Entry and Exit in Treasury Auctions

PRELIMINARY DRAFT: Please do not cite or circulate!

Jason Allen Ali Hortaçsu Eric Richert Milena Wittwer*

July 31, 2023

Abstract

Regulated banks have traditionally dominated Treasury markets. More recently, less regulated institutions, such as hedge funds, have entered these markets. We document steady bank exit and rising, yet volatile hedge fund participation in the Canadian primary market. To understand hedge fund entry and to trade-off the benefits of greater competition against the costs of higher market volatility, we introduce and estimate a model with multi-unit auctions and endogenous entry. A counterfactual analysis suggests that hedge fund entry was partially driven by bank exit, and that competition benefits may be small compared to volatility costs. This trade-off is likely present in other markets with regular and irregular participants, which can be studied in our framework.

JEL: D44, D47, G12, G28

Keywords: Treasury market, multi-unit auction, intermediaries, hedge funds, regulation

*The presented views are those of the authors, not necessarily of the Bank of Canada. All errors are our own. We thank Walter Muiruri for excellent research assistance. We are especially grateful to Philippe Besnier. We thank Mark de Guzman, Jim Hatley, Faizan Maqsood, and Matthieu Truno for sharing their expertise on the evolution of the Canadian Treasuries market. Jason Allen - Bank of Canada, Email: jallen@bankofcanada.ca, Ali Hortaçsu: University of Chicago, Email: Hortacsu@gmail.com; Eric Richert, University of Chicago, Email: richert.eric@gmail.com, Milena Wittwer, Carroll School of Management, Email: wittwer@bc.edu.

1 Introduction

Governments worldwide have traditionally relied on regulated banks, known as primary dealers, to consistently purchase government debt and facilitate trades between investors, such as firms, public entities, and individuals, effectively "making markets." In recent times, however, other institutions have assumed an increasingly important role. For instance, hedge funds doubled their exposure to U.S. Treasuries between early 2018 and February 2020, reaching $1.45 trillion and $0.94 trillion in long and short exposures, respectively (Banegas et al. (2021)). Given that hedge funds and other non-dealer institutions, referred to as customers, are subject to less stringent regulations than dealers and have no obligation to participate in market-making activities, the implications of this observed trend on the functioning of Treasury markets remain uncertain.[1]

The goal of this paper is to better understand the increased participation of customers in Treasury markets and evaluate the consequences for market functioning. Answering these questions is challenging because there is limited data that allows us to study how customers, in particular hedge funds, trade Treasuries. We overcome this challenge by focusing on the Canadian primary market in which the government issues bonds via regularly held auctions. Using twenty-three years of bidding data, we document dealer exit and rising, yet volatile, customer participation. We then introduce and estimate a structural model to assess the role of dealer exit in explaining customer entry, and to quantify the benefits of greater customer competition against the costs of higher volatility.

Our data covers bidding information of all Canadian Treasury auctions from 1999 until 2022. For most of the paper we focus on bond auctions, which are more profitable than bill auctions and have attracted hedge funds. Two types of bidders participate in these auctions: dealers and customers. Only dealers can submit their bids directly to the auctioneer; customers must bid via a dealer. We observe bidder types, unique anonymized bidder identifiers, all submitted bids (including bid-updating, i.e. any bid before auction-closing that refreshes an existing bid), and which security is issued at what amount. We also know the

[1]Unprecedented market turmoil in March 2020 triggered a policy debate on whether to reform Treasury market rules and capital requirements. For statements by the Federal Reserve and newspaper articles see Logan (2020), Ackerman and Hilsenrath (2022), Grossman and Goldfarb (2022); for related academic literature see Falato et al. (2020), Haddad et al. (2020), He et al. (2022), Kargar et al. (2021), Vissing-Jorgensen (2021), Kutai et al. (2022), Ma et al. (2022).

time each bid is placed, and how bids are submitted (to the auctioneer by a dealer or via a dealer by a customer), and where the market cleared.

With these data, we document a series of novel facts. The first set of facts is about entry and exit. We show that dealers have systematically exited the primary market since the very first auctions took place. In contrast, customers, in particular hedge funds, have entered the market. However, unlike dealers, who have an obligation to regularly attend auctions and buy sufficient amounts of debt, hedge fund auction participation is volatile. This suggests that hedge funds select into specific auctions. The second set of facts help us model bidding behavior, conditional on auction participation. We show that dealers who observe an aggressive customer bid systematically adjust their own bids in anticipation of greater competition in the auction. Sophisticated customers, such as hedge funds, should take this effect 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 whether to commit to bidding in all auctions of the upcoming year at a cost. Then, before each auction, customers (such as hedge funds) decide whether or not to enter that specific auction, at a cost. Conditional on participation, all customers and dealers draw a signal about how much they value the bond—representing how much profit they expect to generate post-auction. Bids are then placed, either directly to the auctioneer or via a dealer. The auction clears at the price at which aggregate demand meets supply and each bidder pays their offered prices for all units won.

To solve for equilibrium conditions of this game and estimate dealer and customer values and entry costs, we build on the empirical literature on multi-unit auctions, in particular Hortaçsu and Kastl (2012), Guerre et al. (2000), Hortaçsu (2002) and Kastl (2011). Unlike existing studies, we allow customers to behave strategically in that they anticipate dealer updating, and endogenize bidder participation. This contributes to the empirical auction literature that allows for endogenous entry, but has focused on single-object auctions so far (see Hortacsu and Perrigne (2021) for an overview). For multi-unit auctions we know too little about how the equilibrium bid function behaves to adopt existing estimation approaches. Our innovation is to leverage estimated bounds on bidder-specific surpluses together with a matching procedure to estimate entry costs in multi-unit auctions. This approach could be used to endogenize bidder participation in other settings, including electricity, renewable

energy, carbon allowances, or diamond auctions.

We estimate high participation costs for both dealers and customers, and find that customers are willing to pay more for bonds than dealers. This suggests that (participating) customers expect to execute more profitable trading strategies with the bond post-auction than (participating) dealers. One reason for this is that customers face less stringent regulation than dealers, especially post global financial crisis. Consistent with this, we find that dealer values are not significantly different from customer values prior to 2007-2009, but differ significantly afterwards.

To disentangle whether the rise in customer participation was because dealers exited the market (potentially due to regulatory changes) or because of general changes in market conditions that made it more profitable for customers to buy bonds, we leverage our structural model and conduct a counterfactual exercise. To compute counterfactual bids without adopting the standard simplifying assumption (used in empirical studies of multi-unit auctions) of truthful bidding, we employ the empirical guess-and-verify approach by Richert (2021, 2022). When we add back the dealers that exited the market since 2014 and simulate how many customers would have participated in these counterfactual auctions, we find that an average customer would have been about 56% less likely to participate in an average auction had dealers not exited. The sizable decline in participation suggests that customer entry was meaningfully driven by dealer exit.

Then, we use our model to evaluate market consequences from the rise in customer participation. Since we are able to compute counterfactual bids without having to abstract from bid-shading, we can analyze changes in competition and volatility in market outcomes. On the one hand, stronger participation may increase competition, which would reduce debt funding costs and price distortions due to bid-shading. On the other hand, volatile bidder participation may increase volatility in market outcomes, such as the market price.

To quantify the increase in the degree of competition when adding an extra bidder in the spirit of Bulow and Klemperer (1996), we start with a simplified environment with one type of bidder (dealers) that bids directly to the auctioneer. We we find that the expected price drops by 10% when removing one bidder, because of stronger shading, and because there is a non-zero risk of auction failure due to insufficient demand. When varying the probability of that one extra customer participates in the typical auction, fixing the observed number of dealers, there is no risk of auction failure. However, the expected price decreases by 25%,

expected revenue decreases by 3% (or C$98 million) and bid-shading increases by 176%, when the customer does not participate relative to when she participates.

To quantify the impact of volatile dealer and customer participation on price volatility, we create a scenario where each auction is identical to the median auction, removing any auction-specific differences. We then analyze how the expected price changes when each dealer or customer randomly decides whether to participate in each auction, instead of committing to all of them. The estimated standard deviations of expected prices, 17.3 with volatile dealer participation and 3 with volatile customer participation, are substantially larger than the observed standard deviation of prices within a year, which averages at 3.6 and is mostly driven by auction-heterogeneity. This emphasizes the risk of significant price volatility when there is insufficient commitment from bidders to participate regularly in auctions.

In light of the competition-volatility trade-off, we compare the smooth functioning of primary auctions without obligating dealers to commit to regular participation, to a scenario where both customers and dealers bid directly, but dealers have to commit ex ante to participate (as in the status quo). Despite not obligating dealers to participate regularly, most auctions attract sufficient bidders to ensure full auction coverage. However, about 2% of auctions are at risk of failing without dealer commitment, resulting in a drop of over 20% in the expected price of these auctions. Even in fully covered auctions, expected revenue decreases by 0.04% in the median without dealer commitment.

Taken together, our findings indicate that some form of commitment is crucial for obtaining stable market conditions in primary markets. Absent a commitment, market volatility increases and auction coverage is at risk. Intuitively, these insights generalize to the secondary market. Requiring regular auction participation is equivalent to requiring some traders to act as designated market maker who provides liquidity, even in bad times when it is not profitable to do so.

More broadly, the competition-volatility trade-off we highlight might be present in many settings other than Treasury markets. First, it is common in auctions for financial products to have a set of regular and irregular bidders (e.g., Hendricks et al. (2023) for mortgage securities; Richert (2022) for credit event auctions). Our framework can be easily adjusted to fit these applications. Second, many financial markets are populated by core market makers that consistently handle a fraction of the trade volume, and opportunistic traders (or even retail) that come and go based on individual needs; even non-financial markets

are often populated by regular and irregular participants, e.g., global versus local firms, or loyal versus non-loyal customers. In electricity markets, for example, volatile production from renewable energies competes with stable generation from other sources (Petersen et al. (2022)). While our framework does not fit these applications directly, the economic insights generalize.

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 (see Appendix Figure A1).

Canadian Treasury market. Canada adopted this market structure to distribute it's debt in November 1998.[2] Since then, Treasury auctions are held regularly according to an annual auction schedule. For instance, between 1999 and 2022 there have been about 28 government bond auctions per year, with an average issuance size of C$3.24 billion.

Anyone can participate in Treasury auctions, but only dealers can bid directly.[3] Other bidders, called customers, can only bid indirectly by placing their bids via one of the dealers. This gives dealers access to more information that customers (and other dealers).[4]

Customers include different types of institutions, such as pension, mutual, and ETF funds, insurance companies, sovereigns, or bank treasuries. For over a decade, the biggest customer category are "alternative investments companies" (AIC), which include hedge funds. For simplicity, we use the term hedge fund to describe AICs throughout the paper.

A bidder may submit and update two types of bids from the time the tender call opens until the auction closing. The first type 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. Bids must be stated in multiples of $1,000, and the minimum bid is $100,000. The second type of bidding is a non-competitive bid, which is a quantity order that the

[2]Prior to 1983 government debt was issued via syndication. This involved the sale of Treasuries at negotiated prices, to a group of underwriting dealers who then redistributed the securities to their clients. By 1991, 75% of government debt was distributed by discriminatory price auctions (with many participants) and 25% via syndication (Pellerin (2006)).

[3]Strictly speaking, there are two types of dealers in Canada: primary dealers and government security distributors, which are smaller dealers. We do not distinguish between these groups.

[4]Appendix Figure A2 shows what a dealer sees when bidding.

bidder will win for sure at the average price of all accepted competitive bid prices. However, since non-competitive bids cannot be larger than $10 million for dealers and $5 million for customers, their size is trivial relative to competitive bids.

Market players and regulations. In Canada, as elsewhere, dealers have an obligation to actively buy bonds in the primary market, and to act as market-makers in the secondary cash and repurchase (repo) markets where they provide liquidity to investors who seek to exchange government bonds for cash. In exchange, dealers enjoy benefits. For instance, dealers have privileged access to liquidity facilities, overnight repurchase operations at the central bank, and access to Receiver General auctions (Chapman et al. (2007)).

Given the important role dealers play in the market, they are heavily regulated. In the aftermath of the 2007-2009 financial crisis, regulation tightened for the largest banks. For instance, mandatory trade-reporting, meant to increase price transparency, came into effect in 2015.[5] Banks also saw an increase in capital requirements. One prominent example of that is the Basel III leverage ratio, which was implemented at the end of 2014. It constitutes the most relevant capital constraints for dealers when trading government bonds according to a survey among market participants, and has been shown to affect dealer demand for government debt, as well as market participation (see CGFS (2016); Allen and Wittwer (2022); Favara et al. (2022)).

Traditionally, customers, such as hedge funds, played a negligible role in Treasury markets.[6] Hedge funds are typically exposed to the Treasury market either as part of a broader investment strategy, or as a source of liquidity. Further, they conduct arbitrage trades to exploit mispricing. Examples include "on/off-the-run arbitrage", which involves buying cheaper off-the-run bonds and shorting more expensive on-the-run-bonds; and "cash/future basis trading", which means going long on the bond issued at auction and shorting the future (Banegas et al. (2021)). In addition, Canadian hedge funds conduct arbitrage between U.S. and Canadian government bonds. For instance, when they expect a devaluation of the Canadian dollar relative to the U.S. dollar, they short-sell Canadian bonds and buy U.S. bonds.

[5] All fixed-income transactions are reported to the Investment Industry Regulatory Organization of Canada via the Market Trade Reporting System.

[6] 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)).

During regular times, hedge funds can act as market makers, engaging in trades that counter the positions of other investors (as documented by Sandhu and Vala (2023)). This can increase market competition. However, since hedge funds are not obligated to provide liquidity, they can also contribute to higher market volatility during times of market distress. For instance, in March 2020, hedge funds, along with other investors, sold significant amounts of U.S. Treasuries (e.g., Barth and Kahn (2020); Vissing-Jorgensen (2021)), leading to market imbalances, reduced liquidity, and increased price volatility (see Appendix Figure A4). Increased hedge fund trading may also have implications for systematic risk in the market, given that hedge funds are more likely to employ riskier trading strategies—an effect that we do not consider in our analysis.

Data. Understanding the activities and impact of hedge funds on market functioning poses challenges due to limited data availability.

For one, comprehensive long panels of trade-level data with unique identifiers for all traders are not readily accessible. For instance, the U.S. started collecting trade-level data through TRACE in mid-2017, but customer reporting with unique IDs is not mandatory. Similarly, while the Bank of England and Bank of Canada provide trader IDs, identifying all customers remains difficult (e.g., Kondor and Pintér (2021); Allen and Wittwer (2021); Pintér and Semih (2022); Barth et al. (2022)).

Additionally, customers are not obliged to report their trades, unlike dealers. Consequently, studies analyzing hedge fund trading behavior in the past years, including Banegas et al. (2021), and Sandhu and Vala (2023), face limitations in examining trades between hedge funds and dealers, which could represent a biased sample of hedge fund trades. For instance, Banegas et al. (2021) acknowledge challenges in assessing the extent of hedge funds' Treasury selling during March 2020 due to a lack of detailed data on hedge funds' cash and derivatives positions.

We overcome the data challenge by focusing on the primary market. We observe all winning and losing bids in regular government bond auctions from the beginning of 1999 until the end of January, 2022. This represents the entire history from the start of the auctions until the current day with the exception of a few auctions held in December of 1998. All auction participants are identified by name (which are visible only to certain Bank of Canada staff). This is crucial for studying entry and exit of dealers because it allows us

Table 1: Data Summary of Bond Auctions

	Mean	SD	Min	Max
Issued amount (face value C$ 100, in C$M)	32.42	10.45	10.00	70.00
Revenue (in C$B)	3.25	1.04	0.88	7.00
Number of Dealers	14.46	2.61	11	23
Number of Customers	6.74	2.56	1	15
Comp demand of a dealer (as % of supply)	14.80	7.51	0.00	40
Comp demand of a customer (as % of supply)	5.83	4.70	0.01	25
Non-comp demand of a dealer (as % of supply)	0.09	0.04	0.00	0.30
Non-comp demand of a customer (as % of supply)	0.19	0.16	0.00	0.76
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 (as % of supply)	4.79	5.85	0	35
Amount won by a customer (as % of supply)	4.02	5.90	0	25

Table 1 displays summary statistics of our sample, which goes from 10 February, 1999 until January 27, 2022. There are 645 auctions. The total number of competitive bidding functions (including updates) is 62,813. The total number of non-competitive bids is 10,552.

to capture mergers, acquisitions, and name changes over the long sample period. Table 1 provides key summary statistics for our auction sample.

