by Customers

So far we have considered only the dealer-side of the customer-dealer interaction. One might, however, wonder whether customers recognize and react to the fact that dealers may utilize the information from these interactions. We now describe and discuss the empirical evidence for several possible customer responses.³⁶

A. Customers Use Multiple Dealers

One potential strategic response by customers to counter dealers' informational advantage is to route their bids through multiple dealers. To test this hypothesis, for each customer in our sample, we obtain the average number of dealers through whom she routes his bids in an auction.³⁷ If this average is less than 1.5, we define this customer as a "single-submitter" customer as he routes bids through a single dealer in most but not all auctions in the sample. If this average is 1.5 or more, then we classify this customer as a "multiple-submitter" customer. Table 6 indicates that while "single-submitter" customers are in the majority, 24 (26 in bond auctions) customers in our sample use multiple dealers.³⁸

We now investigate why multiple-submitters are utilized. A first hypothesis is that a customer who needs to purchase a large amount of securities at the auction routes his bid through more than one dealer. Column (1) of Table 7 reports the regression of the number of submitters utilized by a customer on the proportion of this customer's bid amount to the total issue amount in the auction. Customers use an additional dealer to submit tenders for a 36% point increase in the

³⁶The focus is primarily customer-dealer interactions in the primary market as our data does not allow us to identify a customer and her individual transactions with a dealer in the when-issued market.

³⁷The average is across the auctions in which the customer participates (we do this separately for bond and T-bill auctions). Importantly, the number calculated here is not the average number of *unique* dealers utilized by the customer over the sample. This is analyzed in Section 4.C.

³⁸We should note that 82% of the "single submitter" customers participated in less than 20% of the auctions in our data set, and can be regarded as "passive" customers; while only 48% of the "multiple submitter" customers can be regarded as being "passive." So, the relevant breakdown within "active" customers is that 13 distinct customers are "multiple submitters," whereas 17 distinct customers are "single-submitters," i.e. the fraction of active customers who are multiple submitters is much higher.

ratio of bid amount to issue amount.

An additional test is based on our previous observation that order flow plays a more important role in bond auctions than in treasury bill auctions. This should also affect customers' behavior: the tendency to use a larger number of dealers in response to quantity demanded should be much more pronounced for bonds than treasury bills. Thus, we rerun the regression in Column (1), Table 7 by interacting the ratio of bid amount to issue amount with an indicator for bond auctions. Results in Column (2), Table 7 show that this is the case. For bond auctions, customers use an additional dealer to submit tenders for a smaller increase in the ratio of bid amount to issue amount compared with treasury bill auctions, and the difference is statistically significant.³⁹ But this result could also suggest that customers demand is higher in bonds relative to treasury bills, and therefore it is the "submission limit" effect that leads them to submit bids through multiple dealers. We find the reverse: the ratio of bid amount to issue amount is 10% higher for treasury bills than bonds, and this difference is statistically significant.

An alternative explanation to these findings is that customers may be routing bids through several dealers to circumvent a dealer's submission limit. As we explained in Section 2, a dealer is subject to a maximum constraint on the amount he can submit for customers in an auction. Thus, when customers have a large quantity of security to buy, they will tend to use multiple dealers since a single dealer may not have enough capacity to route their bid.

However, there is some reason to believe that this alternative explanation does not fully account for the use of multiple dealers. In particular, we find that not all customers use multiple dealers; some use a single dealer. If the "submission limit" effect was indeed the only driving reason to using multiple dealers to submit one's bids through, then controlling for the quantity demanded by the customer, we should find that the slope coefficient of our regression in Table 7, Column (1) should not be statistically different between "single-submitter" and "multiple-submitter customers," as defined in Table 6. Recall that the slope measures the average change in the number of submitters used by a customer due to a change in the amount demanded by this customer in an auction scaled by the issue amount in that auction. We actually find the reverse. We re-run the regression in Column (1) of Table 7 by interacting the ratio of bid amount to issue amount with a dummy that equals 1 if the tender is submitted by a multiple-submitter customer, and zero otherwise. The sample consists of official tenders of customers in auctions where they use two or more dealers as submitters. Column (3), Table 7 reports the results of this

³⁹For bond auctions, the regression estimates mean that customers use an additional dealer to submit tenders for a 24% increase in the ratio of bid amount to issue amount, and for treasury bill auctions they use an additional dealer for a 40% increase in the ratio of the bid amount to issue amount.

regression. While multiple-submitter customers use an additional submitter for a 55% increase in the proportion of bid amount to issue amount, single-submitter customers actually decrease the number of submitters with an increase in the proportion of bid amount to issue amount, with the difference being statistically significant at the 99% level. These results suggest that while single-submitter customers route bids through an additional dealer when the dealer hits his “submission” limit constraint, multiple-submitter customers route bids through several dealers due to both the “submission limit” effect and the strategic “demand concealing effect.”

B. Long Term Dealer-Customer Relationships

The agency problem between customers and dealers may be resolved through contractual means. Although we are not aware of explicit contracts written between customers and dealers, long-term relationships may be an equally powerful remedy. Our data supports the existence of long-term relationships. Table 8 lists the 7 customers we have identified as having long-term relationships with their dealers. For example, customer A is a single-submitter customer who has a long-term relationship with dealer 1 in treasury-bill auctions. Customer A participates in 25% of the treasury bill auctions held in the sample period. In 78% of these treasury-bill auctions in which customer A participates, she uses dealer 1 as the submitter.⁴⁰

Absent any information on explicit or implicit agreements between customers and dealers, it is difficult to say too much about how these long-term relationships work. However, we can examine our data to see if customers who are in long-term relationships bid differently than those who are not. In Table 9, Columns (1) and (2), we look at the difference between the (quantity-weighted) price bids of customers who are in long-term dealer relationships vs. those who are not. Column (1) reports the results of a regression of the difference of quantity weighted average price and the cutoff price on the long-term dealer indicator.⁴¹ Customers in long-term relationship, when

⁴⁰To identify customers in long-term relationships with their dealers, we concentrated on customers in row 1 in Table 6; we classified these customers as “single-submitter” customers. However, not all of these single-submitter customers are active participants – in fact 82% participated in less than 20% of the auctions. When we eliminated these “passive” customers, we ended up with 17 distinct active “single-submitter” customers in our sample. For each of the single-submitter customers, we constructed the ratio of number of auctions in which a specific submitter was used, to the total number of auctions in which the customer participated, for each distinct submitter used by a single-submitter customer. This is referred to as a customer’s *submitter proportion*. For each customer, we tested if the highest and second-highest submitter proportion is significantly different. Customers for whom this is the case are customers with long-term relationship with a specific dealer in that they submit most of their tenders in the sample period through a distinct dealer.

⁴¹To construct the sample for this regression, we remove all tenders submitted by the customer through the dealer with whom he is not in a long-term relationship. This gives us allotted tenders of customers who are in long-term

submitting tenders through the dealer with whom they are in long-term relationship, bid 0.5 cents more than customers who are not in long-term dealer relationship, and this difference is significant at the 99% level. In Column (2), Table 9 we run the same regression, but now introducing the distinction between treasury bills and bond auctions. In both the treasury bill and bond sector, customers in long-term relationship when they submit tenders through the dealer with whom they are in long-term relationship, pay more compared with customers who are not in long-term relationship. For bonds (treasury bills), the former pay 6 (5) cents more than the latter, and this difference is statistically significant.

