

NBER WORKING PAPER SERIES

ORDER FLOW AND THE FORMATION OF DEALER BIDS:
INFORMATION FLOWS AND STRATEGIC BEHAVIOR
IN THE GOVERNMENT OF CANADA SECURITIES AUCTIONS

Ali Hortacsu
Samita Sareen

Working Paper 11116
<http://www.nber.org/papers/w11116>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2005

The views expressed in this paper are those of the authors and not those of the Bank of Canada. We thank Bruno Biais, Kjell Nyborg, Robert Porter, Rafael Romeu, Leonardo Rezende, Christian Upper, S. Viswanathan and participants at the IIOC 2004, AMD 2004 and MTS 2004 conferences for their valuable comments and suggestions. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2005 by Ali Hortacsu and Samita Sareen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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

Ali Hortacsu and Samita Sareen

NBER Working Paper No. 11116

February 2005

JEL No. G1, L1

ABSTRACT

Is order-flow an important component of private information possessed by traders in government securities markets? Utilizing a detailed data set on Government of Canada securities auctions, we argue that the answer is yes. Direct participation in these auctions is limited to government securities dealers. However, non-dealer customers can also submit bids through dealers. We document patterns of strategic behavior by both sides of the market, dealers and customers, that support the hypothesis that customer bids provide valuable order-flow information to dealers. Dealer bids respond to privately observed customer bids, and dealers observing customer bid can predict the auction cutoff price better. Customers also respond strategically to dealers' use of the information contained in their bids.

Ali Hortacsu
Department of Economics
University of Chicago
1126 East 59th Street
Chicago, IL 60637
and NBER
hortacsu@uchicago.edu

Samita Sareen
Department of Economics
Duke University
ssareen@econ.duke.edu

A large theoretical literature in financial economics studies the question of how asset market equilibrium will be attained when agents possess private information regarding asset values and make strategic use of this information.¹ However, the institutional sources of private information in some asset markets, especially in the vast markets for government securities, are not easy to determine. After all, it is difficult to argue that many participants in government securities markets are privy to insider information regarding the economic fundamentals underlying the valuation of a particular Treasury bill or bond. Most, if not all, larger players in these markets use the same computer screens showing the same financial and political data and news. Hence, aside from possessing heterogeneous priors due to exogenous reasons, the forecasts they will make will differ from each other only to the extent that their forecasting technologies are different.

There is, however, one source of truly private information, identified by the market microstructure literature as “order flow.” Most large players in government securities markets are intermediaries who buy and sell securities to profit from the bid-ask spread. A source of private information for securities dealers is their interactions with customers. Since each dealer interacts with a different set of customers, each, in effect, sees different portions of market demand and supply curves, leading to differing inferences regarding where the equilibrium market price might lie.

In this paper, we will investigate whether and how “order-flow” information matters in the Government of Canada securities auction market. In these auctions, securities dealers authorized by the Bank of Canada place bids on debt issues of the Government of Canada (GoC henceforth). These securities dealers bid for their own accounts, but they also submit bids on behalf of a set of qualified bidders (called “customers”) who can not place bids directly on the GoC’s computer system, but have to route their bids through a dealer. Hence, “order flow” in this context can be defined as the customer bids that a particular dealer sees, but others do not.

The possibility that customer bids provide private information to dealers bidding in Treasury auctions was first suggested by the influential paper of Nyborg and Sundaresan (1996, page 68). Our paper provides direct evidence for this observation through a very detailed bidder-level data set. Similar to other treasury auction data sets used in previous studies, we have access to the entire set of bids submitted for a set of securities offerings.² The unique aspect of our data set is that we are able to observe which bids are submitted by the dealers for themselves and which bids are submitted on behalf of their customers. Another unique aspect of our data is that, along with the final set of bids, we are also able to track how dealers and customers modify their bids

¹See O’Hara (1995) and Brunnermeier (2001) provide comprehensive surveys of this theoretical literature.

²Examples of empirical studies on Treasury auctions that also have access to bidder-level data are Umlauf (1993), Gordy (1994), Nyborg, Rydqvist and Sundaresan (2002), Hortaçsu (2002), Fevrier, Prguet and Visser (2002), and Keloharju, Rydqvist, Nyborg (2004).

on the Bank of Canada’s bid submission database over time.

Given this extremely rich data source, we address the following questions:

1. What is the nature of information flows in this market? Specifically, do dealers respond to the information contained in their customers’ bids? Do dealers benefit from this information?
2. One may expect that customers know that they are revealing valuable information to their dealers, and will react accordingly. Is there any evidence for strategic behavior on the part of the customers?

To begin answering these questions, we first describe a feature of the bidding process that indicates that information flows are important in determining market outcomes. In Section 2.A we document that a vast majority of bid submissions are concentrated within the last 10 minutes preceding the bid submission deadline. Since this is a sealed-bid auction, where bidders can not see each others’ bids, there is no strategic benefit to waiting until the last minute, unless the bidders are waiting for payoff-relevant information to arrive. Moreover, reminiscent of the information aggregation process noted by Biais, Hillion and Spatt (1999), we find that bids that come in later are less dispersed and much more concentrated around the eventual cut-off (market clearing) price of the auction, indicating that later bids are much more informative about the auction outcome than earlier bids. Throughout, we also note that all of the temporal patterns we have discovered are much more pronounced in bond auctions compared to T-bill auctions. This finding is consistent with the intuition that the valuation of long-term securities is subject to much more uncertainty, and hence the value of waiting for new pieces of information to arrive is larger.

Next, in section 2.B, we document that dealers submit their own bids seconds after they submit their customers’ bids. This suggests that an important piece of information that the dealers are “waiting” for is the information contained in their customers’ bids. Consistent with the previous finding, the “waiting” behavior of dealers is much more apparent in bond auctions as opposed to Treasury bill auctions. This once again points out that the value of information provided by customer orders may be much more important in bond auctions as opposed to Treasury bill auctions.

An alternative explanation for the “waiting” behavior of dealers is that both customer and dealer bids reflect responses to releases of public information that are not observable to us. Thus, in Section 3.A, we investigate the cross-sectional variation in dealers’ bids *within* a given auction. We find that customer bids privately observed by a dealer is an important component of

the observed variation of bids across dealers participating in a given auction. Furthermore, the importance of privately observed customer bids in explaining the variation across dealer bids is more pronounced in bond auctions, as opposed to treasury bill auctions. In the same section, we also investigate the causes underlying modifications in dealers' bids. We find, again using a auction-level fixed-effects specification, that the observed modifications in a dealer's bid are positively correlated with modifications in customer bids observed by that dealer. Once again, consistent with previous results, the importance of customer bids is more pronounced in bond auctions, as opposed to treasury bill auctions.

We then examine whether the ability to observe customer bids affects dealers' bidding performance. Indeed, in Section 3.B, we report a systematic difference across dealers' bidding performances over the sample period: dealers who submit bids on behalf of a larger number of customers consistently bid closer to the market clearing price of the auction. Since the auction uses the discriminatory (pay-as-bid) format, bidders have the incentive to bid as close to the (expected) market clearing price as possible, so as not to leave money-on-the-table to the auctioneer. Our result suggests that dealers who submit bids for a larger number of customers are more successful in doing this. Consistent with above, we document that the positive correlation between a dealer's success and the number of customer bids routed by this dealer is much more pronounced in bond auctions, as opposed to Treasury bill auctions.

Although these findings suggest that dealers in this market possess an informational advantage due to their access to customer bids, we still have to consider the alternative possibility that the dealer might not be the only party who is benefitting from the dealer-customer information. That is, there might be information sharing between the dealer and the customer, in a manner that is beneficial for both.

