WORKING PAPER · NO. 2026-06

Entry and Exit in Treasury Auctions

Jason Allen, Ali Hortaçsu, Eric Richert, and Milena Wittwer

JANUARY 2026

THE UNIVERSITY OF CHICAGO
BECKER FRIEDMAN INSTITUTE FOR ECONOMICS

5757 S. University Ave.
Chicago, IL 60637
Main: 773.702.5599
bfi.uchicago.edu

Entry and Exit in Treasury Auctions
Jason Allen, Ali Hortaçsu, Eric Richert, and Milena Wittwer
January 2026
JEL No. D44, D47, G12, G28

ABSTRACT

Hedge funds have become increasingly active in Treasury markets, yet we have little evidence on how this shift has affected the market. We focus on the primary market, where the government issues its debt through auctions. Using over 20 years of data from Canadian Treasury auctions, we document dealers exiting and hedge funds entering. To evaluate the costs and benefits of this trend for debt issuance, we develop an auction framework with endogenous participation. We find that the expected gains from greater hedge fund participation, driven by increased competition, are smaller than the costs that arise because hedge funds—unlike dealers—do not participate regularly. Currently, dealer participation remains sufficiently strong that policies aimed at stabilizing participation have little effect on funding costs.

Jason Allen
Bank of Canada
jallen@bankofcanada.ca

Ali Hortaçsu
University of Chicago
and NBER
hortacsu@uchicago.edu

Eric Richert
University of Chicago
Department of Economics
richert.eric@gmail.com

Milena Wittwer
Columbia University
Columbia Business School
and NBER
mw3941@columbia.edu

1 Introduction

Governments worldwide have traditionally relied on regulated banks, known as primary dealers, to consistently purchase government debt and facilitate trade between investors, such as firms, public entities, and individuals. More recently, other institutions, in particular hedge funds, have assumed an increasingly important role (e.g., Banegas et al. (2021) for the U.S., Bank of England (2025) for the U.K., and Bahceli et al. (2024) for Europe). Yet, due to data limitations, there is no systematic evidence on entry and exit in government bond markets. Moreover, given that hedge funds and other non-dealer institutions—referred to as customers—have no obligation to regularly participate in the market and intermediate trades, the implications of this trend on the functioning of Treasury markets remain uncertain.[1]

In this paper, we focus on the primary market, where governments issue debt through auctions. We document dealer exit alongside rising customer participation before examining the resulting costs and benefits for debt issuance. Greater customer participation can strengthen within-auction competition, but irregular participation can also create a volatility cost. This happens when revenue losses in auctions with few bidders outweigh the revenue gains in auctions with many bidders, leading to lower revenues than in a scenario with stable participation. While auction theory offers some guidance on how to think about competition benefits and volatility costs in stylized settings—typically single-unit environments where bidders seek one object—intuition from such models often breaks down in multi-unit settings (Ausubel et al. (2014)). We develop a framework to evaluate these effects in complex multi-unit auctions such as Treasury auctions.

For our analysis, we collect bidding information on all Canadian Treasury auctions, from 1999 to 2022. We combine this with price information from secondary markets (futures, repo, foreign exchange). The auctions data allows us to observe unique anonymized bidder identifiers and all submitted bids. As is common in Treasury auctions, two types of bidders participate: dealers, who are obligated to bid regularly and submit their orders directly to the auctioneer, and customers, who face no participation requirement and must submit their bids through a dealer—an order that we observe.

With these data, we document a series of facts, which motivate our structural analysis. The first set of facts helps us model market participation. We show that dealers have been

[1] Unprecedented market turmoil in March 2020 triggered a policy debate on whether to reform Treasury market rules (e.g., Logan 2020; Ackerman and Hilsenrath 2022; Grossman and Goldfarb 2022).

systematically exiting the Canadian primary market. In contrast, customers have entered, in particular hedge funds. However, unlike dealers, who have an obligation to regularly attend auctions and buy sufficient amounts of debt, customer participation is irregular. This suggests that customers select specific auctions, depending on market conditions. In line with this idea, we show that customers participate in auctions when secondary-market spreads in the days leading up to the auction are high and when issuance volumes are larger. The second set of facts helps us model bidding behavior, conditional on auction participation. We show that when dealers observe an aggressive customer bid, they systematically adjust their own bids in anticipation of greater competition in the auction. Sophisticated customers, such as hedge funds, should take this adjustment into account when bidding.

Motivated by the empirical evidence and institutional features of the market, we construct a model that mimics a fiscal year. At the start of each year, when the government announces its debt issuance plan, each dealer decides, at a cost, whether to commit to bidding in every auction in the upcoming year and the number of participating dealers is announced publicly. Then, prior to each auction, customers observe the market conditions and decide whether or not to enter that specific auction, at an entry cost. Conditional on participation, all customers and dealers draw private signals about how much they value the bond—representing how much profit they expect to generate post-auction, given their private balance sheet positions—and place their bids. The auction clears at the price at which aggregate demand meets supply and each bidder pays their offered prices for each unit won.

To solve for the equilibrium conditions of this game and to estimate the dealer and customer values and entry costs, we build on the empirical literature on multi-unit auctions, in particular Guerre et al. (2000), Hortaçsu (2002), Kastl (2011), and Hortaçsu and Kastl (2012). Unlike existing studies, we allow customers to behave strategically, in that they anticipate dealers updating their own bids and we endogenize bidder participation. This contributes to the empirical auction literature that allows for endogenous entry but has so far focused on single-object auctions (see Hortaçsu and Perrigne (2021) for an overview). Estimation approaches for entry in single-object auctions require knowledge of how the equilibrium bid function behaves. In multi-unit auctions, this is a complicated object and so these tools are not directly transferable. Instead, we exploit estimated bounds on bidder-specific surpluses together with a matching procedure that pools auctions with similar value distributions to estimate the entry costs of multi-unit auctions. This approach could be used to endogenize bidder participation in other multi-unit auction settings, including those for electricity, renewable energy, or carbon allowances.

We find that both dealers and customers have non-negligible entry costs, and that customers are willing to pay more for bonds than dealers. The latter suggests that, post-auction, participating customers expect to execute more-profitable trading strategies with the bonds than participating dealers. One reason for this may be that customers face less stringent regulation than dealers, especially after the global financial crisis. Consistent with this, we find that prior to 2007-2009, dealer values were not so different from customer values but differ significantly afterwards.

We then use our structural model to conduct two sets of counterfactuals. For this, we utilize the empirical guess-and-verify method introduced by Richert (2021), while allowing for endogenous entry. This approach moves beyond truthful bidding (typically used in multi-unit auction studies), which is crucial for assessing how the number of bidders affects expected prices and auction revenues.

Our first set of counterfactuals aims at shedding light on three potential drivers of increased customer participation. There is growing evidence that regulatory changes in response to the 2007-2009 global financial crisis made it relatively more expensive for dealers to buy large quantities of bonds (e.g., Adrian et al. (2017), Bessembinder et al. (2018), and Cohen et al. (2024)) and allowed hedge funds to take advantage of increasing arbitrage opportunities (Du et al. (2018), Boyarchenko et al. (2020), and Barth and Kahn (2020)). Although the gradual implementation and limited data make the regulatory effect hard to isolate, we can quantify how much the rise in customers' willingness to pay relative to dealers—whether due to higher customer (or lower dealer) profit expectations—contributed to their increased participation: 25 percent, on average.

Customers may have also entered because dealers exited, which reduced competition, making entry more attractive. To quantify this we start in 2014—when Basel III regulation significantly limited bond trading for dealers—and add back the two dealers who exited in order to compare counterfactual entry probabilities to those observed in the data. We find that an average customer is 32 percent less likely to participate in the counterfactual.

Moreover, customers may have entered as the supply of government debt expanded, making it attractive to buy at auction despite low per-dollar margins. We estimate that roughly 29 percent of customer entry is explained by the increase in debt issuance.

Figure 1: Expected Price

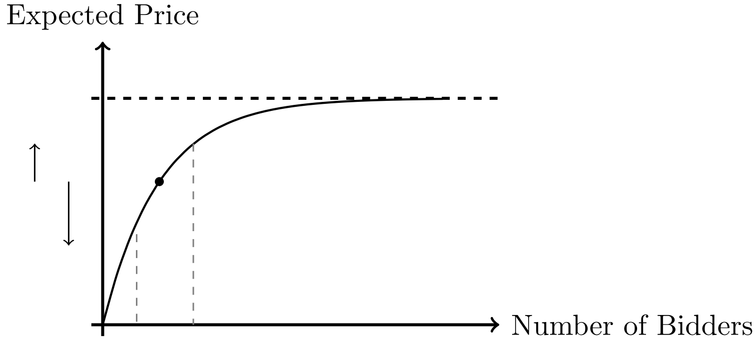

Figure 1 illustrates the relationship between the expected price and the number of bidders, with the gray dashed line representing the expected price under perfect competition. The key characteristic of this relationship is its concavity—an auction feature which is not unique to our setting. Expected revenue is also concave in the number of bidders.

Our second set of counterfactuals evaluates costs and benefits from a rising share of irregular auction participants. On one hand, stronger participation may increase competition, as in Bulow and Klemperer (1996). This reduces debt funding costs and price distortions due to bid shading. On the other hand, irregular bidder participation may introduce volatility across auctions that might destabilize financial markets, and increases the total (annual) cost of issuing government debt. The reason for this is that expected auction prices, and with it auction revenues, fall by more when customer participation is low than they rise when customer participation is high, as illustrated in Figure 1. This is a known theoretical feature in less complex settings, such as single-unit first price auctions (Milgrom (2004)).[2]

To illustrate the effects of participation on competition, we start with the simplest case with only dealers (who bid directly to the auctioneer) and vary their numbers. We find that when we remove one dealer from the status quo, the expected price drops by 6 cents per hundred dollar face value (i.e, 6 bps) because of stronger bid shading. This is sizable compared to the average amount of shading in our data, which is 11 bps. We then add back customers and vary how many customers are expected to bid (through dealers). The effect on expected prices is smaller—0.5 bps (C$0.149 million per auction)—mainly because we are keeping the number of dealers fixed, implying a relatively large number of bidders.

Next, we evaluate the cost of irregular customer participation. Comparing expected

[2]To reduce volatility costs Bhattacharya et al. (2014) suggest an auctioneer can fix the number of entrants using an entry rights auction to increase revenue. Quint and Hendricks (2018) show under what conditions non-binding indicative bids, often used in firm takeover auctions to select eligible bidders, can be informative by helping the seller select higher value bidders with greater probability.

revenues in the status quo to a counterfactual with fixed participation (at the median level) shows a total revenue loss of C$0.488 million per auction (1.63 bps). Conceptually, this loss reflects two main intertwined forces: first, a strategic shading effect, as bidders anticipating random competition shade their bids more aggressively; and second, a concavity effect. Since expected revenue is concave in the expected number of participants (analogous to Figure 1), adding one customer in expectation raises expected revenue less than removing one reduces it. In reality, auction heterogeneity means no two auctions are identical and variation in participation probabilities across auctions alone generates a loss of C$0.223 million per auction (0.74 bps).

Taken together, the costs arising from irregular participation are sizable relative to the benefits of attracting additional market participants, underscoring commitment as an important source of revenue stability. Under current market conditions, with dealer competition remaining sufficiently strong, policies aimed at stabilizing customer participation, however, either generate negligible revenue gains or are difficult to implement. For example, we show that subsidizing participation to sustain stable competition effectively amounts to a zero-sum transfer across bidders, since subsidies cannot be targeted to specific bidders based on their private information. Further, reallocating supply across auctions to smooth participation is difficult to reconcile with the government's objective of predictable issuance that avoids surprising the market. If the global trend of dealer exit continues, however, some of the policies we discuss may warrant reconsideration.

More generally, the competition-volatility effects we highlight is relevant for many countries that rely on a primary dealer model to issue debt. It is also common in auctions for financial products to have a set of regular and irregular bidders; e.g., Hendricks et al. (2023) for mortgage securities and Richert (2022) for credit event auctions. Our framework can be adjusted to fit these applications. Furthermore, the economic insights generalize to other markets that are populated by regular and irregular participants. Examples include market makers versus opportunistic traders in financial markets, global versus local firms in production markets, loyal versus non-loyal customers in consumption markets, and irregular versus stable energy generation in electricity markets (Petersen et al., 2024).

2 Institutional details and data

In most countries, government bonds are issued in primary auctions to a small set of regulated banks, often called primary dealers, and to customers. While specific auction rules may vary slightly, the main market features are common across developed economies.

Primary market. In November 1998, Canada transitioned from issuing debt via a syndication model to a more typical primary dealer market structure with debt issued according to an annual schedule. Between 1999 and 2022 there have been about twenty-eight government bond auctions per year, with an average nominal issuance size of C$3.24 billion. The auction schedule specifies the timing of auctions and the total debt to be issued. One week prior to the auction, the precise issuance size is announced.

Anyone can participate in Treasury auctions, but only dealers can bid directly.[3] Other bidders, called customers, can only participate indirectly by placing their bids via a dealer, who observes their customers' bids (but not the bids of customers bidding through other dealers). Indirect customer bidding is a common feature of Treasury auctions. It is also present, for example, in Japan and the U.S. (Boyarchenko et al., 2021).

Similar to Treasury auctions in other countries, customers include different types of institutions, such as pension, mutual and exchange traded funds, insurance companies, sovereigns, or bank treasuries. For over a decade, the biggest customer category has consisted of alternative investments companies, which includes hedge funds. For simplicity, we use the term hedge fund to describe these investment companies throughout the paper.

A bidder may submit and update two types of bids from the time the tender call opens until the auction closes. This is also a common feature of most Treasury auctions worldwide. The first type of bid is a competitive bid. This is a step function, with at most 7 steps, which specifies how much a bidder offers to pay for specific amounts of the asset for sale.[4] During normal times, primary dealers can demand up to 25 percent of the supply of Treasuries for their own account and 25 percent for their customers (with an aggregated cap of 40 percent), in the form of competitive bids.

[3]Strictly speaking, there are two types of dealers in Canada. Most dealers are primary dealers, but some are government security distributors. These are smaller dealers, who also place bids on behalf of customers, but face fewer market-making requirements. For simplicity, we do not distinguish between groups.

[4]Prices are expressed with three decimal places, e.g., C$ 99.999, and quantities must be stated in multiples of C$1,000. The minimum demand is C$100,000.

The second type of bidding is a non-competitive bid. This is a quantity order that the bidder wins for sure at the average price of all the accepted competitive bid prices. The Bank of Canada utilizes non-competitive bids to passively manage its balance sheet, and hence reduce the announced auction supply (before observing the submitted bids). For example, the Bank buys Treasuries (assets) to match the value of its bank notes (liabilities). For dealers and customers, non-competitive bids are trivial—no more than C$10 million for dealers and C$3-5 million for customers. Our model therefore abstracts from the strategic decision to place non-competitive bids.

All submitted bids are aggregated and the market clears at the price at which aggregate demand equals supply. In case of excess demand, bidders are rationed pro-rata on the margin, which means that the auctioneer proportionally adjusts demand at the clearing price until supply equals demand (see Kastl 2012 for a formal definition). Every bidder wins the units they demanded at bids weakly above the clearing price, and pays the bids for each unit won.[5]

Secondary market. Although our analysis focuses on the primary market, it is useful to briefly discuss the degree of connectedness between the primary and secondary markets. The secondary market is an over-the-counter market where dealers trade continuously with a wide range of investors—either through bilateral negotiations or via electronic platforms (Allen and Wittwer (2023)).

We know from the existing literature that this market is far from frictionless: there are search and bargaining frictions and dealers often face balance sheet constraints that limit their willingness to trade large volumes (see Duffie (2023) or Bessembinder et al. (2020)). This implies that from the perspective of a customer seeking to purchase substantial amounts of debt at favorable prices, the secondary market offers only a limited alternative. Indeed, customer demand in the primary market is, in the median, fives times the size of a median institutional secondary market trade (see Allen and Wittwer (2025)). Moreover, customers who wish to buy bonds but are hesitant to bid through a dealer in the primary market would instead have to purchase them non-anonymously from a dealer within their network, offering no significant advantage in terms of anonymity.

[5]Shortly after the auction clears, the clearing price and some additional aggregate summary statistics about the auction are publicly announced. This implies that no one has the incentive to participate in the auction only to learn about the market price, without wanting to win.

Regulation. In Canada, as elsewhere, dealers have an obligation to actively buy bonds in the primary market. Concretely, in normal times, primary dealers face minimal bidding requirements of roughly 10 percent of the auction supply. The minimum level of bidding must be at reasonable prices, and accepted bids should be approximately equal to a dealer's secondary market share over a specified time period. Primary dealers are also expected to act as market makers in the secondary cash and repo markets, where they provide liquidity to investors who seek to exchange government bonds for cash.[6] In exchange, primary dealers enjoy benefits. For instance, they have privileged access to liquidity facilities, overnight and term repurchase operations, and extract auction rents from observing customer bids.

Given the important role dealers play in the market (in Canada and elsewhere), they are heavily regulated. In particular, in the aftermath of the 2007-2009 financial crisis, regulation tightened for dealers and large banks more broadly. For instance, banks in developed countries faced heightened capital requirements. A notable illustration of this is the Basel III leverage ratio, which was enforced starting in 2014 and represents a significant limitation for bond trading (CGFS 2016; Wittwer and Allen 2022; Favara et al. 2022). In addition, in Canada, starting in 2016, dealers must report all trades to the Investment Industry Regulatory Organization of Canada, while customers are exempt. A similar trend happened in other countries; for example, in the U.S., dealers started reporting trades in 2017.

Traditionally, customers, such as hedge funds, played a negligible role in Treasury markets.[7] This has changed in recent years, but we have a limited understanding of what hedge funds are doing, given that there are only a handful of empirical studies. For example, Sandhu and Vala (2023) argue that hedge funds can act as market makers, engaging in trades that counter the positions of other investors. In this way, they are not very different from dealers. During times of distress, such as March 2020, however, hedge funds can contribute to market imbalances, reduced liquidity, and increased price volatility (e.g., Barth and Kahn 2020; Vissing-Jorgensen 2021). Increased hedge fund trading may also have implications for systemic risk in the market, given that hedge funds are more likely to employ riskier trading strategies (Dixon et al., 2012)—an effect that we do not consider in our analysis.

[6] For the most recent terms and conditions, see `https://www.bankofcanada.ca/wp-content/uploads/2016/08/standard-terms-securities180816.pdf`.

[7] In stock markets, hedge funds have long been active, and their role in these markets has been discussed in the academic literature (e.g., Stein 2009).

Table 1: Data summary of bond auctions

	Mean	SD	Min	Max
(Nominal) amount issued (in C$B)	3.24	1.04	1.00	7.00
Revenue (in C$B)	3.25	1.04	0.88	7.00
Number of participating dealers	14.46	2.61	11	23
Number of participating customers	6.74	2.56	1	15
Comp demand of a dealer (in %)	14.80	7.51	0.00	40
Comp demand of a customer (in %)	5.83	4.70	0.01	25
Non-comp demand of a dealer (in %)	0.09	0.04	0.00	0.30
Non-comp demand of a customer (in %)	0.19	0.16	0.00	0.76
Non-comp demand of the Bank of Canada (in %)	13.84	4.00	0.01	20
Number of submitted steps of a dealer	4.34	1.71	1	7
Number of submitted steps of a customer	1.86	1.02	1	7
Amount won by a dealer (in %)	4.79	5.85	0	35
Amount won by a customer (in %)	4.02	5.90	0	25

Table 1 displays summary statistics of our sample—February 10, 1999 to January 27, 2022. There were 645 auctions. The typical auction issues $3.24 billion worth of debt. The total number of competitive bidding functions (including updates) is 62,813. Competitive bids are step functions with at most 7 steps. The total number of non-competitive bids is 10,552. Demand and amount won are in percentage of supply.

Data. We focuses on the Canadian primary market for government debt, where we can trace market participants and their purchases over a long time series, which enables us to study exit, entry, prices, and allocations.

Our auction data covers all regular government bond auctions from the beginning of 1999 to the end of January 2022. This represents the entire auction history with the exception of three auctions held in December 1998. We observe unique (legal) identifiers for each bidder that account for mergers, acquisitions, and name changes. We also observe all winning and losing bids. We know whether the bid is competitive or non-competitive, and see through which dealer a customer places their orders. Table 1 provides summary statistics.

We augment the auction data with the daily average prices of each security from the secondary market and the futures, foreign exchange and repo markets, from the Canadian Depository for Securities. These data range from the beginning of 2014 to the end of 2021.

3 Empirical evidence: Exit, entry and bid updating

We document a series of stylized facts to motivate our structural model.

Exit and entry. We observe two striking time trends regarding entry and exit in the Canadian primary market. Appendix Figure E1 visualizes similar trends using public data of U.S. Treasury auctions; moreover, policy reports and newspaper articles indicate similar patterns for other countries.[8] This highlights that these trends are not Canadian-specific.

On one hand, dealers have (voluntarily) been exiting the market since the first Treasury auctions took place. While our data shows this fact only for auctions, we know from exit interviews that primary dealers exited the Canadian Treasury market entirely. The total number of dealers declined from 24, in 1999, to 15, in 2021 (see Figure 2A). After an early round of exits, the number of dealers remained stable until 2014, when Deutsche Bank and Morgan Stanley exited. Broker-dealers such as PI Financial Corporation and Ocean Securities followed in 2015 (see Appendix Table E1). The exit of global banks from the Canadian bond market around 2013-2015, during a period of tighter banking regulations and monitoring, suggests that stricter regulations may have pushed some banks out of the market.[9] More recently, two dealers sought buyers, which would have meant exiting via acquisition: RBC purchased HSBC in December 2023 and Laurentian failed to find a buyer.

On the other hand, customers—in particular, hedge funds—have become more active (see Figures 2B and 3A). The number of hedge funds participating in an auction increased from zero to ten in 2021. Furthermore, they have been buying an increasingly larger share of the auction allotments relative to dealers and other customer groups (see Appendix Figures E2 and E3). However, since customers have no obligation to participate regularly, like dealers, their auction participation is highly volatile, as shown in Figure 3.

[8] For example, the OECD raised the concern that primary dealers might not have the capacity to commit capital to debt markets post-global financial crisis and would exit (Blommestein et al. 2010); Bloomberg reports on dealer exit in the U.K., in particular the exit of Credit Suisse Group AG and Societe General SA in 2016. More recently, the Association for Financial Markets highlights the exit of 7 primary dealers in the Eurozone between June 2021 and January 2022; The Economist and the Wall Street Journal discuss the rising role of hedge funds in the U.S.; Reuters highlights the increasingly important role of hedge funds in European Treasury markets. All websites were accessed on 07/05/2024.

[9] This would align with anecdotal evidence. For instance, according to research by Greenwich Associates—a leading financial consultancy—regulations implemented after the 2007-2008 global financial crisis caused a general (voluntary) retreat from Canadian debt markets (primary and secondary) in 2014 (Altstedter, 2014).