Why do customers in long-term relationship pay a higher price for Government of Canada securities when submitting tenders through the long-term dealer, compared with customers who are not in a long-term relationship? Long-term customers and dealers interact with each other in sectors other than Government of Canada securities. Thus, it is conceivable that the payoffs to the long-term customers are being given in these sectors by the long-term dealer. Alternatively, it is conceivable that a customer enters a long-term relationship with a dealer because institutional reasons specific to a customer's business push him to the inelastic part of his demand curve relative to customers who are not in a long-term relationship.⁴² Indeed, there is some evidence that although customers in long-term dealer relationships may be different, they do benefit from established relationships. In column (3), we restrict attention to the winning bids of customers who are in a long-term relationship, only in the sectors (T-bills or bonds) where there is a long term relationship. These tenders could be submitted either through the dealer with whom the customer is in a long term relationship, or through the dealer with whom the customer is not in a long term relationship. The regression indicates that customers in a long-term relationship pay 0.5 cents less when they submit tenders through their long-term dealer compared with other dealers (the difference is significant at the 89% level).

C. Last Minute Bidding by Customers

In Section 2.B we observed that while customer bids lag dealer bids, customer bids come quite close to the submission deadline as well. A possible reason for late bidding by customers is that customers are attempting to conceal their demand curves from dealers, to prevent them from

relationship, submitted through the dealer with whom they are in long-term relationship, in the sector in which the customer is in a long-term relationship. As the control sample, we add allotted tenders submitted by customers who are not in a long-term relationship in treasury bill and bond auctions.

⁴²For example, a customer that is a pension fund could be legally required to hold a certain proportion of its portfolio in the form of Government of Canada securities. This may not be the case for another customer that is a hedge fund.

free-riding on their information. Since dealers are under legal obligation to route customer bids as soon as they receive them, customers may try to bid very close to the deadline in order not to leave enough time for dealers to revise their bids. Moreover, the tendency to delay interaction with the dealer need not be confined to the primary market. Customers may delay transacting with dealers in the when-issued market to prevent a dealer from submitting a revised bid before the bid submission deadline.⁴³

To clarify this intuition, let $[T - \Delta, T], \Delta > 0$ denote the “last time interval” before the close of the auction at time T .⁴⁴ Suppose any bid transmitted within this time interval is accepted by the BoC system with probability $p < 1$. In this situation, a customer who wants to lessen the advantage possessed by her dealer can submit her bid exactly at time $T - \Delta$. At this point, the dealer can try to submit an updated bid, which is accepted with probability p . However, just in case this updated bid is not accepted, the dealer can also place an “insurance” bid anytime before $T - \Delta$. Without any customer information, this “insurance” bid by the dealer will be equal to the dealer’s optimal bid without access to customer information.

We now look at late tenders in our data set to investigate the plausibility of this hypothesis. Throughout, *late bids* will refer to bids that were rejected by the issuer for being submitted after the submission deadline. In case customers bid close to the bid submission deadline to prevent dealers from revising their bids on seeing customer bids, we should find that most of the late bids are dealer bids, especially when compared with bids originating from dealers prior to the bid submission deadline.⁴⁵ We find this to be the case: 77% of the late tenders are submitted by primary dealers and 13% by customers,⁴⁶ with only 48% of the official bids originating from dealers and 33% from customers. Moreover, we have ample evidence that late dealer tenders are immediately preceded by dealer-customer interactions in both the primary and when-issued markets,⁴⁷ suggesting that last-minute when-issued trading and primary market bidding may

⁴³We thank our referee for pointing out this possibility.

⁴⁴If a bid is interpreted generically as a customer transaction with a dealer, this intuition can be extended to the when-issued dealer-customer interactions as well.

⁴⁵We thank a referee for pointing this out to us.

⁴⁶The difference is statistically significant at the 99% level.

⁴⁷30% of late dealer tenders are preceded by a customer tender in the 5-minute interval prior to the late tender. Another 20% of late dealer tenders report a “new” net position compared with the official dealer bid immediately preceding the late tender. The remaining 50% of the late dealer tenders are neither accompanied by a “new” net position nor preceded by late tenders. But we do find that these late tenders originate from dealers with **no** customers in the auction where the late tender is submitted. In conjunction with the requirement by the Bank of Canada to report net position changes only in excess of \$25 million, this suggests that late bids from dealers without customers are likely the result of customers delaying their transactions with dealers in the when-issued market till just before the bid submission deadline.

indeed be a customer response.

We now show that late bids, if accepted, would have had an important impact on the auction outcome. Since a dealer’s late bids are “informed” bids compared with the “insurance” bid placed anytime before $T - \Delta$, one would expect that if the late bids of dealers were accepted rather than rejected, the auction outcome would be different, in a manner that is more profitable to the dealer. To test this, we calculated the counterfactual auction cutoff price and the resulting allocation *had* the late bids been accepted. Notice that a late bid would affect the cutoff price or the allocation only if it is marginal or inframarginal. We found that the late bid had an impact on the auction outcome in 111 auctions where a late bid was received. However, these late bids had very little effect on the revenue of the seller – taking the late bids into account, the GoC would have received, on average, only an additional 767 Canadian dollars per auction!⁴⁸ This is also borne out by the fact that the quantity weighted average price is really unchanged between the real auction and the reallocated auction.⁴⁹

These last set of findings imply that the late bids, when successful, are right at the margin, and end up displacing other marginal bids (if the late bids had been at the top end of the aggregate bid curve, i.e. if they had been inframarginal, we would have expected larger changes in revenue and quantity-weighted average price). Moreover, within each of these 111 auctions, we find that there is a large amount of redistribution of the allocated securities. On average three dealers are affected in an auction, with the number of bidders who win more units in the reallocated auction being almost equal to those who win less. In the auction where we observe the maximum redistribution with the inclusion of late bids, one primary dealer ends up winning 398 million Canadian dollars more of the auctioned security! Of course, this reallocation would have come at the expense of other dealers – had this late bid been accepted in this auction, 3 other dealers would have each won 100 million dollars less of the security. Moreover, the redistribution of rewards is consistent with our intuitive prediction outlined in Section 3.C that more/better order-flow is profitable for the dealer – the winning dealer’s late tender is immediately preceded by a customer bid and when-issued trading, which is not the case for the three dealers who lose!⁵⁰

⁴⁸This amount is also not statistically significant with $Z=0.0002$.

⁴⁹The change in price, averaged across the 111 late auctions is 0.004 cents, and $Z=1.2$.

⁵⁰The winning dealer’s late tender is accompanied by a change in his long position that is 25% lower than that reported with his last accepted tender submitted 3 minutes before the bid submission deadline suggesting that he shorts securities in the when-issued market in the last 3 minutes. In addition, the last accepted bid of this dealer (submitted less than a minute before deadline) is one that he has routed for a customer. In contrast, two of the three losing dealers do not route any customer bid in the auction, while the third dealer route’s a customer bid only 10 seconds before the bid submission deadline circumventing this dealer from even submitting a late bid. We also do not observe significant net position changes for them.