Therefore, in Section 4, we investigate patterns of strategic behavior by customers. If a customer does not want a dealer to profit from the information contained in her bid, she might hold off her bid until the very last seconds of the auction, so that the dealer will not have the time to modify her own bid in response. We find some evidence supporting this hypothesis. In particular, we find that a large number of bids that are submitted too late, and are not taken into account when determining the winners of the auction and their payments, are dealer bids. Customers occasionally also miss the deadline, but dealers are much more likely to submit bids after deadline. Moreover, when we calculate what the auction outcome would have been if the late dealer bid had been taken into account, we find very large changes in the late dealers' outcomes. In one striking example, one dealer would have been allocated 398 million dollars more of the auctioned security had his late bid been taken into account.

Another pattern we observe in customer bidding behavior is that some customers, especially those who are buying large amounts, spread their bids among several dealers. As described in Sections 1 and 4.B, another explanation for this pattern may be the presence of “submission limits” imposed by the Bank of Canada on bidders to prevent short-squeezes in secondary market. However, although submission limits affect both bond auctions and T-bill auctions, we observe that customers are much more likely to use multiple dealers in bond auctions, as opposed to Treasury bill auctions. Combined with the previous evidence that “order-flow” information is much more important in bond auctions, this suggests that the use of multiple dealers is indeed a “strategic” response by customers, employed where it matters the most.

Interestingly, we also find that a number of large customers do not route their bids through multiple dealers. In Section 4.C, we report that these customers have formed long-term relationships with their dealers. However, customers in a long-term relationships, as opposed to those who spread their bids among several dealers, appear to differ in their bidding behavior. In particular, they bid prices that are further above the market-clearing price of the auction, compared to customers who are not in long-term relationships. This might be explained by the fact that these “long-term relationship” customers have more inelastic demand for government securities than others. However, we also find that these “long-term relationship” customers, whenever they use multiple dealers, submit bids that are closer to the market-clearing price through their long-term dealer, and submit their higher-price bids through other dealers. Bids that are closer to the market-clearing price may provide more valuable information to a dealer, hence it is reasonable to conclude that customers reserve this information for their long-term dealers, who may compensate them for this information through other means.

Section 5 provides a brief discussion regarding the nature of the information contained in customer bids. We differentiate between two types of information: the first is “strategic ” information, that allows a dealer to better assess the competition he faces within the auction. This type of information can be present in a purely private value world, in which dealers’ and customers’ valuations are driven purely by idiosyncratic demand shocks. However, even in a purely private value world, a dealer can improve her inferences about the competition she faces in the auction by observing her customer’s bid, and hence change her strategic behavior. The second type of information can be called “valuation ” information. In a common value auction environment, observing a customer bid will enable the dealer to modify, in a Bayesian fashion, her own forecast of the value of the auctioned security. We argue that our empirical evidence is consistent with both kinds of information being present in Government of Canada securities auctions.

We believe that our results shed light into the source and aggregation of private information in

government securities markets, and the formation of customer-dealer relationships as a response to the exchange of valuable information across these two parties. It also points to one of the benefits of being a dealer, that of having private access to order flow information. Of course, as in any industry with fixed costs, some revenues above marginal cost are needed to sustain entry; hence, one may regard the informational advantage possessed by the dealers as this additional revenue component.³

Our paper also contributes to prior empirical literature in market microstructure studying the role of order flow in various other securities markets. These include 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)). We should note that aside from Hansch, Naik and Viswanathan (1998), this literature has used aggregate, market-level data to investigate the importance of aggregate measures of order-flow and trading activity on asset prices. Hence, our study makes an important data contribution by focusing on the trading behavior of individual dealers and their customers.

Empirical work that is closest to ours is by Drudi and Massa (2001), who examine the individual trading behavior of government securities dealers on the primary (auction) market and the interdealer exchange, using a detailed fixed-income transaction database from Italy. Their starkest finding is that dealers behave strategically on the interdealer market to manipulate outcomes in the auction market. Massa and Simonov (2001) use the same data source to investigate strategic trading in the interdealer exchange, and find, consistent with our findings, that information gained from secondary market trading can affect bidding behavior in the primary market. Massa and Simonov (2003) document that long-term interactions between dealers leads to the formation of "dealer reputations," which affect the informational content that counterparties ascribe to the trades originating from these dealers.

Our results complement the findings of Drudi and Massa (2001) and Massa and Simonov (2001, 2003) in delving deeper into the sources of private information and strategic behavior in government securities markets, especially in the primary market. 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

³Aside from setting up the technological infrastructure to participate in these auctions, as discussed in section 1, dealers may be thought of incurring fixed costs due to the regulations they need to comply by. For example, dealers are subject to participation requirements to keep their status, this may be construed as a fixed cost of doing business. Sareen (2002) argues that the dealership system, whereby the issuer makes access to advantages conditional on the dealers satisfying obligations is a means to resolve the agency problem between the issuer and the security dealers.

of the informational and strategic environment (see Binmore and Swierbinski (2003), Sareen (2003) for recent surveys of the theoretical and empirical literature on the design of Treasury auctions).

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.⁴ 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.⁵

The process for issuing a “typical” GoC debt instrument links three markets: the when-issued market, primary market or the auction, and the secondary market. The process begins with government securities distributors (*dealers* henceforth) typically taking short positions in the when-issued market through forward contracts with other participants, for the yet-to-be auctioned security.⁶ 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. In all instances, being profit maximizing agents, they will attempt to buy the securities in which they have a short position at the cheapest price. Details of the three markets are given below.

A. Primary Market

Following the when-issued market is the *primary* market, where the Bank of Canada issues GoC securities. 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. The auctions are sealed-bid auctions.

Potential bidders in the government securities auctions can be classified into three groups: primary dealers, other government securities distributors and customers. “Other government securities distributors” refers to government securities distributors excluding primary dealers. On

⁴Bonds 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.

⁵The 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.

⁶The Bank of Canada designates certain institutions as distributors of government securities. These institutions are obligated to buy and sell securities to individual investors.

an average, 12.7 primary dealers, 4.9 other government securities distributors, and 8.6 customers participate in an auction. Of these, primary dealers account for as large as 90% of the primary market in terms of the volume of securities issued.

The distinction between government securities distributors 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.⁷ 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 section 4.B, we will argue that the use of multiple dealers by customers may be a strategic response by customers who do not want to reveal too much information to dealers who route their bids.

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. The 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.

Bids can be submitted as *competitive* tenders and *noncompetitive* tenders.⁸ Typically, a 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. Net position at a point in time refers to the participant’s net holdings (whether long or short) of the security being auctioned at that point of time. The net positions really capture the forward contracts of a dealer with other dealers or customers, prior to the auction. A 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

⁷Government securities distributors do not charge a fee to route customer bids.

⁸A *noncompetitive* tender comprises a quantity subject to an upper bound of \$3 million, with a participant being allowed to submit a single *noncompetitive* tender.

through dealers. 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.

All tenders have to be submitted before the expiry of the bidding deadline. 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 set captures several aspects of the primary market. For the primary market we have data over the period October 1998 to March 2003. A security will be uniquely identified by its maturity date and coupon rate if it is a bond, and by a maturity date if it is a treasury bill. An auction will refer to the issuance of a security of a specific maturity range (30, 10, 5, 2 years for bonds and 1 year, 6 months, 3 months for treasury bills), held at a specific time; thus, an auction for a security will be uniquely identified by the maturity range and the date on which it is held.

We have 347 treasury-bill auctions and 66 bond auctions in our sample. For each auction in the sample, we have the issue amount; issue date and maturity date of the auctioned security; total amount bid; total amount allotted; cutoff yield; total bid amount at the cutoff yield; coupon rate (if the security is a bond).