Figure 2: Dealer exit and hedge fund entry from 1999 to 2022

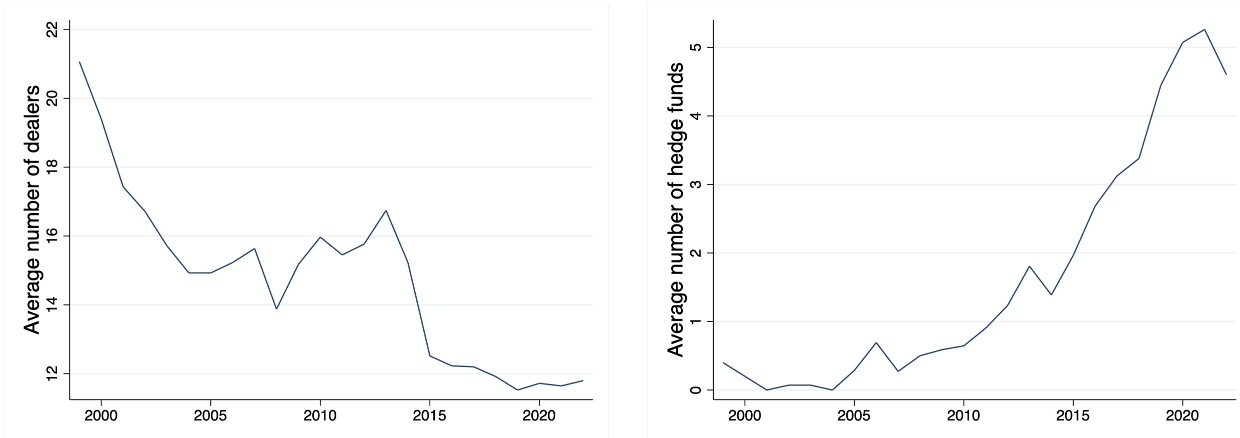

(A) Average number of dealers in a year (B) Average number of hedge funds in a year

Figure 2A shows how many dealers have participated in primary auctions, on average in a year, between 1999 and 2022. Figure 2B shows how many hedge funds participate, on average, each year.

Figure 3: Irregular hedge fund participation

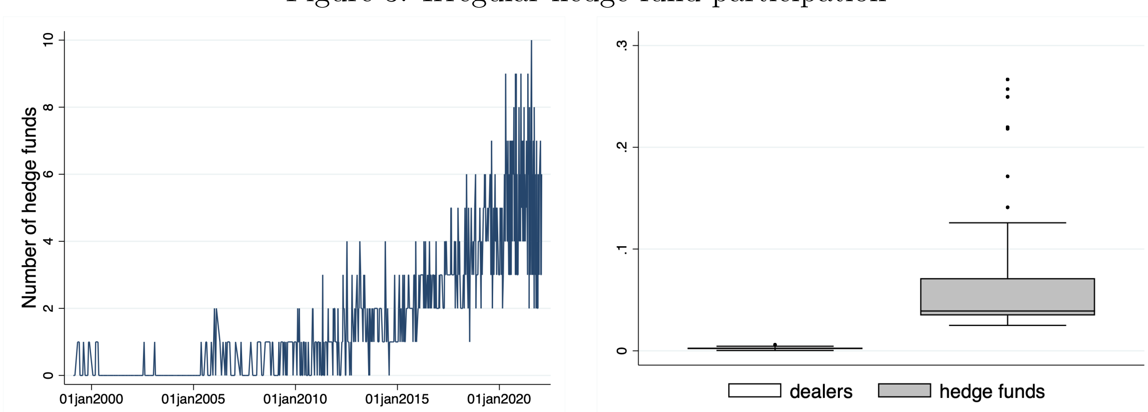

(A) Number of hedge funds per auction (B) Variance in participation

Figure 3A shows the number of participating hedge funds in all auctions. The first box plot of Figure 3B, called dealers, shows the distribution of the variance in the percentage of dealers bidding across auctions in a year out of all active dealers in that year. The second box plot, called hedge funds, shows the analogue for hedge funds. The distribution is similar when including all customers.

Predicting customer participation. In Appendix B.1, we analyze potential predictors of customer auction participation, which are inspired by emerging research into the trading behavior of hedge funds, particularly during 2020, without claiming causality. Across regression specifications, we find that the spread between the highest and lowest price at which a bond trades shortly prior to being auctioned is positive and significant. Since this spread correlates closely with liquidity measures such as the bid-ask spread (Duffie et al. 2023), this finding suggests that customers buy bonds at auction when they can quickly sell them at high prices in the secondary market. Additionally, some coefficients, detailed in the Appendix, indicate that fewer customers participate when there is increased interest rate uncertainty. The issuance amount (supply) is also a significant positive factor, although the estimate is noisy. The R^2 ranges between 37 and 41 percent, highlighting the influence of unobservable factors on customer participation—a feature that our model will incorporate.

Dealer updating. Our auction model, presented in Section 4, builds on Hortaçsu and Kastl (2012), who model the bidding behavior of dealers in Canadian Treasury auctions, in which customers must bid via dealers. Hortaçsu and Kastl (2012) show that dealers in the Canadian Treasury bills market frequently update their bids after observing a customer bid. To provide evidence that dealer updating plays a role in our data, we would like to regress the "change in a dealer's bid conditional on observing a customer's bid" on "the customer's bid." Given that bids are step functions, this is non-trivial.

We compute the change in the dealer's bid as the difference between the average quantity-weighted bid of the dealer's update, having observed a customer's bid, minus the average of the dealer's bid, prior to the update. If no update is placed, we use the last bid before observing the customer bid, since by refusing to update, the dealer decided that that bid was still optimal. We regress this variable on three moments of the customer's bidding function to get a sense of what aspects of the step function are most strongly correlated with updating: the quantity-weighted average bid (that is, the price a bidder is willing to pay per unit of the bond), the number of steps, or the highest amount demanded. We normalize the customer's quantity-weighted average bid by the average of all customer bids within an auction to capture the idea that dealers update their bid when the customer's bid they observe differs from what they expected. In rare cases in which the dealer's update follows multiple customer bids, we average over them.

We report in Table 2 that a dealer bids more aggressively when observing a more ag-

Table 2: Dealer updating

Change in qw-bid of dealer		
Customer's qw-bid - average customer qw-bid	0.047***	(0.009)
Customer's number of steps	0.032	(0.029)
Customer's total demand	+0.001	(0.001)
Auction fixed effects	Yes	
Adjusted R^2	0.21	

Table 2 shows the results from regressing the change in a dealer's quantity-weighted average (qw) bid, calculated as the difference between the qw-bid of their next submitted bid and the average qw-bid of all other steps that they submit in the auction, on the qw average bid of the customer less the average customer qw-bid in that auction (to reflect how aggressive a particular customer's bid is relative to the average customer at that auction), the number of steps, and the customer's total demand. The regression also includes auction fixed effects. We use bidding data from all bond auctions from 1999 to 2022—8,193 observations. Bids are in bps, quantities are in C$ millions. Standard errors are in parentheses.

gressive customer bid. Consistent with Hortaçsu and Kastl (2012), dealers therefore respond to customer bids. Of the moments in the customer's step function that we consider, only the quantity-weighted average is statistically significant—a detail that we exploit in our estimation presented in Section 5.

4 Model of the Canadian primary market

Motivated by the empirical evidence, we construct a model with two main features, which are both novel relative to the existing literature. First, we endogenize the bidders' participation decision. Here, we distinguish between the dealer's decision to exit the market at an annual frequency and the customer's decision to enter specific auctions.[10] This allows us to highlight the benefit of greater competition versus the cost of irregular customer participation. Second, we allow customers to be sophisticated in that they can anticipate that dealers might update their bids when observing a customer bid. We show in Appendix Figure E4 as well as Appendix B.3 that this feature is empirically relevant.

[10]We assume that dealers decide whether to participate in the market at the beginning of a fiscal year to capture the institutional feature that dealers must commit to regularly participate in primary auctions, in combination with the fact that debt financing is planned annually. In counterfactuals in which dealer participation is held fixed, the assumptions governing their participation do not matter.

4.1 Players, timing, and preferences

Our model describes bidding decisions over the course of a fiscal year. We highlight four assumptions in particular because they directly relate to the model primitives that we later estimate. We denote random variables in **bold**.

There are two groups (g) of market participants: dealers (d) and customers (h). The number of dealers who consider remaining as dealers, $\bar{N}^d > 1$, and the number customers who are interested in bidding, $\bar{N}^h \geq 0$, are commonly known.[11] Similarly, all distributions and functional forms are commonly known.

At the beginning of each year a debt issuance plan is announced, stating T auctions. At this point, dealers decide whether they wish to continue being a dealer, that is, whether they commit to participating in all auctions for that year. Whether this is profitable depends on the private annual cost each dealer faces, γ_i^d. One may think of this as an opportunity cost that an institution acting as a dealer suffers because it cannot do other things during the time it fulfills its dealer activities, such as bidding at auction and making markets.

Assumption 1. *At the beginning of the year, dealers' private entry costs for all T auctions of the year, $\boldsymbol{\gamma_i^d}$, are drawn independently from a common atomless distribution, G^d.*

After each dealer makes their decision, the market is informed about the number of bidders who will act as dealers in the upcoming year, N^d. In reality, this information is posted on the website of the auctioneer. The game ends if no dealer enters.

Before each auction, t, customers observe how costly it is for them to enter the auction. Similar to dealers, one may think of these entry costs as opportunity costs, since it takes time to monitor the market and compute competitive bids.[12] Crucially, as for dealers, these

[11] We could distinguish between different types of customers, for instance hedge funds versus other customers, or different types of dealers. Theoretically, such an extension is straightforward. However, given the number of bids placed per auction, differentiating between more bidder groups than dealers and customers would require pooling bids across auctions for estimation purposes. This would imply abstracting from fundamental changes in bond values over time, which are sizable. Moreover, non-hedge fund customers play such a small role that the cost of complicating the model and increasing measurement error (in the resampling procedure described below) outweigh the benefit of having multiple customer types.

[12] In principle, these costs could be constant over time within customer or could be constant for all customers but vary across auctions. However, both these assumptions would produce entry patterns that are inconsistent with that data. Constant entry costs would imply a constant order of entry among the bidders, with the lowest cost bidder always entering first. While auction-specific common entry costs, would lead participation to be monotone in customers' expected surpluses. Instead, we assume that the cost distribution is fixed across auctions. In Appendix B.6 we test this assumption and do not reject it.

entry costs are independent of the auction outcome—they are incurred regardless of whether, or how much, the customer ultimately wins. For instance, a trader spending time preparing and bidding in the auction forgoes the potential profits from trading in other markets. This cost is sunk once the auction concludes.

Assumption 2. *Before each auction t, customers' entry costs for auction t, $\boldsymbol{\gamma_{ti}^h}$, are drawn independently from a common atomless distribution, G^h.*

In addition to their entry cost, customers observe the auction-specific distribution from which they will draw (multi-dimensional) private signals that affect their willingness to pay when entering the auction. This is specified in Assumption 3. The signal distribution captures current market conditions, which can be unobserved by the econometrician. For example, the expected signal may be high when secondary spreads are high (as suggested by the evidence in Appendix Table B1), or when interest rates are expected to fall.

Observing the signal distribution, each customer decides whether to enter an auction before learning their private signal. This timing assumption is standard in the literature on endogenous bidder participation in single-object auctions. In our setting, it reflects the idea that most customers are part of large institutions that tend to allocate tasks (such as bidding at auction) some time in advance, before the bidding process starts.

An alternative would be to assume that customers first observe private signals and then enter; that is, allow for selective entry. While we cannot perfectly test this implication, we can provide some supporting evidence against this hypothesis by comparing the average customer bid (relative to dealer bids) around dealer exits. In a selective entry model, customers with lower signals would enter an auction after a dealer exits, as fewer bidders make the auction less competitive, lowering the entry threshold. Consequently, the average customer value in auctions after a dealer exits should be lower than before they exit. In Appendix B4 we formally test for this and do not find convincing evidence of a change in relative values; the same is true when comparing bids.

Assumption 3. *Dealers' and customers' private signals $\boldsymbol{s_{ti}^d}$ and $\boldsymbol{s_{ti}^h}$ are, for all bidders, i, independently drawn from common atomless distribution functions F_t^d and F_t^h with support $[0,1]$ and strictly positive densities f_t^d and f_t^h.*

Signals are independently drawn from auction- and bidder-group specific distributions, which captures all information that all bidders know when bidding. For example, bidders observe a reference price range provided by the auctioneer and leverage insights from the active secondary market. The existence of this market suggests that most, if not all, price-relevant information is aggregated prior to the auction. For instance, if dealers anticipate a high auction clearing price, the mean of their signal distribution would incorporate this shared prior belief. Any residual heterogeneity in information at the auction stage likely stems from idiosyncratic factors, such as the structure of individual dealer balance sheets, client order flows, or investment opportunities.

Assumptions 1-3 rule out that bidders have an incentive to adopt strategies that connect multiple auctions. Thus, technically, our game consists of dealers' exit decisions plus T separate auction sub-games. We think this is reasonable in our setting, given that the typical dealer sells most of their auction-purchased bonds before the next auction takes place. Furthermore, there is no evidence that dealers are willing to pay more in subsequent auctions when they were allocated fewer bonds than expected. Further, the number of customers participating in two subsequent auctions is uncorrelated (once we control for the upward time trend in customer participation, in Appendix Table B2). In other settings, it can be important to take inter-temporal strategic considerations into account.[13]

How bidders bid in a specific auction, t, depends on how much they value the bond at that time. This, in turn, is driven by their private signals, implying that their values are independent and private. To provide support for this assumption, we follow Hortaçsu and Kastl (2012) and test whether dealers who observe customer's bids only learn about the degree of competition in the auction and not about an uncertain common value of the bond. Our findings, reported in Appendix Table B3, support the assumption of private values.

Assumption 4. *A bidder, i, of group $g \in \{d, h\}$ with signal s^g_{ti} values amount q of the bond by $v^g_t(q, s^g_{ti})$. This value function is non-negative, measurable, bounded, strictly increasing in s^g_{ti} for all q, and weakly decreasing in q for all $s^g_{ti,}$.*

In contrast to entry costs, the value a bidder derives from an auction depends on the amount won and is therefore directly tied to the auction outcome. In practice, a bidder's value—or

[13] Rüdiger et al. (2023), for example, analyze inter-temporal incentives of Argentinian primary dealers to forgo short-term auction gains to fulfill the longer-term minimal bidding requirements necessary to maintain their dealer status. Appendix Figure E5 provides evidence that Canadian primary dealers are above minimal bidding requirements and, therefore, do not face the same inter-temporal trade-off. In addition, if these bidding requirements played an important role, dealers should be willing to pay more in subsequent auctions when winning less than expected, for which we find no evidence in our setting.

willingness to pay—reflects the expected profit from acquiring the bond at auction and using it in subsequent trading strategies. Accordingly, factors that lower these expected profits—such as regulatory costs that vary with the traded amount—affect bidders' values (rather than their entry costs, which are independent of auction outcomes).

Given values, bidders place bids. Each bid is a step function that characterizes the price the bidder would like to pay for each amount offered at auction. Specifically, bidder i has the following action set to place a bid in auction t, given the (maximal) auction supply $\overline{Q}_t$:

$$A_{ti} = \begin{cases} (b,q) : dim\,(b) = dim(q) = K \\ b_k \in [0,\infty) \text{ and } q_k \in [0,\overline{Q}_t] \\ b_k > b_{k+1} \text{ and } q_k < q_{k+1}\ \forall\ k < K. \end{cases} \tag{1}$$

Here, and throughout the paper, we omit the t- and i-subscripts on b, q, and K for notational convenience. K is the exogenous maximal number of steps bidder i can submit in auction t. This number could be determined by auction rules and thus be common across bidders. Alternatively, it could be bidder-specific, for example, because bidders face an additional small cost for computing and placing extra steps (Kastl (2011)).

Bidding evolves in three rounds. First, dealers can place early bids directly with the auctioneer. Second, each participating customer is randomly matched to a dealer and places their bid with this dealer.[14] Third, each dealer observes these customer bids (if any) and may update their own bids.

To rationalize early bidding (which we observe in the data), we follow Hortaçsu and Kastl (2012), and let one dimension of each dealer's signal, $\boldsymbol{s^d_{ti}}$, be a random variable, $\boldsymbol{\Psi_{ti}} \in [0,1]$, which is the mean of another Bernoulli random variable, $\boldsymbol{\Phi_{ti}}$, that determines whether the dealer's later bid will be made before the auction closes. The idea is that, when observing a customer's bid shortly before the auction, the dealer might not have sufficient time to recompute their bid and enter it into the bidding interface. Whether there is sufficient time or not is revealed only in the last stage. Formally, in the last stage, the dealer observes the

[14]The assumption of random matches simplifies the equilibrium conditions and the estimation procedure. In Appendix Figure E6, we provides some evidence that random matching is a reasonable approximation of reality. Further, we provide evidence that since 2014, dealers have not been leaving the market because they systematically observe fewer customers than dealers who remain in the market, which would violate the random matching assumption. We find that on average the total number of customers that an exiting dealer observes in a year (because the customer bids via this dealer) is not statistically different from the average number of customers that the remaining dealers observe. The test's p-value is 0.28 for auctions post-2014.

realization $\omega_{ti} \in \{0, 1\}$ of $\boldsymbol{\Phi_{ti}}$, where $\omega_{ti} = 1$ means that the late bid will make it on time, in addition to their signal, s^d_{ti}, and information that includes the bid(s) of the customer(s) that was (were) matched to this dealer or the fact that no customer bid has arrived.

A pure bidding strategy is a mapping from the information set of a bidder to the action space at each stage of the game. To capture everything that a bidder knows, we introduce a bidder type, labeled $\theta^g_{\tau i}$. This is the private signal of the bidder in the first and second stages of the game, and it includes information about the observed customer bid at the final stage. We use subscript τ to summarize the auction and bidding stages of dealers with $\tau = t1$ in the first stage and $\tau = t2$ when the dealer moves for the second time in the third stage in auction t. For a customer, who only moves one time $\tau = t$. With this, bidding strategies can be represented by bidding functions, labeled $b^d_{\tau i}(\cdot, \theta^d_{\tau i})$, for dealer i with type $\theta^d_{\tau i}$ at time τ in auction t, and $b^h_{ti}(\cdot, \theta^h_{ti})$, for customer i with type $\theta^h_{ti} = s^h_{ti}$ in auction t.

When choosing the bidding function, a customer anticipates that their dealer can update their own bid. This differs from Hortaçsu and Kastl (2012), who only model the bidding decision of dealers. Since the customer does not know the dealer's type, $\theta^d_{\tau i}$, they do not know whether and how the dealer will update their bid. As a result, a customer cannot be sure if the market-clearing price will increase or decrease when they marginally increase their own demand at step k—anything can happen.[15] This makes it complex for the customer to determine their optimal bid.

To render the customers' optimization problem solvable, we assume that dealers only pay attention to finite sets of L_t moments of the customers' bidding function when updating their own bid—motivated by the empirical evidence presented in Table 2.[16] Formally, a moment is a mapping μ^l_t that transforms the bidding function, $b^h_{ti}(\cdot, \theta^h_{ti})$ for type θ^h_{ti}, into a real number, $\mathbb{R}$. We restrict our attention to moments that are differentiable with respect to the quantity at each price. This includes, for example, the intercept with the price or the quantity axis, some smooth approximation of the slope, or the quantity-weighted bid, defined as follows:

$$\mu^1_t(b^h_{ti}(\cdot, \theta^h_{ti})) = \frac{b_1 q_1 + \sum_{k=2}^{K} b_k(q_k - q_{k-1})}{q_K} \tag{2}$$

[15] As a comparison, when dealers do not update their bids, the market-clearing price will weakly increase in all states of the world if the customer increases their quantity q_k at price b_k by a little bit, assuming that all other participants play as in the equilibrium.

[16] Alternatively, we could assume that customers only think that this is the case, even though the dealer responds to the full curve. This is similar to restrictions commonly used in the dynamic games literature to simplify the dimension of the state-space.

where $\{b_k, q_k\}_{k=1}^{K}$ constitutes bidding function $b^h_{ti}(\cdot, \theta^h_{ti})$.

Once all bidders submit their step function, the market clears at the highest price, P^*_t, at which the aggregated submitted demand satisfies supply. The supply, $\boldsymbol{Q_t}$, is unknown to the bidders when they place their bids because a significant fraction of the total allotment goes to the Bank of Canada, which is the largest non-competitive bidder. $\boldsymbol{Q_t}$ is distributed according to an auction-specific distribution on $[0, \overline{Q}_t]$ with strictly positive marginal density conditional on s^g_{ti} $\forall\, i, g = h, d$.

Given all bidding functions, bidder i of group g wins amount q^{g*}_{ti} at market clearing. They pay the amount they offered to win for each unit won. In case there is excess demand at the market-clearing price, each bidder is rationed pro-rata on the margin.

4.2 Equilibrium conditions

We first characterize the equilibrium in auction t, conditional on customer and dealer participation. Then, we determine the customer and dealer entry and exit decisions.

Bidding. To find the optimal bidding strategy, a dealer maximizes their expected total surplus, taking the behavior of other bidders as given. For bidder i in group g of type $\theta^g_{\tau i}$ the expected total surplus in auction t at time τ is

$$TS^g_{\tau i} = \mathbb{E}\left[\int_0^{\boldsymbol{q^{g*}_{ti}}} [v^g_t(x, s^g_{ti}) - b^g_{\tau i}(x, \theta^g_{\tau i})]dx\right]. \tag{3}$$

In all cases, the expectation is taken over the amount the bidder will win at market clearing, $\boldsymbol{q^{g*}_{ti}}$, which depends on all bidders' strategies and types as well as the unknown supply. For customers who bid a single time, $\tau = t$, and therefore $TS^h_{\tau i}$ is equivalent to TS^h_{ti}. For dealers, we have two surpluses: TS^g_{t1i} and TS^g_{t2i}, one for each bidding round.

We focus on the pure-strategy group-symmetric Bayesian Nash equilibrium (BNE) in which all dealers and customers play the same bidding strategy if they are the same type. Formally, a BNE of auction t is a collection of bidding functions, $b^g_{\tau i}(\cdot, \theta^g_{\tau i})$, such that each bidder, i, and almost every type, $\theta^g_{\tau i}$, maximizes their expected total surplus (3) each τ.