We augment the auction data with daily average prices of each security from the secondary market, the future market and the repo-market from the Canadian Depository for Securities (CDS). These data range from beginning of 2014 until the end of 2021.

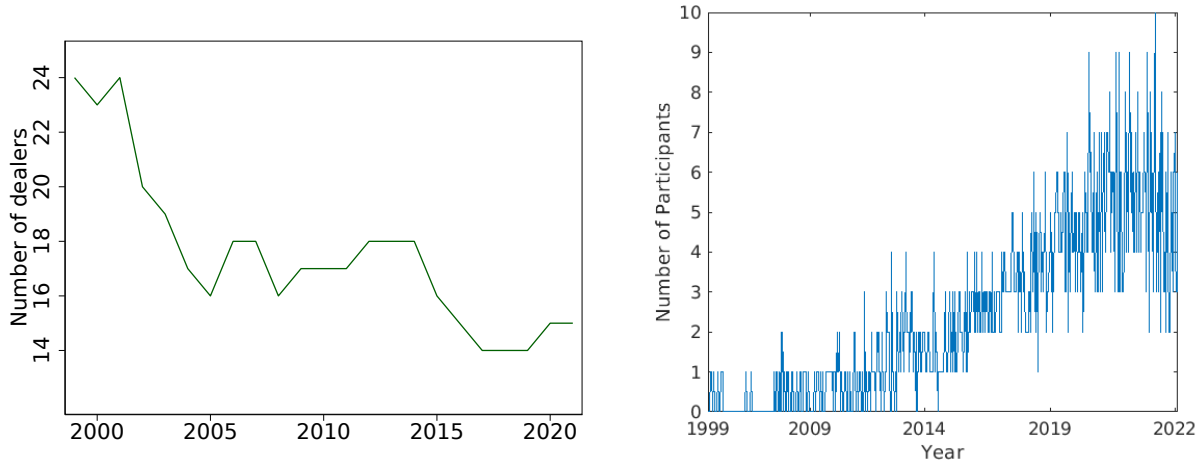
3 Empirical evidence: Exit, entry and bid updating

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 since the adoption of the primary dealer model. Appendix Figure A5 visualizes similar trends using public data of U.S. Treasury auctions to illustrate that thees trends are not Canadian-specific, but rather driven by global trends.

On the one hand, dealers have exited the market. The total number of dealers declined from 24 in 1999 to 15 in 2021 (see Figure 1A). After an initial round of exits in the first 5 years, the number of dealers remained stable until 2014, when global payers—such as Deutsche Bank and Morgan Stanley—exited the auctions for Canadian bonds. Smaller

Figure 1: Dealer exit and hedge fund entry from 1999-2022



(A) Number of dealers (B) Number of hedge funds

Figure 1A shows how many primary dealers and government security distributors have participated in primary auctions since the first auction in 1999 until 2022. Figure 1B shows how many hedge funds participate in each auction from 1999 until 2022.

broker-dealers—PI Financial Corporation and Ocean Securities—followed in 2015.[7]

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. This would align with anecdotal evidence. For instance, according to research by Greenwich Associates—a leading financial consultancy—regulations implemented after the 2008 global credit crises caused general retreat from Canadian debt markets in 2014 (Altstedter (2014)). Providing causal empirical evidence for this conjecture is beyond the scope of this paper, which focuses on customer entry.

On the other hand, customers (in particular hedge funds) have become more active (see Figure 1B). The number of hedge funds has increased from zero participation 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 Figure 2, and Appendix Figure A3). Interestingly, customer entry coincides with dealer exit: the number of hedge funds increased around 2014, and the percentage of supply they purchased increased from below 10% to almost 20% between 2014 and 2015.

[7]Appendix Table A1 lists all dealer entry/exits.

Figure 2: Purchased amount by dealers, customers, and hedge funds

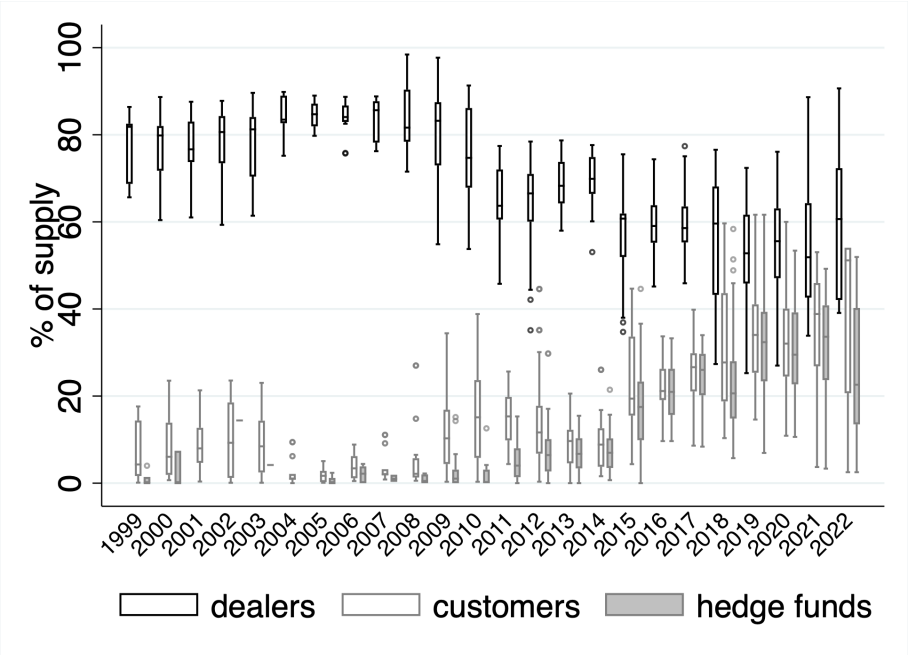


Figure 2 shows the distribution of how much dealers, customers, and hedge funds win (as a group) in percentage of the total amount issued across all bond auctions in our sample for each year from 1999 until 2022.

Drivers of customer participation. To find out whether we can predict customer participation based on institutional about hedge funds, we regress the number of participating customers, N_t^h, on a set of explanatory variables using CDS data from 2014 onward (when customers are almost exclusively hedge funds):

$$N_t^h = \beta_0 + \beta_1 CF_t + \beta_2 BM_t + \beta_3 MP_t + \beta_4 QE_t + \beta_5 FXRate_t + \beta_6 Spread_t + \epsilon_t. \quad (1)$$

The first variable, CF_t, indicates the auction dates on which cash/future basis trade is profitable (inspired by Barth and Kahn (2020); Banegas et al. (2021)).[8] The second indicator variable, BM_t, tells us whether the bond-to-be-issued has benchmark status, which is the Canadian equivalent of being on-the-run (Berger-Soucy et al. (2018)). The third variable, MP_t, indicates whether an auction occurs within two days prior to a monetary policy com-

[8]To determine profitability, we first determine whether the bond being auctioned is the cheapest-to-deliver for the nearest upcoming futures delivery date. Then we check if the rate of return on the basis trade is greater than the cost of borrowing money in the repo market. If both is true, the basis trade is profitable.

mittee (MPC) meeting. This variable is inspired by the findings of Lou et al. (2023), who demonstrate that hedge funds tend to purchase bonds outside of the pre-MPC window to avoid uncertainty associated with potential interest rate changes. The fourth variable, QE_t, flags auctions days on which the government conducts a bond issuance in the morning and engages in quantitative easing by purchasing bonds in the afternoon. The idea is that hedge funds might buy cheap and sell high (similar to evidence by An and Song (2018, 2023)). The fifth variable, $FXrate_t$, is the rate at which Canadian dollars are exchanged to U.S. dollars on auction day, which might affect the desirability to conduct arbitrage between Canadian and U.S. government bonds. The sixth variable, $Spread_t$, measures the difference between the highest and lowest price at which a bond that is about to be issued at auction is traded in the secondary market on a day prior to the auction; we include three prior-auction days and average across days. This variable captures price dispersion and likely correlates with the (unobserved) bid-ask spread that customers can charge their clients when trading the bond.

We are currently waiting for some explanatory variables of this regression, so we cannot show the results yet. However, in all specifications we have analyzed, a moderate size of the R^2 suggest that customer participation in the primary market cannot be easily explained by the factors considered in previous studies. It indicates that there may be unobservable factors that vary across auctions and play a significant role in driving customer participation. This motivates us to employ a non-parametric approach to account for changes in these factors across auctions in our structural model.

Market consequences of customer participation. Independently of what triggers increased customer participation at auction, the change likely has advantages and disadvantages.

On the one hand, their participation increases the number of bidders at auction. This implies more competition, which could result in lower funding costs for the government. On the other hand, auction participation has become more volatile since customers, unlike dealers, do not participate in all auctions. Customers, therefore, are not a perfect substitute for dealers.

To assess the benefit of higher (average) competition and the cost of higher volatility in auction participation, we introduce and estimate a structural model of Treasury auctions

Figure 3: Dealer adjustment and customer bid

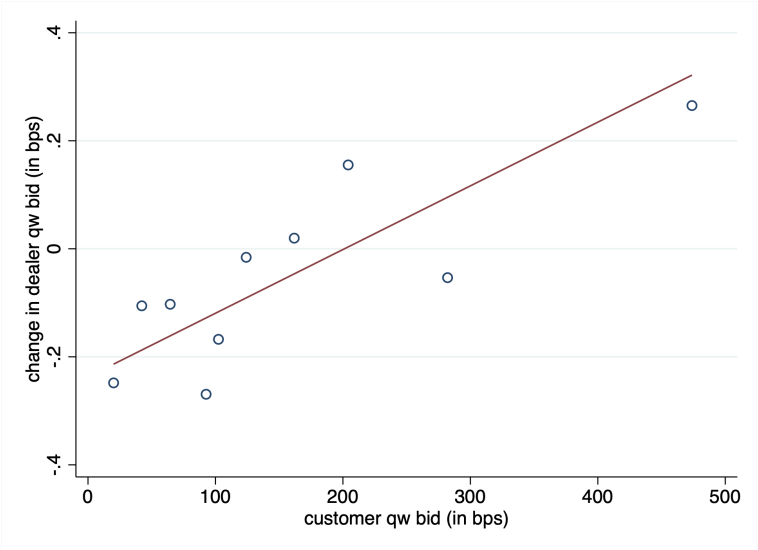


Figure 3 shows the correlation between the change in the dealer's quantity-weighted bid and the quantity-weighted bid of the customer, conditional on the dealer updating her bid. A quantity-weighted bid is the total amount a bidder bid divided by the total amount she demanded. A bid is measured in yields to maturity and expressed in bps.

with endogenous bidder entry. The model allows us to conduct counterfactual simulations in order to analyze why more customers entered the market.

Dealer updating. Our auction model, presented in Section 4, builds on Hortaçsu and Kastl (2012), who model bidding behavior in Canadian Treasury auctions. One interesting feature is that all customers must bid via dealers. This generates an information asymmetry between dealers and customers, and potentially disincentivizes customers participation.

Intuitively, dealers who observe an aggressive customer bid should aggressively update their own bid—the aggressive customer bid signals to the dealer that the auction is going to be competitive.[9] Supporting our conjecture, the correlation between the change in the dealer's (quantity-weighted) average bid, conditional on updating, and the customer's (quantity-weighted) average bid is positive (see Figure 3). Naturally, there might be a selection problem—we only observe how a dealer updates their bid if there is an actual update.

[9]In theory, it could also be that after observing an aggressive customer bid, a dealer re-evaluates their own valuation for the bond upwards. This would generate the same updating behavior, but be problematic for estimation, which relies on valuations being conditionally independent. In our empirical application we reject the hypothesis that dealers learn about fundamentals (see Appendix A.1). This is consistent with Hortaçsu and Kastl (2012).

To overcome the selection concern, we need an instrument that exogenously affects the probability that a dealer updates (conditional on observing a customer's bid) but not the bid itself. Our idea is to use the time until the auction closes as an instrument—dealers are less likely to be able to update their bids the less time they have left in the auction. Concretely, we estimate a Heckman-Selection model, in which the time until the auction closes, time_left_i, affects whether the dealer updates her bid after observing a customer bid, bid_i^d, but not by how much the dealer changes her own bid Δbid_i^d:

$$\Delta bid_i^d = \beta_1 bid_i^h + \epsilon_{1i} \text{ if } \beta_2 \text{time_left}_i + \epsilon_{1i} > 0, \tag{2}$$
$$\text{where } \epsilon_1 \sim N(0, \sigma), \epsilon_2 \sim N(0, 1), \text{corr}(\epsilon_1, \epsilon_2) = \rho.$$

One complication in implementing this identification strategy arises from the fact that bids are step functions, and it is not obvious how to capture the shape and movement of that function. To get a sense of what matters most, we consider three moments of the bidding step function: the (quantity-weighted) average bid, the number of steps, and the highest amount demanded.

Table 2 presents the results. In line with our expectations, we find that a dealer is more likely to update her bid conditional on observing a customer bid when she has more time until the auction closes. A dealer bids more aggressively when observing a more aggressive customer bid. Of the moments in the customer's step function that we consider, only the quantity-weighted average is statistically significant—a detail that we will leverage 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. First, we allow customers to be strategic in that they take into account that dealers will systematically update their bids when observing their bid. Second, we endogenize the decision to participate in the market. Here we distinguish between the dealer's decision to exit the market and the customer's decision to enter specific auctions. This is crucial for studying customer entry and dealer exit in the primary market.

Table 2: Heckman-Selection Model

Δ bid^d		
qw-bid	0.000800***	(0.000215)
number of steps	0.0393	(0.0314)
total demand	0.000264	(0.000163)
constant	-0.905***	(0.166)
$\mathbb{I}$(Update| customer bid)		
time_left	0.000318***	(0.0000169)
qw bid	-0.000422***	(0.0000744)
number of steps	0.0981***	(0.0133)
total demand	0.000346***	(0.0000695)
constant	0.102***	(0.0294)
Inverse mills	1.198***	(0.280)
ρ	0.44449	
σ	2.69551	
Observations	16156	

Table 2 shows the results from estimating Heckman-Selection model (2) using bidding data from all regular Treasury auctions from 1999 until 2022. The model tests whether and in which direction a dealer changes her quantity-weighted bid, Δbid^d, (in bps) after observing a customer's bid—which we summarize by its quantity-weighted bid (in bpd), the number of steps and the total demand (in million C$)—, using the seconds left until auction closure as exogenous driver of selection into updating. Standard errors are in parentheses.

4.1 Players, timing and preferences

In order to mimic the events that take place during a fiscal year, our model has multiple phases. It is played by two groups (g) of market participants: dealers (d) and customers (h). The number of dealers who consider remaining as dealers, $\bar{N}^d$, and the number customers who are interested in bidding, N^h, are commonly known.

Alternatively, we could distinguish between different types of customers, for instance hedge funds versus other customers. Theoretically, such an extension would be straightforward. However, empirically, non-hedge fund customers play such a small role that the cost of complicating the model and increasing measurement error (due to more bidder groups in the resampling procedure described below) outweighs the benefit of separating non-hedge fund customers from other customers.

Phase 0 occurs at the beginning of a year, when the debt issuance plan is announced. At this point, dealers decide whether they wish to continue being a dealer, which means

committing to participate in all auctions in the upcoming 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 that acts as a (primary) dealer suffers because it cannot do other things during the time it fulfills dealer-activities, such as bidding at auction and making markets.

Assumption 1. *At the beginning of the year, dealers' private commitment 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 debt management office.

Phase 1 starts right before the first auction, 1; similarly for phase 2, ...,T, which each begin right before auction $2, ..., T$. All phases are analogous—each phase t starts with customers observing current market conditions and deciding whether to bid in the upcoming auction t.

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

A customer's entry decision is taken after they learn about market conditions (which affects the market value of the bond), but before they learn their private signal about how much their own bond-valuation deviates from the market value—driven by institutional-specific reasons. This timing reflects the idea that most customers are part of large institutions, who tend to allocate tasks (such as bidding at auction or trading other assets) some time in advance. For instance, a trader who works for an insurance company is likely told to participate in the next Treasury auction before the bidding process starts. Similar to dealers, one may think of the entry cost as an opportunity cost.

Conditional on participation in auction t, each customer and dealer draws a private (possibly multi-dimensional) signal $\boldsymbol{s_{it}^g}$. The signals may be drawn from two different distributions for dealers and customers.