Several bidders participate in each auction, and each bidder can submit more than one tender in an auction. For each tender submitted in an auction, we have the following information: tender type (competitive or noncompetitive); time-stamp of the tender indicating the time at which the tender was submitted in the auction; tender status, indicating whether the tender was submitted by the participant, cancelled by the participant, or rejected by the issuer;⁹ identity of the submitter of the tender; identity of the bidder of the tender; participant type of the bidder (primary dealer, other government securities distributor, customer, Bank of Canada) and submitter (primary dealer, other government securities distributor); net position amount indicating a

⁹Tenders received before the auction deadline will have a status of “submitted” or “cancelled”. A participant can cancel his last submitted tender. No bid level information is available for cancelled tenders. Tenders received after the auction deadline will have a status of “pending submission” or “pending cancellation”. These tenders will either be accepted or rejected by the issuer, with a tender beyond the auction deadline usually being accepted by the issuer if it is on account of transmission failure, and rejected otherwise. The former if accepted has a status of “submitted”, and the latter if accepted has the status of “cancelled”. All rejected tenders have the status of “rejected”.

participant’s net holdings of the yet-to-be auctioned security at that time;¹⁰¹¹ 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.

In addition to describing the data in terms of auction-tender-bid level, we also distinguish the *official* vs *history* aspect of the data. As mentioned before, a bidder can revise a submitted tender before the auction deadline. All tender revisions and the constituent bid level revisions, made by each participant in the auction will be referred to as the *history* of the auction. 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.¹² Thus, the official data is a subset of the history data. While several empirical studies of treasury auctions have used official data, this is the first paper that makes use of the history aspect of treasury data to make inferences about the strategic environment under which bidding takes place in these auctions.

B. “When-Issued” and Secondary Markets

Recall that the *when-issued* market precedes the auction by a week. 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 comprise the resale market for Government of Canada securities. The resale market can be decomposed into two: the customer-dealer market and the interdealer market. In 2002, the customer-dealer market accounted for 54% of the total volume traded of Government of Canada bonds, and the interdealer market accounted for the rest.

In the customer-dealer market, institutional investors (for example, pension funds, mutual funds) trade with dealers on a bilateral over-the-counter basis over the telephone, with the result

¹⁰A net position has to be reported with a revised tender only if the change in the net position since the last submitted tender exceeds \$25 million.

¹¹Customers can submit their net positions to the issuer instead of revealing them to the dealer when submitting their tenders through them. This has to be done 30 minutes before the auction deadline unless there is a change in the net position by \$25 million, in which case the latest net position reported is recorded. For customers who submit net positions directly to the issuer, only the last net position amount reported by a customer before the auction deadline is available; the history of revisions in the net position is not available.

¹²An official tender is the last submitted or cancelled tender of a participant.

of these transactions known only to the two counterparties participating in the transaction. Unfortunately, we do not have real-time quote or trade data for this part of the market, thus we can not observe the over-the-counter customer-dealer market directly. However, we do observe this market indirectly through the net positions that bidders report when they submit tenders in the primary market.

This is in contrast to the interdealer market that operates primarily through electronic interdealer brokers. Electronic interdealer brokers post on an electronic screen, bid, offers, and trade outcomes communicated to him by the dealers.¹³ In 2002, 86% of the interdealer market was brokered through these electronic interdealer brokers. 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.¹⁴

We have access to transactions level data from CanPX, allowing us to observe 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 would, at least in principle, allow us to track activity on the when-issued and secondary market for an auctioned security.

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, we subsequently discovered that there is indeed an active market for “when-issued” trading, as evidenced from our observation of the levels and changes in the declared net positions of the bidders. 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.

¹³Typically, the identity of the dealer is not revealed.

¹⁴Prior to the establishment of CanPX, only dealers could observe the trade and quote activity on the electronic interdealer-broker screens.

II. Customer Order Flow and Dealers' Bidding Behavior

In this section, we examine bidding patterns in the data to investigate whether dealers respond to, and make profitable use of information contained in customer bids that they observe. We first document a salient temporal pattern in this market: bids come in during the very last minutes of the auction. This suggests that bidders are waiting for important information to arrive. This is reflected in the “quality” of the bids that come late – these bids are much more concentrated around the market clearing price of the auction. We then document that dealer bids follow customer bids. This suggests that the information that dealers are waiting for is in their customers' bids. Next, we provide more direct evidence for the hypothesis that dealers utilize the information contained in their customers' bids, by showing that dealers modify their bids in response to modifications in their customers' bids. Finally, we show that dealers who route bids for a larger number of customers are systematically more successful in their bidding.

A. Temporal Patterns in Bidding

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. Since our data set includes the time stamps for each tender (official or not) submitted by the dealers, we can analyze the timing of official bid submissions.

In Figure 1, we plot the cumulative distribution function of official bid submission times for different subsamples of our data set. First, observe that bidding activity is very much concentrated within the hour before the submission deadline. 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 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.¹⁵

¹⁵The 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 that the phenomenon we are reporting 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 can not 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.

An observation that is consistent with the “waiting for new information” hypothesis is the difference between the bid submission times for long-term vs. short-term securities seen in Figure 1. The submission timings for long-term securities are much more concentrated at the very last minutes of the auction, with the median official bid for securities with maturity greater than one year coming 2.5 minutes before the auction deadline, as opposed to 9.3 minutes for securities with maturities less than or equal to one year. 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.

We now bolster the information gathering hypothesis by presenting evidence that later bids are less dispersed, and, in particular, tend to be much more clustered around the eventual market clearing price of the auction than early bids. To do this, we first categorized official bids in every auction by their submission times by dividing time 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 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 the absolute difference between each quantity weighted average price bid and the cutoff price of the auction, i.e. lowest price at which the securities were sold. We then average these absolute deviations across each bid in a given submission period. Hence, what we are constructing is a measure of the dispersion (in absolute deviation terms) of the bids that arrive in different periods.¹⁶ Moreover, the dispersion is calculated around the cutoff price, hence the measure of dispersion can be interpreted as a “quality” indicator. Since Bank of Canada securities

¹⁶If K_t bids $\{p_1^{QW}, \dots, p_{K_t}^{QW}\}$ arrive in time period t , we compute $ABS_DISPERSION_t = \frac{1}{K_t} \sum_{k=1}^{K_t} |p_k^{QW} - p^c|$ where p^c is the cutoff price in the auction. Note that this is a dispersion measure, and, in particular, is not equal to $\frac{1}{K_t} \sum_{k=1}^{K_t} (p_k^{QW} - p^c)$, which would mechanically decline with K_t .

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.¹⁷ Alternatively, one can view the cutoff price as reflecting a weighted average of bids submitted by the bidders.

Figure 2 plots, as an average over all auctions in our data set, both the number of bids received in each 5 minute time interval, and the closeness of these bids to the eventual cutoff price, as measured by the dispersion of the bids around the cutoff price. Observe first that, consistent with the discussion above, many more bids arrive late than early. Observe also that the average absolute prediction error of bids declines from 15 cents at the 30 minute mark to 2.5 cents within the last 5 minutes. Thus early bids are, on average, quite far from the equilibrium market clearing price.

In Figure 3, we conduct a similar analysis separately for bond auctions and T-bill auctions. In this figure, we take the absolute prediction error of bids that come within the last 5 minutes as the base case, and report the relative size of the absolute prediction errors of bids that come in other time periods.¹⁸ Observe that this ratio is constant around 1 for short-term securities auctions, suggesting that earlier bids are as good predictors of the auction outcome as later bids. In contrast, for longer term securities, bids that come in earlier than 30 minutes before the auction deadline have absolute prediction errors that are 10 times as large as bids that come within the last 10 minutes of the submissions process.