V. Discussion and Conclusions

We conclude that in Government of Canada securities auctions, where direct access is restricted to authorized security dealers, privately observed order-flow information is an important and valuable driver of bidding behavior, and thus the formation of asset prices. As we have demonstrated, the importance of order-flow information in this market manifests itself in a rich array of strategic responses by both dealers and customers that are difficult to reconcile through other means.

Given our empirical results so far, one may wonder whether the current market rules, which provide an informational advantage to dealers, are undesirable from the point of view of achieving economic objectives such as maximizing revenues or allocational efficiency. In the Appendix to this paper, we use a simple, stylized model of the bidding environment to argue that, although the current mechanism is likely to distort the efficiency of the allocation, its impact on revenues is likely to depend sensitively on the particular economic environment. The main intuition underlying this result is that the informational advantage possessed by (some) dealers introduces asymmetry into the bidding strategies of (ex-ante symmetric) customers and dealers. This intuition follows from the work of Maskin and Riley (2000), who have shown that asymmetric valuation structures in first-price auctions can lead to loss of efficiency, and ambiguous revenue rankings. What we show through our modelling exercise is that a different source of asymmetry, which comes from the ability of some dealers to observe the information of others, leads to similar results.

The modelling exercise in the Appendix only considers the case where the source of informational advantage is “dissipated” by allowing direct access to customers to place bids in the primary market. Another possibility is to impose some form of transparency obligations on the authorized security dealers with respect to their secondary market activity. For example, in Italy, the secondary market is transparent in that it is a centralized, regulated electronic screen-based market. However, only authorized dealers are allowed to place bids in primary issuance. Still, Drudi and Massa (2001) show how authorized dealers use the discrepancy in transparency to obtain government securities in the less transparent primary market at below-market prices. In the U.S., customers are allowed to place bids in the primary issuance, making the primary market transparent. But unlike Italy, the secondary market is largely over-the-counter, with the customer-dealer interaction protocol requiring the customer to reveal his intention to buy or sell when he requests a quote. The Italian and the U.S. comparison, along with several other countries examined by Sareen (2003), highlight that a mechanism for issuing government securities appears to retain privacy of a security dealer’s “order-flow” in at least one of the markets in which the dealer is participating. Which one will be less costly for the issuer? Which one will increase participation in the primary issuance? These questions will be explored in future research.

VI. References

- Bajari, Patrick, and Ali Hortagsu, "Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay," *RAND Journal of Economics*, Summer 2003, pages 329-355.
- Biais, Bruno, Pierre Hillion, and Chester Spatt, "Price Discovery and Learning During the Pre-Opening Period in the Paris Bourse," *Journal of Political Economy*, 1999, 107(4), pages 1218-1247.
- Binmore, Kenneth and Joseph Swierzbinski, "Treasury auctions: Uniform or discriminatory?" *Review of Economic Design*, 2000, vol. 5, issue 4, pages 387-410.
- Bikhchandani, Sushil, Patrick L. Edsparr, and Chi-Fu Huang, "The Treasury bill auction and the when-issued market: Some evidence," UCLA working paper, 2000.
- Brandt, Michael W. and Kenneth A. Kavajecz, "Price Discovery in the U.S. Treasury Market: The Impact of Orderflow and Liquidity on the Yield Curve," *Journal of Finance*, 2004, 59, 2623-2654.
- Brunnermeier, Markus, *Asset Pricing under Asymmetric Information - Bubbles, Crashes, Technical Analysis and Herding*, Oxford University Press, 2001.
- Drudi, Francesco and Massimo Massa. "Asymmetric Information and Trading Strategies: Testing Behavior on the Primary and Secondary T-Bond Markets around Auction Days," working paper, 2001.
- Easley, David, Maureen O'Hara, and P. S. Srinivas. "Option Volume and Stock Prices: Evidence on Where Informed Traders Trade." *Journal of Finance*, Vol. 53, No. 2. (Apr., 1998), pp. 431-465.
- Engelbrecht-Wiggans, Richard, Paul R. Milgrom, and Robert J. Weber, "Competitive Bidding and Proprietary Information," *Journal of Mathematical Economics*, 1983, 11, 161-169.
- Evans, Martin D.D. and Richard K. Lyons, "Order Flow and Exchange Rate Dynamics," *Journal of Political Economy*, February 2002, 170-180.
- Fevrier, Philippe, Raphaelae Preguet, and Michael Visser. "Econometrics of Share Auctions," working paper, 2002.
- Fleming, M.J., "Are Larger Treasury Issues More Liquid? Evidence from Bill Reopenings," *Journal of Money, Credit and Banking*, 2002, Vol 34, No. 3, 709-735.
- Fleming, M.J., "Measuring Treasury Market Liquidity" *Federal Reserve Bank of New York Staff Report No. 133.*, 2001.
- Green, T. Clifton, "Economic News and the Impact of Trading on Bond Prices," *Journal of Finance*, 2004, 59(3), 1201-1234.
- Gordy, Michael B. "Structural Analysis of Reservation Price Policy in a Treasury Bill Auction," working paper, September 1994.
- Hansch, O., N. Naik and S. Viswanathan, 1998, Do Inventories Matter in Dealership Markets? Evidence From The London Stock Exchange, *Journal of Finance*, 53, 1623-1656.
- Hendricks, Kenneth and Robert H. Porter, "An Empirical Study of an Auction with Asymmetric Information," *American Economic Review*, 1988, 78, 865-883.

- Hortaçsu, A. "Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market," working paper, 2002.
- Ito, Takatoshi, Richard K. Lyons, and Michael T. Melvin. "Is There Private Information in the FX Market? The Tokyo Experiment." *Journal of Finance*, June 1998, 1111-1130.
- Keloharju, Matti, Kristian Rydqvist, and Kjell G. Nyborg, "Strategic Behavior and Underpricing in Uniform Price Auctions: Evidence from Finnish Treasury Auctions," *Journal of Finance*, forthcoming.
- Lyons, Richard K., "Tests of Microstructural Hypotheses in the Foreign Exchange Market," *Journal of Financial Economics*, 39, 321-351.
- Madhavan Anand and Smidt, Seymour, "A Bayesian Model of Intraday Specialist Pricing," *Journal of Financial Economics*, 30, 99-134.
- Maskin, Eric and John Riley, "Asymmetric Auctions," *Review of Economic Studies*, 2000, 67(3), 381-411.
- Massa, Massimo and Andrei Simonov, "The Sneaky, the Sleepy and the Skeptic: a Behavioral Model of Market Making: Evidence of Strategic Market Making on the Treasury Bond Market," working paper, December 2001.
- Massa, Massimo and Andrei Simonov. "Reputation and dealers' trading," *Journal of Financial Markets*, vol. 6 (2003), pp. 99-141.
- O'Hara, Maureen, *Market Microstructure Theory*, Blackwell Publishers, 1995.
- Nyborg, Kjell G. and Ilya Strebulaev. "Multiple Unit Auctions and Short Squeezes," *Review of Financial Studies* 17, 849-877, 2004.
- 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.
- Nyborg, Kjell G., Kristian Rydqvist, and Suresh M. Sundaresan. "Bidder Behavior in Multi-Unit Auctions - Evidence from Swedish Treasury Auctions." *Journal of Political Economy*, Vol. 110, April 2002.
- Roth, Alvin E. and Axel Ockenfels. "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet," *American Economic Review*, 92(4), pages 1093-103.
- Sareen, S., Cross-Country Comparison of Models for Issuance of Government Securities, Unpublished manuscript, Bank of Canada, 2003.
- Sareen, S., Commodity Bundling in Government Securities Auctions, Unpublished manuscript, Bank of Canada, 2005.
- Sareen, S., Auctioning Government of Canada Securities: Can we do better, Unpublished manuscript, Bank of Canada, 2001.
- Simon, David P., "Markups, quantity risk, and bidding strategies at Treasury coupon auctions," *Journal of Financial Economics*, 35, 43-62.
- Umlauf, Steven. "An Empirical Study of the Mexican Treasury Bill Auctions," *Journal of Financial Economics*, 1993, 33, 313-340.
- Viswanathan, S. and James Wang. "Auctions With When Issued Trading: A Model of the U.S. Treasury Markets," working paper.