Assumption 3. *Dealers' and customers' private signals $\boldsymbol{s_{it}^d}$ and $\boldsymbol{s_{it}^h}$ 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.*

Within an auction t, a bidder's signal must be independent from all other signals, conditional on anything and everything that everyone knows when bidding. This includes a reference price-range provided by the auctioneer, in addition to all public information that is available in the forward (when-issued) market. Thus, bidders are equally informed about the price at which they can re-sell the bond in the secondary market. The private information arises, for instance, from observing client orders in the secondary market.[10] To further validate the assumption of independent private signals, we follow Hortaçsu and Kastl (2012)'s and conduct a statistical test. Results are reported in Appendix A.1. Supporting the assumption, we fail to reject that dealers who observe customer's bid's only learn about the degree of competition in the auction and not about the fundamental value of the bond.

Assumptions 1-3 rule out that bidders have incentive to adopt strategies that connect multiple auctions, because signals and costs are independent across auctions, and the signal distributions, F_t^g, are the same for all bidders in the group. This means, for instance, that a customer who won a lot in a previously held auction is as likely to participate in the upcoming auction as a customer who didn't win or who didn't participate. Thus, technically, our game consists of dealers' exit decisions in phase 0 plus T ex-ante identical, and separate auction sub-games. This is similar to existing studies of Treasury auctions. We think that this is reasonable given that we observe bidders selling their t-period auction-purchased bonds several days before the $t+1$ auction takes place.[11]

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 signals.

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

Given their values before auction t, bidders place bids. Each bid is a step function that characterizes the price the bidder would like to pay for each amount.

[10] Clients are different than customers. Clients interact with dealers and customers in the secondary market; they do not bid at auction. One reason for their orders orders to be conditionally independent is the fact that clients keep close relationships with a single intermediary (as shown by Allen and Wittwer (2021)).

[11] In the case of Treasury bills, auctions are held simultaneously and in parallel. See Allen et al. (2020) for an analysis of bidding in these auctions.

Assumption 5. *Each bidder has the following action set to place a bid:*

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

Dealers can submit their bids directly to the auctioneer, while customers have to place their bids with a dealer. A dealer who observes a customer's bid might gain additional information (about the competition she faces in the auction) relative to other bidders. To capture this, we define the information that is available to dealer i before placing her bid at time τ in auction t by $Z_{i\tau}$, and call $\theta^d_{i\tau} = (s^d_{it}, Z_{i\tau})$ the dealer's type. The type of a customer is her private signal $\theta^h_{i\tau} = s^h_{it}$—where, for simplicity, we use the same notation for dealers and customers.

The simplified notation allows us to define a pure strategy in an auction as a mapping from the bidder's set of types to the action space: $\Theta^g_{i\tau} \to A$. It is a bidding function, labeled $b^g_{i\tau}(\cdot, \theta^g_{i\tau})$, for bidder i of group g with type $\theta^g_{i\tau}$.

A dealer who observes a customer's bid might choose to update her own bidding function $b^g_{i\tau}(\cdot, \theta^g_{i\tau})$. In principle, even without observing a competing bid, dealers may update their bids at any point over the course of the auction. Customers too can update their bids prior to the auction closing. Rationalizing this type of updating is difficult—bidders can simply wait to place their bid just prior to the auction closing. This appears optimal if calculating and placing bids is costly. Given that we focus on bid-updating after observing customer bids, the anticipation of such updating and customer entry decisions, we abstract from modeling "irrational" updating, and refer to Hortaçsu and Kastl (2012) for ways on how to model these updates.

A customer anticipates that their dealer will update their own bid. However, since the customer doesn't know the dealers type, $\theta^d_{i\tau}$, she doesn't know, with certainty, if the dealer will update their bid. Further, she doesn't know how the dealer will change their own bid. As a result, a customer cannot be sure if the market-clearing price will increase (or decrease) when she marginally increases her own demand, q_k, at step k—anything can happen.[12] This

[12]As a comparison, when dealers don't update bids, the market clearing price will weakly increase in all states of the world if the customer increases her quantity q_k at price b_k by a little bit, assuming

makes it computationally complex for the customer to determine her optimal bid.

We assume that customers cannot solve the full-fledged (auction) optimization problem. To render the problem solvable, we assume that customers think that their dealer only considers a finite set of moments of their bidding function when updating their own bid. For simplicity, we restrict attention to moments that are real numbers and that are differentiable w.r.t. q at each price, for instance, the quantity-weight bid, the intercept with the price or quantity axis, or the steepness of the curve.[13] Formally, a moment is a mapping μ_t^l that transforms the bidding function, $b_{i\tau}^h(\cdot, \theta_{i\tau}^h)$ for type $\theta_{i\tau}^h$, into a real number $\mathbb{R}$, and that is differentiable w.r.t. q at each price.

$$\mu_t^l(b_{i\tau}^h(\cdot, \theta_{i\tau}^h)) = \frac{\sum_k^K b_k q_k}{\sum_k q_k} \text{ when } \{b_k, q_k\}_{k=1}^K \text{ constitute bidding function } b_{i\tau}^h(\cdot, \theta_{i\tau}^h). \quad (3)$$

Once all bidders submit their step function, the market clears at the lowest price, P_t^*, at which the aggregated submitted demand satisfies the total supply. The supply is unknown to bidders when they place bids because a fraction of the total allotment goes to noncompetitive tenders. These are quantity-only bids that win with certainty. The Bank of Canada is the largest non-competitive bidder at the auction.

Assumption 6. *Supply $\boldsymbol{Q_t}$ is a random variable distributed on $[\underline{Q}_t, \overline{Q}_t]$ with strictly positive marginal density conditional on s_{it}^g $\forall i, g = h, d$.*

Given all bidding functions, $b_{i\tau}^g(\cdot, \theta_{i\tau}^g)$ for all g, i, bidder i of group g wins amount q_{it}^{g*} at market clearing. She pays the amount she offered to win for each unit won. In case there is excess demand at the market clearing, each bidder is rationed pro-rata on-the-margin (see Kastl (2011) for details).

4.2 Equilibrium conditions

We solve for equilibrium conditions by backward induction. First, we characterize the equilibrium in auction t conditional on customer and dealer participation. Then, we determine which customers enter auction t. Finally, we analyze the dealers' decision to commit to being

that all other participants play as in the equilibrium. This monotonicity helps the bidder find her best reply.

[13] Lemma 1 in Appendix D.1 formally proves that the quantity-weighted bid function, which we focus on later, is differentiable w.r.t. q.

a dealer before the very first auction takes place.

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

$$TS^g_{it} = \mathbb{E}_t \left[\int_0^{\boldsymbol{q^{g*}_{it}}} [v^g_t(x, s^g_{it}) - b^g_{i\tau}(x, \theta^g_{i\tau})] dx \right]. \tag{4}$$

The expectation is conditional on the bidder's type; it is taken over the amount the bidder will win at market clearing, $\boldsymbol{q^{g*}_{it}}$. This depends on the strategies and types of all bidders, as well as the supply, and thus unknown to the dealer.

We focus on type-symmetric Bayesian Nash equilibrium (BNE) in which all dealers and all customers play the same bidding strategy if they have the same type. Formally, a pure strategy BNE of auction t is a collection of bidding functions $b^g_{i\tau}(\cdot, \theta^g_{i\tau})$ that each bidder i and almost every type $\theta^g_{i\tau}$, $b^g_{i\tau}(\cdot, \theta^g_{i\tau})$ maximizes the bidder's expected total surplus.

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

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

$$\Pr(b_k > \boldsymbol{P^*} > b_{k+1} | \theta^d_{i\tau}) \left[v^d_t(q, s^d_{it}) - b_k \right] = \Pr(b_{k+1} \geq \boldsymbol{P^*_t} | \theta^d_{i\tau})(b_k - b_{k+1}). \tag{5}$$

At the last step, $b_k = v^d_t(\bar{q}(\theta^d_{i\tau}), s^d_{it})$, where $\bar{q}(\theta^d_{i\tau})$ is the maximal amount the dealer may be allocated in the auction equilibrium.

(ii) Every step k in a customer's bid function $b^h_{i\tau}(\cdot, \theta^h_{i\tau})$ that generates moments $m^l_t = \mu^l_t(b^h_{i\tau}(\cdot, \theta^d_{i\tau}))$ for all l has to satisfy:

$$\begin{aligned} &\Pr(b_k > \boldsymbol{P^*} > b_{k+1} | \theta^h_{i\tau}, m^l_t) \left[v^h_t(q, s^h_{it}) - b_k \right] = \\ &\qquad \Pr(b_{k+1} \geq \boldsymbol{P^*} | \theta^h_{i\tau}, m^l_t)(b_k - b_{k+1}) - \sum_l \lambda^l_t \frac{\partial \mu^l_t(b^h_{i\tau}(\cdot, \theta^h_{i\tau}))}{\partial q} + Ties(b^h_{i\tau}(\cdot, \theta^h_{i\tau}) | \theta^h_{i\tau}, m^l_t) \end{aligned} \tag{6}$$

where $Ties(b^h_{i\tau}(\cdot, \theta^h_{i\tau}) | \theta^h_{i\tau}, m^l_t) = \Pr(b_k = \boldsymbol{P^*_t} | \theta^h_{i\tau}, m^l_t) \mathbb{E}[v^h_t(\boldsymbol{q^{h*}_{it}}, s^h_{it}) \frac{\partial \boldsymbol{q^{h*}_{it}})}{\partial q_k} | b_k = \boldsymbol{P^*_t}, \theta^h_{i\tau}, m^l_t] - \Pr(b_{k+1} \geq \boldsymbol{P^*_t} | \theta^h_{i\tau}, m^l_t) \mathbb{E}[v^h_t(\boldsymbol{q^{h*}_{it}}, s^h_{it}) \frac{\partial \boldsymbol{q^{h*}_{it}}}{\partial q_k} | b_{k+1} \geq \boldsymbol{P^*_t}, \theta^h_{i\tau}, m^l_t] + \Pr(b_k = \boldsymbol{P^*_t} | \theta^h_{i\tau}, m^l_t) \mathbb{E}[\frac{\partial \boldsymbol{q^{h*}_{it}}}{\partial q_k} | b_k = \boldsymbol{P^*_t}, \theta^h_{i\tau}, m^l_t] + \Pr(b_{k+1} = \boldsymbol{P^*_t} | \theta^h_{i\tau}, m^l_t) \mathbb{E}[\frac{\partial \boldsymbol{q^{h*}_{it}}}{\partial q_k} | b_{k+1} = \boldsymbol{P^*_t}, \theta^h_{i\tau}, m^l_t]] + \Pr(b_{k+1} < \boldsymbol{P^*_t} | \theta^h_{i\tau}, m^l_t) \mathbb{E}[\frac{\partial \boldsymbol{q^{h*}_{it}}}{\partial q_k} | b_{k+1} <$

$\boldsymbol{P_t^*}, \theta_{i\tau}^h, m_t^l]$, *and*

$$\lambda_t^l[m_t^l - \mu_t^l(b_{i\tau}^h(\cdot, \theta_{i\tau}^h))] = 0 \text{ for all } l \text{ in auction } t. \tag{7}$$

From the existing literature we know how the dealer determines her equilibrium bid. Essentially, she chooses her bid to maximize total surplus, TS_{it}^d, defined in (4), subject to market clearing. Similar to a first-price auction, she trades off the expected surplus on the marginal infinitesimal unit versus the probability of winning it (see Kastl (2017), p. 237 for more details). Given that it is never optimal for a dealer to submit a bid above her true value (Lemma 2 of Kastl (2011)), ties between the dealer's bid and someone elses bid can only occur at the last step of the bidding function (Lemma 1 of Kastl (2011)). As a result, the dealer wins the full amount that she demanded at all steps but the last. The first order condition simplifies to expression (5)—without ties and rationing of winning quantities.

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. To derive an intuition for this, take the perspective of a customer and let all other bidders play the equilibrium strategy. Assume that the customer considers changing her own bidding function $b_{i\tau}^h(\cdot, \theta_{i\tau}^h)$. She thinks that the dealer (who observes the perturbed bidding function) changes her bid only in response to observing her bid if the moments of the bidding function $\{m_t^l\}_{l=1}^L$ change. With this information, the customer first considers perturbation of the bidding function that leave the moments unchanged. Fixing some $\{m_t^l\}_{l=1}^L$, the customer finds the bidding function that maximizes her expected total surplus such that $\mu_t^l(b_{i\tau}^c(\cdot, \theta_{i\tau}^h)) = m_t^l$ for all l. Then, the customer chooses moments $\{m_t^l\}_{l=1}^L$ so that the total surplus is actually maximized.

Optimality condition (6) formalizes the first step of this maximization procedure with $\lambda_t^l \in \mathbb{R}$ as Lagrange multipliers. This condition has two additional terms relative to the condition for dealers (5). The first term, $-\sum_l \lambda_t^l \frac{\partial \mu_t^l(b_{i\tau}^h(\cdot, \theta_{i\tau}^h))}{\partial q}$, comes from the fact that the customer limits her search for bidding functions to those that generate moments $\{m_t^l\}_{l=1}^L$. The second term, $\text{Ties}(b_{i\tau}^h(\cdot, \theta_{i\tau}^h)|\theta_{i\tau}^h, m_t^l)$, includes all the cases in which the customer's bid ties with some other bid and must be rationed. These cases do not drop out for customers, because it can be optimal for a customer to place a bid at a step above her true value. This happens when such a bid changes the moments and with that the dealer's decision in a way that increases the customer's total surplus.

Entry and exit decisions. A customer i enters an auction t if her entry cost for auction t, γ_{it}^h, is smaller than the total surplus she expects to earn from participating in the auction before she observes her private signal, but after observing the signal distributions, F_t^g, and the shape of the value functions, $v_t^g(\cdot,\cdot)$, for $g = \{h, d\}$ which capture market conditions.

Proposition 2. *customer i with entry cost γ_{it}^h enters auction t if*

$$\gamma_{it}^h \leq \mathbb{E}_t[\boldsymbol{TS_{it}^h}|N^d] \text{ with } TS_{it}^h \text{ given by (4).} \tag{8}$$

The expectation is taken over the customer's private signal s_{ti}^h, and conditional on N^d dealers bidding in the auction.

Technically, $\mathbb{E}_t[\boldsymbol{TS_{it}^h}|N^d]$ is a mapping $\mathcal{E}^h$ from the signal distributions, F_t^d, F_t^h, the value functions, v_t^d, v_t^h, the customer's entry cost distribution, G^h, and the number of dealers and customers who participate in phase t of the game, N^d, N^h, into $\mathbb{R}^+$.

Anticipating all auctions, $t = 1, ..., T$ of the upcoming year, dealer i exits the market if her commitment cost is higher than the surplus she expects to earn from bidding in all T auctions of the upcoming year.

Proposition 3. *At the beginning of the year, dealer i with commitment cost γ_i^d exits the market if*

$$\gamma_i^d \geq \sum_{N^d=1}^{\bar{N}^d} \left(\sum_{t=1}^{T} \mathbb{E}_t \left[\boldsymbol{TS_{it}^d} | N^d \right] \right) \Pr(\boldsymbol{N^d} = N^d) \text{ with } TS_{it}^d \text{ given by (4).} \tag{9}$$

As for customers, the expectation is taken over the dealer's type for each auction t. Different to customers, the dealer cares about the aggregate surplus she expects to earn over the entire year, and when taking her decision she doesn't know how many other dealer's will compete in the auctions, N^d. Therefore, the dealer considers all possible realizations of N^d and weights each by how likely it is to occur, $\Pr(\boldsymbol{N^d} = N^d)$.

Technically, $\sum_{t=1}^{T} \mathbb{E}_t \left[\boldsymbol{TS_{it}^d} | N^d \right]$ is a mapping, $\mathcal{E}^d$, from all signal distributions and value functions of the upcoming year, F_t^g, v_t^g, for $g = \{h, d\}$ and all t, the customer's entry cost distribution G^h, and the number of dealers and customers who participate in phases $1 - T$ of the game, N^d, N^h, into $\mathbb{R}^+$. The probability of having N^d dealers, $\Pr(\boldsymbol{N^d} = N^d)$, is a mapping, $\mathcal{P}$, from the same objects as for $\mathcal{E}^d$ plus the dealer's cost distribution, G^d, and the number of maximal dealers, $\bar{N}^d$, into $[0, 1]$.

5 Identification and estimation

The goal is to learn about the unobserved bidder values, $v_t^g(\cdot, s_{ti}^g)$, their signals, s_{ti}^g, and the entry and exit cost distributions, G^h and G^d. For this, we fix the maximal number of dealers to what we observe in our sample, i.e., $\bar{N}^d = 24$. We define the number of potential customers in a year, N^h, as the maximal number of customers we observe bidding in any auction of that year.