This last set of results supports the hypothesis that the late submission of bids reflects an important information gathering process, and that this process is more important for bond auctions than for T-bill auctions.

B. Dealer Bids Follow Customer Bids

We now argue that one piece of information that dealers may be waiting for is the information contained in their customers' bids. Table 1 displays evidence supporting the hypothesis that dealers submit their own bids after seeing customers' bids. To construct this table, we looked at all instances in which a dealer submitted own her own behalf as well as for a number of customers. We then compared the submission time of the latest customer bid to the submission time of the dealer's own 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

¹⁷If 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.

¹⁸Specifically, we are reporting $\frac{|p_t^{QW} - p^c|}{|p_{5mins}^{QW} - p^c|}$ for each time period t , where the ratio is equal to 1 for $t = 5$ minutes.

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 1 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.

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.

Hence, the evidence suggests that dealers react to the information contained in their customers' bids *when and where it matters*. In particular, the second and third columns of Table 1 confirm our findings so far regarding the difference in the importance of “customer information” across longer term vs. shorter term securities auctions. Given that customer bids also come very near the bid submission deadline, one may suspect that dealers will have little time to perform elaborate calculations. However, our data suggests that, on average, a minute or two appears to suffice.

III. Do Dealer Bids Reflect The Informational Content of Customer Bids?

Do dealers utilize the information contained in their customers' bids? We try to answer this question in two ways.

A. Do Customer Bids Drive Differences in Dealer Bid Levels and Dealer Bid Modifications?

We first investigate whether the cross-sectional variation in dealer bids *within* a given auction can be explained by the customer bids seen by each dealer. Under the null hypothesis that customer bids do not provide valuable private information to a dealer, we should expect differences across two dealers' bids within a given auction to be independent of the customer bids each receives.

Note that examining *within* auction variation across dealer bids is crucial. Dealer bids may respond to public as well as private information flows; thus variation in dealer bids across auctions may reflect unobserved variation in public information flows as well as variation in private information flows. We will control for the impact of public information flows by including auction

level fixed effects in our regressions.¹⁹

Another piece of private information possessed by each dealer is their net position. Recall that dealers are actively involved in the (over-the-counter) when-issued market before the auction, accumulating significant short or long positions of the security that is about to be sold. One can also interpret the when-issued market activities of the dealers as a component of the private order flow information possessed by the dealers, since, as noted in Section 1, these activities are conducted on the over-the-counter market, which is not a transparent market. As also noted above, however, these net positions have to be reported to the Bank of Canada, hence we use this information in our regressions.

Thus, the regression equation we will estimate is:

$$DealerBid_{it} = \beta_1 CustBid_{it} + \beta_2 NetPos_{it} + \gamma_i + u_t + \varepsilon_{it}$$

where t indexes auctions in our data set, i indexes dealers within an auction. The dependent variable, $DealerBid_{it}$, is Dealer i 's quantity-weighted price bid in auction i . $CustBid_{it}$ is Dealer i 's customer's bid (averaged over different customers if Dealer i submits bids for multiple customers). $NetPos_{it}$ is Dealer i 's declared net position in auction t , normalized by C\$ 100 million. In many specifications, we also include a dealer fixed-effect term, γ_i , to control for differences across the bidding patterns of different 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).

¹⁹An alternative method to control for public information flows is to use proxy variables generated from secondary data sources. We should note, however, that most central banks, including Bank of Canada, set auction deadline timing so that it does not coincide with major news releases like monetary policy target or income and employment data. However, it is possible that other sources of financial news, such as developments in U.S. financial markets, may affect asset prices in Canada. One could, in principle, use real-time trading prices on the "when-issued" market, or prices of securities that might be close substitutes to the security being auctioned. As noted in Section 1.B, however, the one data source one might expect to be most helpful in this regard, CanPX, is conspicuously silent during the period preceding the auction. As described in Section 1.B, we observed *no* interdealer trading activity for the security about to be auctioned. We did observe some trading activity in other securities during the period preceding the auction. However, our conversations with experienced GoC securities traders pointed out several problems with treating these securities as close substitutes. The traded securities either had very different maturities, or had very different coupon structures. Nyborg and Sundaresan (1996) also note the problems attached to using yield-curve based valuation benchmarks. Even in the cases where there was a maturity and cash-flow structure match, the auctioned security was a benchmark and the security trading on CanPX was a non-benchmark issue, and hence not comparable due to the liquidity premium attached to the benchmark issue. For the U.S. Treasury market, Fleming (2001, 2002) has documented that liquidity of on-the-run ("benchmark" in the Canadian context) issues differ significantly from that of off-the-run ("non-benchmark") issues, with the price of the on-the-run issues embedding a premium for liquidity.

In the above specification, the auction fixed-effect, u_t , serves to capture unobservable changes in dealer bids that are common to all dealers' participating in the auction, presumably due to public information flows. We also conduct our analysis on Treasury bills and bonds separately to relax the assumption that the linear relationships in this specification are constant across different types of securities auctions in our data.

The results of the above regression are reported in Table 2. 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 0.6 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 a 100 million Canadian dollar increase in a dealer's *short* position is correlated with a bid that is 0.4 cents higher than average, i.e. dealers with short positions bid more 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.²⁰

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 0.62 for bonds, as opposed to 0.16 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.

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.

²⁰This finding appears to be 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 a short position bid more aggressively than dealers without a forward position 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. To further investigate this intriguing theoretical result, we ran the regressions in Table 2 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.

We can conclude from these results that there is a robust statistical dependence between dealer and customer bid levels, and that this can explain some of the variation in dealer bids within a given auction. We now strengthen this finding of correlation in bid levels by looking at bid changes. I.e. we investigate whether dealers change their bids in response to the arrival of customer bids, and whether the direction and magnitude of this change is explained by the informational content of the customer order.

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 . However, the dealer can change this bid until the bid submission deadline. Thus, we find the dealer’s “official” bid, i.e. the final bid by the dealer that is accepted by the Bank of Canada. We then calculate the annualized basis point difference between the dealer’s official bid, and her standing bid at T minutes prior to the deadline. Since bids comprise of multiple price-quantity points, we take the quantity-weighted average price as the bid.

We then perform the same calculation for the customers, taking the annualized difference between their official bid and their standing bids at T minutes prior to the bid submission deadline. Thus, the change in a dealer’s information set between time T and the auction deadline includes these modifications in customer bids.²¹

The dealer’s information set also contains changes to his net position. 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.

Table 3 reports the results of the regression of modifications in dealer bids on modifications in customer bids that the dealer observes. 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

²¹Note that the dealer may choose not to change her bid in response to a modification in her customer’s bid. This choice 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.

the entire sample of auctions, the estimates indicate that a 1 cent increase in a customer’s bid translates into a 0.261 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 0.135 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-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.

Overall, the results of this section indicate that customer bids are important drivers of the variation across dealers’ bids. We now investigate whether the information provided by customer bids is useful for the dealers in terms of their bidding performance.

B. Dealer Profitability and Access to Customer Information

Having established in the previous section that dealer bids reflect the orderflow information of their customers, we would expect that dealers who observe “more” orderflow predict the cutoff price more accurately. We find that this is indeed the case with dealers who submit bids on behalf of a larger number of customers consistently bidding closer to the cutoff price at the auction. Moreover this greater accuracy in prediction appears to translate into higher ex post profits for dealers who submit bids for a larger number of customers, since these dealers are able to win the auction more often, but do not appear to pay more for the units they win.

Figure 4 compares the bidding patterns of three classes of bidders: customers, dealers who route customer bids, and dealers who do not route customer bids. In this figure, 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.²²

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 the “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.²³

We now investigate the prediction performance of the dealers using a regression framework. To do this, we grouped auctions into their respective maturity classes, and averaged each dealer’s absolute prediction error (calculated as the absolute difference between her quantity-weighted average price bid and the cutoff price in the auction) across all auctions she participated in within this maturity class.