Proposition 1. *Fix a set of L_t moment functions that map a customer's bid function into a real number, $\mu^l_t : b^h_{ti}(\cdot, \theta^h_{ti}) \to \mathbb{R}$, and consider a group-symmetric BNE with $N^d > 1, N^h > 1$.*

(i) Every step k but the last step in the dealer's bid function, $b^d_{\tau i}(\cdot, \theta^d_{\tau i})$, has to satisfy

$$\Pr(b_k > \boldsymbol{P_t^*} > b_{k+1}|\theta_{\tau i}^d)\left[v_t^d(q_k, s_{ti}^d) - b_k\right] = \Pr(b_{k+1} \geq \boldsymbol{P_t^*}|\theta_{\tau i}^d)(b_k - b_{k+1}). \tag{4}$$

At the last step, $b_K = v_t^d(\bar{q}(\theta_{\tau i}^d), s_{ti}^d)$*, where* $\bar{q}(\theta_{\tau i}^d)$ *is the maximal amount the dealer may be allocated in the auction equilibrium.*

(ii) Every step k *in a customer's bid function,* $b_{ti}^h(\cdot, \theta_{ti}^h)$*, that generates moments* $m_{ti}^l = \mu_t^l(b_{ti}^h(\cdot, \theta_{ti}^h))$ *for all* l *has to satisfy:*

$$\begin{aligned}&\Pr(b_k > \boldsymbol{P_t^*} > b_{k+1}|\theta_{ti}^h)\left[v_t^h(q_k, s_{ti}^h) - b_k\right] = \\ &\Pr(b_{k+1} \geq \boldsymbol{P_t^*}|\theta_{ti}^h)(b_k - b_{k+1}) - \sum_{l=1}^{L_t} \lambda_{ti}^l \frac{\partial \mu_t^l(b_{ti}^h(\cdot, \theta_{ti}^h))}{\partial q_k} + Ties(b_{ti}^h(\cdot, \theta_{ti}^h)),\end{aligned} \tag{5}$$

with $Ties(b_{ti}^h(\cdot, \theta_{ti}^h)) = \Pr(b_k = \boldsymbol{P_t^*})\mathbb{E}[v_t^h(\boldsymbol{q_{ti}^{h*}}, s_{ti}^h)\frac{\partial \boldsymbol{q_{ti}^{h*}})}{\partial q_k}|b_k = \boldsymbol{P_t^*}] - \Pr(b_{k+1} \geq \boldsymbol{P_t^*})\mathbb{E}[v_t^h(\boldsymbol{q_{ti}^{h*}}, s_{ti}^h)\frac{\partial \boldsymbol{q_{ti}^{h*}}}{\partial q_k}|b_{k+1} \geq \boldsymbol{P_t^*}] + \Pr(b_k = \boldsymbol{P_t^*})\mathbb{E}[\frac{\partial \boldsymbol{q_{ti}^{h*}}}{\partial q_k}|b_k = \boldsymbol{P_t^*}] + \Pr(b_{k+1} = \boldsymbol{P_t^*})\mathbb{E}[\frac{\partial \boldsymbol{q_{ti}^{h*}}}{\partial q_k}|b_{k+1} = \boldsymbol{P_t^*}] + \Pr(b_{k+1} < \boldsymbol{P_t^*})\mathbb{E}[\frac{\partial \boldsymbol{q_{ti}^{h*}}}{\partial q_k}|b_{k+1} < \boldsymbol{P_t^*}]$*. Here we omitted the dependence on* θ_{ti}^h*. In addition,*

$$\lambda_{ti}^l[m_{ti}^l - \mu_t^l(b_{ti}^h(\cdot, \theta_{ti}^h))] = 0 \text{ for all } l, \tag{6}$$

with Lagrange multipliers $\lambda_{ti}^l \in \mathbb{R}$ *for all* l*.*

(iii) The moments, $\{m_{ti}^l\}_{l=1}^L$*, are such that the expected total surplus (3) is maximized and* $m_{ti}^l = \mu^l(b_{ti}^h(\cdot, \theta_{ti}^h))$ *for all* l *and all customers.*

When $N^h = 0$*,* $N^d > 1$*, we consider a symmetric BNE for which condition (i) characterizes dealer bidding, and conditions (ii) and (iii) do not apply.*

From the existing literature, we know how dealers determine their equilibrium bid. In each stage dealers choose their bids, $\{q_k, b_k\}_{k=1}^K$, to maximize the total expected surplus, $TS_{\tau i}^d$, defined in (3), subject to market clearing. They trade off the expected surplus on the marginal infinitesimal unit versus the probability of winning it, summarized in Proposition 1 (*i*). Given that it is never optimal for a dealer to submit a bid above their true value, dealer demand is never rationed in equilibrium, except for the last step. At the last step, the dealer submits their true value because it is not possible to increase the winning probability of (non-existing) subsequent steps by shading their bid.

Our innovation is to characterize the equilibrium bidding of a customer. The key difference between the dealer and the customer comes from the fact that customers take into account the dealer's response to observing their bid. For illustration, assume that dealers only observe a single customer bid. Since dealers only change their own bids in response to changes in the L_t moments of the observed customer bidding function, $\{m_{ti}^l\}_{l=1}^L$, the customer's optimality conditions can be decomposed into two parts. This helps to draw the connection to the dealer's equilibrium condition. First, for each fixed set of moments, $\{m_{ti}^l\}_{l=1}^L$, the customer's equilibrium bidding function must achieve the highest expected surplus among all of the functions that induce the same dealer updating:

$$\max_{\{q_k,b_k\}_{k=1}^K} TS_{ti}^h \text{ subject to market clearing, and } m_{ti}^l = \mu_t^l(b_{ti}^h(\cdot,\theta_{ti}^h)) \text{ for all } l \in L_t. \qquad (7)$$

The corresponding optimality conditions are summarized in Proposition 1 (*ii*). Second, among those partially optimal functions, the customer chooses the optimal one by choosing moments $\{m_{ti}^l\}_{l=1}^L$ so that the expected total surplus, TS_{ti}^h, is maximized—Proposition 1 (*iii*).[17] Both conditions generalize to the case where dealers observe more than one customer. The difference is that the customer forms a different expectation over where the market will clear, taking into account the possibility that the dealer might observe other customer bids.

Dealer updating implies that it may be optimal for a customer to place a bid above their true value, in which case they might prefer to be rationed and win less than they bid for. As a consequence, we cannot rule out that ties may occur at any step, with positive probability. Formally, in Proposition 1 (*ii*) the term $Ties(b_{ti}^h(\cdot,\theta_{ti}^h))$ includes all of the cases in which the customer ties with some other bid and the bid must be rationed.

To illustrate why it might be optimal for customers to bid above their value and tie at non-final steps with positive probability, assume customer i places their bids via dealer j. They could submit a step function with $b_k \leq v_k$ for all k. Alternatively, they could deviate and bid above their value, for example, at the fourth step, $b_4 > v_4$, as depicted in the solid and dotted lines, respectively in Figure 4A. This would increase their quantity-weighted bid, and therefore trigger dealer j to update their own bid—going from the solid line to the dotted line in Figure 4B. Note that the dealer's new bid is more aggressive on most but not all parts of the bidding function.

[17] Solving this maximization is challenging (even when we restrict our attention to moments that are real numbers) because the objective function, i.e., the expected auction surplus, is not differentiable w.r.t. these moments. To see this, consider a change in moment m_i^l. The dealer who observes the corresponding bid function updates their own bid because a change in m_i^l (weakly) changes the dealer's information set and, with that, its type, θ_i^d. Therefore, the dealer submits a different bid function. This changes the customer's beliefs about the price at which the auction will clear. Formally, the distribution of the clearing price changes when the customer fixes their own bid function—since bidding functions are step functions, a change in such a function easily leads to non-continuous jumps that render the objective function, TS_i^h, non-differentiable.

Figure 4: Bidding example

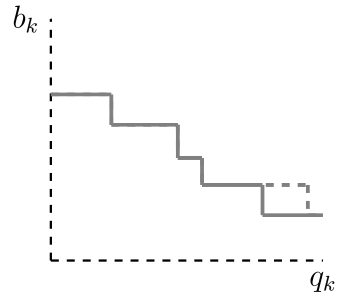

(A) Initial and new customer bid

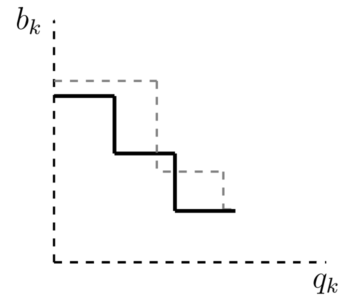

(B) Initial and updated dealer bid

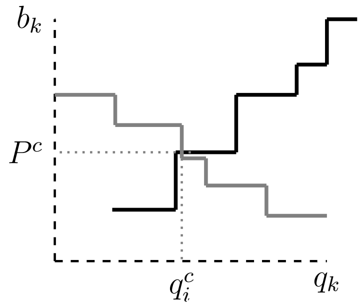

(C) Initial market clearing

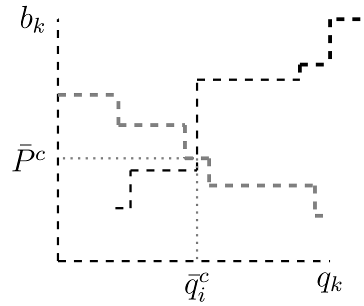

(D) New market clearing

Figure 4 provides an example for why it can be optimal for a customer to deviate from a bidding function with $b_k \leq v_k$ (shown in solid lines) and bid above their value, here, at step $b_4 > v_4$ (the dashed lines in panel A). This is true even though this triggers a more aggressive dealer bid, as shown in the dashed line of panel B. Panels C and D display one realization of market clearing for both customer bids; the downward-sloping step function is the customer's bid and the upward-sloping function is one realization of the residual supply curve that the customer faces.

How the dealer updates their bid depends on their private beliefs about where the market will clear—beliefs the customer does not observe. The customer can reason through potential market-clearing outcomes, defined by the point where their own bid intersects the residual supply curve, which in turn depends on the bids (and types) of all bidders other than the customer. Figures 4C–4D illustrate one possible market-clearing outcome before and after the customer deviates. We see that the customer finds it profitable to bid above their value in this event: rather than clearing at the second step of the customer's bidding function, the auction now clears at the third step, so that the customer wins more units at prices below their values. In the event that the residual supply curve crosses the customer's bid at the forth step, the customer would want to be rationed to win fewer units at prices above value.

Entry and exit decisions. A customer, i, enters an auction, t, if their entry cost, γ^h_{ti}, is smaller than the total surplus they expect to earn from participating in the auction before they observe their private signal. When deciding, they know the current market conditions, which are captured by the signal distributions, F^g_t, and the shape of the value functions, $v^g_t(\cdot,\cdot)$, for $g = \{h, d\}$.

Proposition 2. *Customer i with entry cost γ^h_{ti} enters auction t if*

$$\gamma^h_{ti} \leq \mathbb{E}[\boldsymbol{TS^h_{ti}}|N^d] \text{ with } TS^h_{ti} \text{ given by (3).} \tag{8}$$

Note that, relative to expression (3), we relabeled the time subscript τ to t to highlight that TS^h_{ti} is the surplus that customer i expects from participating in auction t. The expectation is taken over the customer's private signal, $\boldsymbol{s^h_{ti}}$, and is conditional on N^d dealers bidding in the auction.

Anticipating all auctions, $t = 1, ..., T$, of the upcoming year, dealer i exits the market if their entry cost is higher than the surplus they expect to earn from bidding in all T auctions of the upcoming year.

Proposition 3. *At the beginning of the year, dealer i with entry cost γ^d_i participates in the market if*

$$\gamma^d_i \leq \sum_{N^d=1}^{\bar{N}^d} \left(\sum_{t=1}^{T} \mathbb{E}\left[\boldsymbol{TS^d_{ti}}|N^d\right] \right) \Pr(\boldsymbol{N^d} = N^d), \tag{9}$$

where $TS^d_{ti} = [1 - \Psi_{ti} + \Psi_{ti} \Pr(\textit{no customer})]TS^d_{t1i} + \Psi_{ti} \Pr(\textit{at least 1 customer})TS^d_{t2i}$ *with* TS^d_{t1i} *and* TS^d_{t2i} *given by (3) for* $\tau = t1, t2$.

The dealer cares about the aggregate surplus they expect to earn over the entire year and, when making their decision, they do not know how many other dealers will compete in the auctions, N^d. Therefore, the dealer considers all possible realizations of $\boldsymbol{N^d}$ and weights each by how likely it is to occur, $\Pr(\boldsymbol{N^d} = N^d)$. Furthermore, for each auction, t, the dealer must take an expectation over bidding rounds, τ, to determine how much surplus they expect to earn in the auction, TS^d_{ti}. This is because ex-ante the dealer does not know whether they will place a final bid and earn a surplus of TS_{t2i}, or just an early bid which leaves them with a surplus of TS_{t1i}. This depends on the probability that the final bid will make it on time (Ψ_{ti}) and the probability that at least one customer is matched to the dealer.

5 Identification and estimation

The goal is to learn about the unobserved bidder values, $v_t^g(\cdot, s_{ti}^g)$, and the participation cost distributions, G^h and G^d. To achieve this, we assume that all bidders act according to our equilibrium conditions and are aware of the equilibrium they will play when making their participation decisions. This allows us to align the predicted equilibrium bidding and participation behavior with the observed data.

Identifying and estimating dealer values. For dealers, other than our object of interest—the dealer values—we can observe all of the elements in Proposition 1. This allows us to point identify dealer values at all submitted steps, q_k (Kastl (2012)). We then leverage monotonicity of the value function to construct an upper and a lower bound for the values at the intermediate quantities, $q \in (q_k, q_{k+1})$, where steps are not submitted.

To estimate dealer values, at the end-points of each step, we first need an estimate of the probabilities of where the market will clear. We estimate these probabilities by extending the resampling procedure in Hortaçsu and Kastl (2012) (see details in Appendix C). Similar to Hortaçsu and Kastl (2012), we account for differences across auctions (including those arising from different market conditions, such as secondary market spreads), by resampling within, rather than across, auctions.

Identifying and estimating customer values. It is more difficult to learn about the customer values because the customer's equilibrium condition (5) involves ties, which implies that the values at all quantities $q \in [0, q_K]$ enter the system of equations, not only those at submitted steps, q_k. This leaves us with K equations but infinitely many unknowns, so that the customer values cannot be point identified.

We can, however, construct informative bounds on the customer values, $v_t^h(q, s_{ti}^h)$, for each customer i in auction t at all q, that are consistent with the observed bids (Proposition 4 in Appendix A). For this, we assume that dealers only pay attention to the quantity-weighted bid when updating their own bid—motivated by the empirical evidence in Table 2. Condition (5) then simplifies to an equation with a single moment, the quantity-weighted bid (2). For notational convenience, we drop the superscript $l = 1$.

Before describing how we identify the customer bounds, let us take a step back and ask what variation in the data identifies these bounds. As in the existing literature, the first-order conditions for a quantity change at each submitted step provide information about the value of the bidder at that step. However, in our setting the customers' optimality condition (5) contains two additional terms, $\lambda\frac{\partial\mu}{\partial q}$, and a term that comes from ties. While $\frac{\partial\mu}{\partial q}$ is directly observed in the bidding data, λ is an additional free parameter. To pin it down, we use the additional optimality restriction of Proposition 1, according to which the chosen moment must be optimal, together with data on how a dealer changes their own bid when observing a higher or lower quantity-weighted customer bid. This additional information is not leveraged to estimate values in the standard approach.

We proceed in three steps to identify customer bounds. First, we guess a Lagrange multiplier, $\lambda_{ti} \in \mathbb{R}$, and replace the system of equations from Proposition 1 (*ii*) with a system that eliminates the infinitely many unknown values due to rationing. To do this, we utilize boundedness and monotonicity to replace the values at quantities where a step is not submitted, with a bound. For example, the upper bound on the value at quantity $q \in (q_k, q_{k+1})$ is $\overline{v}_t^h(q_k, s_{ti}^h)$ and the lower bound is $\underline{v}_t(q_{k+1}, s_{ti}^h)$. This results in a system of $2K$ equalities, which are linear in the unobserved values, with $2K$ unknown.

Second, we simplify this system of equations by showing that, at a subset of steps, rationing never occurs in equilibrium, which implies that we can cancel out all of the terms involving rationing at these steps. We do this by constructing profitable deviations at these steps in Lemma 1 in Appendix A.[18] With these simplifications, it is straightforward to express the system of equations in matrix format and show that the matrix has full rank, which proves that the system is identified.

Third, we check that the guessed Lagrange multiplier, λ_{ti}, is valid in equilibrium, by using Proposition 1 (*iii*). To do this we rely on the fact that we observe the distribution of dealer-bid updates following any customer bid. This allows us to calculate what residual supply curves a customer would have faced had they chosen any other quantity-weighted bid. With this, and the implied marginal values from λ_{ti}, we can compute bounds on the expected surplus that the customer could have achieved at any alternative quantity-weighted bid, given λ_{ti}. We then check whether the guessed λ_{ti} is consistent with the changes in the

[18] To provide an intuition for how Lemma 1 works, consider a customer with $\lambda_{ti} > 0$. If the dealer did not update their bid ($\lambda_{ti} = 0$), this customer would submit a lower quantity-weighted bid. With dealer updating, the customer instead choose a step-function that reaches an inflated quantity-weighted bid. At steps other than the last step, this may involve rationing, that is, demanding larger quantities than in the unconstrained bid so as to inflate the quantity-weighted bid without winning all of the extra units. However, at the last step, demanding a larger amount causes the quantity-weighted bid to fall and, therefore, the customer does not find it optimal to tie at the last step. Terms involving rationing at the last step are therefore zero. A similar logic applies when $\lambda_{ti} < 0$.

surplus between the submitted quantity-weighted bid and the alternative quantity-weighted bids. For example, if the guessed λ_{ti} suggests that the customer is forgoing profitable changes in their bid in order to reduce their quantity-weighted bid, then the residual supply curves that they expect to face after submitting a larger quantity-weighted bid must imply large enough losses in their surplus. However, these losses should not be so big that the customer would rather choose a bid function that resulted in an even lower quantity-weighted bid.

To estimate the model, we apply the same algorithm, replacing win probabilities by estimates obtained via resampling.[19]

Identifying and estimating cost distributions. To learn about the cost distributions, G^h and G^d, we rely on Propositions 2 and 3, respectively. For identification, we need to fix the maximal number of dealers and customers that can show up to an auction. We set these to what we observe in our sample—$\bar{N}^d = 24$ for dealers, and $\bar{N}^h = 12$ for customers.

Entry costs distributions are identified from the variation in entry behavior across auctions together with differences in bounds on expected surpluses conditional on participation, $\overline{TS}^g_{ti}$, and $\underline{TS}^g_{ti}$ (which can be constructed directly from the value bounds).[20] For customers, relying on the fact that

$$\bar{N}^h \Pr(\boldsymbol{\gamma^h_{ti}} \leq \mathbb{E}[\underline{\boldsymbol{TS}}\boldsymbol{^h_{ti}}|N^d]) \leq \mathbb{E}[\boldsymbol{N^h_t}] \leq \bar{N}^h \Pr(\boldsymbol{\gamma^h_{ti}} \leq \mathbb{E}[\overline{\boldsymbol{TS}}\boldsymbol{^h_{ti}}|N^d]) \tag{10}$$

we could identify bounds for the customer's cost distribution non-parametrically as long as we can construct sufficiently many different surpluses, $\mathbb{E}[\overline{\boldsymbol{TS}}\boldsymbol{^h_{ti}}|N^d])$ and $\mathbb{E}[\underline{\boldsymbol{TS}}\boldsymbol{^h_{ti}}|N^d])$, to cover the full support of the cost distribution. A similar argument applies for dealers but requires summing across auctions and over N^d.

In order to estimate the distribution of entry costs, we first compute bounds on how much the customers and dealers expect to gain from participating in the game—the RHS of (9) and (8), respectively. For the customers, this is straightforward. We compute each customer's auction surplus (3) at the upper and lower bounds of their estimated values, $\overline{TS}^h_{ti}$, and $\underline{TS}^h_{ti}$, conditional on participation. We then average $\overline{TS}^h_{ti}$ across all participating customers in an auction to obtain the expected surplus prior to auction entry, $\mathbb{E}[\overline{\boldsymbol{TS}}^h_{ti}|N^d]$,

[19]Concretely, we replace the probabilities in equations (16) and (17) of the proof of Proposition 4 in the Appendix.

[20]For example, $\overline{TS}^h_{ti}$ is the area between the upper bound on the customer's value function and the submitted bid function, where each quantity is weighted by its win probability at market clearance.

and similarly for the lower bound.

For dealers, the estimation is more difficult. This is because dealers make their entry decisions before knowing the number of dealers participating in the upcoming year. This implies that we need to compute bounds on their expected annual auction surplus (given estimated values). Thus, in addition to the expected total surplus, given the observed number of dealers (which we can construct), we need to compute expected total surplus when there is a different number of dealers, and how likely it is that each possible number of dealers will participate.[21] To do this, we could compute the counterfactual surpluses following a similar approach presented in our counterfactual exercises. However, this is computationally intensive, as it involves resolving equilibria for all auctions under all possible numbers of participating dealers.

Therefore, we estimate these surpluses using a matching estimator. We find auctions that are similar in terms of their average (quantity-weighted average) values of bidders, but with different numbers of participating dealers. For example, to obtain the counterfactual surplus for an auction in which a different number, $N_1^d \neq N^d$, dealers participate, we find a similar auction in which N_1^d dealers participate. For that similar auction, we compute the expected surplus for dealers: $\mathbb{E}[\overline{\boldsymbol{TS}}_{\boldsymbol{ti}}^{\boldsymbol{d}} | N_1^d]$. Repeating this exercise for other similar auctions with different numbers of dealers provides an estimate for each counterfactual surplus.

With the bounds on how much dealers and customers expect to gain from auction participation, we estimate the bounds on their cost distributions by matching the predicted participation probability of a customer and a dealer (according to Propositions 2 and 3) to what we observe in the data. For example, the predicted expected number of customers entering an auction, $\bar{N}^h \Pr(\boldsymbol{\gamma}_{\boldsymbol{ti}}^{\boldsymbol{h}} \leq \mathbb{E}[\boldsymbol{TS}_{\boldsymbol{ti}}^{\boldsymbol{h}} | N^d])$, must equal the observed expected number. With our data, we impose an exponential distribution and estimate parameter β^h for customers and parameter β^d for dealers.[22]

[21]Formally: $\Pr(\boldsymbol{N^d} = N^d) = \binom{\bar{N}^d}{N^d} (N^d/\bar{N}^d)^{N^d} (1 - (N^d/\bar{N}^d)^{\bar{N}^d - N^d}$.

[22]Concretely, for customers we estimate a set of β^h using the following criterion function: $Q'(\beta^h) = Q(\beta^h) - \inf_{\beta'} Q(\beta')$ with $Q(\beta^h) = (\frac{N_t^h}{\bar{N}^h} - H(\mathbb{E}[\underline{\boldsymbol{TS}}_{ti}^h | N^d]; \beta^h))_+)^2 + (\frac{N_t^h}{\bar{N}^h} - H(\mathbb{E}[\overline{\boldsymbol{TS}}_{ti}^h | N^d]; \beta^h))_-)^2$, where H is the CDF of an exponential distribution with parameter β^h. As the sample size grows, all points in the identified set produce criterion values of zero. To account for finite sample errors, we define a contour set of level $c_n = \log(645)/645$, based on the number of auctions in our sample, and estimate the parameter set $\{\beta^h | Q'(\beta^h) \leq c_n\}$.

6 Estimated values and costs

For each auction, t, we estimate the dealer values, $\hat{v}_t^d(q_k, s_{ti}^d)$, which we call $\hat{v}_{tik}$, and the bounds for the customer values, $\underline{v}_{tik}$ and $\overline{v}_{tik}$, at each submitted quantity step, k, of final bids. In addition, we obtain the upper and lower bounds for the parameters of the exponential cost distributions of both bidder groups, G^h and G^d.

Customer and dealer values. From Figure 5 we see that customers are typically willing to pay more than dealers per unit of the bond. This aligns with expectations, as willingness to pay declines with quantity, and customers purchase smaller quantities than dealers, who are obligated to acquire larger volumes to maintain their dealer status. We estimate that, in the median, customers are willing to pay between 4.6 and 7.3 bps more per unit of the bond than dealers. This difference is economically meaningful compared to the median market yield-to-maturity of the bonds, which is 161 bps, and is sizable compared to the median difference between the quantity-weighted average winning bid and the quantity-weighted average secondary market price on auction day (one day after the auction), which is 0.4 bps (2 bps). Further, the difference is statistically significant at the 5 percent level (see Appendix Table B3) and cannot be driven by customer selection into more-profitable auctions, given that we are considering value differences conditional on auction participation. This is true when using the lower and upper bounds of the customer values.

Over time, customer values have increased relative to dealer values, as shown in Figure 5B.[23] Since these values reflect how much an auction participant expects to earn from trading bonds in the secondary market, this finding suggests that customers anticipate increasingly larger per-unit returns from buying bonds, relative to dealers. This is in line with growing evidence that more stringent regulations for dealers since the 2007-2009 financial crisis have decreased profit margins for dealers relative to customers. This finding is also in agreement with evidence documented by Sandhu and Vala (2023), who argue that hedge funds that entered the market after 2009 are able to obtain higher than average returns in the secondary market for Canadian government bonds because they have more flexibility to employ complex and risky trading strategies.

As validation, we test whether our estimated distribution of customer values correlates

[23] This trend could come from the same customers becoming more profitable over time. Alternatively, if different customer types (e.g., hedge funds vs. pension funds) have systematically different value distributions, this could come from a change in the composition of customer types over time.

with the observed market conditions that predict customer participation. This provides a validity check for our value estimates since we have not used any information on these market conditions in our estimation. Specifically, we regress the average quantity-weighted customer value per auction, in addition to other moments of the customer value distribution, such as the standard deviation, on the explanatory variables that we used to predict customer participation shown in Appendix Table B1. Our findings, reported in Appendix Tables E2 and E3, confirm our prior that customer values are higher (and more disperse) when secondary market spreads are wide.