Identifying and estimating values. The idea behind identifying and estimating bidder values from bidding data is to infer how much each bidder is truly willing to pay from equilibrium conditions under the assumption that everyone plays this equilibrium, here specified in Proposition 1. To estimate the probabilities that enter the equilibrium conditions, we extend Hortaçsu and Kastl (2012)'s resampling procedure to estimate the distributions of the market clearing price by allowing information of a bidder to be correlated over the course of the auction (details in Appendix B).

For dealers, we know that we can point-identify values (at submitted steps) given condition (5): Kastl (2012) shows that there is a unique $v_t^d(q_k, s_{it}^d)$ at each submitted point q_k in bidder i's bidding function. This implies that we can construct an upper and a lower bound for the entire value curve, i.e. at all quantity points, as illustrated in Figure 4.

Our contribution is to learn about customer values (at submitted steps) from their equilibrium condition (6). Due to dealer updating, we cannot point-identify customer values. However, we can construct sets of informative bounds on customer values that are consistent with the observed bids. For this, we assume that customers think that their dealer only cares about their qw-bid when updating their own bids—motivated by the empirical evidence in Table 2. The customer's optimality condition (6) then simplifies to an equation with a single moment, the qw-bid (3). For simplicity, we drop the super- and subscript $l = 1$.

To identify value bounds, we proceed in two steps. First, we determine K_{it} upper and K_{it} lower values for $v_t^h(q_k, s_{it}^h)$ for a customer i who submits K_{it} steps in a fixed auction, for a guessed value of the Lagrange multiplier λ_t. Given λ_t, we determine bounds for how much a customer values winning a step at which they tie and thus gets rationed, by relying on boundedness and monotonicity of $v_t^h(\cdot, s_{it}^h)$ (recall Figure 4). For example, we replace $v_t^h(\boldsymbol{q_{it}^{h*}}, s_{it}^h)$ by $\overline{v}_t^h(q_k, s_{it}^h)$ to determine an upper bound, and by $\underline{v}_t(q_{k+1}, s_{it}^h)$ to determine a lower bound. This leaves us with a system of $2K_{it}$ linear equations.To simplify the system

Figure 4: Illustration of what we can infer about values from bids

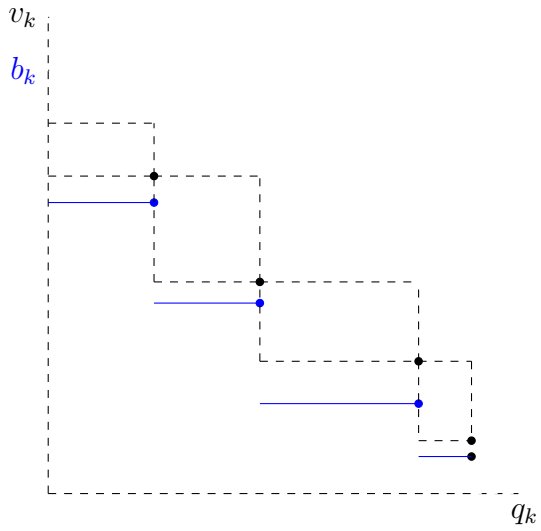


Figure 4 illustrates what can infer about values of a dealer who submits a bid function with 4 steps (shown in blue). We can point-identify her values at q_1, q_2, q_3 and q_4—the black circles. Further, since the value function is decreasing in quantity and bounded, we know that the dealer's values for quantities in between must lie in the dashed boxes.

of equations and reduce estimation noise, we show that at a subset of steps, rationing never occurs in equilibrium (see Lemma 2). For these steps, we can construct profitable deviations and set all terms in (6) involving rationing to zero. For instance, consider a step $k < k_{it}$ for the case in which $\lambda < 0$. Slightly increasing the bid at step k can break the tie and either generate a positive surplus or decrease the loss from placing a bid to favorably influence dealer updating.

Second, we check which values for λ_t are valid in equilibrium. For this, we rely on the fact that the equilibrium bidding function must generate a moment m_t^l that is optimal. To see how this works, assume for an instance that we could point-identify the customer's entire value curve. Then, moving away from m_t^l must decrease the customer's expected total surplus, so that a marginal upward and downward ϵ perturbation of m_t^l would lead to two inequality conditions $+TS_{it}^h(m_t^l) - TS_{it}^h(m_t^l + \epsilon) \geq \lambda_t \epsilon$ and $-TS_{it}^h(m_t^l) + TS_{it}^h(m_t^l - \epsilon) \leq \lambda_t \epsilon$ with $\epsilon > 0$. TS_t^h's as defined by (4)—here we highlight the dependence on a chosen moment, m_t^l, but drop this for simplicity elsewhere—, and $q_0 = b_{K+1} = 0$. With these inequalities, we can find the set of λ_t's that correspond to a valid equilibrium. The same logic applies when using bounds on value functions (shown in Figure 4), just that the RHS of the above expressions are derived from bounds; see Appendix D.2 for details.

Identifying and estimating cost distributions. To learn about the cost distributions, G^h and G^d, we rely on Propositions 2 and 3, respectively.

We first compute bounds on how much customers and dealers expect to gain from participating in the game, i.e., the RHS of (9) and (8), respectively. For customers, we construct upper and lower bounds of the surplus (4) a customer (who submitted bid function i) expects to gain after entering auction t, $\overline{TS}^h_{it}$, and $\underline{TS}^h_{it}$. Visually, $\overline{TS}^h_{it}$ is the area between the upper bound on the value function and the bid function in Figure (4), where each quantity is weighted by the probability that it is won at market clearance. To approximate bounds on the surplus customers expect to gain before entering the auction, $\mathbb{E}_t[\overline{\boldsymbol{TS}}^h_{it}|N^d]$ and $\mathbb{E}_t[\underline{\boldsymbol{TS}}^h_{it}|N^d]$, we take the averages of $\overline{TS}^h_{it}$ and $\underline{TS}^h_{it}$ across customers i in auction t. Here we rely on the assumption that customers signals are drawn *iid* from the same distribution.

For dealers, constructing bounds on the expected annual auction surplus, given in (8), is more difficult, since dealers make their entry decision before knowing how many other dealers will participate in the upcoming year. This implies that we need to compute two additional objects besides the expected total surplus given the observed number of dealers (which we obtain as for customers).

First we get the expected total surpluses when different numbers of dealers participates. We could exactly compute these counterfactual surpluses using a similar approach presented in our counterfactual exercises. However, this is computationally intensive. Therefore, we approximate the counterfactual surpluses by finding auctions that are similar to an auction in question with respect to all other model primitives but with a different number of dealers participating. Within this set of auctions, we compute $\overline{TS}^d_{it}$ for all dealers and take an average in order to obtain a $\mathbb{E}_t[\overline{\boldsymbol{TS}}^{\boldsymbol{d}}_{\boldsymbol{it}}|N^d]$ for a fixed counterfactual N^d.

Second we compute how likely it is that N^d dealers participate when $\bar{N}^d$ is the maximal number of dealers from the empirical probability that a dealer participates in a fixed year: $\Pr(\boldsymbol{N^d} = N^d) = \binom{\bar{N}^d}{N^d}(N^d/\bar{N}^d)^{N^d}(1-(N^d/\bar{N}^d)^{\bar{N}^d-N^d}$.

With the bounds on how much dealers and customers expect to gain from auction participation, we identify bounds on the 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 probability for a customer of entering an auction, $\Pr(\boldsymbol{\gamma^h_{it}} \leq \mathbb{E}_t[\boldsymbol{TS^h_{it}}|N^d])$, must equal the observed entry probability (the number of customers who bid in an auction, N^h_t, relative to the number of customers who consider bidding, N^h).

Relying on the fact that

$$\Pr(\boldsymbol{\gamma}^{\boldsymbol{h}}_{\boldsymbol{it}} \leq \mathbb{E}_t[\underline{\boldsymbol{TS}}^{\boldsymbol{h}}_{\boldsymbol{it}}|N^d]) \leq \frac{N^h_t}{N^h} \leq \Pr(\boldsymbol{\gamma}^{\boldsymbol{h}}_{\boldsymbol{it}} \leq \mathbb{E}_t[\overline{\boldsymbol{TS}}^{\boldsymbol{h}}_{\boldsymbol{it}}|N^d]) \tag{10}$$

we could identify a lower and upper bound for the cost distribution non-parametrically as long as there are sufficiently many different surpluses, $\mathbb{E}_t[\overline{\boldsymbol{TS}}^{\boldsymbol{h}}_{\boldsymbol{it}}|N^d])$ and $\mathbb{E}_t[\underline{\boldsymbol{TS}}^{\boldsymbol{h}}_{\boldsymbol{it}}|N^d])$, from different auctions, to cover the full support of the cost distribution. With our data, we impose an exponential distribution with parameter β^h for customers, and β^d for dealers.[14]

6 Estimated values and costs

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

Dealer and customer values. We find that customers who participate in the auctions tend to have higher values than dealers, per average unit of a bond (see Figure 5). A customer is willing to pay about C$0.069-0.077 (or 6.9-7.7 bps) more than a dealer in the median (per unit of a bond with face-value C$100). This is true when using customer values at the lower bound and upper bound, given that bounds are relatively tight. For the rest of the paper, we rely on the lower-bounded values to be conservative.

A 6.9 bps difference in willingness to pay is economically meaningful compared to the median (average) market yield-to-maturity of the bonds, which is 161 bps (214 bps), and compared to secondary bid-ask spreads, which tend to be less than 1 bps (e.g., Allen and Wittwer (2022)). Roughly speaking, it implies that a customer is willing to pay C$9.8 million more per auction than a dealer. Further, we show in Appendix A.3 that the difference is statistically significant at the 5% level, where standard errors are computed via bootstrap.

[14] 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 $Q(\beta^h) = (\frac{N^h_t}{N^h} - H(\mathbb{E}_t[\underline{\boldsymbol{TS}}^h_{it}|N^d];\beta^h))_+)^2 + (\frac{N^h_t}{N^h} - H(\mathbb{E}_t[\overline{\boldsymbol{TS}}^h_{it}|N^d];\beta^h))_-)^2$, 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, and estimate the parameter set $\{\beta^h | Q'(\beta^h) \leq c_n\}$. We choose the cutoff c_n proportionally to the number of auctions in our sample $c_n = \log(645)/645$, inspired by Chernozhukov et al. (2007).

Since values reflect how much an auction participant expects to earn from trading bonds in the secondary market post-auction, this finding suggests that (participating) customers anticipate larger average returns (per unit) form buying and using bonds than dealers.[15] This is in line with evidence documented by Sandhu and Vala (2023), who argue that hedge funds 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 (Ontario Securities Commission (2007)).

Different to customers, dealers have been facing more and more stringent regulations since the global financial crisis in 2007-2009. In line with this fact, Figure 6 shows that dealer values have decreased over time relative to customer values, suggesting, that the average customer has become more profitable over time relative to the average dealer.[16] Below, we assess to what extent this trend explains customer entry.

Before doing so, we verify that higher customer values correlate with the observable factors that predict stronger customer participation to validate our value estimates (shown in Table **??**). For this we regress the average quantity-weighted customer value per auction on the same explanatory variables that we used in regression (1). Our findings, reported in Appendix Table A3, are consistent with the descriptive evidence.

Participation costs. Participation costs capture the opportunity cost of profits that customers and dealers could generate outside the Treasury market. These costs differ from values in that they are independent of how much an institution wins at auction. The cost must be payed when an institution spends time bidding at auction (and in the case of dealers, conduct other market making-activities), even if the institution doesn't buy any bonds.

On average, the annual cost of being a dealer is between C$3.198 and C$4.310 million with (bootstrapped) confidence interval [C$2.837M, C$4.551M]. This is sizable compared to the average annual profit a bank (at the company holding level) generates from its market-making activities in all financial markets combined, which is roughly C$413 million (Allen and Wittwer (2021)).

[15]Since we are considering value differences conditional on auction participation, this pattern cannot be driven by customer selection into more profitable auctions.

[16]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, it could come from a change in the composition of customer types over time.

Figure 5: Difference between dealer and customers values

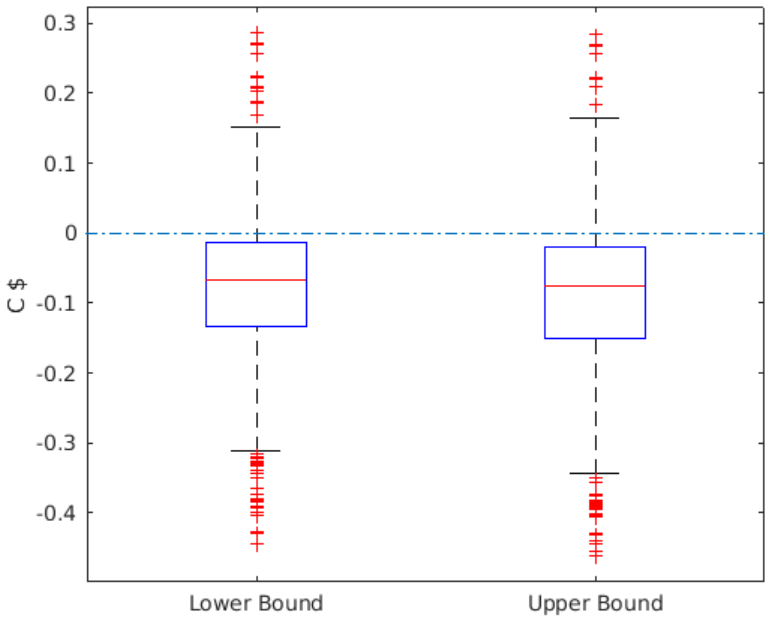


Figure 5 shows two box plots. The LHS shows the distribution of the difference in quantity-weighted average values between participating dealers and customers, using customer values at the lower bound. The RHS shows the analogous distribution, but uses the customer values at the upper bound. In both cases, we remove outliers, defined as a value that is more than three scaled median absolute deviations from the median. Prices are in C$ with a face value of 100.

The average entry cost of a customer is between C$429,500 and C$459,300 per auction with a (bootstrapped) confidence interval [C$402,000, C$490,710]. Given that there are about 30 auctions per year, the customer cost is larger than the dealer's cost (which is about C$106,600 - C$143,670 per auction). This 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 (participation) cost.

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

7 Drivers and consequences of customer participation

With the estimated model, we conduct counterfactuals to understand why customers, in particular, hedge funds, entered the market in 2014-2015, and evaluate consequences for market functioning. Ideally, we would compute counterfactual market outcomes for all auctions in

Figure 6: Time trend in value difference between dealers and customers

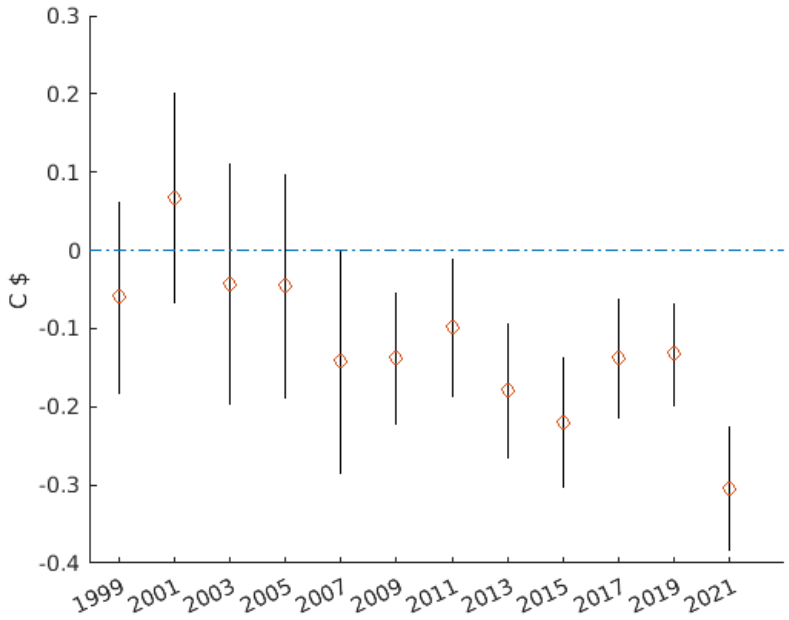


Figure 6 plots the point estimates and 95% confidence intervals from regressing the difference between the average dealer value and the average customer value in each auction on a set of time indicator variables. The first indicator is 1 for 1999-2000, the second indicator is 1 for 2001-2002, and so forth, until 2021-2022. The confidence intervals are computed using the point-estimates of dealer and customer values at the lower bounds, and therefore do not account for noise in the estimation of these values.

our sample. However, this is computationally intense. Therefore, we consider every 3rd auction for the past 9 years, i.e., over the period when hedge funds became the dominating customer group (recall Figure 2).