In Table 4, 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 4 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

²²The 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 out M days from now. Then, if we see a bid of P_{bid} for this bond, the daily, 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.

²³Similar patterns were observable when we plotted the corresponding bid distributions for bond and bill auctions.

served by a dealer is correlated with price bids that are 7.7 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 interact the two independent variables with an indicator for T-bill auctions (hence the base case estimates are for bond auctions). We see that the negative correlation between the number of customers served and the absolute deviation of a dealer’s bid from the cut-off price is much stronger for bond auctions. That is, the marginal effect of a customer on a dealer’s bidding performance is much larger for bond auctions than T-bill auctions. This finding is consistent with our earlier findings, and the hypothesis that customer information is more valuable in bond auctions than in T-bill auctions. However, also interestingly, the positive correlation between dealer size and absolute bid deviation is larger for bond auctions (in fact the coefficient on size for T-bill auctions is not significantly different from zero).

In the third column, we use a different interaction term: we look at the difference between the performance drivers of primary dealers and non-primary dealers. It appears that the marginal impact of a customer is not significantly different for a non-primary dealer when compared to a primary dealer (though the coefficient is negative). It also appears that the anomalous positive correlation between dealer size and absolute bid deviation is driven mostly by non-primary dealers.

Going across the table, we see that the positive correlation between number of customers served and the prediction performance of a dealer appears to be larger for bond auctions than for bill auctions.

The next set of results reported in the table replicate the three 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.”²⁴ We note that the correlation and explanatory power of the independent variables is low for this performance measure (none of the estimated coefficients appear statistically or economically significant. The reported R^2 also includes maturity fixed effects, hence the reported magnitude is due to these variables).

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 4, these last two sets of results indicate that customer

²⁴We set this variable equal to zero when the dealer bids below the cutoff price.

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 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. Customers Respond to Dealer’s Use of Order-Flow Information

In the previous sections, by looking at the dealer-side of the primary market, we found considerable support for the hypothesis that an important source of private information for dealers is customer bids, and that dealers use customer bids to revise their opinion about where the cutoff price in the auction might be. In this section we present evidence to support this hypothesis from the customer side of the primary market. Customers realize that dealers revise their opinion about where the cutoff price will be, on the basis of customer bids. We investigate three hypotheses regarding how customers may react strategically to dealers’ use of their information.

In section 4.A, we investigate the hypothesis that by holding-off bid submission till just before the auction deadline, customers can try to prevent dealers from making use of their order flow information. Since a dealer wants to condition on as much customer information as possible, customers may try to push a dealer “beyond the deadline” by submitting their bids at the last minute, and forcing the dealer to risk submitting his own bid later than the deadline.²⁵ In Section 4.B, customers attempt to conceal their entire demand schedule from a single dealer by using multiple dealers to submit bids. In Section 4.C we show that some customers are in a long-term relationship with a dealer in that they use a distinct dealer to submit their bids across all auctions. Presumably these repeated interactions provide dealers with valuable information about the cutoff price in the auction, for which they are likely to compensate the customers. We find that the payoff to the customer takes the form of the customer paying a lower price for the securities at the auction when he submits bids through his long-term dealer compared with other dealers. Finally, in parallel with the difference in the “information gathering” hypothesis between

²⁵This strategy hinges crucially on the fact that dealers are obligated by law to route customer bids immediately.

bonds and treasury bills in Section 2, we find that the strategic responses of the customer to bid revision by the dealer are much more pronounced in bond auctions compared with treasury bill auctions.

A. Last Minute Bidding by Customers Leads to Late Bids by Dealers

In Section 2.B we observe that while customer bids lag dealer bids, customer bids come quite close to the submission deadline as well. A possible reason for this could be that customers attempt to conceal their demand curves from dealers as long as they can, as they realize that dealers will use these to revise their own bids. Moreover, 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 send in their own, revised bids.

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 deadline²⁶. We found that 2% of the official bids submitted by the dealers were late. However, the fact that we see so few late tenders does not necessarily contradict our hypothesis that customers bid close to the auction deadline to conceal their demand curve from dealers. This is because a customer knows that a bid that she submits too close to the submission deadline has a higher likelihood of not being transmitted by the electronic auction system, as the dealer through whom this bid has been submitted may have hit his maximum submission limit.²⁷ In this instance bidding too close to the submission deadline would mean that the customer does not participate in the auction, either. If she has a large order, which is when she will attempt to bid as late as possible, the consequence of not participating can be particularly severe as this customer could be potentially “squeezed” in the secondary market.

Evidence that is more in line with the hypothesis that customers bid close to the auction deadline to prevent dealers from revising their bids emerges, when we look at the bidders from whom the late bids originate, and the auctions where the late bids are submitted. 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. We find that this is the case. 77%

²⁶In few instances, a bid that comes after the submission deadline is accepted as an official bid in case it is due to a genuine transmission error.

²⁷As described in Section 1, dealers can not submit bids for more than 40% of the total issue net of the dealer’s long position. When a dealer submits a bid on behalf of a customer on the electronic auction system, and including this bid the dealer has bid in excess of the amount he is permitted as a submitter, the auction system automatically does not transmit this bid and flashes this message on the dealer’s screen. Since these bids are not entered in the auction system, they are not recorded, and we could not analyze them.

of the late tenders were submitted by the primary dealers and 13% by the customers.²⁸

Do these late bids affect the auction outcome? If customer bids indeed provide valuable information to the dealers, then we should 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 examine the manner in which the auction outcome changes, 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 issuer would have received, on average, only an additional 767 Canadian dollars per auction!²⁹ 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.³⁰

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.

Finally, late bidding should be more pronounced for bond auctions rather than treasury bill auctions as there is greater uncertainty in pricing the former, and hence the tendency by customers to conceal their bids till the end should be more pronounced for bond auctions than treasury bill auctions. This is supported by the data as well. We find that the proportion of late tenders in bond auctions is more than double that in treasury bill auctions.³¹ In addition late tenders

²⁸The difference is statistically significant at the 99% level with the test statistic, $Z=7.9$.

²⁹this amount is also not statistically significant with $Z=0.0002$.

³⁰The change in price, averaged across the 111 late auctions is 0.004 cents, and $Z=1.2$.

³¹For bonds, late tenders as a proportion of the number of auctions in the sample was 0.64; for treasury bills the corresponding proportion was 0.31. The difference is significant at 99% with the test statistic, $Z=4.3$.

were submitted in twice as many bond auctions compared with treasury bill auctions.³² If the value of orderflow information is greater in bond compared with treasury bill auctions as we have hypothesized, then we should find that the difference in the amount paid by the bidders when the status of the late tenders is changed from rejected to accepted is significantly greater for bond rather than treasury bill. We find this to be the case with the former being thrice the latter.³³

B. Customers Use Multiple Dealers

In this section we document another possible strategic response by customers to dealer’s using the information contained in customer bids submitted through them. Table 5 reports a breakdown of the customers by the average number of dealers they utilize per auction. Specifically, for each customer in our sample, we obtain the average number of dealers through whom she routes his bids in an auction; the average is across the auctions in which the customer participates (we do this separately for bond and T-bill auctions).³⁴ If this average is less than 1.5, we define this customer as a “single-submitter” customer who he uses one dealer to route bids 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 5 indicates that while “single-submitter” customers are in the majority, 24 (26 in bond auctions) customers in our sample use multiple dealers.³⁵ 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 6 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% increase in the ratio of bid amount to issue amount.