A. Appendix: Revenue and Efficiency Consequences of Current Market Rules

We now analyze the simplest possible example where there is a single dealer, D , and her customer, C , who are competing for a single unit of the security, and where both of these agents have symmetric independent private valuations for the auctioned security.⁵¹ Both the dealer and the customer have symmetric, independent private values drawn iid from the uniform $[0, 1]$ distribution. However, as in the GoC auction market, the dealer can observe the customer's bid.

Call the customer's valuation v_c and her corresponding bid, b_c . Similarly, call the dealer's valuation v_d and her corresponding bid b_d . The unique symmetric Bayesian Nash equilibrium of this game when the dealer does not observe the customer's bid is $b_d(v) = b_c(v) = \frac{v}{2}$. When the dealer can see the customer's bid, the dealer will bid $b_d = b_c + \varepsilon$ whenever $v_d > b_c$, where ε is a small increment to outbid the customer. If $v_d < b_c$ the dealer knows he does not want to win, so he can place any undominated bid (i.e. any bid below v_d).

Knowing the dealer's strategy, the customer's best response problem is to maximize:

$$\max_{b_c} (v_c - b_c) \Pr(b_c > v_d)$$

since the customer only wins the auction when $b_c > v_d$. Taking the first-order condition, we get $b_c = \frac{v_c}{2}$, i.e. the customer's strategy is the same as in the case where the dealer does not benefit from the customer.

Let us now see how the asymmetry between the dealer and the customer affects the revenues to the auctioneer. Since valuations are drawn iid uniform $[0, 1]$, the expected revenue from the first-price auction without the dealer advantage is $\frac{1}{3}$, the expected value of the second highest valuation (by the revenue equivalence theorem). When the dealer has informational advantage, however, the auctioneer's revenue is $E[\frac{v_c}{2}]$, since the winning bid will always equal $\frac{v_c}{2}$ (disregarding the ε , which can be arbitrarily small). But with uniform $[0, 1]$ valuations, the expected revenue will be $\frac{1}{4}$, i.e. the auctioneer loses revenue from requiring customers to bid through dealers.

Moreover, the informational advantage given to the dealer introduces allocational inefficiency. Efficiency requires that the good be allocated to the agent with the highest valuation. However, with the dealer-customer arrangement, there exists cases where $b_C^* < v_D < v_C$ i.e. cases where the dealer wins the auction even though the customer has the higher valuation.

A. Two Dealers, One Customer

The second example we examine, however, shows that the revenue loss prediction is not robust (though the efficiency loss prediction is). In this example, we introduce a second dealer to the model, with the twist

⁵¹In a common value framework, there would be an additional informational advantage to observing a customer's bid – the dealer could alleviate her winner's curse by combining her own private information and the informative content of her customer. Modelling the strategic interaction in a common value environment fully is much more difficult. Therefore, we will restrict attention to the private values environment. Note, however, that the private value environment is able to generate many interesting predictions that we see in the data.

that this dealer does not have a customer whose bid she can observe (however, the bid of this dealer is not observed by the other dealer). In this example, the “advantaged” dealer with the customer is once again the most profitable party, followed by the non-advantaged dealer, and finally by the customer. However, the auctioneer’s revenue actually increases in this example.

This slightly more complicated setup is as follows: suppose, in addition to the dealer and customer, there is another dealer in this auction, A , but this dealer does not have a customer. Let v_a be the (iid) valuation draw of this dealer, and b_a her corresponding bid.

In this case, D faces some uncertainty regarding her chances of winning the auction even upon seeing the bid of C , since she still has to outbid A . Therefore, D ’s expected profit from winning the auction, upon seeing b_c is:

$$\max_{b_d} (v_d - b_d) \Pr(b_d > b_a | b_d > b_c) \quad \text{if } v_d > b_c \quad (8)$$

We now assume that A follows a linear strategy, $b_a = \alpha v_a$ (we will verify this assumption later when solving A ’s problem). The optimal bidding strategy of D is given by:

$$b_d = \frac{v_d + b_c}{2} > b_c \quad \text{if } v_d > b_c$$

and any bid $b_d \leq v_d$ if $v_d < b_c$.

The optimal bidding problem of the customer is:

$$\max_{b_c} (v_c - b_c) \Pr(\{b_c > b_d\} \cap \{b_c > b_a\})$$

which is

$$\max_{b_c} (v_c - b_c) \Pr(b_c > v_d) \Pr(b_c > b_a)$$

by independence and D ’s bidding strategy above, and:

$$\max_{b_c} (v_c - b_c) \frac{b_c^2}{\alpha}$$

under the linear strategy assumption for A . Thus, the customer’s bidding function is given by:

$$b_c = \frac{2v_c}{3}$$

which is the Bayesian Nash strategy with 3 symmetric bidders.

Finally, A ’s optimal bidding problem is given by:

$$\max_{b_a} (v_a - b_a) \Pr(b_a > \max\{\frac{2v_c}{3}, \frac{v_c}{3} + \frac{v_d}{2}\})$$

when we substitute in D and C ’s strategies. Graphical inspection yields:

$$\Pr(b_a > \max\{\frac{2v_c}{3}, \frac{v_c}{3} + \frac{v_d}{2}\}) = \frac{9b_a^2}{4}$$

and A ’s optimal bidding strategy can be solved to be:

$$b_a = \frac{2v_a}{3}$$

which is once again the Bayesian Nash strategy with 3 symmetric bidders. Notice that in this last step, we did not impose the linearity of A 's bid function, hence we have verified the linearity assumption made in the earlier steps.

The fact that the customer C and customer-less dealer A have the same bidding strategies may appear counterintuitive at first; however, these strategies are not outcome equivalent (since D acts upon C 's bid). Moreover, it appears that D benefits the most from this arrangement. A and C are at a disadvantage compared to the setting with 3 symmetric bidders, with C being the most disadvantaged party. In particular, our Monte Carlo simulations (with 1,000,000 draws of iid bidder valuations) show that D wins the auction 45% of the time, A wins the auction 33% of the time, and C wins the auction 22% of the time (instead of each winning with probability $\frac{1}{3}$). D 's expected profit is 0.0868 compared to 0.0833 in the symmetric case, A 's profit is 0.0810, and C 's profit is 0.056. Interestingly, the auctioneer appears to gain slightly from this arrangement in terms of revenues: the expected revenue is 0.5089, compared to 0.5 in the 3 symmetric bidders case. However, this revenue gain is a result of an efficiency loss, which happens when D wins the auction in instances where C had the highest valuation.