How to compute counterfactuals? Conducting counterfactuals (for multi-unit auctions in which bidders have multi-unit demand) is challenging because it is impossible to solve for an equilibrium analytically. We can only characterize necessary equilibrium conditions, which implies that it is difficult to compute counterfactual bids. Following Richert (2021, 2022), we adopt an empirical guess-and-verify approach.

We assume that all model primitives, such as the distributions of values and costs, remain unchanged when changing market rules and use the lower bound estimates of values and costs. For the value distribution, this means that we fix our estimated values for all auctions, bidders i and submitted steps, in addition to the position of the steps (i.e., the demanded quantities).

To change the auction environment, we want to take a stance on how much each bidder is

willing to pay for larger quantities than the last submitted step to account for the possibility that she might demand more when competing against fewer bidders. For this, we construct a distribution of maximal demand from the data. For each bidder, we determine the maximal quantity (in percentage of supply) that this bidder ever demanded. If this maximal demand is below the bidding limit, which is 25% of total supply, we use the observed maximal demand, else we take the bidding limit. Given these maximal quantities for each bidder, we obtain an empirical distribution of maximal quantities for dealers and customers, respectively.

In our main specification, we assume that the counterfactual distribution of maximal quantities within the auction is first order stochastically dominated by the empirical distribution of maximal quantities across all auctions, and solve for the distributions of bids under the counterfactual such that the distribution of values implied by these bids, and optimal bidding, are indistinguishable from the distribution of the estimated values associated with all submitted bids. We explain all details in Appendix C.

It is important to note that our approach flexibly accounts for local changes to the observed equilibrium, but might not fully capture the size of demand in counterfactual auctions that are far from any observed auction. We therefore focus our discussion on local changes where the 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 the expected price at which an auction clears under current market rules, relative to counterfactual market rules.

Why did customers enter? We first aim at understanding whether customers entered the market because of dealer exit or because of changes in market conditions. To answer this question, we compare the status quo, in which two dealers left, with a counterfactual in which we add back two dealers (fixing dealer participation but allowing customers to select into auctions).

Our findings suggest that customer participation was partially driven by dealer exit, and partially because changes in the market made it more profitable for customers to buy bonds (as the time-trend in Figure 6 suggested). Adding the two dealers reduces a customer's participation probability by roughly 15.47 percentage points, or 55.56%, on average. In

Figure 7: Why did customers enter?

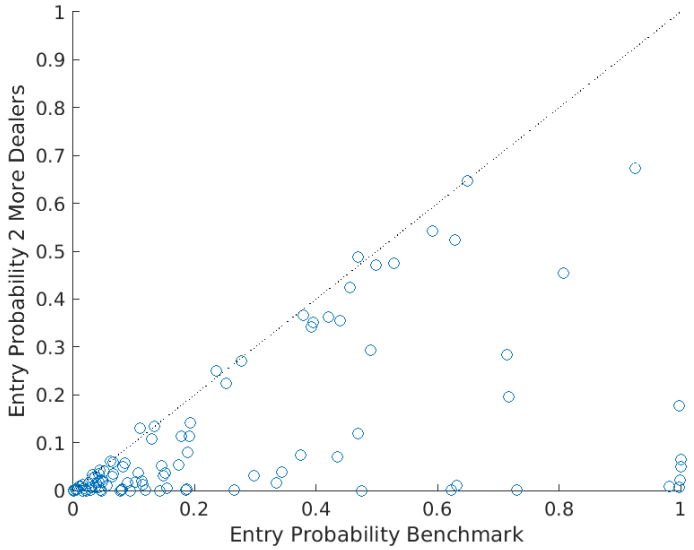


Figure A6 shows the probability for a customer to participate in every 3rd auction from 2014 onward in the status quo (on the x-axis) and the counterfactual in which we add back the two dealers who left (on the y-axis). Probability is in percentage points.

some auctions, the participation probability drops to almost zero, as shown in Figure 7.[17] This finding complements an empirical analysis by Favara et al. (2022) who argue that increasingly stringent regulation post-global financial crisis distorted dealer participation in the U.S. Treasury market.

Competition-volatility trade-off. Next, we illustrate the competition-volatility trade-off that arises when adding market participants who don't participate with regularity.

For this, we first consider a hypothetical auction environment with only dealers to eliminate effects coming from changes in the composition of bidders, or from dealer bid-updating. In this simplified auction environment, we ask by how much the expected auction price varies in the number of dealers who compete in the auction. This quantifies the increase in the degree of competition when adding an extra bidder in the spirit of Bulow and Klemperer (1996), and tells us by how much prices and auction coverage are expected to vary when dealer participation is volatile.

[17]Appendix Figure A6 shows by how much the expected revenue and expected price changes in the counterfactual relative to the status quo. Given that all auctions are relatively competitive given the observed number of bidders, as we explain next, these effects are relatively small.

Consider a typical auction, shown in Figure 8: If 14 dealers compete, the market clears at a competitive price, which is similar to the observed one. If only 13 or 12 dealers compete, the expected price drops by close to 10%, and 20% respectively.[18] Further reducing the number of dealers to 7, increases the likelihood of auction failure to close approximately 95%, leading to a dramatically low expected price of C$5, and a massive expected revenue loss of 34% or C$1.1 billion.

The expected price, and revenue drop for two reasons. First, with fewer bidders, the auction is less competitive, so that bidders shade their values more strongly. Second, auction failure, because total demand is insufficient, becomes more likely (as shown in Figure 8 (B)). To disentangle these two effects, we can compare compare these price drops to the price change that would result if bids were unchanged, but the same portion of auctions failed to clear. This price change is roughly equal to the portion of auctions that fail to clear (since values are close to 100). For example, if only 13 dealers compete, the failure probability is 5%, which means that the total price drop (of 10%) splits equally into the first effect coming from competition and the second effect coming from auction failure.

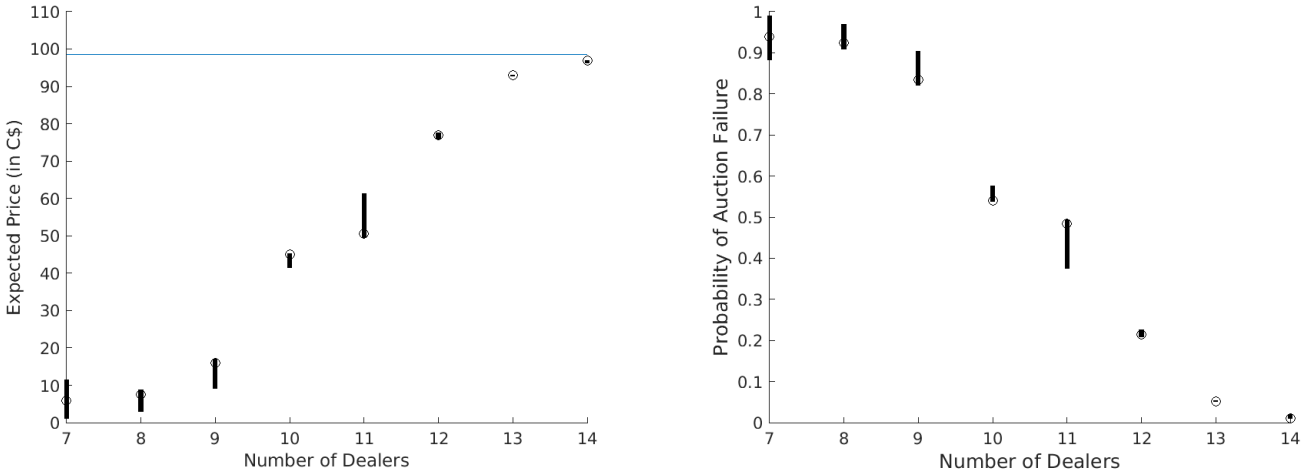
Next, we analyze the effect of varying the probability that a customer participates in the typical auction, fixing the observed number of dealers (see Figure 9 and Appendix Figure A8). While there is no risk of auction failure (in the typical auction), the expected price decreases by 25%, expected revenue decreases by 3% (or C$98 million) and bid-shading increases by 176%, when the customer does not participate relative to when she participates.

To quantify price volatility due to volatile dealer and customer participation, respectively, we eliminate auction heterogeneity and let each auction be a copy of the median auction. In this environment we ask by how much the expected price moves if each dealer (customer) randomly decided whether to participate in each auction, instead of committing to all of them. We estimate a standard deviation of expected prices of 17.3 with volatile participation of a dealer, and 3 for customers. Both numbers are large compared to the observed standard deviation of prices within a year (3.6 on average).

Taken together, these findings highlight the significant impact on market competition when bidders are removed and the potential risk of introducing substantial volatility in auction coverage and clearing prices when bidder participation becomes unstable. The results also indicate that losing more dealers could have detrimental effects unless an adequate

[18] Appendix Figure A7 shows the analogues results for other auctions.

Figure 8: Typical auction: Varying number of dealers



(A) Expected Price (B) Auction Failure

Figure 8 shows how the range of expected price (in C$) at which an auction clears varies as the number of dealers increases from 7 to 14 in an auction, which issues the average supply of all auctions with medium customer participation—by that we mean that 3 customers participate, where 3 is the medium number of participating customers in our sample. In theory, there is one expected price for each fixed number of dealers. In practice, we determine a range of prices (marked in black) 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) for details). The blue horizontal line shows the average observed bid which is close to the observed market clearing price.

number of new customers enter the market. This is not a hypothetical concern, as two of the dealers are currently contemplating leaving the market.

Policy. Finally, we assess the extent to which primary auctions run smoothly when allowing for market volatility, that is, without forcing dealers to commit to participate regularly. For this we compare a market in which customers and dealers directly place bids with the auctioneer but dealers are not obligated to participate, to a market in which both bid directly, but dealers have to commit ex ante to participate (like in the status quo). To compute the counterfactual without dealer commitment, we assume that a dealer's cost to enter one auction equals her 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 participate regularly (see Figure 10). However, both dealer and customer participation is highly volatile. Moreover, two out of one-hundred

Figure 9: Typical auction: Varying customer participation probability

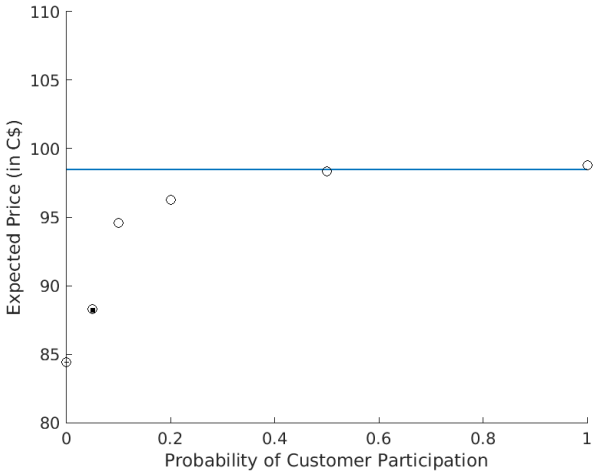


Figure 9 shows how the range of expected price (in C$) at which an auction clears varies as the customer participation probability varies between 0 and 1 in an auction, which issues the average supply of all auctions with medium customer participation—by that we mean that 3 customers participate, where 3 is the medium number of participating customers in our sample. In theory, there is one expected price for each counterfactual. In practice, we determine a range of prices (marked in black) 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) for details). The blue horizontal line shows the average observed bid which is close to the observed market clearing price.

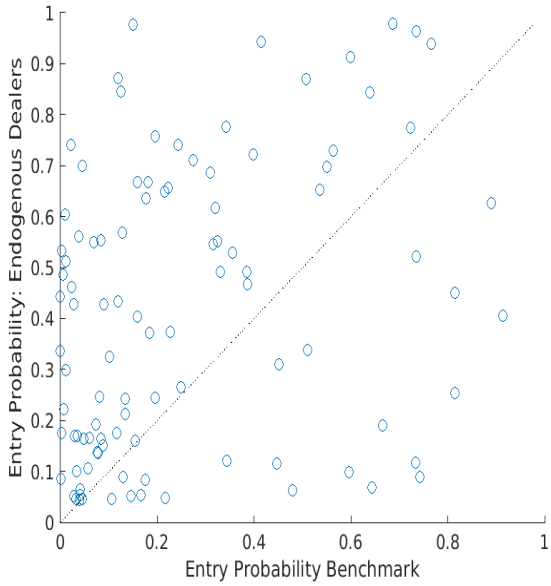
and one auctions risk to fail without dealer commitment. The expected price of these auctions at risk drops by more than 20%. Further, even in fully covered auctions, expected revenue decreases without dealer commitment, by 0.04% in the median.

Taken together, our counterfactual exercises help us understand why more customers entered the market, and evaluate the consequences for market functioning. Our findings indicate that it might be detrimental for the Canadian governments to loose more primary dealers who commit to regularly buy debt from the government.

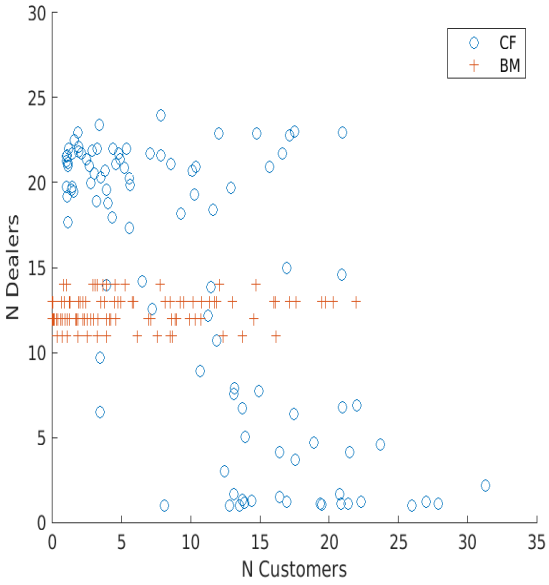
8 Conclusion

We analyze dealer exit and customer entry in the primary market for Canadian government debt, and analyze consequences for the functioning of the market. We show that hedge fund participation has increased in around 2014, but remains highly volatile. Then we introduce and estimate a structural model to trade-off the benefits of higher competition from customer entry with the costs of higher market volatility. Our framework could be used in

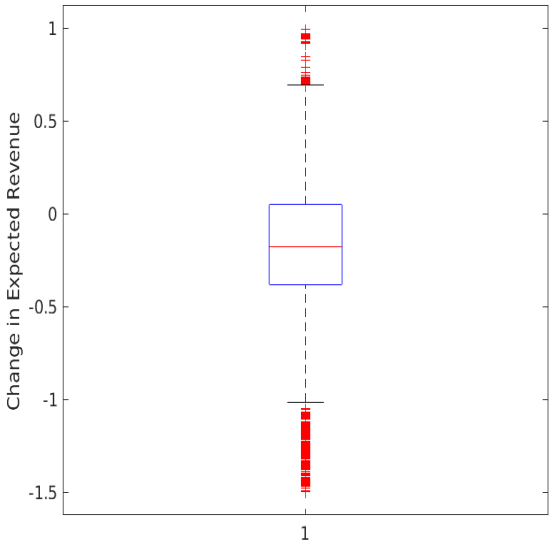
Figure 10: No dealer commitment



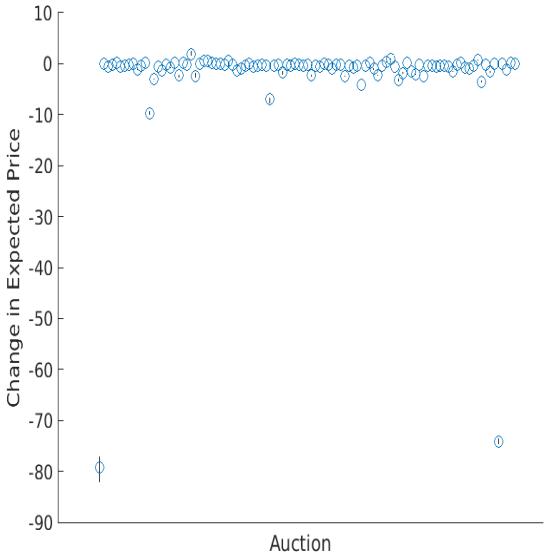
(A) Hedge Entry Probability (in %)



(B) Expected Number of Dealers/Customers



(C) Percentage Change in Expected Revenue



(D) Percentage Change in Expected Revenue

Figure 10A shows the probability for a customer to participate in every 3rd auction from 2014 onward 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). Probability is in percentage points. Figure 10B shows the expected number of dealers and customers that participate in each auction in the status quo and the counterfactual. Figure 10C shows the distribution of the percentage change in the expected auction revenue. Figure 10E a 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.