Finally, in Appendix Figure E4 we evaluate whether it is quantitatively important to account for the fact that sophisticated customers, such as hedge funds, anticipate that a dealer might update their own bid after observing the customer's bid. An alternative assumption would be to neglect the customer bidding incentives that are triggered by dealer updating. In that case, we could simply follow the estimation procedure of the existing literature and back out the customer values in the same way we back out the dealer values. Average values of such naive customers are lower than those of our sophisticated customers. This is because a high bid is less attractive to sophisticated customers who understand that this bid could trigger an aggressive response by their dealer. Therefore, values that rationalize high bids by sophisticated customers should be higher than by naive customers. This has implications for our counterfactuals, where we would predict less entry if we ignored customer sophistication.

Entry costs. The entry costs capture the opportunity cost of the profits that customers and dealers forgo by participating in the primary auctions. These costs differ from values in that they are independent of the amount purchased. They must be paid when an institution spends time bidding at auction (and in the case of dealers, conducting other market- making activities), even if the institution does not buy any bonds.

On average, we estimate an annual cost of being a dealer of between C$3.263 and C$4.117 million, with a (bootstrapped) confidence interval of [C$3.111M, C$4.296M]. Since there are 28 auctions in an average year, the per-auction dealer cost is between C$3.263M/28 = C$116,535 and C$4.117M/28 = C$147,036. In comparison, Lu and Wallen (2024) report that a typical trader of a U.S. bank trading desk generates $84,000 in profits per day, which is equivalent to approximately C$118,000. Thus, if a trader spends a day preparing for and bidding in a Treasury auction, the forgone profit is close to our cost estimate.

The average entry cost of a customer is between C$471,505 and C$492,605 per auction,

Figure 5: Difference between customers and dealers' values

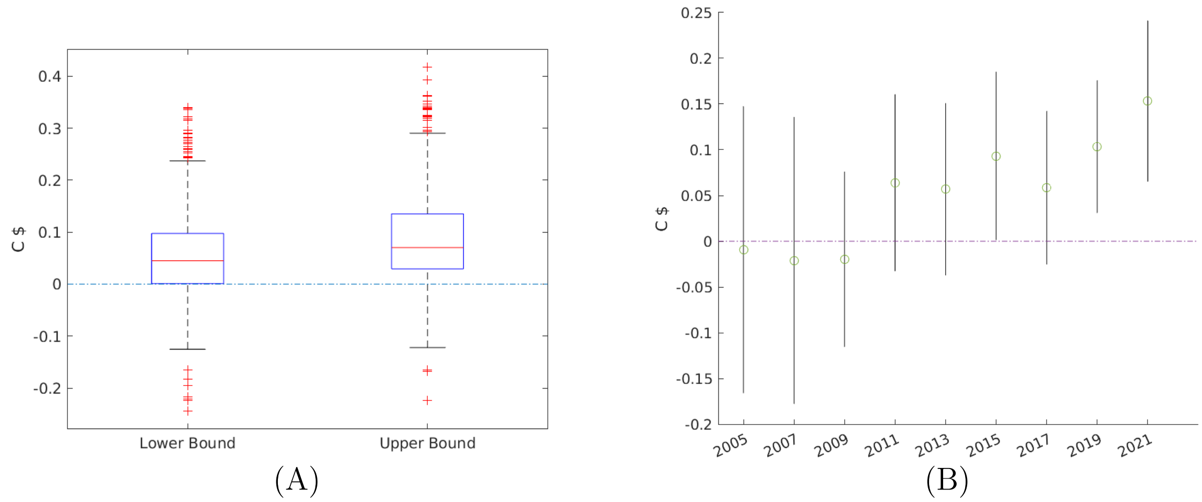

Figure 5A shows the distribution of the difference in the quantity-weighted average values between participating customers' and dealers, using customer lower bound values on the LHS and upper bound values on the RHS. We remove outliers, defined as those with a value that is more than three scaled median absolute deviations from the median. Figure 5B plots the point estimates and the 95% confidence intervals from regressing the difference between the lower bound of the average customer value and the lower bound of the average dealer value, winsorized at 1 and 99 percentiles, in each auction on a set of time indicator variables. The first indicator is 1 for 1999-2000, the second indicator is 1 for 2001-2002, and so forth, until 2021-2022. The confidence intervals are computed using the point estimates of the dealer and customer values at the lower bounds and, therefore, do not account for noise in the estimation of these values. Prices are in C$ with a face value of $100.

with a (bootstrapped) confidence interval of [C$401,204, C$562,905]. The customer's cost is larger than the dealer's cost, which is in line with the idea that customers are better at executing profitable trading strategies—here, outside of the Treasury market, driving up the opportunity cost.

Summarizing, our findings highlight systematic differences between dealers and customers, both in terms of their values and their entry costs.

7 Drivers and consequences of customer participation

Using the estimated model, we conduct counterfactuals to clarify the drivers of customer entry and to assess the associated costs and benefits for debt issuance. For this, we assume that all model primitives, such as the distributions of values and costs, remain fixed when we change the market rules. We use final bids only and the lower bound estimates of the values and costs.

Computing counterfactuals. Conducting counterfactuals for multi-unit auctions in which bidders have multi-unit demand is challenging since it is analytically impossible to solve for an equilibrium. We proceed using a numerical solution approach, detailed in Appendix D.1.

We extend the empirical guess-and-verify approach of Richert (2021) to allow bidders to expand their auction-specific demands in response to changes in the auction environment. We limit the amount of expansion to the quantity that would result if the bidder submitted the greater of (*i*) their largest ever bid and (*ii*) a bid for 10 percent of the supply. The bidder-specific limit accounts for the unobservable constraints that restrict the amount a bidder is able to buy, such as balance sheet constraints; the 10 percent bound approximates the minimal bidding requirement of dealers in normal times.[24]

Concretely, we solve for the counterfactual bid distributions such that two conditions are satisfied. First, the distribution of the values implied by these bids and optimal bidding must be indistinguishable from the distribution of the estimated values. Second, the counterfactual distribution of the maximal amount each bidder demands in an auction must be first-order stochastically dominated by the observed distribution of maximal demands. We obtain this distribution by first fixing each bidder and computing the maximum between the largest fraction of the bond supply this bidder bids on in any auction and 10 percent of this supply. We then take the empirical distribution of these maximal quantities across bidders. For robustness, we alternatively impose constraints such that when each dealer is removed the coverage ratio—the ratio of the sum of all bids over supply—falls by the same amount as the average decrease in coverage when a dealer exits, as estimated using the 2014-2015 exits.

Our approach does not account for endogenous changes in the value or cost distributions, given that we need to fix the estimated model primitives to compute counterfactual bids and entry decisions. For example, since bidders at least partially derive value from the money they make from reselling bonds, we might be concerned that their values change when their competitors win more at auction—since this impacts the degree of competition the bidder faces in the secondary market. Sufficiently small changes in the auction allocation, however, should result in similar degrees of competition in the secondary market. We therefore focus our discussion on local changes where our approach provides more-reliable predictions.

To avoid focusing on any particular draw from the cost or value distribution, we take the ex ante perspective and compare expected market outcomes. For instance, we analyze

[24]For small changes in the number of dealers, this approach may be conservative, given that dealers do not increase their maximum quantity demanded when a dealer exits the market—see Appendix Table E4.

the expected price at which an auction clears under the current market rules, relative to counterfactual market rules.

Why did customers enter? We start by examining why customers, specifically hedge funds, entered the primary market (Figure 2B). Although it is difficult to pinpoint the exact causes—given many potential, endogenous, and largely unobserved factors that may drive the observed time trends—we can use our structural model to try and separate potential forces into three broad categories. Ideally, we would compute counterfactual outcomes for all auctions in our sample; however, this is computationally intense. We therefore only consider every third auction from 2014, when hedge funds became the dominating customer group, up to 2020, when debt issuance spiked as a result of the pandemic.

There is growing evidence that regulatory changes following the global financial crisis made it relatively more expensive for dealers to intermediate in bond markets (e.g., Adrian et al. (2017), Bessembinder et al. (2018), Duffie (2023), and Duffie et al. (2023)). Through the lens of our model, this is reflected in customers being willing to pay more for Canadian bonds compared to dealers (Figure 5B)—our first potential driving factor. Using our auction data and structural model, we quantify the extent to which the rise in customers' willingness to pay relative to dealers contributed to their increased participation. We do this by conducting a counterfactual where we take the relative values from Figure 5B and remove the increasing trend. Operationally what this means is that for every auction where customer values are larger than dealer values, we shift down the customer-value distribution by the mean difference in the quantity-weighted average values of the two groups. In doing so, we eliminate the higher customers' willingness to pay. We find that a customer's participation probability falls, on average, by 8.23 percentage points, or 25.07 percent.

The second driving force we assess, also likely due to regulation, is dealer exit. To assess whether dealer exit created room for entry, we compare the status quo, in which two dealers left in 2014, with a counterfactual having two additional dealers (fixing dealer participation but allowing customers to select into auctions). Figure 6 plots the entry probabilities on the x-axis and the counterfactual entry probabilities on the y-axis. Blue circles denote entry in the counterfactual with two additional dealers. On average, adding two dealers reduces a customer's participation probability by 10.70 percentage points, or 31.99 percent. Therefore, without dealer exit some customers might not have chosen to enter the primary market.

The third potential driving force we assess is increasing auction sizes over our sample

Figure 6: Why did customers enter?

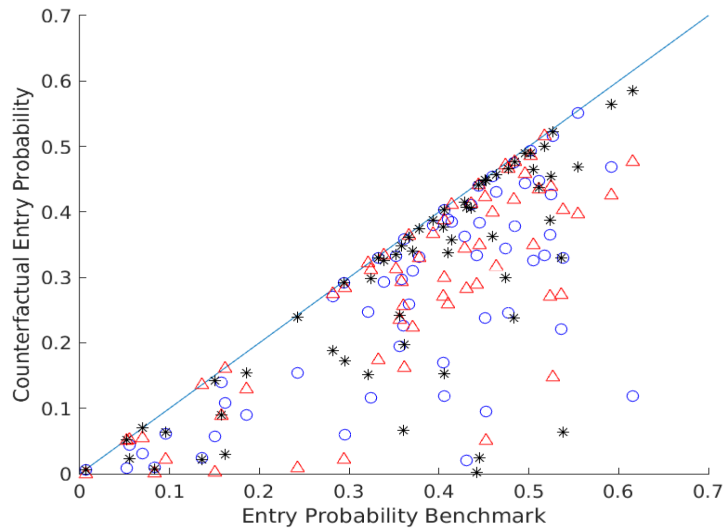

Figure 6 shows the customer's participation probability in percentage points for every 3rd auction from 2014 onward in the status quo (on the x-axis) and the three counterfactual entry probabilities (on the y-axis). Blue circles denote entry probabilities adding back two dealers. Red triangles denote entry probabilities with reduced supply. Black stars denote entry probabilities with low customer values.

period that were needed to accommodate increasing government debt—a common feature across many countries post-financial crisis and the COVID-19 pandemic. Margins on bond purchases are low, so that sufficient gains can be generated only when quantities are large enough. In the counterfactual we scale down the post-2014 issuance by the ratio of the mean issuance pre-2014 to the mean issuance post-2014. Figure 6 (red triangles) plots the counterfactual entry probabilities against the benchmark entry probabilities. Similar to our other counterfactuals, there is dispersion across auctions. On average, a customer's participation probability falls by 9.41 percentage points, or 28.81 percent.

In summary, our counterfactuals attribute sizable and fairly similar effects of changes in relative values, dealer exit, and changes in issuance size as drivers of larger customer participation. Looking ahead, we expect this trend to continue given the substantial upcoming debt issuance and the prospect of further dealer exits. The latter is not only a hypothetical concern—in Canada as one dealer (RBC) recently acquired another dealer (HSBC), while Europe continues to face a decline in the number of primary dealers (Khadbai, 2023).

Competition benefits and volatility costs. Next, we quantify the competition benefits and volatility costs that arises when adding market participants who do not participate with regularity.

To build an intuition for competition benefits, we first consider a hypothetical auction environment with only dealers. This eliminates the effects coming from changes in the bidder composition or from dealer bid updating. In this simplified auction environment, we ask by how much the expected auction price varies in the number of competing dealers. If 14 dealers compete in a typical auction, shown in Figure 7, the market clears at a competitive price, which is similar to the observed one. If the number of dealers competing drops to 13 or 12, the expected price drops by close to 6 basis points and 41 basis points, respectively.[25]

Then, we assess how adding customers (in expectation) to a typical auction affects competition while holding the number of dealers fixed. Because customers bid through dealers, the bidding dynamics become more complex than in an environment with only dealers. Even so, adding bidders should raise prices and revenues in expectation. But because customers typically submit smaller bids than dealers, and because we keep the number of dealers fixed—so the auction is already crowded—we expect the competition effect from adding customers to be smaller than the effect from varying dealers.

Our estimates confirm our prior: Figure 8A plots by how the expected price (red diamonds) increases when we increase the customer's participation probability (implying a larger expected number of participating customers). Adding one customer in expectation increases the expected clearing price by 0.4 basis points. Expected revenue increases by 0.5 basis points (or C$0.149 million), and bid shading decreases by 2.72 percent when the expected number of customers increases by one relative to the status quo. With fewer dealers the effects are larger, as shown in Figure 8B (red diamonds). This indicates that losing additional dealers could have sizable effects unless an adequate number of new customers enters the market.

Next, we examine how irregular participation affects auction revenues. Conceptually—in a world without auction heterogeneity—volatility losses stem from two intertwined forces: (*i*) a strategic shading effect (within an auction), as bidders anticipating random competition shade their bids more aggressively; and (*ii*) a concavity effect (operating across auctions), since expected revenue is concave in the number of participants, implying that adding one customer raises revenue less than removing one reduces it. These two effects are not additively separable because strategic bid shading responds in a non-additive way whenever the auction environment changes—for example, when participation shifts from random to fixed

[25]These large changes rely on the assumption that the dealers are limited in their increases in demand around an exit. This is consistent with evidence when a dealer exits the market—see Appendix Table E4.

Figure 7: Typical auction: Varying number of dealers

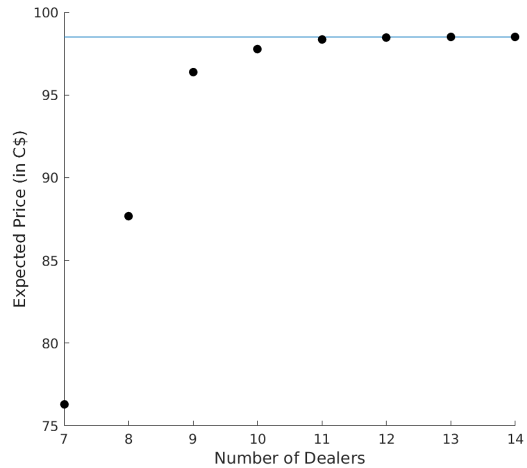

Figure 7 shows how the expected prices at which an auction clears varies as the number of dealers increases in an auction that issues the average supply with medium customer participation, i.e., 4 customers. Given that our numerical procedure to determine counterfactual bids allows for small differences between the true value distribution and the one implied by the counterfactual bids (Richert 2021), we plot the midpoint of the expected clearing price ranges at each number of dealers. The blue horizontal line shows the average observed bid, which is close to the observed market-clearing price.

or when the number of bidders varies.

To isolate the within-auction effect stemming from bid-shading, we eliminate the randomness bidders face about their number of competitors within an auction by fixing the number of customers. In Figure 8A (blue dots), we plot expected clearing prices fixing the number of participating customers rather than the expected number of participating customers. As expected, this causes bidders to shade less. The effect is starker when fewer bidders participate, as shown in Figure 8B (blue dots), where we remove two dealers.

To illustrate the concavity effect (which operates across auctions), we compare the expected revenue gain when adding a customer in expectation in a typical auction to the loss when removing one. The net expected revenue loss is C$0.16M per auction, or 0.496 bps of bonds issued and the net expected gain from adding one customer is C$0.149M per auction, or 0.53 bps of bonds issued. While this comparison helps build intuition, in reality no two auctions are identical.

Auction heterogeneity implies that each auction has its own expected revenue–participation curve. Consequently, we are never moving along a single curve, making it difficult to illustrate the loss graphically. Nonetheless, the key insight remains: variation in participation probabilities across auctions leads to cross-auction revenue losses. To compute the actual across-auction volatility loss, we compare total revenues from fixing participation probabilities (at the median) across auctions to expected revenues with auction-specific fixed, yet auction-specific participation probabilities. On average, revenues are 0.74 bps higher or C$0.223M on C$3B of bonds issued with fixed participation probabilities than with across-auction variation in participation probabilities.

To assess the total volatility cost and compare it to the benefit from increased competition, we eliminate uncertainty about opponents' participation by fixing customer participation at its median level across all auctions and computing the expected revenue change.[26] This increases expected revenue by C$0.488M on C$3B issued relative to the status quo. In comparison, adding one customer in expectation had a smaller effect (C$0.149M), implying that in the current environment volatility costs are larger than competitive benefits.

To get a better sense of how large the total volatility cost is, we benchmark it against the compensation the auctioneer would need to offer customers—who face an additional entry cost when forced to participate—to make them indifferent to participating: C$0.215M. While this is a useful conceptual exercise, it is not an implementable policy. The reason is that the auctioneer cannot target their transfers based on the customer's private information about their entry costs. We discuss implementable policy ideas next.

Alternative policies. Canada, like many countries, have a debt-management objective to raise stable and low-cost funding and maintain a well-functioning market for Government of Canada securities. Within this objective there is some flexibility, so long as changes do not surprise markets (Garbade (2015)). We therefore ask whether there are policies that could both reduce volatility and increase competition.[27] Given the existing evidence that the efficiency gains from moving to first-best from the status are small—Kastl (2011)—we do not expect large gains from these interventions. Nonetheless, thinking through these options is useful should auction conditions change, for instance, because more dealers exit.

The first—admittedly blunt—idea is to subsidize customer participation. Without the ability to target specific customers (based on customers' private information), the auctioneer

[26] An alternative way to assess volatility costs is to compute how many committed customers would be required to generate the same revenue as under the status quo with irregular customer participation. We find that if the government could find 1-2 committed customers, it would generate the same expected revenue as the status quo with a median of 4 irregular customers, depending on the year.

[27] Although not included, our framework also allows us to evaluate the role of direct versus indirect bidding. For instance, letting customers bid directly to the auctioneer instead of via dealers would decrease dealer information rents, and could increase revenues.

Figure 8: Typical auction: varying customer participation probability

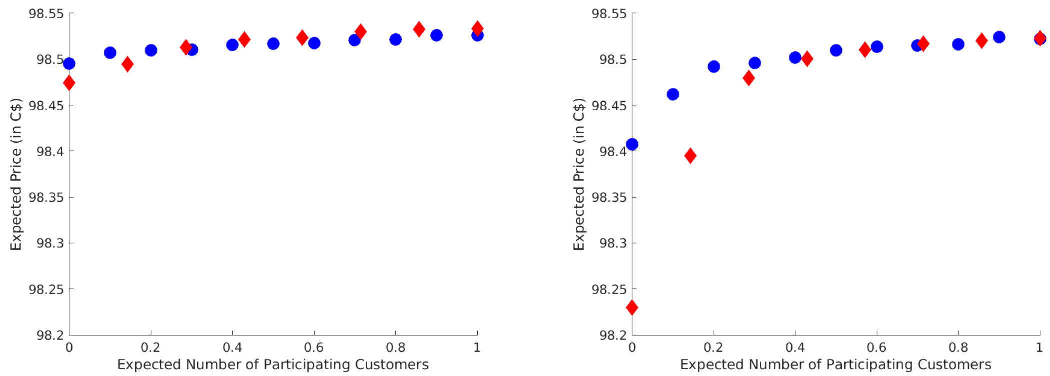

(A) Customer participation (B) Customer participation less dealers

Figure 8 shows how the expected price (in C$) at which an auction clears varies as the participation probability of (all) customers varies between 0 and 1, computed at an evenly spaced grid point in an auction that issues the average supply with medium customer participation, i.e., 5 customers. The blue dots denote the outcomes under deterministic participation of opponents while the red diamonds denote random participation of opponents. The right panel, shows the same auction but with 2 less dealers. We determine a range of prices given that our numerical procedure to determine counterfactual bids allows for small differences between the value distribution that is implied by the counterfactual bids and the true value distribution (see Richert 2021) and we plot the midpoint of the expected clearing price ranges at each number of dealers. The blue horizontal line shows the average observed bid, which is close to the observed market-clearing price.

could offer a uniform transfer to all of them. The total cost would be C$0.446M per auction. Because this amount is close to the revenue gain (C$0.488M), this policy effectively amounts to a zero-sum transfer across bidders.

The second idea is to modify commitment requirements. We first eliminate dealer commitment, allowing them to freely decide whether to enter each auction, in the hope that this increases competition without harming volatility. Then, we require customers to commit in the same way as dealers, in the hope that this decreases volatility without harming competition. Due to endogenous bidder participation, in both cases, it is theoretically possible to increase competition and decrease volatility relative to the status quo. Both policies would be implementable and have been previously considered by the Bank of Canada. Empirically, we find that neither of these alternative policies achieve the goal of greater expected revenue (see Appendix D.2).

The final idea is to strategically reshuffle supply across auctions to incentivize and stabilize customer participation. The idea is that we can, based on market observables, predict how many customers each auction would attract under the current supply schedule, and then shift a small fraction of the supply from the attractive to unattractive auctions.[28] We implement this idea in Appendix D.3. A key challenge is that, within a fiscal year, supply (to competitive bidders) can only be adjusted using information revealed up to that point—we cannot condition on future outcomes. As the year progresses, the number of remaining auctions shrinks, and with a fixed annual debt-issuance target, earlier adjustments constrain what is feasible later. This creates the risk of having to issue too much or too little toward the end of the year, depending on earlier choices. While this policy does help stabilize customer participation, it likely conflicts with the guidelines that debt issuance be regular and predictable. For this reason, we do not view it as a viable option.

8 Conclusion

We study dealer exit and customer entry in the primary market for Canadian government debt and analyze some consequences for the functioning of the market. We show that customer participation has increased, but remains highly irregular. We introduce and estimate a structural model to compare the benefits of higher competition from customer entry with the costs of irregular customer participation. Our framework could be used in other settings with regular and irregular market participants.

References

Ackerman, A. and Hilsenrath, J. (2022). Regulators look to lessen Treasury market reliance on big bank dealers. The Wall Street Journal.

Adrian, T., Boyarchenko, N., and Shachar, O. (2017). Dealer balance sheets and bond liquidity provision. *Journal of Monetary Economics*, 89:92–109. Carnegie-Rochester-NYU Conference Series on the Macroeconomics of Liquidity in Capital Markets and the Corporate Sector.

Allen, J. and Wittwer, M. (2023). Centralizing over-the-counter markets? *Journal of Political Economy*, 131(12):3310–3351.

Allen, J. and Wittwer, M. (2025). Bundling trades in over-the-counter markets. mimeo.

Altstedter, A. (2014). Auction bids falling as foreign banks retreat: Canada credit. Bloomberg.

[28]In practice, this amounts to the central bank purchasing more or less of the newly issued debt, rather than the government actually changing the total amount of debt issued per auction.

An, Y. and Song, Z. (2018). Quantitative easing auctions of Treasury bonds. *Journal of Financial Economics*, 128:203–124.

An, Y. and Song, Z. (2023). Does the Federal Reserve obtain competitive and appropriate prices in monetary policy implementation? *Review of Financial Studies*, 36(10):4113–4157.

Ausubel, L. M., Cramton, P., Pycia, M., Rostek, M., and Weretka, M. (2014). Demand reduction and inefficiency in multi-unit auctions. *The Review of Economic Studies*, 81(4):1366–1400.