Figure 1: CDF of Bid Submission Times

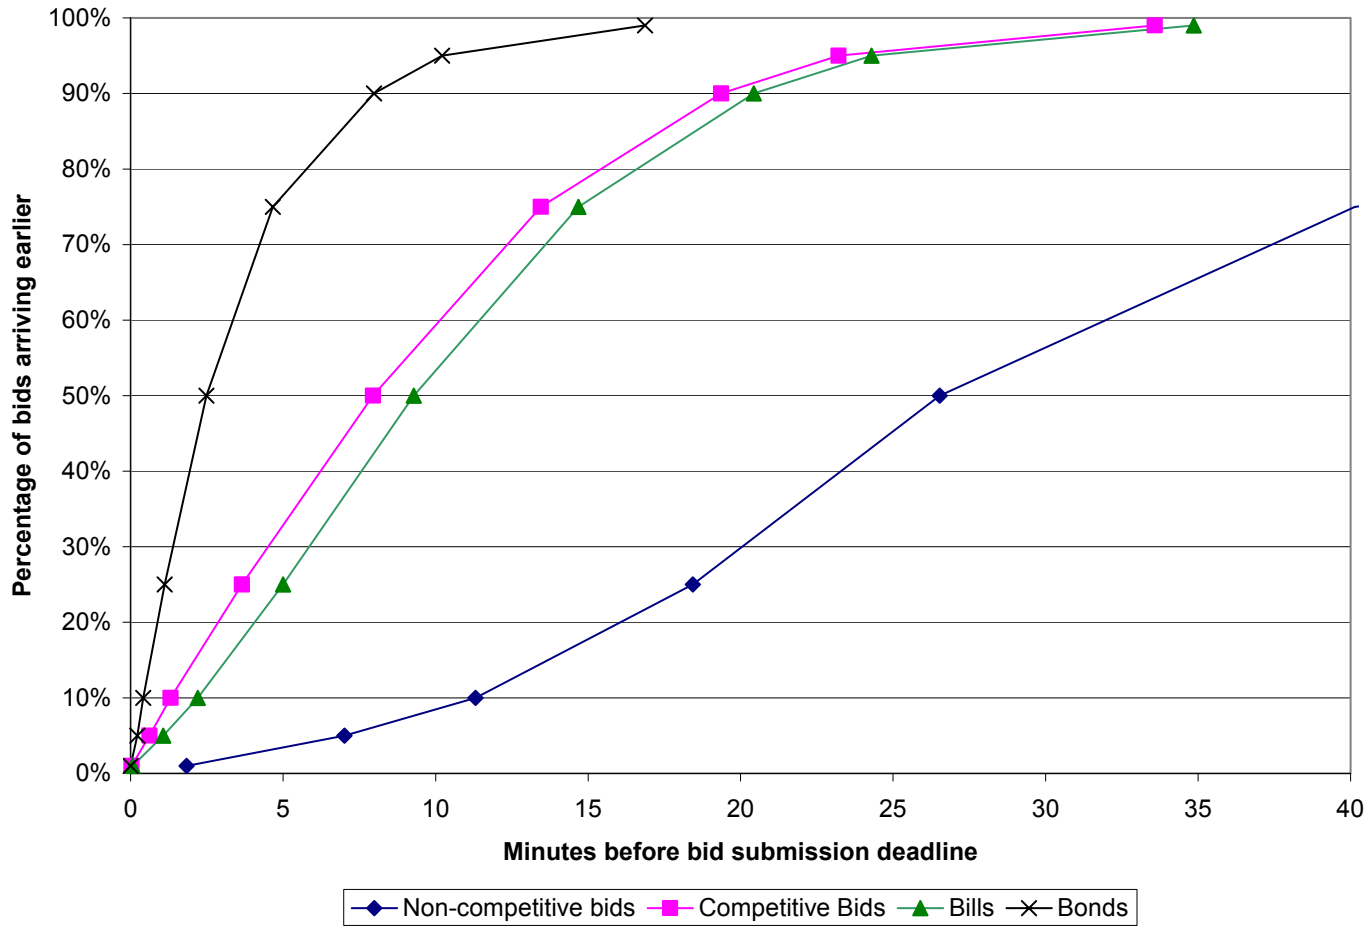
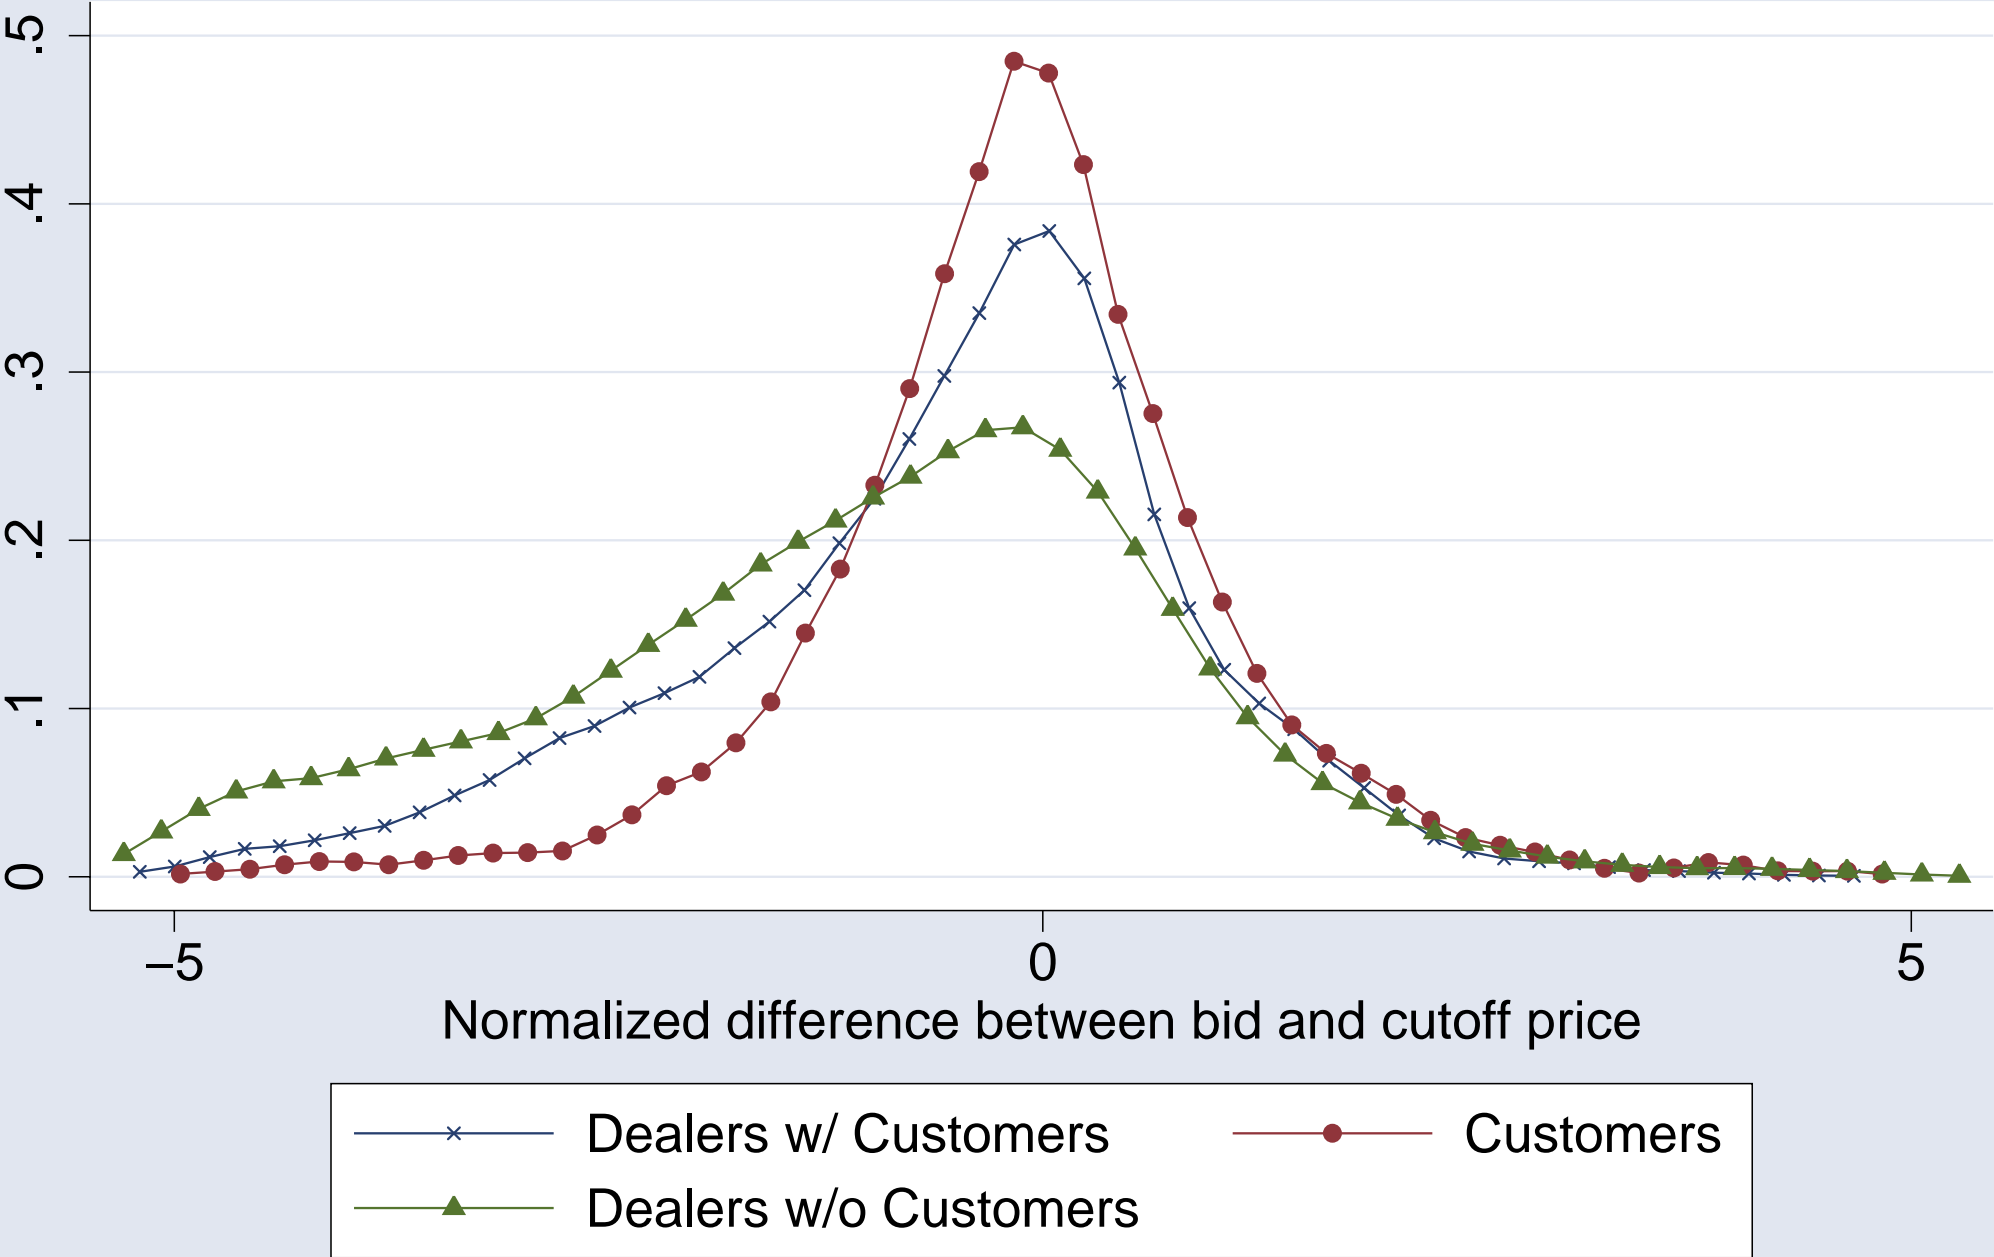