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.

Allen, J., Kastl, J., and Wittwer, M. (2020). Estimaging demand systems for treasuries. Bank of Canada working paper No. 2020-29.

Allen, J. and Wittwer, M. (2021). Centralizing over-the-counter markets? Forthcoming, Journal of Political Economy.

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

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

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*, page forthcoming.

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

Barth, D., Coen, J. K., and Coen, P. (2022). A structural model of liquidity in over-the-counter markets. Working paper.

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

Berger-Soucy, L., Garriott, C., and Usche, A. (2018). Government of canada fixed-income market ecology. Bank of Canada discussion paper 2018-10.

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*.

Chapman, J. T. E., McAdams, D., and Paarsch, H. J. (2007). Bounding revenue comparisons across multi-unit auction formats under ϵ-best response. *American Economic Review*, 97(2):455–458.

Chernozhukov, V., Hong, H., and Tamer, E. (2007). Exchange design and efficiency. *Econometrica*, 75:1243–1284.

Falato, A., Goldstein, I., and Hortaçsu, A. (2020). Financial fragility in the COVID-19 crisis: The case of investment funds in corporate bond markets. NBER working paper.

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

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.

Haddad, V., Moreira, A., and Muir, T. (2020). When selling becomes viral: Disruptions in debt markets in the COVID-19 crisis and the Fed's response. NBER working paper.

He, Z., Nagel, S., and Song, Z. (2022). Treasury inconvenience yields during the COVID-19 crisis. *Journal of Financial Economics*, 143(1):57–79.

Hendricks, K., Jean-Franccois, 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.

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

Kargar, M., Lester, B., Lindsay, D., Liu, S., Weill, P.-O., and Zuniga, D. (2021). Corporate bond liquidity during the covid-19 crisis. Working paper.

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.

Kastl, J. (2017). Recent advances in empirical analysis of financial markets: Industrial organization meets finance. In Honoré, B., Pakes, A., Piazzesi, M., and Samuelson, L., editors, *Advances in Economics and Econometrics: Eleventh World Congress (Econometric Society Monographs)*, pages 231–270. Cambridge: Cambridge University Press.

Kondor, P. and Pintér, G. (2021). Clients' connections: Measuring the role of private information in decentralized markets. *Journal of Finance*, forthcoming.

Kutai, A., Nathan, D., and Wittwer, M. (2022). Exchanges for government bonds? Evidence during COVID-19. Working paper.

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.

Ma, Y., Xiao, K., and Zeng, Y. (2022). Mutual fund liquidity transformation and reverse flight to liquidity. *Review of Financial Studies*, 35(10):4674–4711.

Muller, P. (2019). Primary dealer system and market making in Canada. Bank of Canada presentation.

Ontario Securities Commission (2007). Hedge funds. CSA Staff Notice 81-316.

Pellerin, M. (2006). The evolution of the government of Canada's debt distribution framework. Bank of Canada Review.

Petersen, C., Reguant, M., and Segura, L. (2022). Measuring the impact of wind power and intermittency. Working paper.

Pintér, G. and Semih, Ü. (2022). Comparing search and intermediation frictions across markets. Working paper.

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.

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

ONLINE APPENDIX

Entry and Exit in Treasury Auctions

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

Appendix **A** presents additional empirical test.

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

Appendix **C** provides details about how to compute counterfactuals.

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

A Empirical tests

A.1 Testing independent private values

We perform the test introduced in Hortaçsu and Kastl (2012) for the auctions of bonds in our sample and obtain a p-value of 0.384, which does not provide evidence of common values. The test checks for equality of the estimated marginal values of dealers 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 marginal values for acquiring the bond being sold. Under independent private values, the customers bid reveals information about the expected level of competition, but should not affect their marginal value. Similar to the findings of Hortaçsu and Kastl (2012) for bills, in our sample we fail to find evidence that dealer marginal values are shifted by the information learned through the customers bids in the bond market.

A.2 Testing bounds

Here we test whether customers take dealer updating into account. Formally, we want to know if $\lambda = 0$. To do this, we construct measures $T_i = |v_i(q, s_i; \lambda = 0) - v_i(q, s_i)|$ and construct a test statistic analagous to SSQ_T from Hortaçsu and Kastl (2012). The test rejects the null hypothesis with p-value of 0.00, where the p-value is computed using the bootstrap.

A.3 Testing value differences

Here 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. Rejecting the null would imply that customer values at the lower bound are strictly greater than dealer values at the upper bound.

We consider four different tests of this hypothesis. First, we compute three aggregate test statistics following Hortaçsu and Kastl (2012), but because we are interested in a one sided null hypothesis (are customer values larger than dealer values), we drop the absolute value. The first is in the spirit of a chi-squared test, the second based on the 95th percentile of across auction differences and the third based on the max difference in values across auctions. In all cases we omit the subset of auctions where no customers showed up. We compute these test statistics for the differences in average quantity weighted value, the average max value, and the average min value of dealers/customers. We also provide confidence intervals for a set estimate of the mean difference. Results are reported in Table A2.

B Resampling procedure to estimate values

To back out 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, $\boldsymbol{P_t^*}$.

To build an intuition for how to estimate the price distribution, consider a simpler setting with N ex-ante symmetric bidders. Here, we fix a bidder in an auction, and draw $N-1$ bid functions with replacement from all observed bids in that auction.[19] This simulates one possible market outcome for our fixed bidder. Repeating this many times, we obtain the distribution of the market clearing price $\boldsymbol{P_t^*}$ for this bidder.

Our setting is more complicated because there are dealers and customers, and customers must bid via dealers. Hortaçsu and Kastl (2012) introduce a resampling procedure to back out dealer values for this setting. We extend their method in two ways. First, we allow signals within a bidder to be correlated over the course of an auction to avoid estimation

[19] In order to allow the auctions to differ in some characteristics which are unobserved to the econometrician but are observed by bidders we focus on resampling schemes that use the empirical distribution of bids within a given auction.

bias coming from the fact that we observe some bidders updating their bids without observing a bid of a customer. For simplicity, our model does not rationalize such updates, but an extended model based on Hortaçsu and Kastl (2012) could. Second, we obtain all elements in the customer's FOC needed to bound customer values.

To simulate a market outcome in our setting, we first construct the residual supply curve that a bidder i faces in auction t. For this, we start by randomly drawing a customer bid from the set of N^h potential customer bids. If the customer did not participate in the auction, her bid is 0; if she updated her bid, we randomly select one of her bids. Next, we find a dealer that saw a similar bid to the bid of that customer.[20] If the selected dealer made multiple bids, we select at random one of the bids 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 her 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 N^h times if i is a dealer and $N^h - 1$ times if i is a customer.

Next, we resample the uninformed dealers. Starting with a list of uninformed dealers, draw one such dealer.[21] If they submitted more than one bid, randomly select one bid and drop the others. Continue drawing from the set of uninformed dealers so that there are $N^h + N^d - 1$ included bidding curves. This is one realization of the residual supply curve that i expects to face. 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 the market clearing at each step is strictly bounded away from zero. However, in our finite sample this event may occur, leading the estimated marginal value to explode. For steps with win probabilities close to or equal to zero we therefore mix the estimated clearing price distribution with a uniform distribution over the range of the bids made in order to give all the bids a small win probability. In addition to reducing the sensitivity of the analysis to these small probabilities, we truncate the estimated marginal values by assuming that

[20] Ideally we would choose a dealer that saw the exact same bid, however given the limited sample size this event is extremely unlikely. To reflect the uncertainty that the customer faces about the value of the dealer observing their bid, we set a bandwidth and define similar bids using the quantity weighted average bids.

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

marginal values are below the maximum bid ever made in the year of the auction for a bond with the same maturity as the bond being sold plus C$1,000/months to maturity, which is roughly equivalent to 10 basis points.

With the estimated price distributions, we can solve for the value that rationalizes a dealer's bid at each step, using condition (5). To obtain customer value bounds, we implement an estimation procedure that follows our identification argument; for details see Appendix D.2.

C Details regarding the counterfactual

We construct a criterion function which allows us to compare the implied and true value distributions along a number of dimensions. We are interested in the set of bid distributions for which the value of this criterion is close enough to zero that remaining violations are likely due to errors from the number of simulations used in estimation of the implied value distribution. The criterion involves three main components.

First, we evaluate the marginal distributions of 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% s of the total supply. At each of these quantity levels we construct bounds on the marginal distribution of values in auction j for bidder type g along a grid of evenly spaced points running from the value at the 5^{th} percentile in auction j for a bidder of type g to the value at 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 marginal distribution of marginal values at each of the quantity levels evaluated on the grid points and F_U the corresponding upper bound on the marginal distribution known from the data.

The second component in the criterion measures the distribution of slopes across the same fixed grid of quantity levels. Given the bounds on bidders' values across these points we calculate the largest and smallest slope that fit within these bounds (or minimize the violations of the bounds). We then compare the marginal distribution of the slopes, along an evenly spaced grid of slopes again adding to the criterion function $max(F_L^{IM} - F_U, 0)^2$ and $min(F_U^{IM} - F_L, 0)^2$ where F now represents the marginal distribution of the slopes (and the notation for bounds, and implied values is as above).

Third, we want the distribution across bidders of type g of the largest quantity bid on in auction j to be smaller than the distribution of the largest quantities (within a bidder) ever purchased. The restriction is designed to capture the fact that not all bidders have the capacity to purchase up to the 25% limit and even in a counterfactual with lower prices may not be interested in purchasing up to this limit. Therefore we require the distribution of largest quantity bid by each bidder of type g in auction j be first order stochastically dominated by the distribution of max (across auctions) quantities ever bid. Evaluating this along a set of grid points (evenly spaced in quantity-space from 0 to 25%), this results in an additional contribution to the criterion function of 0 if the counterfactual distribution of max quantities at a given grid point is above the max-quantity distribution in the data and the squared difference of the probabilities of quantities less than that grid point when the counterfactual distribution is below the distribution from the data.

In an alternative specification we instead use in-sample variation in the average (across bidders) of the within-bidder maximum quantity submitted for both bidder groups to estimate how the average max quantity varies as a function of the number of customers, number of dealers, and total supply being auctioned. Using this we then generate predictions (both in and out of sample) for the average max-quantity submitted in each auction j. In the counterfactuals with this restriction we require that the average max quantity submitted be smaller than this number and add to the criterion function the squared difference of any violations where the counterfactual average max-quantity exceeds the predicted level.

Finally, recall that we specified the customers' optimization problem using a two-step optimization problem with the exterior optimization over moments represented in the bidders problem via a constraint with multiplier λ. Since these λ will change in the counterfactual world, we don't have a primitive distribution to match for this parameter and it appears as though this introduces additional degree of freedom which might make informative counterfactuals impossible. However, fixing any proposal distribution the λ are not free parameters, but are pinned down by the FOC (in particular, by the condition for optimality of the moment chosen). Therefore, given a bid distribution, the λ cannot be changed to map into different value distributions. In practice both the λ and the marginal values of a bidder are set-identified, which will lead to some information loss. This loss will result in wider counterfactual bounds in the empirical analysis.

D Mathematical appendix

D.1 Auxiliary statements and proofs

Lemma 1. *The moment function, which computes the quantity weighted bid,*

$$\mu^l(b^h_{i\tau}(q,\theta^h_{i\tau})) = \frac{\sum_k^K b_k q_k}{\sum_k q_k} \text{ with } \{b_k,q_k\}_{k=1}^K \text{ constituting bidding function } b^h_{i\tau}(\cdot,\theta^h_{i\tau}) \tag{3}$$

is differentiable w.r.t. q.

Proof of Lemma 1. We want to show that $\mu^l(b^h_{i\tau}(q,\theta^h_{i\tau}))$ is differentiable at any q, meaning that $\lim_{h\to 0}\frac{\mu^l(b^h_{i\tau}(q+h,\theta^h_{i\tau}))-\mu^l(b^h_{i\tau}(q,\theta^h_{i\tau}))}{h}$ exists. Given (3) there are K different cases, depending on which q we consider. The first case is:

$$\lim_{h\to 0}\frac{\frac{b_1(q_1+h)+\sum_{k=2}^K b_k q_k}{q_1+h+\sum_{k=2}^K q_k} - \frac{\sum_k^K b_k q_k}{\sum_k q_k}}{h} = \frac{b_1\sum_{k=2}^K q_k - \sum_{k=2}^K b_k q_k}{(\sum_{k=1}^K q_k)^2} \geq 0 \text{ since } b_1 \geq b_2 \geq b_3.$$

The limit exists. An analogous argument applies to all other cases, only that the limit becomes negative for steps towards the last step. At the last step, the limit is non-positive. Therefore, the moment function is differentiable w.r.t. q. □

Proof of Proposition 1. The proof of statement (i) is analogous to the proof of Kastl (2011)'s Proposition 1. To show statement (ii), take the perspective of customer i and assume all other bidders play an equilibrium. Fix L moment functions μ^l. Given that the customer thinks that the dealer only updates its own function when the customer submits a function with at least one different moment $\{m^l\}_{l=1}^L$, we can solve a two step problem to characterize the necessary condition for the customer's best reply—which is an equilibrium bid function given that all other bidders play an equilibrium strategy by assumption.

First, we fix some set of moments $\{m^l\}_{l=1}^L$, and find $b^c_{i\tau}(\cdot,\theta^h_{i\tau})$ that maximize her expected total surplus $TS^h(b^c_{i\tau}(\cdot,\theta^h_{i\tau}),\{m^l\}_{l=1}^L)$ such that $\mu^l(b^c_{i\tau}(\cdot,\theta^h_{i\tau})) = m^l$ for all l. The objective function is:

$$TS^h(b^h_{i\tau}(\cdot,\theta^h_{i\tau}),\{m^l\}_{l=1}^L) - \sum_{l=1}^L \lambda_l\left(m^l - \mu^l(b^h_{i\tau}(\cdot,\theta^h_{i\tau})\right) \tag{11}$$

with $\{b_k, q_k\}_{k=1}^K$ making up bidding function $b_{i\tau}^h(\cdot, \theta_{i\tau}^h)$. Given this objective function, we can follow Kastl (2011)'s perturbation argument in the original proof step by step (shown in his Appendix A.2.). There are two differences. First, there is an additional term that comes from the constraints that all moments l of chosen function, $\mu^l(b_{i\tau}^h(\cdot, \theta_{i\tau}^h)$ must equal the fixed moments m^l. However, since all moment functions are differentiable w.r.t. q by assumption, no complication arises from this extra term. Second, with dealer updating it can be more difficult to predict where the market will clear when the customer demands a different amount than q. However, with the assumption that the customer thinks that dealers only update their own bids when a moment of the customer's bidding function changes, Kastl (2011)'s original arguments go through and we obtain equation A2, A3 and the additional term, $\sum -\lambda \frac{\partial m u^l(b_{i\tau}^h(\cdot, \theta_{i\tau}^h)}{\partial q_k}$. Note that unlike in Kastl (2011), these expressions do not simplify because it may be optimal for a bidder to tie. A tie allows the bidder to request additional quantity in order to move their moments, without necessarily winning all the quantity bid. A necessary condition for optimality is that the sum of these three components is zero.

The above characterizes an equilibrium bidding function for a fixed set of moments $\{m^l\}_{l=1}^L$. For our estimation these conditions are sufficient. The reason is that we assume that the observed bids are equilibrium bids and that we observe not only the bids but also the moments that these bid functions generate (for fixed moment functions).

However, for completeness of the argument, in equilibrium, the customer chooses the moments $\{m^l\}_{l=1}^L$ so that $b_{i\tau}^c(\cdot, \theta_{i\tau}^h)$ actually maximizes the total surplus:

$$\max_{\{m^l\}_{l=1}^L} TS^h(b_{i\tau}^h(\cdot, \theta_{i\tau}^h), \{m^l\}_{l=1}^L) \text{ subject to: } \mu^l(b_{i\tau}^h(\cdot, \theta_{i\tau}^h)) = m^l \text{ for all } l. \tag{12}$$

Solving this maximization problem is challenging (even when we restrict attention to moments that are real numbers) because the objective function is not differentiable. To see this consider a change in in moment m^l. The dealer who observes the corresponding bid function updates her own bid because a change in m^l (weakly) changes the dealer's information set and with that its type $\theta_{i\tau}^d$. Therefore, the dealer submits a different bid function, i.e., a step function with steps at different points. This changes the customer's believes about the price at which the auction will clear. Formally, the distribution of the clearing price, when fixing the customer's own bid function, changes, and since biding functions are step functions, a change in such a function easily leads to non-continuous jumps that render the objective

function TS^h non-differentiable.