One possible reason for customers using multiple dealers when they wish to purchase a large amount of the auctioned securities could be due to what we shall call the “bid concealing” effect.

³²Late tenders were submitted in 50% of the bond auctions in the sample, and only 25% of the treasury bill auctions. The difference was significant at 99% with the test statistic, $Z=3.4$.

³³The difference is significant at the 95% level.

³⁴The number calculated here is not the average number of *unique* dealers utilized by the customer over the sample. This is analyzed in the Section 4.C.

³⁵We 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.

When a customer has a large quantity of security to buy and is routing his bid through a single dealer, he is in fact revealing a much larger proportion of the market demand to the dealer, compared to when he has a smaller quantity to buy. This suggests that a customer could be routing his bids through several dealers when he has a large amount to buy, precisely because he wants to conceal his full demand schedule from a single dealer. Routing bids through more than one dealer is simply a customer’s strategic response to a dealer who will use her customers’ bid information to her own strategic advantage.

The following test lends considerable support to the “bid concealing” hypothesis. The test is based on a comparative static that exploits the differences in the temporal bidding patterns established for the dealer-side of the market in the previous sections. Our results in previous sections suggested that customer information played a more important role for bond auctions compared with treasury bill auctions. This effect 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.

To perform this test, we rerun the regression in Column (1), Table 6 by interacting the ratio of bid amount to issue amount with an indicator for bond auctions. Results in Column (2), Table 6 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.³⁶ 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.

There is, however, a second compelling reason for why customers might route bids through several dealers, or why the slope coefficient in the regression in Table 6 Column (1) has a positive sign. Customers could be routing bids through several dealers to circumvent a dealer’s “submission” limit. As we explained in Section 1, 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 due to the “submission limit” of a dealer. Clearly, unlike the “concealing effect”, there is nothing strategic in this effect. Rather, dealers routing their bids through multiple submitters might simply be a quirk of the Canadian

³⁶For 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.

mechanism for selling government securities.

There is, however, a second piece of evidence in our data that suggests that the submission limit can not fully account for the use of multiple dealers. In particular, we find that the use of multiple dealers is a strategy that is not employed by all customers in our data.

However, 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 6, Column (1) should not be statistically different between "single-submitter" and "multiple-submitter customers," as defined in Table 5. 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 6 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 6 reports the results of this 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 lead us to conjecture 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 "concealing effect".

Thus, all customers will tend to route bids through more than one dealer when they have a large amount of security they wish to purchase at an auction. The existence of single-submitter customers shows that it is the "submission limit" imposed on a dealer as a submitter that is at work. But the existence of multiple-submitter customers shows that there is a strategic effect at work as well: using multiple submitters could be a customer's best response to a dealer making strategic use of this customer's order flow information.

C. Customers in Long Term Dealer Relationships

If customers use multiple dealers to conceal their demand schedules from customers, why do we not see the single-submitter customers doing the same? After all they too know that dealers through whom they submit bids, revise their own bids on observing their customer bids. We find

that some of the single-submitter customers are in long-term relationship with their dealers.

To show the existence of a long-term relationship, we want to show that single-submitter customers submit a large proportion of their bids through a distinct submitter throughout the sample. Thus we are going to concentrate on customers in row 1, Table 5; these customers are single-submitter customers.³⁷ However Table 5 does not indicate that when on an average one submitter was used, whether that submitter was kept the same *across* auctions. Hence, for each of the single-submitter customers, we construct 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. Table 7 lists the 7 customers for whom the fraction of bids that they route through their favorite dealer is significantly different from the fraction they route through their second favorite dealer. For example, customer A is a single-submitter active 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.

Having established the existence of long-term dealer relationships, we now explore what kind of payoffs a customer gets from being in this long-term relationship with a dealer. We test three hypothesis about whether the average price paid by a customer differs depending on whether a customer is in a long-term dealer relationship or not.

The first hypothesis we test is whether a customer pays a lower price when she submits tenders through the dealer with whom she is in a long-term relationship, compared with the dealers with whom she is not in a long term relationship. Table 8, columns (1)-(3) report the results of these tests. The sample for testing this hypothesis comprises of winning bids of only those 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. For each winning tender of the a customer in an auction,³⁸ we construct the quantity weighted average price paid on the units won. We then calculate the difference in

³⁷As noted in the previous section, not all of these single-submitter customers are active participants – in fact 82% participated in less than 20% of the auctions. When we eliminate these “passive” customers, we ended up with 17 distinct active “single-submitter” customers in our sample.

³⁸A customer has the option of submitting his tenders through more than one dealer.

quantity weighted average price and the cutoff price of the auction. This is our dependent variable. For the regression in Column (1), Table 8, the explanatory variable is an indicator variable “Use long-term dealer” that takes the value 1 if the tender allotted to the customer was submitted through the dealer with whom he is in a long-term relationship, and 0 if it is submitted through any other dealer. Customers in a long-term relationship pay 0.5 cents less when they submit tenders through their long-term dealer compared with other dealers, with this difference being significant at the 89% level.

The sample that we have used in this regression has two kinds of auctions. In the first set of auctions, customers who are in long-term relationship route bids only through the dealers with whom they are in a long-term relationship. In the second set of auctions, customers who are in a long-term relationship route bids through both the dealer with whom they are in a long-term relationship, and other dealers. Given the “submission” limit effect, it could well be that the second set of auctions are really high demand auctions. But then the negative coefficient of the long-term dealer variable in Column (1), Table 8 could well be the result of this: the higher price paid when the bid is routed through a dealer other than the long-term dealer is because the need to go to other dealers arises only in auctions where demand is high.

Our results in columns (2)-(3), Table 8 rule out this concern. In column (2), Table 8 we introduce an indicator variable “Use multiple dealers” that takes value 1 for auctions in which the customer is submitting tenders through *both* the dealer with whom he is in a long-term relationship, and through other dealers. Presumably these are high-demand auctions, where the customer’s long-term dealer has hit a submission limit. The variable takes a value zero for all other auctions in the sample; presumably these are low-demand auctions. We now re-run the regression in Column (2), Table 8 introducing this indicator variable as well as interacting it with the dummy variable for the use of a long-term dealer.

The first thing to note is that auctions where both the long-term dealer and other dealer are being used are indeed high demand auctions. The average price paid in these auctions is 3 cents higher compared with auctions in which the customers route bids only through the dealer with whom they are in a long-term relationship; the difference is also statistically significant at the 99% level. However, this appears to make no difference to the payoff hypothesis. Customers pay 1 cent less when they route bids through the dealer with whom they are in a long-term relationship, whether or not the auction is a high-demand auction or a low-demand auction.³⁹ But note that the difference in the average price paid increases from 0.5 cents when we do not

³⁹We find a difference of 0.1 cents between high-demand and low-demand auctions. But the coefficient on the interaction of the bond auction indicator and the indicator for the use of a long term dealer is not significant.

control for high-demand, low-demand auctions, to 1 cent when we control for the high-demand, low-demand auctions.

We have established the differences in bidding patterns between treasury bill and bond auctions for the dealer-side of the market in the previous sections. These results suggest that the payoff to a customer for being in a long-term relationship with a dealer should be much more in bond auctions compared with treasury bill auctions. We do not find this to be the case. In Column (3), Table 8, we define an additional indicator variable that takes the value 1 if the long-term relationship is in the bond sector, and zero if the long-term relationship is in the treasury bill sector. The regression in column (2), Table 8 is re-run taking into account the sectoral difference. We find that the coefficient on the interaction term is not statistically significant; the difference in the average price paid by the customer when he routes his bid through the dealer with whom he is in a long-term relationship compared with other dealers, does not vary across the sectors in which there is long-term relationship. These results hold when we control for between-customer differences in the quantity weighted average price and cutoff price between customers, irrespective of whether they submit tenders through long-term dealers or other dealers.⁴⁰ But we do find that the customers who are in a long-term relationship in the treasury bill sector are being compensated by their dealer in the bond sector. When they participate in the bond auction, these customers pay 9 cents less when they route bids through the dealer with whom they are in a long-term relationship in the treasury bill sector!⁴¹ This cross-sectoral compensation supports the bond-bill bidding patterns reported in the paper so far.