Figure 2: Distribution of Bids Around Cutoff Price



"Normalized" difference can be interpreted as the (negative) basis point difference. See fn. 32, pg.18.

Table 1: Dispersion of Bids Across Time

This table documents the dispersion of bids submitted at different time intervals. The dependent variables are (a) the absolute dispersion of (quantity-weighted) price bids around the market clearing price, and (b) the standard deviation of price bids that are submitted within a given time interval. (Thus, the unit of observation is auction/time-interval). As independent variables, we include indicators for 5 minute time-intervals, with bids submitted earlier than 30 minutes to the deadline as the excluded group. We also control for bid volume in the time interval, and include auction fixed effects to isolate within auction variation. Standard errors are reported below the coefficients. Significance at 10-percent, 5-percent and 1-percent levels are denoted by (*),(**), and (***)

	(a) Absolute Dispersion of Price Bids Around Market Clearing Price		(b) Standard Deviation of Price Bids	
	Bills	Bonds	Bills	Bonds
25-30 min to deadline	0.0001 (0.0009)	-0.22 (0.22)	0.002 (0.003)	-0.14 (0.24)
20-25 min to deadline	-0.0002 (0.0008)	-0.65 (0.19)***	0.003 (0.002)	-0.58 (0.19)***
15-20 min to deadline	0.0007 (0.0008)	-1.09 (0.19)***	0.004 (0.002)**	-0.48 (0.17)***
10-15 min to deadline	0.0018 (0.0008)**	-1.20 (0.19)***	0.006 (0.002)***	-0.55 (0.16)***
5-10 min to deadline	0.0025 (0.0008)**	-1.43 (0.23)***	0.009 (0.002)***	-0.71 (0.19)***
0-5 min to deadline	0.0015 (0.0008)*	-1.67 (0.35)***	0.008 (0.002)***	-0.74 (0.26)***
Bid volume in interval	-0.0001 (0.00002)***	0.003 (0.002)	-0.0001 (0.00004)***	0.0003 (0.001)
Auction Fixed Effects	Y	Y	Y	Y
N	2434	375	2029	300
R-sq	0.12	0.22	0.05	0.27

Table 2 (a): Comparison of the Timing of Dealer vs. Customer Bids

“Submission time of dealer bid” reports the time, in minutes, before the auction deadline at which the average official dealer bid is submitted. The standard deviation is reported in parantheses below. “Submission time of Customer Bid” measures the time, in minutes, before the auction deadline at which the official customer bid is submitted. The “paired-difference” is the result of a pairwise t-test between the bid submission time of a dealer and a customer she serves. We report the p-value and standard errors for this test below. (***) indicates significance at the 1% level.