Lemma 2. *(i) For a customer s_i^h: Ties occur with zero probability for a.e. s_i^h in any equilibrium, for either all steps except the last step, or at the last step. (ii) For a dealer s_i^d, Kastl (2011) Lemma 1 applies.*

Proof of Lemma 2. Case 1: $\lambda > 0$. Consider a step $k = K_i$. Suppose bidder i ties on a step $k = K_i$. Take some $q_m = sup\{q|v(q, s_i) \leq b_k\}$ and let $\bar{q} = max\{q_k - \delta, q_m\}$ with δ some strictly positive step size bounded above by $q_k - q_{k-1}$. ie. you step either to where you get positive surplus from the amount you purchase if possible or if not at least you buy less negative. Take the deviation to bid $b'_{ik} = b_{ik} + \epsilon$ where ϵ is sufficiently small, at $\bar{q}$ (and the associated step). Bidder i gets less units than they would in a tie and each one of these they longer have to pay for. There is no loss to them from not winning these units.

Then the constraint part gives $-\lambda(m^l - \mu(\bar{q}))$. Because $\mu(\bar{q}, b') > \mu(q_k, b) = m^l$ the bracketed term is negative, λ is positive (by assumption) and there is one more negative sign so this term is positive overall. This represents an additional strictly positive profit from the deviation.

Case 2: $\lambda < 0$. Suppose bidder i ties on a step $k < K_i$. Take some $q_m = sup\{q|v(q, s_i) \leq b_k\}$ and let $\bar{q} = max\{q_k - \delta, q_m\}$ with δ some strictly positive step size bounded above by $q_k - q_{k-1}$. ie. you step either to where you get positive surplus from the amount you purchase if possible or if not at least you buy less negative. Take the deviation to bid $b'_{ik} = b_{ik} + \epsilon$ where ϵ is sufficiently small, at $\bar{q}$ (and the associated step). Bidder i gets less units than they would in a tie and each one of these they longer have to pay for. There is no loss to them from not winning these units.

Then the constraint part gives $-\lambda(m^l - \mu(\bar{q}))$. Because $\mu(\bar{q}, b') < \mu(q_k, b) = m^l$ the bracketed term is positive, λ is negative (by assumption) and there is one more negative sign so this term is positive overall. This represents an additional strictly positive profit from the deviation.

Case 3: $\lambda = 0$. This is the world of Kastl (2011). Follows from that proof. □

D.2 Mathematical details on bounding customer values

In this section, we provide formal details on how to identify bounds on the customer's value distribution from equilibrium condition (6) of Proposition 1. For this, we fix an auction t and

a customer i, and drop the auction t, time τ, customer h, bidder i subscripts for simplicity. Further, we assume that only one moment, m^l, matters, for instance, the quantity-weighted bid, consistent with our estimation, and drop the l-subscript.

The identification argument is complicated by the fact that condition (6)—here slightly rearranged—not only contains marginal values at submitted steps, $v(q_k, s)$, but also values at some intermediate quantities between submitted steps, $v(\boldsymbol{q^*}, s)$:

$$
\begin{aligned}
0 = &\Pr(b_k > \boldsymbol{P^*} > b_{k+1}|\theta, m)v(q, s)+ \\
&\Pr(b_k = \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[v(\boldsymbol{q^*}, s)\frac{\partial \boldsymbol{q^*})}{\partial q_k}|b_k = \boldsymbol{P^*}, \theta, m\right] + \\
&\Pr(b_{k+1} \geq \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[v(\boldsymbol{q^*}, s)\frac{\partial \boldsymbol{q^*}}{\partial q_k}|b_{k+1} \geq \boldsymbol{P^*}, \theta, m\right] - \\
&\Pr(b_k > \boldsymbol{P^*} > b_{k+1}|\theta, m)(b_k)- \\
&\Pr(b_{k+1} \geq \boldsymbol{P^*}|\theta, m)(b_k - b_{k+1})- \\
&\Pr(b_k = \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_k}|b_k = \boldsymbol{P^*}, \theta, m\right] - \\
&\Pr(b_{k+1} = \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_k}|b_{k+1} = \boldsymbol{P^*}, \theta, m]\right] - \\
&\Pr(b_{k+1} < \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_k}|b_{k+1} < \boldsymbol{P^*}, \theta, m\right] + \lambda\frac{\partial \mu(b(\cdot, \theta))}{\partial q_k}.
\end{aligned} \tag{13}
$$

However, we can still obtain set-identification of these marginal values, by using the monotonicity and boundedness of the step function. That is, for all $q_{k-1} \leq q \leq q_k$ we know that $v(q_{k-1}, s) \geq v(q, s) \geq v(q_k, s)$. We also are able to sign the derivatives of the rationed quantity: 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 $k + 1$.

We obtain an upper bound on the marginal value at q_k by making the terms involving the marginal 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 max possible value at intermediate quantities along the next step for terms involving rationed quantities at the next step (e.g. line 3 of equation (13)) and $\underline{v}(q_k, s)$ 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 (13)). 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 (13)) and

the minimum value at the next step, $\underline{v}(q_{k+1}, s)$ (in line 3 of equation (13)). By substituting these values into the first order conditions from each step, we are left with a system of $2K$ linear equations with $2K$ unknowns to solve.

We can make use of some additional information to simplify the resulting system, by adding zeros into the matrix of coefficients for the linear system. In particular, at the last step, all terms involving b_K drop out. When the current λ is negative, at 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. To see why, suppose that a bidder wanted to tie with positive probability. If the step is above the last, raising the quantity increases the quantity-weighted average bid while at the last step, it decreases it. Lemma 2 formalizes this idea and summarizes all possible simplifications.

With these simplifications, we obtain the following system of equations for two cases $\lambda > 0$, and $\lambda < 0$. Below, we express each equation to guarantee readability. It can help to transform the system of equation into a single matrix, one for each case, to get a better sense for whether the system is identified.

(i) When $\lambda < 0$:

$$
\begin{aligned}
A_1 &= \Pr(b_1 > \boldsymbol{P^*} > b_2|\theta, m)\overline{v}(q_1, s) \\
A_1 &= \Pr(b_1 > \boldsymbol{P^*} > b_2|\theta, m)\underline{v}(q_1, s) \\
&... \\
A_{K-1} &= \Pr(b_{K-1} > \boldsymbol{P^*} > b_K|\theta, m)\overline{v}(q_{K-1}, s) + \Pr(b_K \geq \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_{K-1}}|b_K \geq \boldsymbol{P^*}, \theta, m\right]\overline{v}(q_{K-1}, s) \\
A_{K-1} &= \Pr(b_{K-1} > \boldsymbol{P^*} > b_K|\theta, m)\underline{v}(q_{K-1}, s) + \Pr(b_K \geq \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_{K-1}}|b_K \geq \boldsymbol{P^*}, \theta, m\right]\underline{v}(q_K, s) \\
A_K &= \Pr(b_K > \boldsymbol{P^*} > 0|\theta, m)\overline{v}(q_K, s) + \Pr(b_K = \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_K}|b_K = \boldsymbol{P^*}, \theta, m\right]\underline{v}(q_K, s) \\
A_K &= \Pr(b_K = \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_K}|b_K = \boldsymbol{P^*}, \theta, m\right]\overline{v}(q_{K-1}, s) + \Pr(b_K > \boldsymbol{P^*} > 0|\theta, m)\underline{v}(q_K, s),
\end{aligned}
$$

where

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

for all $k = 1, ..., K_1$, and

$$A_k = \Pr(\boldsymbol{P^*} > b_K|\theta, m)(b_K) + \Pr(b_K = \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_K}|b_K = \boldsymbol{P^*}, \theta, m\right] - \lambda\frac{\partial \mu(b(\cdot, \theta))}{\partial q_K}$$

at the last step K.

(ii) When $\lambda > 0$:

$$\begin{aligned}
B_1 &= \Pr(b_1 > \boldsymbol{P^*} > b_2|\theta, m)\overline{v}(q_1, s) + \Pr(b_2 \geq \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_1}|b_1 = \boldsymbol{P^*}, \theta, m\right]\underline{v}(q_1, s)\\
B_1 &= \Pr(b_1 > \boldsymbol{P^*} > b_2|\theta, m)\underline{v}(q_1, s) + \Pr(b_2 \geq \boldsymbol{P}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_1}|b_2 \geq \boldsymbol{P^*}, \theta, m\right]\underline{v}(q_2, s)\\
B_2 &= \Pr(b_2 > \boldsymbol{P^*} > b_3|\theta, m)\overline{v}(q_2, s) + \Pr(b_3 \geq \boldsymbol{P^*}|\theta, m)\mathbb{E}[\frac{\partial \boldsymbol{q^*}}{\partial q_2}|b_3 \geq \boldsymbol{P^*}, \theta, m]\overline{v}(q_2, s)\\
B_2 &= \Pr(b_3 = \boldsymbol{P^*}|\theta, m)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_2}|b_2 = \boldsymbol{P^*}, \theta, m\right]\overline{v}(q_1, s) + \Pr(b_2 > \boldsymbol{P^*} > b_3|\theta, m)\underline{v}(q_2, s)\\
&\quad + \Pr(b_3 \geq \boldsymbol{P}|\theta, m)\mathbb{E}[\frac{\partial \boldsymbol{q^*}}{\partial q_2}|b_3 \geq \boldsymbol{P^*}, \theta, m]\overline{v}(q_4, s)\\
&...\\
B_{K-1} &= \Pr(b_{K-1} > \boldsymbol{P^*} > b_K)\overline{v}(q_{K-1}, s) + \Pr(b_{K-1} = \boldsymbol{P^*})\underline{v}(q_{K-1}, s)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_{K-1}}|b_{K-1} = \boldsymbol{P^*}, \theta, m\right]\\
B_{K-1} &= \Pr(b_{K-1} = \boldsymbol{P^*})\overline{v}(q_{K-1}, s)\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_{K-1}}|b_{K-1} = \boldsymbol{P^*}, \theta, m\right] + \Pr(b_{K-1} > \boldsymbol{P^*} > b_K)\underline{v}(q_{K-1}, s)\\
B_K &= \Pr(b_K > \boldsymbol{P^*} > 0)\overline{v}(q_K, s)\\
B_K &= \Pr(b_K > \boldsymbol{P^*} > 0)\underline{v}(q_K, s)
\end{aligned}$$

were

$$\begin{aligned}
B_1 &= A_1 - \overline{v}(q_1, s)\Pr(b_1 = \boldsymbol{P^*}|\theta, m))\mathbb{E}\left[\frac{\partial \boldsymbol{q^*}}{\partial q_1}|b_1 = \boldsymbol{P^*}, \theta, m\right]\\
B_k &= A_k \text{ for } k = 2, ...K
\end{aligned}$$

In either case the system of equations is still underidentified as λ, which appears in the terms A_k and B_k and determines which set of equations apply, is unknown. However, in addition to this system of equations we know that any perturbation in m must not be optimal: we consider two such possible perturbations to $m + \epsilon$ and $m - \epsilon$ respectively and argue this can be used to construct bounds on λ.

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

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

where

$$TS(b(\cdot,\theta),m) = \sum_{k=1}^{K}\Big[\Pr(b_k > \boldsymbol{P^*} > b_{k+1}|\theta,m)V(q_k,s) - \Pr(b_k > \boldsymbol{P^*}|\theta,m)b_k(q_k - q_{k-1})\Big] + \sum_{k=1}^{K}\Pr(b_k = \boldsymbol{P^*}|\theta,m)\mathbb{E}\Big[V(\boldsymbol{q^*},s) - b_k(\boldsymbol{q^*} - q_{k-1})|b_k = \boldsymbol{P^*},\theta,m\Big]$$

with $q_0 = b_{K+1} = 0$. $TS(b(\cdot,\theta),m+\epsilon)$ and $TS(b(\cdot,\theta),m-\epsilon)$ are defined analogously.

We can use (17) to find bounds for (14) and (15). Consider (14) first, and omit conditioning on θ for simplicity. We obtain:

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

where $\Delta\Pr(\cdot)$ indicates taking a difference between $\Pr(\cdot|..,m)$ and $\Pr(\cdot|..,m+\epsilon)$ and $v(q_{k-1},s) = \overline{v}(q_1,s)$ for k=1.

Similarly for (17):

$$\sum_{k=1}^{K}[\max(0,\Delta\Pr(b_k \geq \boldsymbol{P}^{\star} \geq b_{k+1})(\sum_{j=1}^{k}\underline{v}(q_j,s)(q_j-q_{j-1}))+ \tag{17}$$

$$\min(0,\Delta\Pr(b_k \geq \boldsymbol{P}^{\star} \geq b_{k+1})(\sum_{j=1}^{k}\overline{v}(q_{j-1},s)(q_j-q_{j-1})) + \Delta\Pr(b_k > \boldsymbol{P}^{\star})b_k(q_k-q_{k-1})]+$$

$$\sum_{k=1}^{K}[\max(0,\Delta\Pr(b_k = \boldsymbol{P}^{\star}))(\sum_{j=1}^{k-1}(\underline{v}(q_j,s)(q_j-q_{j-1}))) + \min(0,\Delta\Pr(b_k = \boldsymbol{P}^{\star}))(\sum_{j=1}^{k-1}(\overline{v}(q_{j-1},s)(q_j-q_{j-1})))$$

$$+ [\Pr(b_k = \boldsymbol{P}^{\star}|m-\epsilon)\underline{v}(q_j)(\mathbb{E}[\boldsymbol{q}^{\star}|b_k = \boldsymbol{P}^{\star}, m-\epsilon] - q_{k-1})-$$

$$\Pr(b_k = \boldsymbol{P}^{\star}|m)\underline{v}(q_j)(\mathbb{E}[\boldsymbol{q}^{\star}|b_k = \boldsymbol{P}^{\star}, m] - q_{k-1})]$$

$$1(\Pr(b_k = \boldsymbol{P}^{\star}|m-\epsilon)(\mathbb{E}[\boldsymbol{q}^{\star}|b_k = \boldsymbol{P}^{\star}, m-\epsilon] - q_{k-1}) - \Pr(b_k = \boldsymbol{P}^{\star}|m)\mathbb{E}[\boldsymbol{q}^{\star}|b_k = \boldsymbol{P}^{\star}, m] - q_{k-1}) > 0)$$

$$+ [\Pr(b_k = \boldsymbol{P}^{\star}|m-\epsilon)\overline{v}(q_{j-1})(\mathbb{E}[\boldsymbol{q}^{\star}|b_k = \boldsymbol{P}^{\star}, m-\epsilon] - q_{k-1})-$$

$$\Pr(b_k = \boldsymbol{P}^{\star}|m)\overline{v}(q_{j-1})(\mathbb{E}[\boldsymbol{q}^{\star}|b_k = \boldsymbol{P}^{\star}, m] - q_{k-1})]$$

$$1(\Pr(b_k = \boldsymbol{P}^{\star}|m)(\mathbb{E}[\boldsymbol{q}^{\star}|b_k = \boldsymbol{P}^{\star}, m-\epsilon] - q_{k-1}) - \Pr(b_k = \boldsymbol{P}^{\star}|m+\epsilon)\mathbb{E}[\boldsymbol{q}^{\star}|b_k = \boldsymbol{P}^{\star}, m] - q_{k-1}) < 0)$$

$$- b_k(\mathbb{E}[\boldsymbol{q}^{\star} - q_k|b_k = \boldsymbol{P}^{\star}, m-\epsilon] - \mathbb{E}[\boldsymbol{q}^{\star} - q_k|b_k = \boldsymbol{P}^{\star}, m])] \leq \lambda\epsilon$$

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 bidder i.

We implement these ideas via the following procedure to construct the identified set of $v(q,s)$. To begin, we search over the (one-dimensional) set of λ. For each feasible λ, there is a unique set of lower and upper bounds for v at each q_k where that customer submitted a step that satisfy equation (6). Using these implied v, together with the definition of the total surplus function allows us to obtain upper and lower bounds on the λ based on (17). This range of (λ_L, λ_U) are the set of feasible λ which are consistent with the observed choices and values. To obtain the maximum and minimum we find the marginal value at each point that maximizes (minimizes) the value of the change in the total surplus function. Whether this is the upper or lower envelope of the set of marginal values consistent with the observed bids, depends only on the sign of the change in clearing probabilities under the increased m. If the initial λ is in the range (λ_L, λ_U), than the associated marginal value curve is part of the identified set of values that can rationalize the behavior of the given customer. When it is outside the range, the bid is not consistent with equilibrium behavior for that set of values. To trace out the identified set, we repeat this exercise along a grid of possible λ.