The second pay-off hypothesis we test is whether a customer in a long-term dealer relationship pays a lower price compared with customers who are not in long-term dealer relationship, in the maturity sector where the customer is in a long-term relationship. Table 8, columns (4) and (5), report the results of this test. The sample to test this hypothesis is built as follows. From the previous sample, 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 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. We add allotted tenders submitted by customers who are not in a long-term relationship in treasury bill and bond auctions. These customers either use more than one submitter in an auction (multiple-submitter customers with a submitter proportion greater than 1), or they use one submitter in an auction, but this submitter

⁴⁰This is done by re-running the regression in column(2), Table 8 after adding a dummy variable for each of the customers listed in Table 7 to the set of explanatory variables in column (3), Table 8.

⁴¹This difference is significant at the 90%level; $t=-1.49$, $t(25, 0.90)=1.32$.

is not a distinct submitter.

Column (4) reports the results of the regression of the difference of quantity weighted average price and the cutoff price on the long-term dealer indicator.⁴² In the sector where there is a long-term relationship, customers in long-term relationship when submitting tenders through the dealer with whom they are in long-term relationship, pay 0.5 cents more than customers who are not in long-term dealer relationship, and this difference is significant at the 99% level.⁴³ In Column (5), Table 8 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.

This result appears counterintuitive at first in that 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. What might explain this result? Given the fact that the long-term customer and his dealer interact with each other in sectors other than Government of Canada securities, it is conceivable that the payoffs to the long-term customers are being given in these sectors by the long-term dealer. The second hypothesis tested above suggests that this could be the case. 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.⁴⁴

Thus, customers who are in a long-term relationship with a dealer, are adequately compensated by the dealer for the information provided by the customers in repeated interactions.

⁴²The constant in the regression in column (3) is the difference in the quantity weighted average price and cutoff price of the customers who are not in a long-term relationship.

⁴³ $t = 2.72$, $P(t > 2.72) = 0.997$.

⁴⁴For 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.

V. Discussion: What is the information contained in customer bids?

The empirical results presented so far support the hypothesis that the ability to see customer bids provides an informational advantage to dealers. But what is the nature of the information being revealed by the customer to the dealer? As we shall argue now, part of the answer to this question depends on whether the auction environment is one with idiosyncratic “private values,” or one where all bidders possess a “common value.”

Although the “common value” assumption is standard in the literature on securities auctions, even a simple independent private values environment can also generate many of the empirical patterns seen in the data. In the Appendix, we provide a simple model based on a private value single-unit first-price auction framework to support this claim. In the model considered in the Appendix, the strategic uncertainty faced by a dealer is the distribution of the bids of her opponents. The dealer’s optimization problem is to find the lowest bid that will allow her to win the auction with high enough probability. Seeing customers’ bids allows a dealer to refine her beliefs about the distribution of competing bids (since customers are also competitors), and this allows the dealer to make a more profitable decision. Thus, in a private values setting, the information contained in customer bids is “strategic” information, that allows the dealer to form a better response against the competition she faces.

As an example, consider the simplest case where there is a single unit of the security being sold, and the customer is the dealer’s only competitor in the auction. In this case, 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. The model in the Appendix shows that this very simple intuition persists when there are other dealers that the informed dealer is competing against. The model also clarifies the manner in which the distribution of surplus is skewed towards the “informed” dealers who see customer bids, and away from customers and “uninformed” dealers. The model also provides a rationale for the late-bidding behavior displayed by customers, and hence the dealers’ incentive to want to wait until the last seconds of the auction to update their bids.

The fact that a strictly private value model can generate several of the predictions does not mean that the presence of a common component in bidder valuations is ruled out, however. The same strategic motives seen in the private values setting are also present in a common value auction. However, in a common value auction, a dealer gains an additional benefit from seeing her customers’ bids. Along with the reduction in competitive uncertainty, the dealer can refine

her estimate of the liquidation value of the security by incorporating the information contained in customers' bids into her expected value calculation. We can call this "valuation" information component of customer bids, as distinct from the "strategic" information.

Based on our empirical results, we can not rule out the hypothesis that customer bids carry "valuation" information; hence the presence of a common value element along with a strictly private values environment. What we can conclude from findings, however, is that customer bids provide "strategic" information to dealers, since this effect is present in both a private value environment, and a common value environment.

VI. Conclusion

An important conclusion of this paper is that in Government of Canada securities auctions, where direct access is restricted to authorized security dealers, "order-flow" information is an important source of private information for these security dealers. "Order-flow" information is revealed to a security dealer through his interactions with customers, who can place bids in the auctions only through the authorized security dealer.

What if this source of private information was "dissipated" through a change in the mechanism to issue government securities? This could happen either by allowing direct access to customers to place bids in the primary issuance, and/or imposing 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, and as we point out in the Introduction, 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 even 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 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.

VII. APPENDIX: A Simple Private Value Model of the Dealer-Customer Interaction

We will now try to characterize the informational advantage possessed by a dealer who can see her customers' bids. The main feature of the following stylized model is that the dealer gains information from her customers regarding the competition she is facing in the auction, and can update her bid accordingly.

We should note that this model is built entirely around a symmetric, independent private value framework. In 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.

We will start by analyzing the simplest possible example where there is a single dealer and single customer in the auction, and where both of these agents have symmetric independent private valuations for the auctioned security. The only twist in this model is the fact that the dealer can observe the customer's bid. We will demonstrate, as we found in our empirical results in Section 3.B of the paper, that the dealer benefits from this arrangement, in that she is able to earn higher profits compared to the situation where she does not observe her customer's bid, and has to compete with her in a "blind" fashion. The result also has a policy implication: we show that this dealer-customer arrangement is costly to the auctioneer, in that expected revenues are lower. The revenue loss is also accompanied by an efficiency loss.

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 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, we can once again match our empirical result in Section 3.B: 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.

The third and final example we consider shows how this simple model can match the temporal patterns in bidding described in Section 2. Here, we allow the customer to wait until the last seconds of the auction to submit her bid through her dealer. This opens up the possibility that the dealer can not submit her updated bid in time. We show in this situation that the "advantaged"

dealer will place the symmetric Nash equilibrium bid as an insurance bid, but will update her bid once the customer's bid arrives, and hope that the updated bid will go through.

A. One Dealer, One Customer

The simplest possible setting is one where there are two bidders, the dealer, D , and her customer, C , who are competing for a single unit of the security. Both the dealer and the customer have symmetric, independent private values drawn iid from the uniform $[0, 1]$ distribution. 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, however, 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. In the case where valuations are 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 the 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 an 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.

B. Two dealers and One Customer

We now consider a slightly more complicated setting: 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:

$$\max_{b_d} (v_d - b_d) \Pr(b_d > b_a | b_d > b_c) \quad \text{if } v_d > b_c \quad (2)$$

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 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.

C. Why Might Customers Want to Delay Their Bids?

Now consider the possibility that the dealer does not always have time to utilize the information in customer bids to update her own bids. This could happen if customers wait until the last seconds of the auction to send in their bids. Since the dealer is under legal obligation to route through customer bids, there might not be enough time for the dealer to send her own updated bids through.