Bahceli, Y., Mackenzie, N., Robertson, H., and Ranasinghe, D. (2024). Hedge funds shake up the Euro zone's $10 trillion government bond market. *Reuters*. Accessed: 13/10/2025.

Banegas, A., Monin, P. J., and Petrasek, L. (2021). Sizing hedge funds' Treasury market activities and holdings. FEDS Notes.

Bank of England (2025). Financial stability report. Technical report, Bank of England.

Barth, D. and Kahn, J. (2020). Basis trades and Treasury market illiquidity. OFR Brief Series.

Bessembinder, H., Jacobsen, S., Maxwell, W., and Venkataraman, K. (2018). Capital commitment and illiquidity in corporate bonds. *The Journal of Finance*, 73(4):1615–1661.

Bessembinder, H., Spatt, C., and Venkataraman, K. (2020). A survey of the microstructure of fixed-income markets. *Journal of Financial and Quantitative Analysis*, 55(1):1–45.

Bhattacharya, V., Roberts, J., and Sweeting, A. (2014). Regulating bidder participation in auctions. *RAND Journal of Economics*, 45:675–704.

Blommestein, H., Guzzo, V., Holland, A., and My, Y. (2010). Debt markets: Policy challenges in the post-crisis landscape. OECD-Financial Market Trends.

Boyarchenko, N., Eisenbach, T. M., Gupta, P., Shachar, O., and Tassel, P. V. (2020). Bank-intermediated arbitrage. Working paper.

Boyarchenko, N., Lucca, D. O., and Veldkamp, L. (2021). Taking orders and taking notes: Dealer information sharing in Treasury auctions. *Journal of Political Economy*, 129(2):607–645.

Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *American Economic Review*, 86(1):180–194.

CGFS (2016). CGFS papers no 55: Fixed income market liquidity. *Bank of International Settlements*.

Cohen, A., Kargar, M., Lester, B., and Weill, P.-O. (2024). Inventory, market making, and liquidity in OTC markets. *Journal of Economic Theory*, 222:105917.

Dixon, L., Clancy, N., and Kumar, K. B. (2012). Hedge funds and systemic risk. RAND Corporation ISBN 978-0-8330-7684-7.

Du, W., Tepper, A., and Verdelhan, A. (2018). Deviations from covered interest rate parity. *Journal of Finance*, 73(3):915–957.

Duffie, D. (2023). Resilience redux in the US Treasury market. Jackson Hole Symposium.

Duffie, D., Fleming, M. J., Keane, F. M., Nelson, C., Shachar, O., and Tassel, P. V. (2023). Dealer capacity and U.S. Treasury market functionality. *FRB of New York Staff Report*, (1070).

Favara, G., Infante, S., and Rezende, M. (2022). Leverage regulations and Treasury market participation: Evidence from credit line drawdowns. Working paper.

Garbade, K. D. (2015). Treasury debt management under the rubric of regular and predictable issuance: 1983–2012. Technical report, Federal Reserve Bank of New York. Available at FRASER: https://fraser.stlouisfed.org/title/9325.

Grossman, M. and Goldfarb, S. (2022). Rocky Treasury-Market trading rattles Wall Street. The Wall Street Journal.

Guerre, E., Perrigne, I., and Vuong, Q. (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68(3):525–574.

Hazelkorn, T. M., Moskowitz, T. J., and Vasudevan, K. (2022). Beyond basis basics: Liquidity demand and deviations from the law of one price. *Journal of Finance*, 78:301–345.

Hendricks, K., Houde, J.-F., and Raisingh, D. (2023). Asymmetric information in the wholesale market for mortgages: The case of Ginnie Mae loans. Working paper.

Hortaçsu, A. (2002). Bidding behavior in divisible good auctions: Theory and evidence from the Turkish Treasury auction market. Working Paper.

Hortaçsu, A. and Kastl, J. (2012). Valuing dealers' informational advantage: A study of Canadian Treasury auctions. *Econometrica*, 80(6):2511–2542.

Hortaçsu, A. and Perrigne, I. (2021). Empirical perspectives on auctions. In: Handbook of Industrial Organization, Chapter 11.

Kastl, J. (2011). Discrete bids and empirical inference in divisible good auctions. *Review of Economic Studies*, 78:974–1014.

Kastl, J. (2012). On the properties of equilibria in private value divisible good auctions with constrained bidding. *Journal of Mathematical Economics*, 48:339–352.

Khadbai, B. (2023). EU primary dealer exits underline pressure facing smaller banks. *Global Capital*. Accessed: 13/10/2025.

Logan, L. K. (2020). Speech: Treasury market liquidity and early lessons from the pandemic shock. Accessed on 12/08/2022 at: https://www.newyorkfed.org/newsevents/speeches/2020/log201023.

Lou, D., Pinter, G., and Uslu, S. (2023). Fiscal-monetary interactions: Pre-announcement liquidity effects after bond issuance. Working Paper.

Lu, L. and Wallen, J. (2024). What do bank trading desks do? Working paper.

Milgrom, P. (2004). *Putting Auction Theory to Work*. Cambridge University Press, Cambridge.

Petersen, C., Reguant, M., and Segura, L. (2024). Measuring the impact of wind power and intermittency. *Energy Economics*, 129.

Quint, D. and Hendricks, K. (2018). A theory of indicative bidding. *American Economic Journal: Microeconomics*, 10(2):118–151.

Richert, E. (2021). Indirect inference techniques as an alternative to direct computation of equilibrium bid distributions. Working paper.

Richert, E. (2022). Quantity commitments in multiunit auctions: Evidence from credit event auctions. Working paper.

Rüdiger, J., Kastl, J., and Gonzalez-Eiras, M. (2023). How much is being a primary dealer worth? Evidence from Argentinian Treasury auctions. Working paper.

Sandhu, J. and Vala, R. (2023). Do hedge funds support liquidity in the government of Canada bond market? Bank of Canada Staff Analytical Note.

Stein, J. (2009). Presidential address: Sophisticated investors and market efficiency. *Journal of Finance*, 64:1517–1548.

Vissing-Jorgensen, A. (2021). The Treasury market in spring 2020 and the response of the Federal Reserve. *Journal of Monetary Economics*, 124:19–47.

Wittwer, M. and Allen, J. (2022). Capital constraints and market power. Working paper.

ONLINE APPENDIX

Entry and Exit in Treasury Auctions

By Jason Allen, Ali Hortaçsu, Eric Richert, and Milena Wittwer

Appendix **A** presents mathematical details, including formal proofs.

Appendix **B** presents additional empirical tests supporting our model.

Appendix **C** explains how we estimate dealer and customer values.

Appendix **D** provides details about counterfactuals, and presents additional findings.

Appendix **E** provides additional tables and figures referenced in the main text.

A Mathematical appendix

Here we prove Proposition 1. The proof relies on Lemma 1, stated and proved below. Finally, we prove that customer value bounds are identified, which is summarized in Proposition 4.

Proof of Proposition 1. The proof of statement (i) and the case when $N^d > 1$ but $N^h = 0$ is analogous to the proof of Proposition 1 in Kastl (2011). To derive the conditions for the case when $N^d > 1, N^h > 1$, take the perspective of customer i and assume that all other bidders play an equilibrium. For ease of notation, we drop the auction and time subscripts, t and τ, and fix L moment functions, μ^l. Given that dealers only update their own function when the customer submits a function with at least one different moment $\{m_i^l\}_{l=1}^L$, we can decompose the conditions that characterize the customer's best reply (and, with that, an equilibrium, given that the other bidders play, by assumption, an equilibrium) into two parts. First, we fix some set of moments, $\{m_i^l\}_{l=1}^L$, and find $b_i^h(\cdot, \theta_i^h)$ that maximizes the customer's expected total surplus (3) such that such that $\mu^l(b_i^h(\cdot, \theta_i^h)) = m_i^l$ for all l. Statement (ii) provides the first-order conditions of this problem. Second, in equilibrium, a customer must choose a bidding function that gives rise to moments, $\{m^l\}_{l=1}^L$, so that the total expected surplus is maximized globally. This is specified in statement (iii).

To derive the first-order conditions, denote Lagrange multipliers of the constraints by λ_i^l. The objective function is $TS_i^h - \sum_{l=1}^L \lambda_i^l \left(m_i^l - \mu^l(b_i^h(\cdot, \theta_i^h)\right)$. Given this, we can follow the perturbation argument in the original proof in Appendix A.2 of Kastl (2011), step by step, with one difference. There is an additional term that comes from the constraints that all moments, l, of the chosen function, $\mu^l(b_i^h(\cdot, \theta_i^h)$, must equal the fixed moments, m_i^l.

This term does not create issues when taking derivatives since the moment functions are differentiable w.r.t. quantity. Further, since dealers only update their own bids if a moment of the customer's bidding function changes, and we are keeping these moments fixed, no complications arise when predicting which states of the world the market will clear relative to the original proof. Following the steps in Kastl (2011), we obtain equations A.2 and A.3 in Kastl (2011) plus the term $\sum_l^L \lambda_i^l \frac{\partial \mu^l(b_i^h(\cdot,\theta_i^h)}{\partial q_k}$. Combining these equations gives condition (5) in Proposition 1. Unlike for dealers, these expressions do not simplify further since it may be optimal for customers to tie in steps other than the last one—see Lemma 1.

Lemma 1. *(i) For a customer, ties occur with zero probability for a.e. s_i^h in any equilibrium for either all steps except the last step (if $\lambda_i \leq 0$), or at the last step (if $\lambda_i > 0$). (ii) For a dealer, Lemma 1 from Kastl (2011) applies.*

Proof of Lemma 1 (i)**.** For ease of notation we eliminate the auction and time subscripts, t and τ, as well as the customer superscript, h.

Case 1: $\lambda_i > 0$. Consider the last step, $k = K$. Suppose bidder i ties on step $k = K$. Take some $q_m = sup\{q|v(q, s_i) > b_k\}$ and let $\bar{q} = max\{q_k - \delta, q_m\}$ with δ some strictly positive number bounded above by $q_k - q_{k-1}$. By moving from q_k to $\bar{q}$, the bidder either moves to where they get a positive surplus from the amount purchased or they buy fewer units at a negative surplus. Take the deviation to bid $b'_{ik} = b_{ik} + \epsilon$, where $\epsilon > 0$ is sufficiently small at $\bar{q}$, the associated step, either bidder i gets fewer units than they would in a tie and no longer pays for those units or they get more units but on each of these units earns an additional positive surplus.

Given dealer updating, this deviation is strictly profitable. To see this, let us abbreviate the moment, that is, the quantity-weighted bid of the deviated bidding function by $\mu(\bar{q}, b')$ and, similarly, the original bidding function. Deviating to b'_{ik} comes at a constraint penalty of $-\lambda_i(m_i - \mu(\bar{q}, b'))$, since the targeted moment, $m_i = \mu(q_k, b)$, is no longer met. Because $\mu(\bar{q}, b') > \mu(q_k, b) = m_i$ and $\lambda > 0$, by assumption, $-\lambda_i(m_i - \mu(\bar{q}, b')) > 0$, representing a strictly positive profit from deviating.

Case 2: $\lambda_i < 0$. Suppose bidder i ties on a step $k < K$. Take some $q_m = sup\{q|v(q, s_i) > b_k\}$ and let $\bar{q} = max\{q_k - \delta, q_m\}$, with some strictly positive step size, δ, that is bounded above by $q_k - q_{k-1}$, that is, the bidder steps either toward where they get a positive surplus from the amount purchased, if possible, or, if not, at least they buy fewer units at prices above their values (reducing losses). Take the deviation to bid $b'_{ik} = b_{ik} + \epsilon$, where $\epsilon > 0$

is sufficiently small, at $\bar{q}$ (the associated step), either bidder i gets fewer units than they would in a tie and no longer pays for those units, or they get more units but on each of these units earns an additional positive surplus. Similar to the first case, there is a positive gain that arises due to dealer updating. The deviation comes at a constraint penalty of $-\lambda_i(m_i - \mu(\bar{q}, b'))$. Because $\mu(\bar{q}, b') < \mu(q_k, b) = m_i$, the bracketed term is positive. By assumption, λ_i is negative, so that $-\lambda_i(m_i - \mu(\bar{q}, b')) > 0$.

Case 3: $\lambda_i = 0$. The original proof from Kastl (2011) applies. □

Proposition 4. *Given customer s^h_{ti} behaves according to Proposition 1 (ii) and (iii), and the dealer only pays attention to one moment of the customer's bidding function, that is, $L = 1$, the upper and lower bounds on the customer's values, $\overline{v}^h_t(\cdot, s^h_{ti})$ and $\underline{v}^h_t(\cdot, s^h_{ti})$, are identified.*

Proof of Proposition 4. We fix an auction, t, and a customer, i, and drop the auction, t, time, τ, customer, h, bidder, i, and moment l subscripts and superscripts, for simplicity. The identification argument is complicated by the fact that condition (5)—here, slightly rearranged—not only contains values at the submitted steps, $v(q_k, s)$, but also the values at some intermediate quantities between the submitted steps, $v(\boldsymbol{q^*}, s)$:

$$\begin{aligned} 0 = &\Pr(b_k > \boldsymbol{P^*} > b_{k+1}|\theta, m)v(q_k, s) + \\ &\Pr(b_k = \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[v(\boldsymbol{q^*}, s)\frac{\partial \boldsymbol{q^*}}{\partial q_k}\Big| b_k = \boldsymbol{P^*}, \theta, m\right] + \\ &\Pr(b_{k+1} \geq \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[v(\boldsymbol{q^*}, s)\frac{\partial \boldsymbol{q^*}}{\partial q_k}\Big| b_{k+1} \geq \boldsymbol{P^*}, \theta, m\right] - \\ &\Pr(b_k > \boldsymbol{P^*} > b_{k+1}|\theta, m)b_k - \\ &\Pr(b_{k+1} \geq \boldsymbol{P^*}\Big|\theta, m)(b_k - b_{k+1}) - \\ &\Pr(b_k = \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_k}\Big| b_k = \boldsymbol{P^*}, \theta, m\right] b_k - \\ &\Pr(b_{k+1} = \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_k}\Big| b_{k+1} = \boldsymbol{P^*}, \theta, m]\right] b_{k+1} + \\ &\lambda\frac{\partial \mu(b(\cdot, \theta))}{\partial q_k}. \end{aligned} \tag{11}$$

Here, and for all other expressions in this section, we include the fixed moment, m, as a condition alongside the bidder's type. Dependence on intermediate quantities implies that the customer values cannot be point identified.

To identify the customer's value bounds, we first guess a Lagrange multiplier, $\lambda \in \mathbb{R}$, and simplify the system of equations (11) to obtain a set of conditions for $\lambda > 0$ and one for $\lambda < 0$. We start by relying on monotonicity and boundedness of the value function: for all $q_{k-1} \leq q \leq q_k$, we know that $v(q_{k-1}, s) \geq v(q, s) \geq v(q_k, s)$. In addition, we sign the derivatives of the rationed quantity as follows: increasing the bid at step k increases the quantity rationed in the event of a tie at step k and decreases the quantity rationed in the event of a tie at step $k+1$.

Next, we eliminate terms in condition (11) to obtain a system of $2K$ linear equations with $2K$ unknowns. To do this, we obtain an upper bound on the value at step q_k by making the terms involving the value and the derivative of the rationed quantity in the event of a tie as small as possible (and making them as big as possible for the lower bound). To do this, we plug in $\overline{v}(q_k, s)$, the maximum possible value at intermediate quantities along the next step, for the terms involving rationed quantities at the next step (e.g., line 3 of equation (11)), since the (conditional) expectation of the derivative of the random quantity won at the next step, $\mathbb{E}\left[\frac{\partial \boldsymbol{q}^*}{\partial q_k}\middle| b_{k+1} \geq \boldsymbol{P}^*, \theta, m\right]$, is negative. This is because increasing q_k decreases $(q_{k+1} - q_k)$, which, due to pro-rata-rationing, decreases the amount the bidder wins in case of a tie at step $k+1$. In addition, we plug in $\underline{v}(q_k, s)$, the smallest possible value at intermediate quantities along the current step, for terms involving ties at the current step (e.g., line 2 of equation (11)), since the (conditional) expectation of the derivative of the random quantity won at the step, $\mathbb{E}\left[\frac{\partial \boldsymbol{q}^*}{\partial q_k}\middle| b_k = \boldsymbol{P}^*, \theta, m\right]$, is positive. This is because increasing q_k increases $(q_k - q_{k-1})$, which, due to pro-rata-rationing, increases the amount the bidder wins in case of a tie at step k.

To obtain a lower bound, we instead substitute the maximum value at the current step, $\overline{v}(q_{k-1}, s)$ (in line 2 of equation (11)) and the minimum value at the next step, $\underline{v}(q_{k+1}, s)$ (in line 3 of equation (11)) .

To further simplify the system of equations, we rely on Lemma 1 according to which ties never occur in equilibrium at a subset of steps. This allows us to cancel terms involving ties at these steps. There are two cases, depending on the sign of λ. When the current λ is negative, at the steps before the last step ties cannot be optimal and the terms involving the derivative of the rationed quantity drop out except for at the second-to-last step. When the current λ is positive, the only simplification occurs at the last step, where the rationing terms all drop out. Furthermore, at the last step, all terms involving b_K drop out.

With these simplifications, we obtain the following system of equations for two cases, $\lambda > 0$, and $\lambda < 0$, where the expected payment is on the LHS and the expected benefit is

Appendix Figure C1: Identification graphically

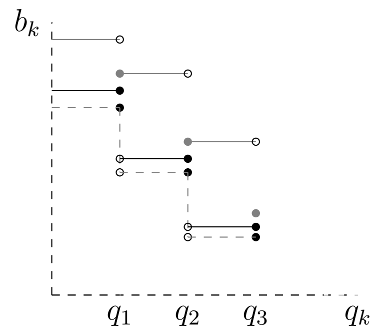

Appendix Figure C1 shows an example of a step function with three steps (in dashed lines) and the corresponding lower bound (in black) and upper bound (in gray) values at each step. In the example, shading is positive at all steps, which isn't crucial for identification. Instead, the main assumption we rely on to create the bounds is that the value curve is monotonically decreasing in quantity.

on the RHS. It can help to transform this system of equations into a single matrix, one for each case, to see that the system is indeed identified (conditional on knowing λ).

Define

$$A_k = \Pr(b_k > \boldsymbol{P^*} > b_{k+1}|\theta, m)b_k + \Pr(b_{k+1} \geq \boldsymbol{P^*}|\theta, m)(b_k - b_{k+1})$$
$$+ \Pr(b_k = \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_k}\Big|b_k = \boldsymbol{P^*}, \theta, m\right] b_k + \Pr(b_{k+1} = \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_k}\Big|b_{k+1} = \boldsymbol{P^*}, \theta, m]\right] b_{k+1}$$
$$- \lambda\frac{\partial \mu(b(\cdot, \theta))}{\partial q_k} \text{ for all } k = 1, ..., K-1, \tag{12}$$
$$A_K = \Pr(\boldsymbol{P^*} > b_K|\theta, m)b_K + \Pr(b_K = \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_K}\Big|b_K = \boldsymbol{P^*}, \theta, m\right] b_k - \lambda\frac{\partial \mu(b(\cdot, \theta))}{\partial q_K}. \tag{13}$$

(i) When $\lambda < 0$, the system of equations equalizing the expected payment with the expected benefit at the upper and lower bound is

$$A_1 = \Pr(b_1 > \boldsymbol{P^*} > b_2|\theta, m)\overline{v}(q_1, s)$$

$$A_1 = \Pr(b_1 > \boldsymbol{P^*} > b_2|\theta, m)\underline{v}(q_1, s), \text{ and analogously for } k = 2, ...K-2.$$

Here, we rely on the fact that when $\lambda < 0$, by Lemma 1, ties cannot occur at non-final steps.

$$A_{K-1} = \Pr(b_{K-1} > \boldsymbol{P^*} > b_K|\theta, m)\overline{v}(q_{K-1}, s) + \Pr(b_K \geq \boldsymbol{P^*}|\theta, m)\underbrace{\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_{K-1}}\Big|b_K \geq \boldsymbol{P^*}, \theta, m\right]}_{<0}\overline{v}(q_{K-1}, s)$$

Here, we again rely on the fact that there is no tie at $K-1$, so that we can cancel the second line in equation (11). We make the third line as small as possible to maximize $\overline{v}(q_{K-1}, s)$. All other lines are part of the expected payment and, thus, are in A_{K-1}. A similar logic applies when determining the lower bound at $K-1$:

$$A_{K-1} = \Pr(b_{K-1} > \boldsymbol{P^*} > b_K|\theta, m)\underline{v}(q_{K-1}, s) + \Pr(b_K \geq \boldsymbol{P^*}|\theta, m)\underbrace{\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_{K-1}}\Big|b_K \geq \boldsymbol{P^*}, \theta, m\right]}_{<0}\underline{v}(q_K, s)$$

For the last step, we apply the same logic from above, just that now the third line of equation (11) drops out because we are at the last step. In addition, we can replace the term $\underline{v}(q_K, s)$ with $\overline{v}(q_K, s)$ since we know the value curve takes one value at that point. We obtain

$$A_K = \Pr(b_K > \boldsymbol{P^*} > 0|\theta, m)\overline{v}(q_K, s) + \Pr(b_K = \boldsymbol{P^*}|\theta, m)\underbrace{\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_K}\Big|b_K = \boldsymbol{P^*}, \theta, m\right]}_{>0}\overline{v}(q_K, s)$$

$$A_K = \Pr(b_K > \boldsymbol{P^*} > 0|\theta, m)\underline{v}(q_K, s) + \Pr(b_K = \boldsymbol{P^*}|\theta, m)\underbrace{\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_K}\Big|b_K = \boldsymbol{P^*}, \theta, m\right]}_{>0}\overline{v}(q_{K-1}, s),$$

where A_k and A_K are given by (12) and (13), respectively.

(*ii*) Similarly, when $\lambda > 0$, we rely on Lemma 1 to eliminate ties at the last step. For earlier steps, we follow the analogous point-wise argument as above to obtain

$$\begin{aligned}
A_1 &= \Pr(b_1 > \boldsymbol{P^*} > b_2|\theta, m)\overline{v}(q_1, s) + \Pr(b_1 = \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_1}\Big|b_1 = \boldsymbol{P^*}, \theta, m\right]\underline{v}(q_1, s) \\
&\quad + \Pr(b_2 \geq \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_1}\Big|b_2 \geq \boldsymbol{P^*}, \theta, m\right]\overline{v}(q_1, s) \\
A_1 &= \Pr(b_1 > \boldsymbol{P^*} > b_2|\theta, m)\underline{v}(q_1, s) + \Pr(b_1 = \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_1}\Big|b_1 = \boldsymbol{P^*}, \theta, m\right]\overline{v}(0, s) \\
&\quad + \Pr(b_2 \geq \boldsymbol{P}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_1}\Big|b_2 \geq \boldsymbol{P^*}, \theta, m\right]\underline{v}(q_2, s)
\end{aligned}$$

Here, and in all of the equations for the upper bound that follow, we replace the $\underline{v}(q_k, s)$ in the terms involving ties with $\overline{v}(q_k, s)$, since we know that the values at that point are a single number. This makes the bounds more informative.

$$A_2 = \Pr(b_2 > \boldsymbol{P^*} > b_3|\theta, m)\overline{v}(q_2, s) + \Pr(b_2 = \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_2}\Big| b_2 = \boldsymbol{P^*}, \theta, m\right]\underline{v}(q_2, s)$$
$$\Pr(b_3 \geq \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_2}\Big| b_3 \geq \boldsymbol{P^*}, \theta, m\right]\overline{v}(q_2, s)$$
$$A_2 = \Pr(b_2 > \boldsymbol{P^*} > b_3|\theta, m)\underline{v}(q_2, s) + \Pr(b_2 = \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_2}\Big| b_2 = \boldsymbol{P^*}, \theta, m\right]\overline{v}(q_1, s) +$$
$$+ \Pr(b_3 \geq \boldsymbol{P}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_2}\Big| b_3 \geq \boldsymbol{P^*}, \theta, m\right]\underline{v}(q_3, s), \text{ and analogously for } k = 3, ...K-2,$$
$$A_{K-1} = \Pr(b_{K-1} > \boldsymbol{P^*} > b_K)\overline{v}(q_{K-1}, s) + \Pr(b_{K-1} = \boldsymbol{P^*})\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_{K-1}}\Big| b_{K-1} = \boldsymbol{P^*}, \theta, m\right]\underline{v}(q_{K-1}, s)$$

In this case, the third line of equation (11) drops out because there is no tie. And similarly for the lower bound equation:

$$A_{K-1} = \Pr(b_{K-1} > \boldsymbol{P^*} > b_K)\underline{v}(q_{K-1}, s) + \Pr(b_{K-1} = \boldsymbol{P^*})\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_{K-1}}\Big| b_{K-1} = \boldsymbol{P^*}, \theta, m\right]\overline{v}(q_{K-2}, s)$$
$$A_K = \Pr(b_K > \boldsymbol{P^*} > 0)\overline{v}(q_K, s)$$
$$A_K = \Pr(b_K > \boldsymbol{P^*} > 0)\underline{v}(q_K, s),$$

where A_k, and A_K are given by (12) and (13), respectively. This system of equations would be identified if λ was known. Since this isn't the case, we rely on Proposition 1 (*iii*) to obtain identification of the value bounds.