	All Maturities	Bills	Bonds
Average submission time of dealer bids (min. from deadline)	8.64 (7.92)	9.59 (8.03)	2.27 (2.15)
Average submission time of customer bid (min. from deadline)	9.08 (6.81)	9.71 (6.87)	4.81 (4.53)
Paired difference between dealer-customer bid submission time (min.)	-0.43***	-0.12	-2.54***
Std.err.	0.20	0.23	0.29
P-value	0.017	0.3	0
% of times dealer bid Precedes customer bid	55.29%	52.39%	74.90%
Number of comparisons	2042	1779	263

Table 2 (b): Comparison of Bid Submission Timing of Customers, Dealers without Customers, and Dealer with Customers

The regressions reported in this table analyze the timing of all official, competitive bids. The dependent variable in the regressions is the time, in minutes, before the auction deadline at which an official bid is submitted. We categorize each official bid as being (i) the “personal” bid of a dealer who routes customer bids, (ii) the “personal” bid of a dealer who does not route customer bids, and (iii) a customer bid. The first category is excluded from the regression. We control for bid size and auction fixed effects.

	T- Bill Auctions	Bond Auctions
Dealers without customers	-2.06 (1.84)	1.32 (0.55)**
Customers	0.24 (0.70)	2.18 (0.28)***
Bid Quantity	-0.92 (0.23)***	-0.34 (0.07)***
Auction Fixed Effects	Y	Y
N	6490	1360
R-sq	0.07	0.18

Table 3: Do Customer Bids Drive Variation Across Dealers' Bids?

Each column reports the estimated coefficients from a regression in which the dependent variable is the difference between the dealer's official (quantity weighted average price) bid and the cut-off price of the auction. The first independent variable is the difference between this dealer's *customers'* bids (averaged over customers if the dealer has multiple customers) and the cut-off price of the auction – these bids are observed by the dealer prior to submitting her official bid. We also control for the dealers' net positions, their maximum submission limits, and how much of this submission limit they are using for their own bids vs. customer bids, as declared in their official tenders. To purge the effect of the public information content in dealer and customer bids, we control for auction and maturity-level fixed effects; thus our estimates reflect within-auction variation across dealers. In the right hand panel, we also control for dealer-fixed effects to control for systematic differences in dealers' bidding behavior. Standard errors are reported below the coefficients. Significance at 10-percent, 5-percent and 1-percent levels are denoted by (*),(**), and (***)

	(1)	(2)	(3)	(4)	(5)	(6)
	Entire Sample	T-Bills	Bonds	Entire Sample	T-Bills	Bonds
Customers' Bid (avg. over customers)	0.590 (0.058)***	0.158 (0.032)***	0.598 (0.1168)***	0.479 (0.048)***	0.194 (0.029)***	0.466 (0.146)**
Net Position/Bid Limit	-0.033 (0.007)***	-0.006 (0.0008)***	-0.200 (0.059)***	-0.026 (0.006)***	-0.004 (0.0001)***	-0.134 (0.053)**
Own Bid Size/Bid Limit	-0.030 (0.009)***	0.008 (0.001)***	-0.357 (0.084)***	-0.017 (0.010)*	0.007 (0.001)***	-0.288 (0.104)**
Customer Bid Size/Bid Limit	0.019 (0.012)	0.001 (0.001)	0.228 (0.103)	0.019 (0.010)*	0.0018 (0.0013)	0.035 (0.096)
Bid Limit	-0.004 (0.003)	-0.0004 (0.0003)	0.026 (0.013)	-0.007 (0.002)**	-0.0003 (0.0004)	-0.031 (0.025)
Auction Fixed Effects	Y	Y	Y	Y	Y	Y
Maturity Fixed Effects	Y	Y	Y	Y	Y	Y
Dealer Fixed Effects	N	N	N	Y	Y	Y
Observations	2042	1779	263	2042	1779	263
No. of Auctions	213	153	60	213	153	60
R-sq overall	0.05	0.15	0.16	0.38	0.25	0.34
R-sq within auction	0.07	0.12	0.18	0.44	0.27	0.52

Table 4: Do Modifications in Customer Bids Drive Modifications in Dealer Bids?

Each column reports the estimated coefficients from a regression in which the dependent variable is the change in a dealer's quantity weighted average bid price within the last 10 or 30 minutes before the bid submission deadline. The sample pools modifications in the bids of two types of dealers: those who route customer bids, and those who do not, which is coded in the first independent variable. The second independent variable interacts with this indicator the change in this dealer's customers' bid prices during the last 10 or 30 minutes in the auction. The rest independent variable is the change in the dealer's reported net position, bid size and customer's bid size during the last 10 or 30 minutes of the auction. To purge the effect of public information sources in comovements across customer and dealer bid changes, we control for auction and maturity level fixed effects in the regressions. We also control for dealer fixed-effects in all specifications. Standard errors are reported below the coefficients. Significance at 10-percent, 5-percent and 1-percent levels are denoted by (*), (**), and (***)

	Last 10 Minutes of the Auction			Last 30 Minutes of the Auction		
	(1)	(2)	(3)	(4)	(5)	(6)
	Entire Sample	T-Bills	Bonds	Entire Sample	T-Bills	Bonds
Dealer Routes Customer Bids	0.004 (0.009)	0.0004 (0.0002)**	0.09 (0.05)*	0.005 (0.006)	0.0002 (0.0001)**	0.02 (0.03)
Routes Customer Bids* Change in Customer Bids	0.24 (0.02)***	-0.00 (0.01)	0.28 (0.06)***	0.06 (0.02)***	0.004 (0.04)	0.11 (0.05)**
Change in Net Position/Bid Limit	0.24 (0.13)*	-0.01 (0.003)***	1.78 (0.65)**	-0.89 (0.06)***	-0.01 (0.002)***	-1.07 (0.18)***
Change in Own Bid Size/Bid Limit	-0.37 (0.03)***	-0.00 (0.00)	-0.71 (0.09)***	-0.02 (0.02)	0.004 (0.0005)***	-0.12 (0.09)
(Routes Customer Bids* Change in Customer Bid Size)/Bid Limit	-0.32 (0.10)***	0.00 (0.00)	-1.27 (0.48)**	-0.13 (0.07)*	-0.0025 (0.0013)*	-0.74 (0.35)**
Auction Fixed Effects	Y	Y	Y	Y	Y	Y
Maturity Fixed Effects	Y	Y	Y	Y	Y	Y
Dealer Fixed Effects	Y	Y	Y	Y	Y	Y
Observations	5616	4576	1044	5616	4576	263
No. of Auctions	223	157	66	223	157	60
R-sq overall	0.27	0.01	0.30	0.06	0.03	0.34
R-sq within auction	0.27	0.01	0.31	0.06	0.03	0.52