To formalize this intuition, let us revisit the simple case with one dealer and one customer considered above, and let $[T - \Delta, T], \Delta > 0$ denote the "last time interval" before the close of the auction at time T . Suppose any bid transmitted within this time interval is accepted by the BoC system with probability $p < 1$. In this situation, a customer has the incentive to 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 bid is not accepted, the dealer can also

place a bid anytime before $T - \Delta$. Without any customer information, however, this bid will be equal to the symmetric Bayesian Nash strategy of the dealer, i.e. $b_d = \frac{v_d}{2}$.

Thus, with probability $1 - p$, the auction outcome will be determined by the symmetric first-price auction strategies, and with probability p , the outcome will be the same as in the case considered in section B above. Notice that this same logic applies to the 2 dealer, 1 customer model considered above – once again, the "advantaged" dealer can post the symmetric Nash equilibrium bid as an insurance policy before seeing the customer bid, and try to sneak in her updated bid, hoping that it will go through. However, the customer will try to fight this by delaying her bid as long as possible.

D. Summary

We now summarize the findings of the above model in relation to the empirical findings reported elsewhere in the paper.

1. Being informed about customer bids gives an advantage to a dealer. This is confirmed by our empirical findings in Sections 3.A and 3.B. In Section 3.A we establish that dealers modify their bids on the basis of customer orderflow information. Section 3.B shows the advantage this gives to the dealers in terms of the accuracy in predicting the primary market clearing price and ex post profits.
2. Customers can try to defend themselves against dealers' use of their information by delaying their bids. This can generate the type "last-minute bidding" and "bid-modification" behavior we documented in Sections 2, 3 and 4.
3. There are policy implications of this simple model, but these implications should be interpreted cautiously. The main and unambiguous policy implication of the model is that forcing customers to bid through dealers distorts the efficiency of the allocation. One may hope that any such efficiency losses will be remedied in the secondary market; but this hope hinges on the lack of transaction costs in the secondary market. The second policy implication of our simple model is that the effect of the current bidding rules on the revenues of the Government of Canada may depend sensitively on the exact specification of the model: note that in the first example considered above, revenues decreased, but in the second example, the revenue increased. An accurate characterization of the revenue impact of the current bidding rules will require the use of a model that provides a much closer approximation to institutional reality.

VIII. References

- 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.
- 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.
- 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., 2002. Are Larger Treasury Issues More Liquid? Evidence from Bill Reopenings. *Journal of Money, Credit and Banking*, Vol 34, No. 3, 709-735.
- Fleming, M.J., 2001. Measuring Treasury Market Liquidity. *Federal Reserve Bank of New York Staff Report No. 133*.
- 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.
- Hortaçsu, A. "Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market," working paper, 2002.
- Hortaçsu, A. and Steven L. Puller "Understanding Strategic Bidding in Restructured Electricity Markets: A Case Study of ERCOT," working paper, 2004.
- 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., *The Microstructure Approach to Exchange Rates*, MIT Press, December 2001.

- 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.
- 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, 2002.
- 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.

Figure 1: CDF of Bid Submission Times




Figure 2: Dispersion of Bids around the Auction Cut-Off Price




Figure 3: Temporal Patterns in the Dispersion of Bids in T-Bill and Bond Auctions






Figure 4: Distribution of Bids Around Cutoff Price



"Normalized" difference is equivalent to the (negative) difference of annual yields. See fn. 22, pg.20.

Table 1: 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: 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, 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)	0.601 (0.058)***	0.162 (0.033)***	0.620 (0.174)***	0.476 (0.048)***	0.191 (0.031)***	0.418 (0.150)***
Dealer's Net Position (CAS 100 million)	-0.004 (0.001)***	-0.001 (0.0001)***	-0.025 (0.011)**	-0.003 (0.001)***	-0.0007 (0.0001)***	-0.016 (0.009)
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	0.07	0.06	0.38	0.23	0.39
R-sq within auction	0.06	0.10	0.08	0.43	0.25	0.49

Table 3: 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 first independent variable is the change in this dealer's customers' bid prices during the last 10 or 30 minutes in the auction. The second independent variable is the change in the dealer's reported net position 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
Change in Customers' Bids During Period (basis points, averaged over customers)	0.261 (0.029)**	0.003 (0.008)	0.268 (0.089)***	0.135 (0.019)***	0.004 (0.051)	0.113 (0.057)**
Change in Dealer's Net Position During Period	0.110 (0.046)**	0.00001 (0.0006)	1.187 (0.465)**	-0.147 (0.011)***	-0.0034 (0.0007)***	-0.157 (0.037)***
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	1918	1672	246	1909	1667	242
No. of Auctions	213	153	60	213	153	60
R-sq: overall	0.28	0.03	0.31	0.14	0.01	0.22
R-sq: within auction	0.36	0.02	0.41	0.19	0.04	0.28

Table 4: 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. The second independent variable is the log of the dealer's size, measured by the amount of securities demanded by the bidder across auctions within the maturity class. These two independent variables are interacted with two indicator variables. The first is an indicator for T-bill auctions. The second is an indicator for dealers who are not primary dealers. 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"		
Average # of customers	-0.077 (0.031)**	-0.147 (0.062)**	-0.040 (0.026)	0.002 (0.008)	0.001 (0.019)	0.005 (0.007)	-0.079 (0.032)**	-0.149 (0.064)**	-0.045 (0.030)
log(Dealer size)	0.026 (0.008)***	0.033 (0.010)***	0.020 (0.016)	-0.002 (0.001)	-0.002 (0.002)	-0.004 (0.007)	0.028 (0.008)***	0.035 (0.010)***	0.024 (0.019)
#customers*T-bill		0.144 (0.062)**			-0.001 (0.019)			0.145 (0.064)**	
log(Dealer size)*T-bill		-0.035 (0.010)***			0.002 (0.002)			-0.037 (0.010)***	
#customers*Non-primary dealer			-0.052 (0.093)			-0.004 (0.028)			-0.048 (0.104)
log(Dealer size)*Non-primary dealer			0.028 (0.014)**			0.004 (0.007)			0.024 (0.016)
Non-primary dealer			-0.435 (0.266)			-0.069 (0.135)			-0.366 (0.317)
Observations	172	172	172	172	172	172	172	172	172
R-sq (including maturity f. e.)	0.54	0.56	0.57	0.25	0.25	0.26	0.47	0.49	0.50

Table 5: 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 6: 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 7: 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 8: Payoffs to 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 in columns (1)-(3) comprises of allotted tenders of customers who are in a long term relationship, in the sector (bonds or t-bills) where there is a long term relationship. The sample for the regressions in columns (4)-(5) 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. “Use multiple dealers” is a dummy variable that takes value 1 for auctions in which the customer is submitting tenders through more than one dealer, possibly including the long-term dealer of the customer. “Bond auction” is a dummy variable that takes value 1 if the auction is a bond auction and zero in treasury bill auctions. Standard errors are in parentheses above the double lines. (***) significant at 99%; (**) significant at 95%; (*) significant at 90%.

	(1)	(2)	(3)	(4)	(5)
Use long-term dealer	-0.005 (0.004)	-0.010 (0.005)**	0.08 (0.02)	0.005 (0.0016)**	0.053 (0.007)***
Use multiple dealers		0.028 (0.008)***	-0.03 (0.01)		
Use long-term dealer * Use multiple dealers		0.001 (0.004)	0.03 (0.001)***		
Bond auction * Use long-term dealer * Use multiple dealers			0.001 (0.004)		
Bond auction * Use long-term dealer					0.0019 (0.0016)
Constant	0.013 (0.004)	0.007 (0.004)	0.007 (0.004)	0.007 (0.001)	0.007 (0.007)
Observations	196	196	196	896	896
R-squared	0.08	0.1	0.2	0.01	0.08