Specifically, we know that perturbing m cannot be optimal. Formally, the total expected surplus must decrease when increasing and decreasing m by $\epsilon > 0$:

$$TS(b(\cdot, \theta), m) - TS(b(\cdot, \theta), m + \epsilon) \geq \lambda\epsilon \tag{14}$$
$$TS(b(\cdot, \theta), m - \epsilon) - TS(b(\cdot, \theta), m) \leq \lambda\epsilon, \tag{15}$$

where $TS(b(\cdot, \theta), m) = \sum_{k=1}^{K}\Big[\Pr(b_k > \boldsymbol{P^*} > b_{k+1}|\theta, m)V(q_k, s) - \Pr(b_k > \boldsymbol{P^*}\Big|\theta, m)b_k(q_k - q_{k-1})\Big] + \sum_{k=1}^{K}\Pr(b_k = \boldsymbol{P^*}|\theta, m)\mathbb{E}\Big[V(\boldsymbol{q^*}, s) - b_k(\boldsymbol{q^*} - q_{k-1})\Big| b_k = \boldsymbol{P^*}, \theta, m\Big]$ with $q_0 = b_{K+1} = 0$. This expression is equivalent to equation (3) for $g = h$. $TS(b(\cdot, \theta), m+\epsilon)$ and $TS(b(\cdot, \theta), m-\epsilon)$ are defined analogously.

We can use (16) and (17) to find the bounds for (14) and (15), respectively. Consider (14) first and, for simplicity omit conditioning on θ; let $1(\cdot)$ denote the indicator function to obtain

$$
\begin{aligned}
&\sum_{k=1}^{K}\Big[\max\Big\{0,\Delta\Pr(b_k>\boldsymbol{P^*}>b_{k+1})\Big\}\Big(\sum_{j=1}^{k}\overline{v}(q_{j-1},s)(q_j-q_{j-1})\Big)\\
&\quad+\min\Big\{0,\Delta\Pr(b_k>\boldsymbol{P^*}>b_{k+1})\Big\}\Big(\sum_{j=1}^{k}\underline{v}(q_j,s)(q_j-q_{j-1})\Big)-\Delta\Pr(b_k>\boldsymbol{P^*})b_k(q_k-q_{k-1})\Big]+\\
&\sum_{k=1}^{K}\Big[\max\Big\{0,\Delta\Pr(b_k=\boldsymbol{P^*})\Big\}\Big(\sum_{j=1}^{k-1}\overline{v}(q_{j-1},s)(q_j-q_{j-1})\Big)+\min\Big\{0,\Delta\Pr(b_k=\boldsymbol{P^*})\Big\}\Big(\sum_{j=1}^{k-1}\underline{v}(q_j,s)(q_j-q_{j-1})\Big)\\
&\quad+\Big[\Big(\Pr(b_k=\boldsymbol{P^*}|m)(\mathbb{E}[\boldsymbol{q^*}|b_k=\boldsymbol{P^*},m]-q_{k-1})-\Pr(b_k=\boldsymbol{P^*}|m+\epsilon))(\mathbb{E}[\boldsymbol{q^*}|b_k=\boldsymbol{P^*},m+\epsilon]-q_{k-1})\Big)\Big]\overline{v}(q_{k-1},s)\\
&\quad\times 1\Big(\Pr(b_k=\boldsymbol{P^*}|m)(\mathbb{E}[\boldsymbol{q^*}|b_k=\boldsymbol{P^*},m]-q_{k-1})-\Pr(b_k=\boldsymbol{P^*}|m+\epsilon)(\mathbb{E}[\boldsymbol{q^*}|b_k=\boldsymbol{P^*},m+\epsilon]-q_{k-1})>0\Big)\\
&\quad+\Big[\Big(\Pr(b_k=\boldsymbol{P^*}|m)(\mathbb{E}[\boldsymbol{q^*}|b_k=\boldsymbol{P^*},m]-q_{k-1})-\Pr(b_k=\boldsymbol{P^*}|m+\epsilon)(\mathbb{E}[\boldsymbol{q^*}|b_k=\boldsymbol{P^*},m+\epsilon]-q_{k-1})\Big)\Big]\underline{v}(q_k,s)\\
&\quad\times 1\Big(\Pr(b_k=\boldsymbol{P^*}|m)(\mathbb{E}[\boldsymbol{q^*}|b_k=\boldsymbol{P^*},m]-q_{k-1})-\Pr(b_k=\boldsymbol{P^*}|m+\epsilon)(\mathbb{E}[\boldsymbol{q^*}|b_k=\boldsymbol{P^*},m+\epsilon]-q_{k-1})<0\Big)\\
&\quad-b_k\Big(\mathbb{E}[\boldsymbol{q^*}-q_k|b_k=\boldsymbol{P^*},m]-\mathbb{E}[\boldsymbol{q^*}-q_k|b_k=\boldsymbol{P^*},m+\epsilon]\Big)\Big]\geq\lambda\epsilon \qquad (16)
\end{aligned}
$$

where $\Delta\Pr(\cdot)$ indicates taking a difference between $\Pr(\cdot|..,m)$ and $\Pr(\cdot|..,m{+}\epsilon)$ and $v(q_0,s)=\overline{v}(q_1,s)$.

Similarly for (15):

$$
\begin{aligned}
&\sum_{k=1}^{K}\Big[\max\Big\{0,\Delta\Pr(b_k>\boldsymbol{P^*}>b_{k+1})\Big\}\Big(\sum_{j=1}^{k}\underline{v}(q_k,s)(q_j-q_{j-1})\Big)\\
&\quad+\min\Big\{0,\Delta\Pr(b_k>\boldsymbol{P^*}>b_{k+1})\Big\}\Big(\sum_{j=1}^{k}\overline{v}(q_{j-1},s)(q_j-q_{j-1})\Big)-\Delta\Pr(b_k>\boldsymbol{P^*})b_k(q_k-q_{k-1})\Big]+\\
&\sum_{k=1}^{K}\Big[\max\Big\{0,\Delta\Pr(b_k=\boldsymbol{P^*})\Big\}\Big(\sum_{j=1}^{k-1}\underline{v}(q_k,s)(q_j-q_{j-1})\Big)+\min\Big\{0,\Delta\Pr(b_k=\boldsymbol{P^*})\Big\}\Big(\sum_{j=1}^{k-1}\overline{v}(q_{j-1},s)(q_j-q_{j-1})\Big)\\
&\quad+\Big[\Big(\Pr(b_k=\boldsymbol{P^*}|m-\epsilon)(\mathbb{E}[\boldsymbol{q^*}|b_k=\boldsymbol{P^*},m-\epsilon]-q_{k-1})-\Pr(b_k=\boldsymbol{P^*}|m)(\mathbb{E}[\boldsymbol{q^*}|b_k=\boldsymbol{P^*},m]-q_{k-1})\Big)\Big]\underline{v}(q_k)\\
&\quad\times 1\Big(\Pr(b_k=\boldsymbol{P^*}|m-\epsilon)(\mathbb{E}[\boldsymbol{q^*}|b_k=\boldsymbol{P^*},m-\epsilon]-q_{k-1})-\Pr(b_k=\boldsymbol{P^*}|m)(\mathbb{E}[\boldsymbol{q^*}|b_k=\boldsymbol{P^*},m]-q_{k-1})>0\Big)\\
&\quad+\Big[\Big(\Pr(b_k=\boldsymbol{P^*}|m-\epsilon)(\mathbb{E}[\boldsymbol{q^*}|b_k=\boldsymbol{P^*},m-\epsilon]-q_{k-1})-\Pr(b_k=\boldsymbol{P^*}|m)(\mathbb{E}[\boldsymbol{q^*}|b_k=\boldsymbol{P^*},m]-q_{k-1})\Big)\Big]\overline{v}(q_{k-1})\\
&\quad\times 1\Big(\Pr(b_k=\boldsymbol{P^*}|m)(\mathbb{E}[\boldsymbol{q^*}|b_k=\boldsymbol{P^*},m-\epsilon]-q_{k-1})-\Pr(b_k=\boldsymbol{P^*}|m+\epsilon)(\mathbb{E}[\boldsymbol{q^*}|b_k=\boldsymbol{P^*},m]-q_{k-1})<0\Big)\\
&\quad-b_k\Big(\mathbb{E}[\boldsymbol{q^*}-q_k|b_k=\boldsymbol{P^*},m-\epsilon]-\mathbb{E}[\boldsymbol{q^*}-q_k|b_k=\boldsymbol{P^*},m]\Big)\Big]\leq\lambda\epsilon, \qquad (17)
\end{aligned}
$$

where $\Delta\Pr(\cdot)$ instead now indicates taking a difference between $\Pr(\cdot|..,m{-}\epsilon)$ and $\Pr(\cdot|..,m)$.

Summarizing, we now have a system of $2K$ linear equations, $2K+1$ unknowns and 2 inequalities from the optimality of the K steps submitted by customer i. Therefore, the customer value bounds are identified. □

B Empirical tests

In Appendix B.1, we analyze what factors predict customer participation. In Appendix B.2, we provide evidence in favor of independent private values. In Appendix B.3, we test whether customers take dealer updating into account and, in Appendix B.4, we test whether customer values are different from dealer values.

B.1 Predicting customer participation

To better understand what predicts customer participation, we regress the number of participating customers on a set of explanatory variables, using data from 2014 onward, a period when customers are almost exclusively hedge funds. Appendix Table B2 predicts customer participation at the bidder level and the results are similar. Importantly, we do not try to estimate the causal effects of customer participation. Instead, we focus on prediction, which we use in Section 7 to propose a policy rule that stabilizes customer participation at a sufficiently high average.

The first predictor of customer participation we include indicates the auction dates for which we estimate that a basis trade, which means buying a bond and shorting the future, could be profitable, an idea inspired by Barth and Kahn (2020) and Banegas et al. (2021).[29] The second indicator variable, denoted CIP, is the ten-year cross-currency basis swap with the U.S. dollar. A non-zero basis indicates a violation of covered interest parity and an opportunity for arbitrage (e.g., Du et al. (2018)). The third variable tells us whether the to-be-issued bond has benchmark status, which is the Canadian equivalent of being on the run. It could be, for example, that hedge funds buy more-liquid on-the-run bonds because they are easier to sell.

The fourth and fifth predictors are indicator variables that capture the importance of monetary policy committee meetings (MPC) and quantitative easing (QE). Including an indicator for MPC meetings is inspired by the findings of Lou et al. (2023), which demonstrate

[29] We calculate the basis as in Hazelkorn et al. (2022). Specifically, to determine profitability of buying bonds at auction and shorting the corresponding futures contract, we approximate the bond's value as the quantity-weighted average price of winning bids (by customers) plus the accrued interest between the auction date and the futures' expiration date. If this price is below the price of the futures contract multiplied by a conversion rate that is determined by the Bank of Canada, we say that a basis trade is profitable. The conversion rates are published here: https://www.m-x.ca/en/markets/interest-rate-derivatives/bond-futures-conversion-factor, accessed on 08/23/2023.

that hedge funds tend to purchase bonds outside of the pre-MPC window to avoid interest rate uncertainty. This is why we also include the bond coupon rate—higher coupon rates are correlated with lower interest rate risk. Including an indicator for the auction days on which the central bank conducts a bond issuance in the morning and engages in QE in the afternoon, where hedge funds have an opportunity to sell bonds back, is inspired by An and Song (2018, 2023).

The sixth variable measures the buy-sell spread at which a to-be-issued bond is traded prior to the auction. We approximate this spread by the average difference between the highest and lowest price within a day during which a bond is to be issued is traded in the secondary market three days prior to the auction. The seventh and eight variables count the number of dealers who participate in the auction and the number of customers who participated in the previous auction. Finally, we control for supply in the auction. Anecdotally, we know that margins are thin in Treasury markets; therefore, countries with large supplies can attract more participants.

In addition, we add controls that capture the interest rate environment and the expectations about the stance of monetary policy and, therefore, future bond prices. Specifically, we construct an overnight index swap (OIS) curve that captures the market expectations of the Bank of Canada's interest rate target for the overnight lending rate over 12 months.[30]

Our estimation findings, reported in Appendix Table B1, indicate that customers are more likely to participate in auctions when the secondary-market spread is high. This suggests that they buy bonds at auction when they can quickly sell them at high prices. In addition, we find some support for the idea put forth in Lou et al. (2023) that hedge funds avoid buying bonds prior to monetary policy announcements. The coefficient on the coupon rate is positive and this is consistent with hedge funds disliking interest rate risk. The other explanatory variables are statistically insignificant when including year-fixed effects. As for the number of dealers, this is because there is little variation within a year—a feature we will incorporate in our model. For some of the other explanatory variables, this might be because of low statistical power.

Low statistical power, in addition to a moderately sized R^2, indicates that there are

[30]Furthermore, even though transactions data only starts in 2016, we did experiment with including interdealer repo rates and repo spreads to capture the cost of overnight borrowing, but the coefficients are not statistically significant. The same is true for 1- and 3-month Canadian Dollar Offered Rates, which are important interest rate benchmarks, but are not statistically significant.

unobservable factors that play a significant role in driving customer participation—a feature our model will incorporate.

Appendix Table B1: Predictors of customer participation

Number of customers	(OLS1)		(OLS2)		(Year-FE)	
β_1: Cash-futures basis trade	0.169	(0.692)	-0.0584	(0.699)	-0.326	(0.697)
β_2: CIP basis trade	-0.0878***	(0.0192)	-0.108***	(0.0204)	-0.0115	(0.0374)
β_3: Benchmark status	0.157	(0.314)	0.0129	(0.319)	-0.138	(0.318)
β_4: MPC	-1.709**	(0.766)	-1.906**	(0.771)	-1.886**	(0.760)
β_5: QE	-0.140	(0.379)	-0.348	(0.403)	-0.359	(0.409)
β_6: Spread	0.259***	(0.0578)	0.263***	(0.0587)	0.226***	(0.0598)
β_7: Number of dealer	-0.187**	(0.0898)	-0.189**	(0.0946)	0.0723	(0.137)
β_8: Lagged number of customers	0.157***	(0.0508)	0.123**	(0.0516)	0.0686	(0.0528)
β_9: Coupon	0.946***	(0.177)	0.976***	(0.186)	0.986***	(0.188)
β_{10}: Supply	1.043*	(0.550)	0.845	(0.668)	0.0571	(0.711)
Extra controls	–		✓		✓	
Adjusted R^2	0.376		0.389		0.411	
Observations	326		326		326	

Appendix Table B1 shows the estimation results of regressing the observed number of participating customers in an auction on a series of explanatory variables using data from the beginning of 2014 to the end of 2021 in column (OLS1). "Cash-futures basis trade" is an indicator variable equal to 1 if buying a bond at auction and shorting the future is profitable (calculated as in Hazelkorn et al. 2022). "CIP basis trade" captures deviations from covered interest parity using the 10-year cross-currency swap basis with the U.S. dollar. "Benchmark status" is an indicator equal to 1 if the issued bond is on the run and 0 otherwise. "MPC" and "QE"" capture conventional and unconventional monetary policy, respectively. "Spread" is the high-minus the low-trading price for the bond being auctioned. "Number of dealers" and "Lagged number of customers" are the number of dealers who participate at auction and the number of customers who participated in the previous auction. "Coupon" is the coupon rate on the bond being issued. Supply is the residual from regressing the log-supply on the bond maturity at issuance. In column (OLS2) we add additional controls that capture the interest rate environment and expectations about the stance of monetary policy, using eight points on the OIS curve. In column (Year-FE) we include year fixed effects, in addition. Standard errors are in parenthesis.

B.2 Testing independent private values

We perform a formal test for independent private values, introduced in Hortaçsu and Kastl (2012), for auctions of the bonds in our sample. The test checks for equality of the dealers' estimated values before and after observing a customer bid. In a common values environment, a customers' bid would provide the dealer with information that changes their expected values for acquiring the bond being sold. Under independent private values, this bid reveals information about the expected level of competition but it should not affect their value. The

Appendix Table B2: Predictors of customer participation—Individual level

Participation$_i$	(OLS)		(Bidder-FE)		(Bidder-Year-FE)	
β_1: Cash-futures basis trade	0.0189	(0.0304)	-0.00307	(0.0251)	-0.0108	(0.0226)
β_2: CIP	-0.00259**	(0.000861)	-0.000521	(0.00141)	-0.00120	(0.00179)
β_3: Benchmark status	-0.0176	(0.0137)	-0.0197	(0.0135)	-0.0180	(0.0130)
β_4: MPC	-0.0413	(0.0329)	-0.0619	(0.0314)	-0.0611	(0.0330)
β_5: QE	-0.0218	(0.0161)	-0.0312	(0.0178)	-0.0366*	(0.0161)
β_6: Spread	0.0170***	(0.00251)	0.0127**	(0.00441)	0.00996**	(0.00369)
β_7: Number of dealers$_d$	0.00986*	(0.00483)	0.00760	(0.0115)	0.0136*	(0.00663)
β_8: Lagged-participation$_i$	0.532***	(0.0106)	0.230***	(0.0521)	0.0354	(0.0300)
β_9: Coupon	0.0462***	(0.00791)	0.0522***	(0.0140)	0.0549***	(0.0144)
β_{10}: Supply	0.0528*	(0.0236)	-0.00198	(0.0367)	0.00887	(0.0296)
Extra controls	✓		✓		✓	
Adjusted R^2	0.289		0.427		0.510	
Observations	6,567		6,551		6,548	

Appendix Table B2 is analogous to Appendix Table B1, but zooms in on customer bidding participation at the individual level. We regress an indicator for whether customer i participated in an auction ("Participation") on the same explanatory variables as in Appendix Table B1 only that we replace the number of customers who participated in the previous auctions by whether customer i participated in the previous auction, called "Lagged-participation." To take the time variation in the set of potential customers into account, we use data from all auctions between the first and last time we observe the customer bidding at auction to construct all customer-specific participation and lagged participation indicators. In column (Bidder-FE) we include a bidder-fixed effect and, in column (Bidder-Year-FE), we include a year-bidder fixed effect. Standard errors are in parenthesis. They are clustered at the bidder level in columns (Bidder-FE) and (Bidder-Year-FE). Our preferred specification includes bidder-year fixed effects, analogous to column (Year-FE) in Appendix Table B1. As in Appendix Table B1, Spread and Coupon are the only significant predictors among the market-level explanatory variables. The coefficient of Lagged-participation is positive without controlling for the upward time trend in customer participation. However, when accounting for this trend, this coefficient becomes statistically insignificant, suggesting that customers are not more likely to participate in an auction based on their participation in the previous auction. The number of dealers is weakly statistically significant, which likely arises from the fact that a fiscal year does not start in January.

test compares marginal values at the set of q_k common to the two dealers bids. Hortaçsu and Kastl (2012) denote the absolute value of the dealers' values before and after updating as T_i, and calculate the test-statistic for the sample as $SSQ_{T,IPV} = \sum(\frac{T_i}{\sigma_{T_i}})^2$. Similar to the findings of Hortaçsu and Kastl (2012) for Treasury bills, we fail to find evidence that dealer values are shifted by the information learned through the customers' bids in the bond market. We calculate a p-value of 0.48 and therefore do not reject the null hypothesis of no learning about fundamentals.

B.3 Testing bounds

We test whether customers take dealer updating into account when placing a bid. Formally, we want to know if $\lambda_{ti} = 0$ in Proposition 1 (ii). To do this, we fix an auction (and therefore omit the auction t subscript). We construct measures $T_i = |v^h(q, s_i^h; \lambda_i = 0) - v^h(q, s_i^h)|$ for each customer i. Here, $v^h(q, s_i^h; \lambda = 0)$ denotes the customer's value for amount q, assuming that $\lambda_i = 0$ and $v^h(q, s_i^h)$ is the value if $\lambda_i \neq 0$. With this, we construct a test statistic analogous to SSQ_T from Hortaçsu and Kastl (2012). The test rejects the null hypothesis with a p-value of 0, where the p-value is computed via bootstrap.

B.4 Testing value differences

We test whether customer values are significantly above dealer values. The null hypothesis is that customer values at the lower bound are weakly smaller than dealer values at the upper bound. Following Hortaçsu and Kastl (2012), we compute three aggregate test statistics: the first test is in the spirit of a Chi-squared test, the second is based on the 95th percentile of the across-auction differences, and the third is based on the maximum difference in these values across auctions. Since we are interested in a one-sided null hypothesis (are customer values larger than dealer values), we drop the absolute value, which differs from Hortaçsu and Kastl (2012). In all cases we omit the subset of auctions where not a single customer participated. We compute these test statistics for the differences in the average quantity-weighted value, the average maximum value, and the average minimum value of dealers/customers. In addition, we compute the confidence intervals for a set estimate of the mean difference. This allows us to evaluate the size of the expected difference and to compare it to 0.

The results are reported in Table B3. For all measures, the customer values appear to be above the dealer values, however, the differences in the average maximum values are less precise, with some of the test statistics being insignificant.

B.5 Testing selective entry

In a model with selective entry by customers (i.e., customers know their values before making their participation decision), the exit of a dealer should lead to entry by low-value customers who were previously unwilling to participate but become willing given the larger rents available after the dealer exits. In this world, the average value of customers would fall relative

Appendix Table B3: Differences in customer and dealer values

	P95	Sum	Max	CI
QWA value	0.00	0.00	0.00	[825, 2906]
Max value	0.06	0.00	0.93	[-836,2406]
Min value	0.00	0.00	0.00	[615, 1976]

Appendix Table B3 shows the results from testing whether customer values are above dealer values. Columns P95, Sum, and Max present the p-values for the test statistics that take the 95th percentile, the sum of squared standardized differences, and the maximum difference across auctions of the within auction average quantity-weighted average values of dealers less the within auction average lower bound of customer quantity-weighted average values. P-values are computed using the bootstrap. The confidence intervals (CI) are for the interval estimates of the mean difference. The QWA value is the average (within customers and dealers) of the individual participant's quantity-weighted average values. The Max value row compares the within group average values of the individual bidders' maximum value (at their first submitted step). The Min value row compares the within group average values of the individual bidder's minimum value (at their last submitted step).

to dealers around a dealer exit.