Table 5: Access to Customers and Dealers' Bidding Performance

The dependent variable in the first three columns is computed by first calculating the absolute deviation of a dealer's (quantity-weighted average) price bid from an auction's cutoff price. This absolute deviation is then average across securities auctions within a maturity class that the dealer participated in. The first independent variable is the average number of customers served by a dealer across auctions within the maturity class. Next, we control for the log of the dealer's size, measured by the amount of securities demanded by the bidder across auctions within the maturity class. We also control for the absolute value of the dealer's net position across these auctions, and the size of customer bids. The coefficients reported are from an OLS regression at the dealer and maturity level, where we also controlled for maturity fixed effects. Standard errors are reported in parentheses and significance at 10-percent, 5-percent and 1-percent levels are denoted by (*),(**), and (***) . The dependent variable in the next three columns is a measure of the amount of "overbidding" by the dealer. This variable is set equal to the difference between the dealer's bid and the auction's cutoff price when the dealer's bid exceeds the cutoff price. It is set equal to zero in auctions where the dealer's bid is below the cutoff price. The variable is then averaged across auctions within a maturity class. The dependent variable for the last three columns of the table is a measure of the amount of "underbidding" by the dealer. This variable is set equal to the difference between the dealer's bid and the auction's cutoff price when the dealer's bid is below the cutoff price. It is set equal to zero in auctions where the dealer's bid is above the cutoff price.

	Abs. bid deviation from cutoff price				"Overbids"		"Underbids"	
	All maturities	All maturities	T-bill	Bond	T-bill	Bond	T-bill	Bond
Average # of customers	-0.08 (0.03)**	-0.08 (0.04)**	-0.004 (0.003)	-0.17 (0.08)**	0.0001 (0.0006)	0.005 (0.02)	-0.005 (0.003)	-0.18 (0.09)**
log(Dealer's Bid Size)	0.03 (0.01)***	0.03 (0.01)***	-0.003 (0.001)*	0.04 (0.01)***	-0.0004 (0.0003)	-0.006 (0.003)**	-0.003 (0.002)	0.05 (0.01)***
log(Dealer's Net Position)		-0.002 (0.002)	0.00 (0.00)	-0.005 (0.004)	-0.0001 (0.0001)	0.0014 (0.0007)**	0.00 (0.00)	-0.006 (0.004)
log(Customer Bid Size)		-0.00 (0.00)	0.00 (0.00)	0.001 (0.003)	0.0001 (0.0001)	-0.00 (0.00)	0.002 (0.002)	0.001 (0.003)
Maturity fixed effects	Y	Y	Y	Y	Y	Y	Y	Y
Observations	172	172	47	125	172	125	47	172
R-sq (including maturity f.e.)	0.54	0.56	0.45	0.53	0.47	0.22	0.37	0.49

Table 6: Frequency Distribution of Customers in terms of Average Number of Submitters

The numbers reported in the table are obtained as follows. Suppose customer ABC participates in 30 year bond auctions, and 1 year treasury bill auctions in the sample. We obtain the average number of submitters used by this customer in the 30 year bond auction and the 1 year treasury bill auction. Suppose this average 3 for the former and 1.25 for the latter auctions. Then customer ABC will appear in row (1) under the "Treasury Bill" column, and row (3) under the "Bonds" column. He is a multiple-submitter customer in bond auctions, and a single-submitter customer in treasury bill auctions.

Avg. # of submitters used	Bonds	Treasury Bills	Customer Type
<1.5	74	95	Single-submitter
>1.5 and < 2.5	18	18	Multiple-submitter
>2.5 and <3.5	6	8	Multiple-submitter

Table 7: Customers Use Multiple Dealers

The dependent variable in the reported regressions is the number of submitters used by a customer in an auction. "Bid amount/Issue Amount" is the ratio of the total amount bid by a customer in an auction over the amount issued in the auction. "Multiple-Submitter" is a dummy variable that takes the value of 1 if the customer is a multiple-submitter in an auction, and zero otherwise. The sample in columns (1)-(2) consists of official tenders of all customers. The sample in column (3) consists of official tenders of all customers in auctions where they submit tenders through two or more dealers. Standard errors are given in parentheses, (***) significant at 99%; (**) significant at 95%.

	(1)	(2)	(3)
Bid Amount/Issue Amount	2.60	2.53	-3.71
	(0.40)***	(0.41)***	(1.03)***
(Bid Amount/Issue Amount)*Bond Auction		1.69	
		(1.01)**	
(Bid Amount/Issue Amount)*Multiple Submitter Customer			5.52
			(1.12)***
Constant	1.32	1.32	2.60
	(0.03)***	(0.03)***	(0.06)***
Observations	1413	1413	385
R-squared	0.03	0.03	0.08

Table 8: Customers with Long Term Dealer Relationships

This table identifies customers who are in long term relationships with dealers. “Maturity-range” identifies the sector (T-bills or Bonds) the relationship is in. We also report the percent of bids that a customer routed through her long-term dealer, and the percent of auctions that the customer participated in.

Customer	Maturity Range	Dealer	% of bids customer routed through dealer	% of auctions customer participated in
A	T-bill	1	0.782	0.245
B	T-bill	2	0.662	0.217
C	T-bill	3	0.851	0.566
D	T-bill	4	0.778	0.248
E	Bond	2	0.790	0.330
F	Bond	2	0.883	0.286
G	T-bill	4	1	0.229

Table 9: Bidding Behavior of Customers with Long Term Dealer Relationships

The dependent variable in the reported regression is the difference in the quantity weighted average price of an allotted tender (winning bid) and the cutoff price in Canadian dollars. Both the quantity weighted average price and the cutoff price are quoted in terms of \$CD 100 of the security allotted. The sample for the regressions is made up of allotted tenders of customers who are in a long-term relationship, submitted through the dealer with whom they are in long-term relationship, in the sector in which the customer is in a long-term relationship. As a control group, we also include allotted tenders submitted by customers who are not in a long-term relationship. “Use long-term dealer” takes the value 1 if a customer in a long-term relationship submits the tender through the long-term dealer in the sector where there is a long-term relationship, and 0 otherwise. “Bond auction” is a dummy variable that takes value if the auction is a bond auction and zero in treasury bill auctions. Standard errors are in parentheses above the double lines. (***) denotes significance at 99%.

	(1)	(2)	(3)
Use long-term dealer	0.005 (0.0016)**	0.053 (0.007)***	-0.005 (0.004)
Bond auction * Use long-term dealer		0.0019 (0.0016)	
Constant	0.007 (0.001)	0.007 (0.007)	0.013 (0.004)
Observations	896	896	196
R-squared	0.01	0.08	0.08