To test whether this is the case, we compute the differences between customer and dealer average values for auction windows, W_t, around each dealer exit event, e:

$$\frac{1}{W_t}\sum_{t=e}^{e+W_t}\left(\frac{1}{N_t^h}\sum_{i\in h}\underline{\mu}_{it}^h-\frac{1}{d}\sum_{i\in d}\overline{\mu}_{it}^d\right)-\frac{1}{W_t}\sum_{t=e-W_t}^{e}\left(\frac{1}{N_t^h}\sum_{i\in h}\underline{\mu}_{it}^h-\frac{1}{N^d}\sum_{i\in d}\overline{\mu}_{it}^d\right) \quad (18)$$

Here $\underline{\mu}_{it}^g$ denotes the lower bound on the quantity-weighted average value for participant i in group g in auction t, and $\overline{\mu}_{it}^g$ the corresponding upper bound. Since we have bounds and not point-estimates, we use customer lower bounds on values and dealer upper bounds—consistent with our approach throughout the paper. There is no reason, however, to suspect a systematic change in the informativeness of the bounds around exit events so this choice is inconsequential. We compute standard errors using the bootstrapped value estimates.

The results in Table B4 suggest that we cannot reject the model with non-selective entry. In auctions following exit, customers' values are not systematically lower than prior to exit. The findings are noisy, but quantities demanded are stable across these events. Results are similar using bids instead of values or using minimum or maximum values instead of the quantity weighted averages.

Appendix Table B4: Testing for selective customer entry

Window	Deutsche Bank	Morgan Stanley	PI Financial
2 Auctions	0.0987	-0.092	-0.316
	[-0.09,0.20]	[-0.30,0.002]	[-0.49,0.38]
5 Auctions	0.014	0.005	-0.168
	[-0.026,0.14]	[-0.097,0.153]	[-0.17,0.24]

Appendix Table B4 provides evidence that entry is not selective. The top panel reports the average change in values of customers relative to dealers using a window of 2 auctions around the exit dates for Deutsche Bank, Morgan Stanley, and PI Financial. The second row uses a window of 5 auctions around the exit dates.

B.6 Test for unobserved heterogeneity in entry cost distributions

An important modeling assumption is that the distribution of entry costs is constant over time. If entry costs were varying across auctions, then conditional on the expected surplus level there would be many different entry probabilities driven by changes in the distribution.

We test for unobserved heterogeneity in entry cost distributions by analyzing the model-implied (conditional) variance of the number of participating customers. Specifically, in our model without unobserved heterogeneity, the number of customers in an auction, N_t^h, is a binomial random variable with parameters $\overline{N}^h$ (the number of potential customers) and $G^h(TS_t^h)$ (the time-invariant customer entry distribution, evaluated at the auction-specific expected surplus, TS_t^h). Therefore, conditional on TS_t^h the variance of the number of participating customers is:

$$Var(N_t^h|TS_t^h) = G^h(TS_t^h)(1 - G^h(TS_t^h))\overline{N}^h. \tag{19}$$

Instead, if there were important unobserved heterogeneity across auctions, then the number of participating customers in the data would follow a mixture of binomial distributions with different probabilities $G_t^h(TS_t^h)$ for each auction, where the entry cost distribution is now time varying according to some unobservable random variable U. In this case, the variance would be:

$$\begin{aligned} Var(N_t^h|TS_t^h) &= \mathbb{E}[Var(N_t^h|TS_t^h, U)] + Var(\mathbb{E}[N_t|TS_t^h, U]) \\ &= \mathbb{E}[G_t^h(TS_t^h)(1 - G_t^h(TS_t^h))\overline{N}^h] + (\overline{N}^h)^2 Var(G_t^h(TS_t^h)). \end{aligned} \tag{20}$$

Here the expectation and the variance are taken with respect to the unobserved variable shifting the entry cost distribution, U.

Comparing the two variances, expression (20) is larger than expression (19) when $\mathbb{E}[G_t^h(\cdot)] = G(\cdot)$, where the expectation is taken over U, which we enforce in our estimation. To see why note that the difference in the variances, (20)-(19), can be written as

$$\begin{aligned}&\mathbb{E}[G_t^h(TS_t^h)(1-G_t^h(TS_t^h))\overline{N}^h] + (\overline{N}^h)^2 Var(G_t^h(TS_t^h)) - G^h(TS_t^h)(1-G^h(TS_t^h))\overline{N}^h \\ =&\overline{N}^h[G^h(TS_t^h)) - \mathbb{E}[G_t^h(TS_t^h)^2]] + (\overline{N}^h)^2[\mathbb{E}[G_t^h(TS_t^h)^2] - (G^h(TS_t^h))^2] - \overline{N}^h[G^h(TS_t^h) - G^h(TS_t^h)^2],\end{aligned}$$

where we have used that $G^h(TS_t^h)) = \mathbb{E}[G_t^h(TS_t^h))]$. Now grouping together the terms multiplied by $\overline{N}^h$ and $(\overline{N}^h)^2$, and simplifying, we obtain:

$$\begin{aligned}=&\overline{N}^h[-\mathbb{E}[(G_t^h(TS_t^h))^2] + (G^h(TS_t^h))^2] + (\overline{N}^h)^2[\mathbb{E}[G_t^h(TS_t^h)^2] - G^h(TS_t^h)^2] \\ =&((\overline{N}^h)^2 - \overline{N}^h)[\mathbb{E}[G_t^h(TS_t^h)^2] - G^h(TS_t^h)^2] = \overline{N}^h(\overline{N}^h - 1)Var(G_t^h(TS_t^h)) \geq 0.\end{aligned}$$

The key idea for our test is that when there is unobserved heterogeneity, so that $Var(G_t^h(TS_t^h)) > 0$, the variance difference is strictly positive (given that $\overline{N}^h > 1$). Since we have not targeted the variance in our estimation, comparing these two variances allows us to test the null hypothesis that entry costs are stable over time.

To conduct a test based on this idea, we need to estimate the variance of N_t^h conditional on TS_t^h in the data, and calculate the distribution of that estimator under the null hypothesis that there is no heterogeneity. We divide the space of TS^h into 5 quantile bins and estimate the $Var(N_t^h|TS_t^h \in bin)$ inside each bin, because there are no auctions in our finite sample with identical values of TS_t^h. We choose 5 to balance the following trade-off. On the one hand, with a larger number of bins the sample estimate of the conditional variance of N_t^h becomes noisy. On the other hand, with fewer bins the across-auction variance in $G(TS_t^h)$, due to changes in TS_t^h gets bigger, limiting the scale of the unobserved heterogeneity that one can detect. In practice, we only estimate bounds on expected surpluses and thus some of the auctions cannot be uniquely assigned a bin. In such cases, we include them in both bins, which will tend to overstate the within-bin variance estimate in the sample, which should reduce the power of the test. From the data we can estimate the variance of the number of participating customers in each bin, $Var(N_t^h|bin)$, based on a sample of T_{bin} auctions within the bin.

To calculate the distribution of $Var(N_t^h|bin)$ under the null hypothesis, assuming no

Appendix Table B5: Test for unobserved heterogeneity in entry cost distributions

Sample	p5 lower bound	p95 upper bound
6.496	0.072	13.726
6.624	0.505	7.531
7.266	0.522	8.006
6.345	1.055	7.190
5.904	4.139	10.608

Appendix Table B5 presents results for our test for unobserved heterogeneity in entry cost distributions. Column (1) is the sample variance while column (2) is the fifth percentile of the lower bound and column (3) is the ninety-fifth percentile of the upper bound. The sample variance lies in between the two bounds, so we do not reject the null that the entry costs are stable over time.

across-auction heterogeneity in entry cost distributions, we use simulation. In particular, we draw at random T_{bin} bounds on expected auction surpluses for customers from their empirical distribution. For each draw we compute upper and lower bounds on entry probabilities using the surplus bounds, and the estimated bounds on entry costs.[31] For each draw, we also simulate $\overline{N}^h$=12 random draws from the uniform distribution. We use these to simulate a lower bound on the number of participating customers (equal to the number of random draws below the entry probability lower bound), and an upper bound on the number of participating customers (equal to the number below the entry probability upper bound). The maximum within-bin variance in number of participating customers consistent with these bounds can be found by maximizing $Var(N_t^h|bin)$ subject to the constraint that the auction-specific number of participants lies between the lower and upper bound constructed for each of the T_{bin} simulated auctions. We repeat this calculation many times to evaluate the distribution of the sample variance estimator based on T_{bin} observations under the null hypothesis.

The results of the test are presented in Table B5. We would reject the null-hypothesis that there is no cross-auction heterogeneity in entry cost distributions if the sample estimator of the variance in column (1) exceeded the 95th percentile of the distribution of upper bound estimates in column (3). However, in all the bins, the sample variance is between the fifth percentile of the lower bound and 95th percentile of the upper bound. Therefore, we do not reject the null that the entry costs are stable over time.

[31] A low surplus bound and a high entry cost bound combine to make a low entry probability bound while a high surplus bound with a low entry cost bound give an upper bound on entry probability.

C Details regarding the estimation of values

To back out the dealer values and bounds on customer values from the equilibrium conditions of Proposition 1, we need to estimate the probabilities that enter these conditions, which are determined by the distribution of the market-clearing price, $\boldsymbol{P_t^*}$. For customers, we also need to estimate the Lagrange multiplier, λ_{ti}, and the terms that capture the ties in condition (5). For simplicity, we assume that the number of customers matched to each dealer is 1. This is the case in the data, except for rare cases.

We estimate the market-clearing price distributions by simulating market clearing. If all bidders were ex-ante symmetric and bid directly to the auctioneer, we would fix a bidder in an auction and draw $N-1$ bid functions, with replacement, from all of the observed bids in that auction. This would simulate one possible market outcome for the fixed bidder. Repeating this many times, we would obtain the distribution of the market-clearing price, $\boldsymbol{P_t^*}$, for this bidder. Our setting is more complicated because there are both dealers and customers and customers must bid via dealers. Hortaçsu and Kastl (2012) introduce a resampling procedure to estimate the price distribution from the dealer's perspective. We extend their method to learn about the customers. Further, we allow signals within a bidder to be correlated over the course of an auction. This is to avoid estimation bias arising from the fact that we observe some bidders updating their bids without observing a customer bid.[32]

Specifically, we resample as follows: We first construct the residual supply curve that bidder i faces in auction t. For this, we start by randomly drawing a customer bid from the set of $\overline{N}^h$ potential customer bids. If the customer did not participate in the auction, their bid is 0; if they updated their bid, we randomly select one of their bids. Next, we find a dealer that observed a similar bid to that customers' bid.[33] If the selected dealer made multiple bids, we select a random bid from the set of bids submitted by that dealer after they observed a bid similar to the current customer's bid. If the dealer did not update their bid after learning the customer's bid, we choose the last bid before learning it. Once a bid is selected, we drop all other bids from that dealer. We repeat this procedure $\overline{N}^h$ times if i is a dealer and $\overline{N}^h - 1$ times if i is a customer. Next, we resample the dealers that did not observe

[32]For simplicity, our model does not rationalize such updates, but an extended model based on Hortaçsu and Kastl (2012) could.

[33]Ideally, we would choose a dealer that observed an identical bid. Given our limited sample size, however, this event is extremely unlikely. To reflect customer uncertainty about the value of the dealer observing their bid, we set a bandwidth and define similar bids using the quantity-weighted average bids.

customer bids. Starting with a list of "uninformed" dealers, we draw one such dealer.[34] If they submitted more than one bid, we randomly select one bid and drop the others. We continue drawing from the set of uninformed dealers so that there are $N^h + N^d - 1$ bidding curves. This is one realization of the residual supply curve that i expects to face. We repeat the process many times to estimate the distribution of the clearing price, each time starting with the full set of bids made in the auction.

As in Kastl (2011), consistency of the estimator requires that the probability of market clearing at each step is strictly bounded away from zero. However, in our finite sample this event may occur. For steps with estimated win probabilities close to or equal to zero, we mix the estimated clearing-price distribution with a uniform distribution over the range of placed bids in order to give all bids a small win probability. In addition to reducing the sensitivity of the analysis to these small probabilities, we truncate less than 10 percent of the estimated values by assuming that these values are below the maximum bid plus C$0.1 times the maturity length in months divided by 12, which is roughly equivalent to 10 bps in terms of yield-to-maturity.

With the estimated price distributions, using condition (4), we can solve for the value that rationalizes a dealer's bid at each step. To obtain the customer value bounds, we implement an estimation procedure that follows our identification argument that is presented in the main text, and formally explained in the proof of Proposition 4 in Appendix A.

To begin, we search over the (one-dimensional) set of λ_{ti}. For each feasible λ_{ti}, there is a unique set of lower and upper bounds for the value at each step, q_k, where that customer submitted a step that satisfies equation (5). Using these implied values together with the definition of the total surplus allows us to obtain upper and lower bounds on the λ_{ti} based on expressions (16 and 17). This range of (λ_L, λ_U) is the set of feasible λ_{ti} that is consistent with the observed choices and values. To obtain the maximum and minimum, we find the value at each point that maximizes (minimizes) the change in the total surplus. Whether this is the upper or lower envelope of the set of values consistent with the observed bids depends only on the sign of the change in the clearing probabilities under the increased moment, m. If the initial λ_{ti} is within the range (λ_L, λ_U), then the associated value curve is part of the identified set of values that can rationalize the behavior of a given customer. When it is outside the range, the bid is not consistent with equilibrium behavior for that set of values.

[34] This includes sampling bids from dealers that later become informed and placed a later bid but that were not selected in the simulated residual supply curve in the customer resampling step.

To trace out the identified set, we repeat this exercise along a grid of possible λ_{ti}.

D Details regarding counterfactuals

In Appendix D.1 we explain how we compute counterfactual equilibria. In Appendix D.2 we present what happens when we make changes to the rules of bidder commitment. In Appendix D.3 we show how to reshuffle supply to stabilize participation.

D.1 Computational details

We are interested in finding a set of bid distributions that implies a value distribution that is similar to the true (in our case estimated) value distribution. We therefore construct a criterion function that compares these distributions along several dimensions. The criterion has three components.

First, we evaluate the distribution of the values at quantiles of the quantity-bid distribution, corresponding to orders for 1.1, 1.7, 2.3, 3.4, 3.6, 4.5, 5.6 and 25 percent of the total supply. To reduce the width of the predicted set of solutions, we add to this set the marginal distribution of the values corresponding to quantities of 10 and 15 percent. For each auction and bidder group, g, we construct bounds on the value distribution at each discrete level of quantity, using an evenly spaced grid running from the 5^{th} percentile to the 95^{th} percentile. At each point, we compare the bounds on the implied values from the guess of the bid distribution to the true values and add to the criterion function $max(F_L^{IM} - F_U, 0)^2$ and $min(F_U^{IM} - F_L, 0)^2$, where F_L^{IM} denotes the implied value distribution at each of the quantity levels evaluated on the grid points and F_U denotes the corresponding upper bound on the distribution known from the data. F_U^{IM} and F_L are defined analogously.

In addition, for each bidder group, we want the (across-bidder) distribution of the largest quantities demanded to be smaller than the observed distribution of the largest quantities a bidder ever purchased in all auctions. The restriction is designed to capture the fact that some bidders might be capacity constrained below the regulatory 25 percent maximum and, therefore, even if the counterfactual price is low, bidders may not be interested in purchasing up to 25 percent. Therefore, we require the distribution of the largest quantity bid by bidders of group g in auction t to be first-order stochastically dominated by the distribution of the largest quantities ever bid by bidders of group g in any auction t. We add to the criterion function the squared difference in the probabilities any time the implied maximum quantity distribution falls below the distribution from the data. These violations are evaluated along a set of grid points (evenly spaced by quantity, from 0 to 25 percent).

In an alternative specification, we require the predicted average coverage, that is, the sum of all bids over the total supply, of the auction to match the observed one, after accounting for the change in coverage that is induced by a change in the number of bidders. Specifically, we estimate the change in coverage caused by dealer exits (pooling 10 auctions before and after each exit) and reduce the coverage by the difference in the number of dealers between the counterfactual and the factual, multiplied by the average reduction in coverage. If we allow each bidder to demand up to the bidding limit (25 percent), no auction will fail. However, it is unrealistic to assume that all bidders have the capacity to buy Treasuries worth more than C$810 million (which is 25 percent of the average amount issued) in each auction. This would be possible only if a bidder received an extraordinary number of client orders or had sufficient balance sheet space despite stringent regulatory constraints.

D.2 Evaluating the importance of commitment

We evaluate two alternative policy regimes regarding bidder commitment. First, we make changes to assess the extent to which primary auctions run smoothly without forcing regular dealer participation. Second, we attempt to minimize volatility by requiring customers to commit in the same way as dealers. In both cases, it is theoretically ambiguous whether competition and volatility increase or decrease relative to the status quo, due to endogenous bidder participation.

To analyze the importance of dealer commitment, we compare two settings—in the first, dealers commit as in the status quo but we allow customers to place bids directly with the auctioneer; in the second, both bidder groups bid directly with the auctioneer but we do not impose obligatory participation on dealers. To compute the counterfactual without dealer commitment, we assume that a dealer's cost of entering one auction equals their estimated annual cost divided by the average number of auctions in a year.

We find that most auctions attract sufficiently many bidders to guarantee full auction coverage, even without obligating dealers to regularly participate (see Appendix Figure D1). However, both dealer and customer participation is highly irregular. Moreover, revenues fall substantially with an expected decrease of 1 percent.

Appendix Table D1: Customer commitment

Year	No. of customers		Expected revenue		SD[Expected revenue]	
	Com	SQ	Com	SQ	Com	SQ
2015	[2, 3]	2.310	[3.576, 3.576]	[3.576, 3.576]	0.209	0.208
2016	[2, 2]	2.256	[3.595, 3.595]	[3.594, 3.595]	0.177	0.176
2017	[2, 3]	2.825	[3.613, 3.613]	[3.613, 3.613]	0.171	0.171
2018	[3, 4]	3.816	[3.617, 3.617]	[3.617, 3.617]	0.322	0.321
2019	[4, 4]	4.290	[3.611, 3.610]	[3.610, 3.611]	0.392	0.392
2020	[5, 5]	5.00	[3.459, 3.459]	[3.458, 3.459]	0.471	0.471
2021	[4, 5]	4.519	[3.504, 3.504]	[3.504, 3.505]	0.420	0.420
2022	[5, 5]	4.378	[3.382, 3.382]	[3.381, 3.382]	0.393	0.393

Appendix Table D1 compares the counterfactual with customer commitment (Com) to the status quo (SQ), where customers make per-auction entry decisions. Dealer participation is fixed. Note that the equilibrium number of customer entrants depends on auction-specific profits for each of the 30 auctions per year across 8 years. To avoid computing the equilibrium in all of these auctions for each possible number of customers, we utilize a selected sample of five auctions. These auctions are strategically chosen to align the number of customers with percentiles (5th, 25th, 50th, 75th, and 95th) of the customer participation distribution since 2014, while the quantity sold approximates the average amount. When calculating profits, and surpluses for each year, we re-weight the predictions from these five auctions to match the composition of auctions in that specific year. Expected revenues are in C$ billions.

To assess the impact of customer commitment, we force customers to decide at the beginning of each year whether or not to commit to participating in all auctions of the upcoming year. They enter the market if their annual entry cost (approximated by the estimated auction-specific cost scaled by the average number of auctions in a year) is larger than the total surplus they expect from participating in all auctions of that year. Dealer participation is fixed.

Forcing customers to commit at the beginning of each fiscal year has only small effects on the number of customers who actually participate relative to the status quo. See Appendix Table D1. For example, in 2015 the average number of customers is 2.31. In the counterfactual the number of customers who choose to commit for the entire year is either 2 or 3 (we cannot be definitive because the set of solutions that we compute includes both outcomes). In contrast, the average number of customers in 2022 is 4.378 but the counterfactual number of customers would actually increase to 5. Auctions remain relatively competitive in the counterfactual since sufficiently many bidders stay in the market, so expected revenues are similar.

Appendix Figure D1: No dealer commitment

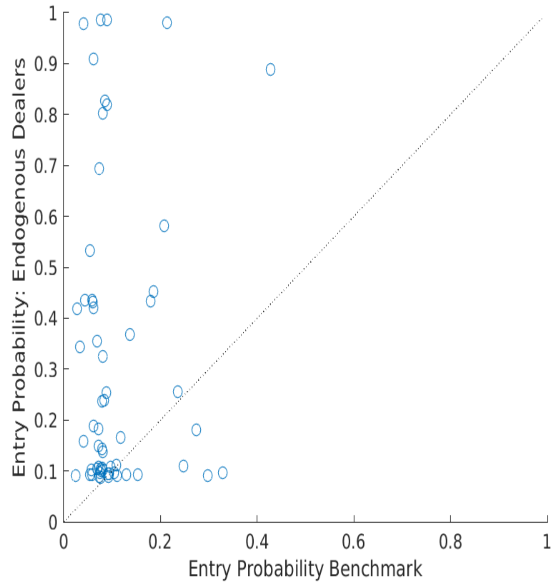

(A) Hedge entry probability (in %)

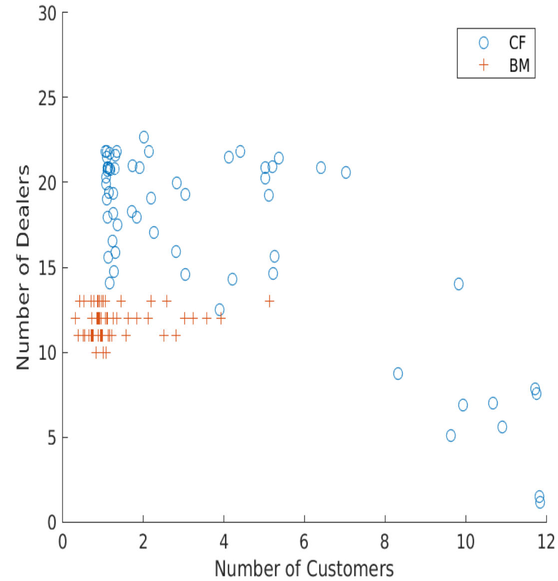

(B) Expected number of dealers/customers

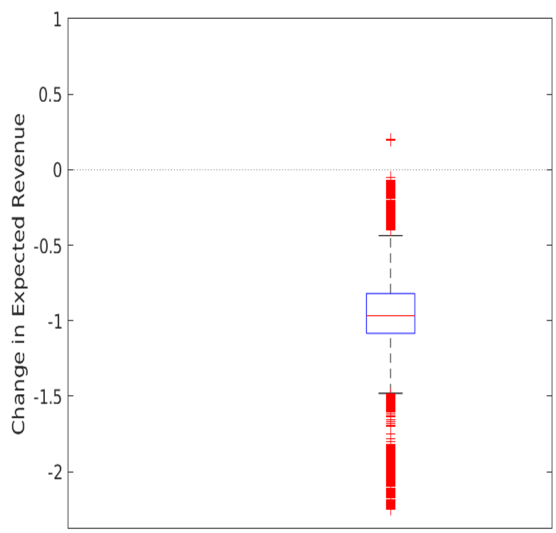

(C) Percentage change in expected revenue

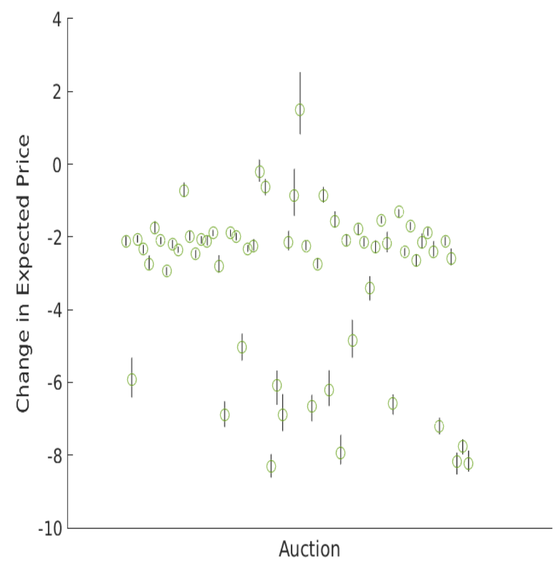

(D) Percentage change in expected revenue

For every third auction starting in 2014, Figure D1A shows the probability, in percentage points, that a customer participates in the status quo (on the x-axis) and the counterfactual in which we add back the two dealers who left (on the y-axis). Figure D1B shows the expected number of dealers and customers that participate in each auction in the status quo and the counterfactual. Figure D1C shows the distribution of the percentage change in the expected auction revenue. Figure D1D is the time series of the percentage change in the expected price when going from the status quo to the counterfactual. Prices are in C$ with a face value of 100.

D.3 Reshuffling supply to incentivize customer participation

We run this counterfactual for all auctions in 2018 and 2019, which are all over-subscribed in the status quo. Concretely, we predict the number of participating customers given the current supply schedule and observable market conditions by relying on the OLS regression from Appendix Table B1. We exclude the lagged number of participating customers from the regression, because it changes endogenously in the counterfactual. We think that it is reasonable to assume that all other predictors of participation, such as the secondary market spread, are not materially affected by the small proposed supply changes. Within each year and maturity bucket (<3 years or over), we rank the predicted number of participating customers, $\hat{N}_t^c$, based on their quantile rank, $s \in [0, 1]$, using the distribution of the number of predicted customers in all auctions before 2018, which captures the intensity of the predicted customer participation. At each auction we add or remove supply of $(0.5 - s)0.2\bar{Q}$ bonds. We then add or remove the negative of this amount evenly from all remaining auctions in the same year, so that at the end of the year the same total volume of bonds are issued.

As we can see from Appendix Figure D2, implementing these supply adjustments in our sample does not have a detectable impact on competition or volatility. The customer participation probability per auction stabilizes around a median of 40 percent and the median revenue per auction is between C$[-1.03, 1.58] million more than the status quo. This simple rule also ignores some of the factors that influence the complicated decision on how to issue government debt. For example, it abstracts from the term structure of bonds and, therefore, ignores the complications that arise from rolling over debt.

Appendix Figure D2: Reshuffling supply to incentivize customer participation

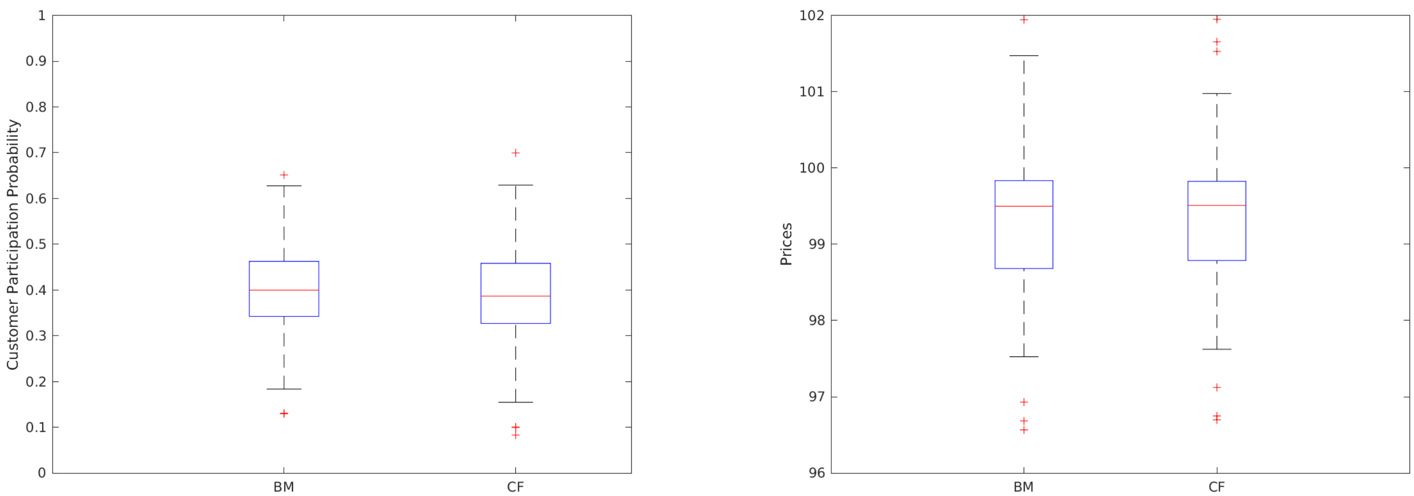

(A) Customer participation (B) Expected price

Figure D2A shows the distribution of customer participation probabilities across auctions in the status quo and the counterfactual in which we strategically reshuffle supply to incentivize stable customer participation (in pp). Figure D2B displays the corresponding distributions of expected auction prices (in C$). Reshuffling could be implemented by the Bank of Canada using it's discretion in placing non-competitive bids.

E Additional Tables and Figures

Appendix Table E1: List of dealers that exited and entered auctions, and year

Dealer name	Bill auctions	Bond auctions
CT Securities	Entry 1999, exit 2000	Entry 1999, exit 1999
Salomon Brothers	Entry 1999 exit 2000	Entry 1999, exit 2001
Goldman Sachs	Entry 1999, exit 2001	Entry 1999, exit 2002
SG Valeurs	Entry 1999, exit 2004	Entry 1999, exit 2004
JP Morgan	Entry 1999, exit 2007	Entry 1999, exit 2007
Deutsche Bank	Entry 1999, exit 2014	Entry 1999, exit 2014
Morgan Stanley	Entry 1999, exit 2001	Entry 1999, exit 2014
PI Financial Corp	Entry 2009, exit 2015	Entry 2009, exit 2015
Ocean Securities	Entry 2006, exit 2008	Entry 1999, exit 2015
Sherbrooke SSC	Entry 2020	Entry 2020

Appendix Table E1 lists all entries and exits of dealers in bill and bond auctions from 1999 to 2022. We only list years, even though we do observe the exact dates of entry and exit.

Appendix Table E2: Predictors of customer participation, customer values and bids

	Values (OLS)		Bids (OLS)		Values (Year-FE)		Bids (Year-FE)	
β_1: Cash-futures basis trade	-1.038	(1.340)	-0.185	(1.352)	-1.554	(1.301)	-0.729	(1.308)
β_2: CIP	0.0426	(0.0390)	0.0387	(0.0394)	-0.0101	(0.0698)	-0.0117	(0.0701)
β_3: Benchmark status	0.476	(0.612)	0.223	(0.617)	0.612	(0.593)	0.389	(0.595)
β_4: MPC	-1.942	(1.478)	-2.017	(1.491)	-1.889	(1.418)	-1.984	(1.425)
β_5: QE	2.896***	(0.771)	2.935***	(0.778)	2.345***	(0.762)	2.350***	(0.766)
β_6: Spread	0.428***	(0.113)	0.457***	(0.114)	0.480***	(0.112)	0.519***	(0.112)
β_7: Coupon	3.831***	(0.356)	3.679***	(0.359)	3.796***	(0.351)	3.612***	(0.353)
β_8: Number of dealers	-0.467**	(0.181)	-0.470**	(0.183)	-0.218	(0.256)	-0.265	(0.258)
β_9: Lagged number of customers	-0.0309	(0.1000)	-0.0223	(0.101)	0.0491	(0.0997)	0.0595	(0.100)
β_{10}: Supply	1.984	(1.281)	1.958	(1.293)	2.413*	(1.326)	2.472*	(1.332)
Extra controls	✓		✓		✓		✓	
Observations	326		326		326		326	
Adjusted R^2	0.481		0.468		0.491		0.477	

Appendix Table E2 is similar to Appendix Table B1. In the "Values" and (OLS) column, we regress our estimated quantity-weighted average values of customers on all of the explanatory variables we used in Appendix Table B1 to predict customer participation. We add a year-fixed effect in the (Year-FE) column. In the "Bids" columns, we replace the value estimates by the observed quantity-weighted bids of customers. The data ranges from the beginning of 2014 to the end of 2021. Standard errors are in parenthesis.

Appendix Table E3: Predictors of customer participation and moments of the customer value distribution

	(1) Average	(2) Median	(3) 5-Percentile	(4) 95-Percentile	(5) Std	(6) Range
Cash-futures basis trade	-1.554	-1.376	-2.737*	-1.047	0.763**	1.690*
	(1.301)	(1.317)	(1.400)	(1.345)	(0.320)	(0.888)
CIP	-0.0101	-0.0111	0.0343	-0.0328	-0.0268	-0.0671
	(0.0698)	(0.0707)	(0.0751)	(0.0721)	(0.0172)	(0.0476)
Benchmark status	0.612	0.465	1.270**	0.542	-0.261*	-0.728*
	(0.593)	(0.600)	(0.637)	(0.612)	(0.146)	(0.404)
MPC	-1.889	-1.980	-0.820	-2.361	-0.531	-1.542
	(1.418)	(1.435)	(1.525)	(1.465)	(0.348)	(0.968)
QE	2.345***	2.335***	2.309***	2.406***	-0.0162	0.0976
	(0.762)	(0.772)	(0.820)	(0.788)	(0.187)	(0.520)
Spread	0.480***	0.527***	0.0957	0.613***	0.174***	0.517***
	(0.112)	(0.113)	(0.120)	(0.115)	(0.0274)	(0.0762)
Number of dealers	-0.218	-0.229	-0.0920	-0.250	-0.0534	-0.159
	(0.256)	(0.259)	(0.276)	(0.265)	(0.0630)	(0.175)
Lagged number of customers	0.0491	0.0508	0.0443	0.0493	0.00298	0.00500
	(0.0997)	(0.101)	(0.107)	(0.103)	(0.0245)	(0.0680)
Coupon	3.796***	3.718***	3.671***	4.113***	0.169*	0.442*
	(0.351)	(0.356)	(0.378)	(0.363)	(0.0864)	(0.240)
Supply	2.413*	2.298*	2.939**	2.280*	-0.298	-0.658
	(1.326)	(1.342)	(1.426)	(1.370)	(0.326)	(0.905)
Extra controls	✓	✓	✓	✓	✓	✓
Year fixed effect	✓	✓	✓	✓	✓	✓
Adjusted R^2	0.525	0.521	0.420	0.557	0.232	0.247
Observations	326	326	326	326	326	326

Appendix Table E3 regresses moments of the estimated customer value distribution (at the lower bound) on all explanatory variables that we include to predict customer participation in Appendix Table B1, plus year-fixed effects. "Average" stands for the quantity-weighted average value, which approximates the quantity-weighted expected value. "Median" considers the median value, "5-" and "95-percentile" show the 5th and 95th percentiles of the quantity-weighted average values, "Std" is the standard deviation, and "Range" is the difference between the 95th and 5th percentiles. The data ranges from the beginning of 2014 to the end of 2021. Standard errors are in parenthesis.

Appendix Table E4: Participating dealers' demand does not increase when a dealer exits

	Demand in C$		Demand in % of supply	
exit	-12.22	(28.28)	-0.288	(0.929)
Adjusted R^2	0.0221		0.0027	
Observations	286		286	

Appendix Table E4 provides evidence that (participating) dealers do not significantly adjust their auction demands when a dealer exits the market. Concretely, we regress the maximal amount any participating dealer demands in the closest auctions around a dealer exit (which we observe, but cannot display, in Appendix Table E1) on an indicator exit variable that is one post-exit and also exit-event fixed effects. We report the estimated coefficients and standard errors in parenthesis, with demands expressed in millions C$ and in percentages of supply. In both cases, the exit coefficient is statistically insignificant at 10%. This is also the case when estimating separate regressions for each of the eight exit events.

Appendix Figure E1: Auction allotment by investor class for U.S. government bond auctions

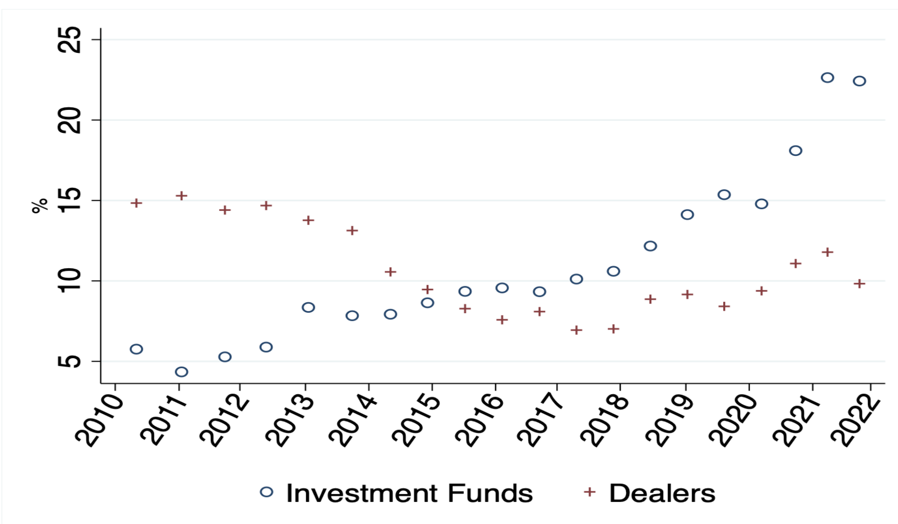

Appendix Figure E1 shows the auction allotment in percentage of supply in U.S. government bond auctions from the beginning of 2010 to the end of January 2022 for broker/dealers (plus) and for investment funds (circle). Broker/dealers include primary dealers, other commercial bank dealer departments, and other non-bank dealers and brokers; investment funds include mutual funds, money market funds, hedge funds, money managers, and investment advisors. To create this graph, we use public data from TreasuryDirect.org, available at https://home.treasury.gov/data/investor-class-auction-allotments, accessed on July 19, 2023.

Appendix Figure E2: Purchased amount by dealers, customers, and hedge funds

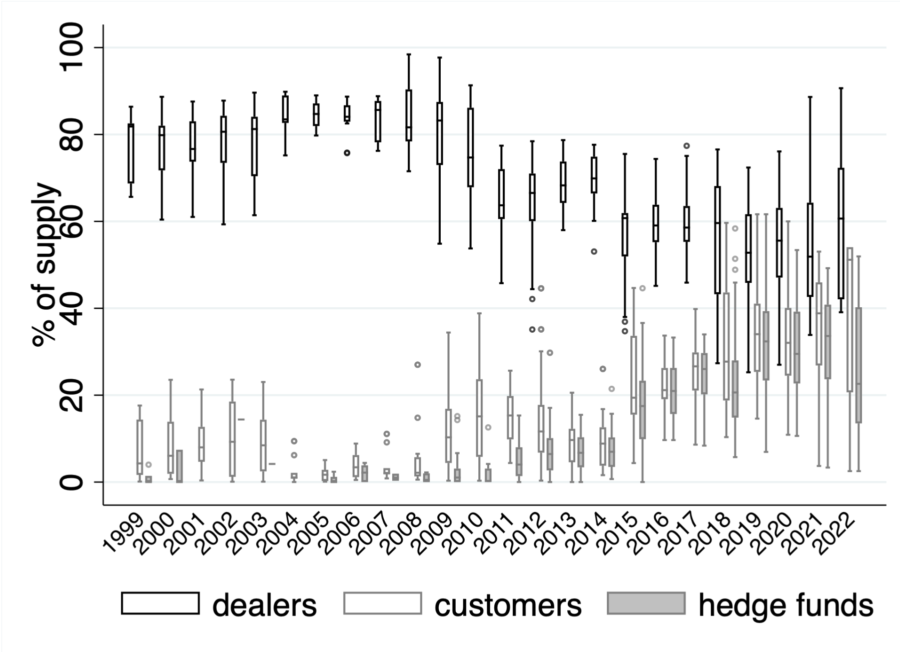

Appendix Figure E2 shows the distribution of how much dealers, customers, and hedge funds win (as a group) in percentage of the total amount issued across all bond auctions in our sample for each year from 1999 to 2022.

Appendix Figure E3: Purchased amount by investor groups

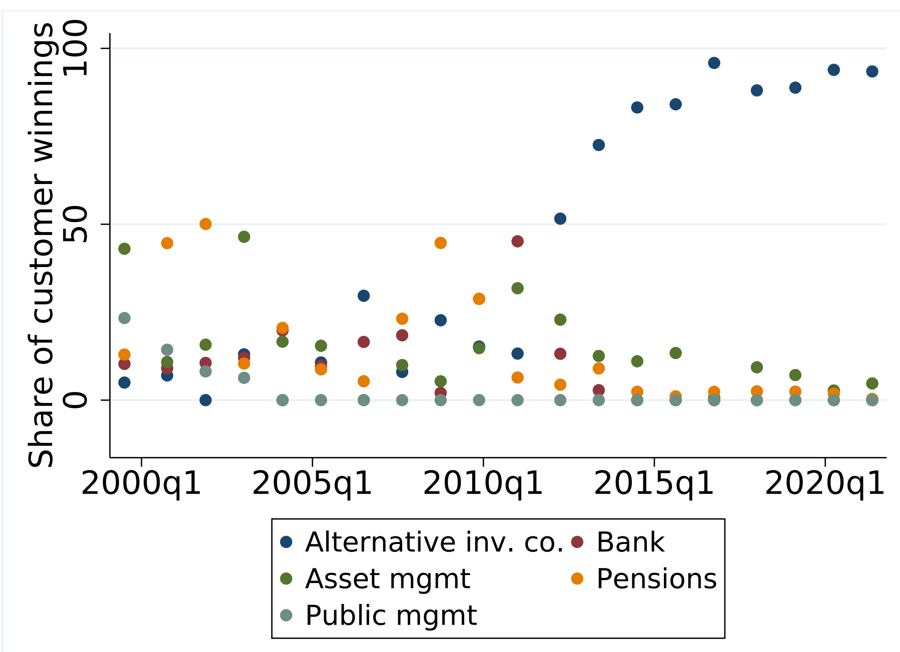

Appendix Figure E3 shows a binned scatter plot of how much each investor group wins in percentage of the total supply bought by non-dealers from 1999 to 2022.

Appendix Figure E4: Customer bid shading under different bidding assumptions

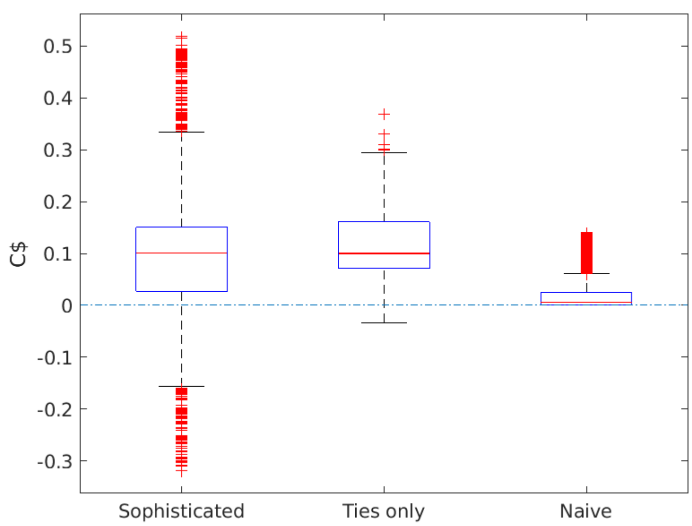

Appendix Figure E4 shows the distribution of the customer's average shading factors in C$, defined as the customer's quantity-weighted average value minus the quantity-weighted average bid, under different assumptions regarding the customer's degree of sophistication. The box plot called "Sophisticated" shows the distribution of the lower bound quantity-weighted customer value estimates according to our model. 'Ties only' estimates the values under the assumption that customers bid according to Proposition 1 (*ii*) but set $\lambda_{ti} = 0$. The 'Naive' box uses the dealers' optimality condition, i.e., forcing the customer to bid without accounting for the impact of the information their bid might have on the dealers' behavior. Prices are in C$ with a face value of $100.

Appendix Figure E5: Primary dealers are above minimal bidding requirements

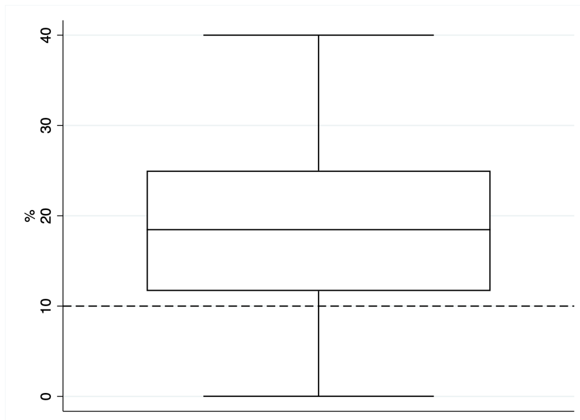

Appendix Figure E5 provides evidence that primary dealers are, with rare exception, above the minimal bidding limit of 10%, which is required to maintain their primary dealer status, conditional on market participation in a given year. In addition, supervisory data show that there have been extremely few violations in the past decade. This suggests that Canadian dealers do not face the dynamic trade-off, noted by Rüdiger et al. (2023), according to which dealers forgo one-shot auction surpluses in order to fulfill the minimal bidding requirements that must be met over a longer horizon. Concretely, the figure shows the distribution of the maximal amount an active primary dealer demands in an auction (as percentage of supply), where a primary dealer is active if they place at least one bid over the course of an entire year and the maximal demand is zero if the dealer does not participate in an auction. The distribution is taken over auctions and primary dealers. Outliers are excluded.

Appendix Figure E6: Random matching of customers to dealers

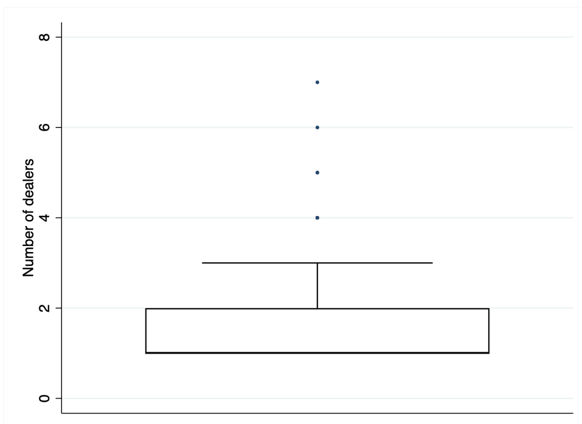

(A) Within auction

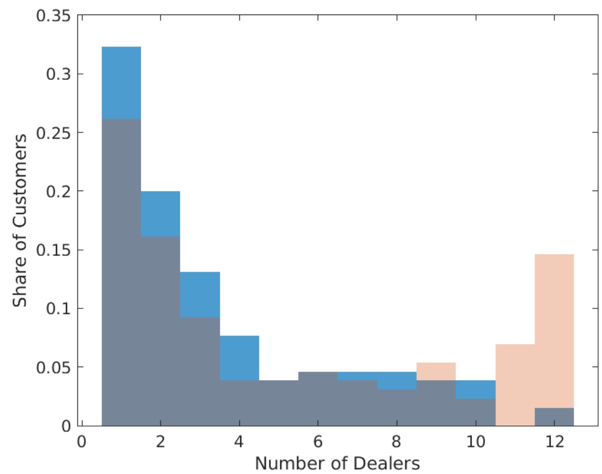

(B) Across auctions

Appendix Figure E6A shows the distribution of how many dealers a customer places a bid through within an auction. The median is 1. Figure E6B plots the distribution of the number of unique dealers used by a customer in all auctions in the data (in pink) and the number of unique dealers that would be predicted for each customer under random matching (in blue). The model prediction fixes the maximum number of dealers at the median number of dealers across years (12). The predicted distribution of the number of dealers used by each customer contains predictions from a single simulation, drawing independently one of the 12 dealers with equal probability for each customer each time they bid. The histogram plots the total number of unique dealers matched to each customer in the simulated sample. The distributions are broadly similar, but the model-predicted distribution somewhat overestimates the probability that a customer sometimes uses all of the possible dealers.