Appendix Table A1: List of dealers that exited and entered auctions, and when

Dealer name	Bill auctions	Bond auctions
CT Securities	Entry 05jan1999, exit 18jan2000	Entry 24feb1999, exit 24nov1999
Salomon	Entry 13apr1999 exit 19dec2000	Entry 10feb1999, exit 18apr2001
Goldman Sachs	Entry 05jan1999, exit 13feb2001	Entry 10feb1999, exit 29aug2002
SG Valeurs	Entry 13apr1999, exit 02nov2004	Entry 10feb1999, exit 10mar2004
JP Morgan	Entry 05jan1999, exit 21dec2007	Entry 10feb1999, exit 28nov2007
Deutsche Bank	Entry 05jan1999, exit 02jul2014	Entry 10feb1999, exit 09jul2014
Morgan Stanley	Entry 05jan1999, exit 10apr2001	Entry 10feb1999, exit 28may2014
PI Financial Corp	Entry 12may2009, exit 10feb2015	Entry 04feb2009, exit 18feb2015
Ocean Securities	Entry 21feb2006, exit 04mar2008	Entry 24feb1999, exit 23sep2015
Sherbrooke SSC	Entry Apr2020	Entry Apr2020

Appendix Table A1 lists all entries and exits of primary dealers in bill and bond auctions from 1999 until 2022.

Appendix Table A2: 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 A2 shows the results from testing whether customer values are above dealer values. Columns P95, Sum, and Max present p-values for the test-statistics which take the 95th percentile, the sum of squared standardized differences, and the maximum difference across auctions of the average values of dealers less the lower bound of customer values. P-values are computed using the bootstrap. The confidence intervals are for interval estimates of the mean difference. The QWA-value is the average (within customers and dealers) of the individual participants 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 bidders' minimum value (at their last submitted step).

Appendix Table A3: Drivers of Customer Participation and Customer Values

	(OLS)	(FE)
β_1: Basis trade	()	()
β_2: Benchmark status	()	()
β_3: MPC	()	()
β_4: QE	()	()
β_5: Exchange rate	()	()
β_6: Spread	()	()
β_0 : constant	()	()
Adjusted R^2		
N	327	327

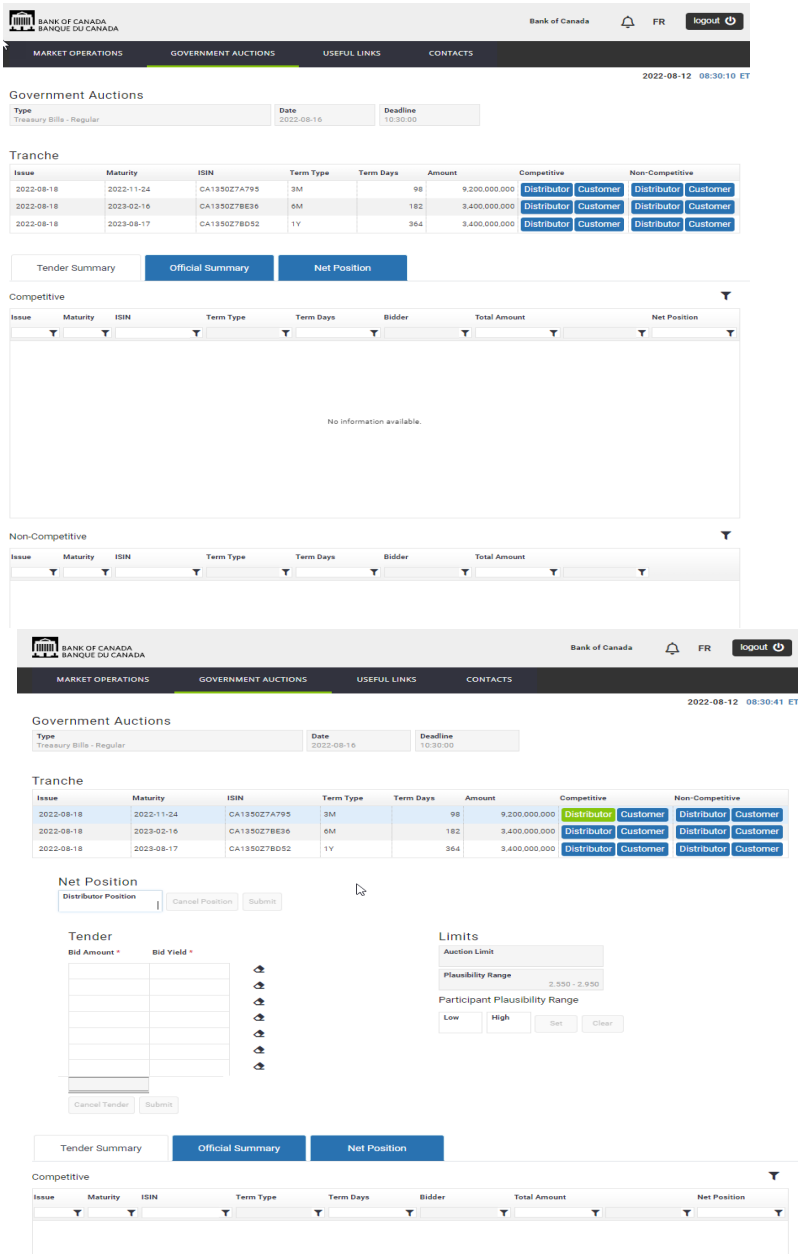
Appendix Table A3 uses the estimated average of the quantity-weighted values of customers instead of the observed number of participating customers as independent variable. The explanatory variables are constructed with price information from the CDS data sample, which ranges from the beginning of 2014 until the end of 2021. Work in progress.

Appendix Figure A1: Primary auctions in different countries



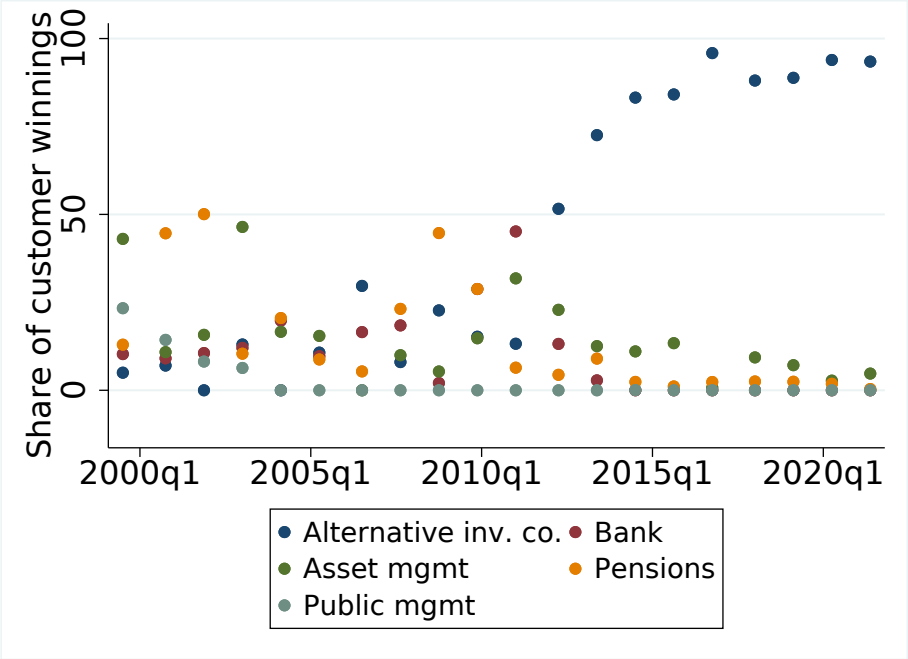
Appendix Figure A1, taken from Muller (2019), shows an overview of how different countries issue debt. Towards the left are countries like Canada who heavily rely on dealers to make markets. Towards the right are countries like the U.S. who let anyone participate in primary auctions.

Appendix Figure A2: What a dealer sees when bidding



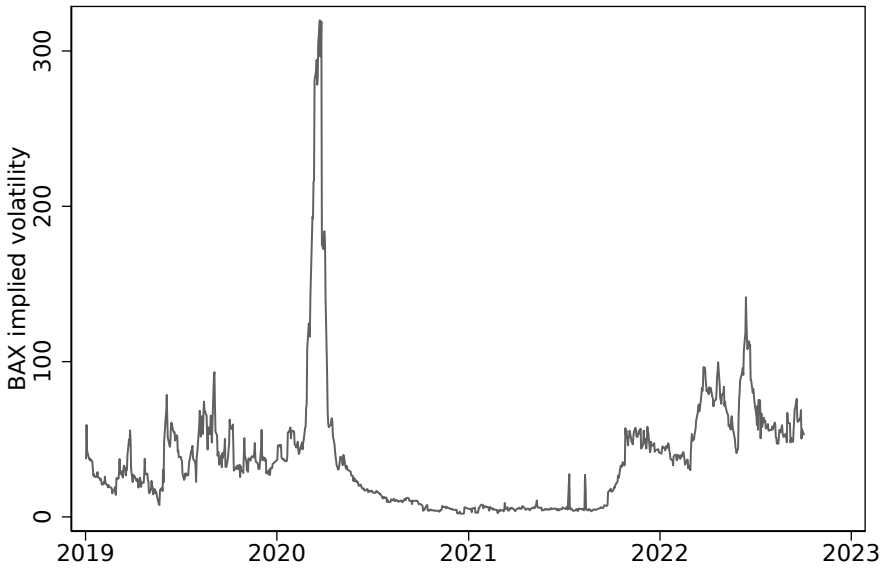
Appendix Figure A2 shows a screenshot of what a dealer sees when placing its bids, either for its own account or on behalf of a customer.

Appendix Figure A3: Purchased amount by investor groups

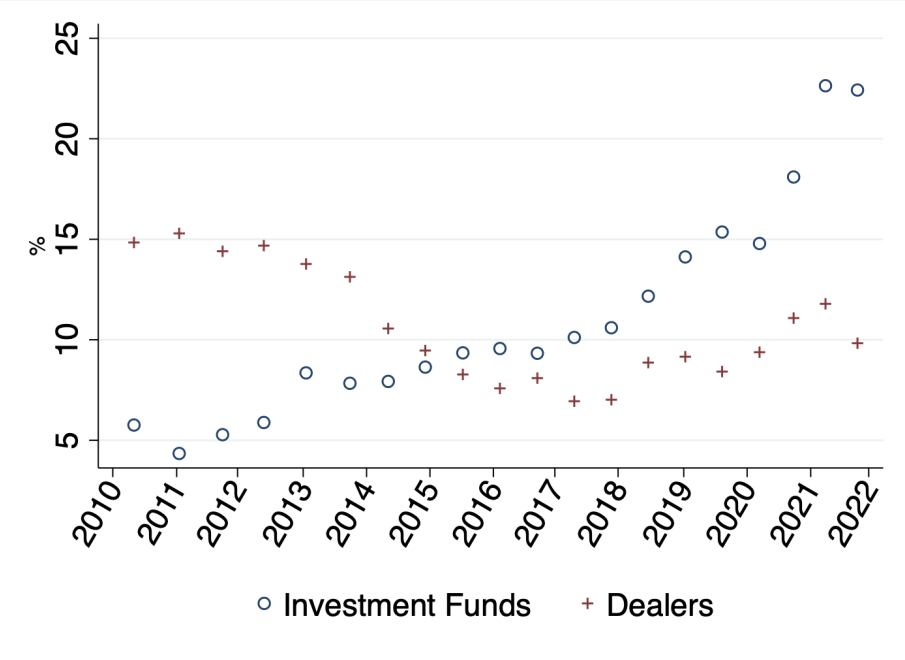


Appendix Figure A3 shows a binned scatter plot of how much each investor group wins in percentage of the total supply bought by non-dealers from 1999 until 2022.

Appendix Figure A4: Volatility in the Canadian Treasury Market

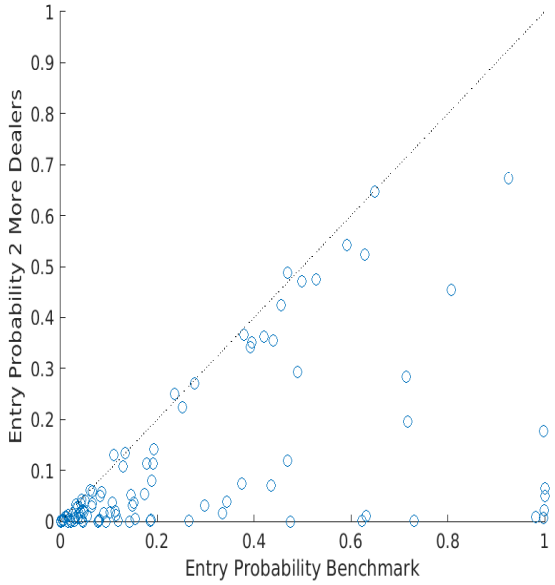


Appendix Figure A5: Auction allotment by investor class for U.S. government bond auctions

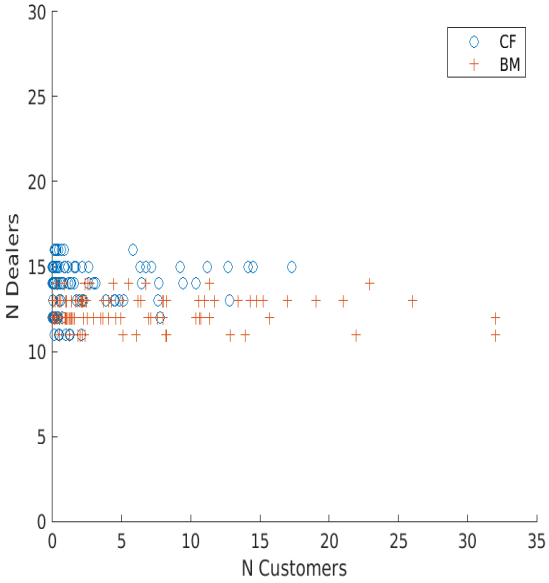


Appendix Figure A5 shows the auction allotment in percentage of supply in U.S. government bond auctions from beginning of 2010 until end of January 2022 for broker/dealers (plus) and for investment funds (circle). Broker/dealers include Includes primary dealers, other commercial bank dealer departments, 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.

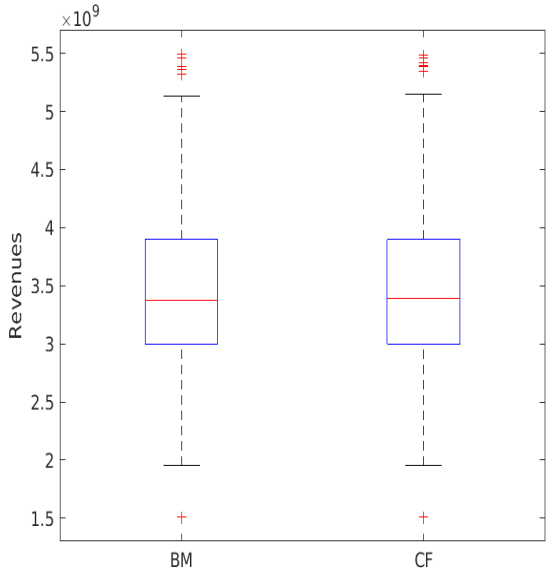
Appendix Figure A6: Revenue and Price Effect—Adding Back Dealers



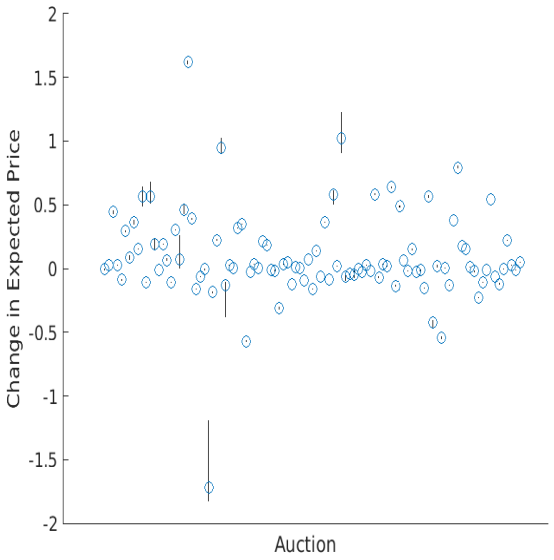
(A) Hedge Entry Probability (in pp.)



(B) Exp. Number of Dealers/Hedge Funds



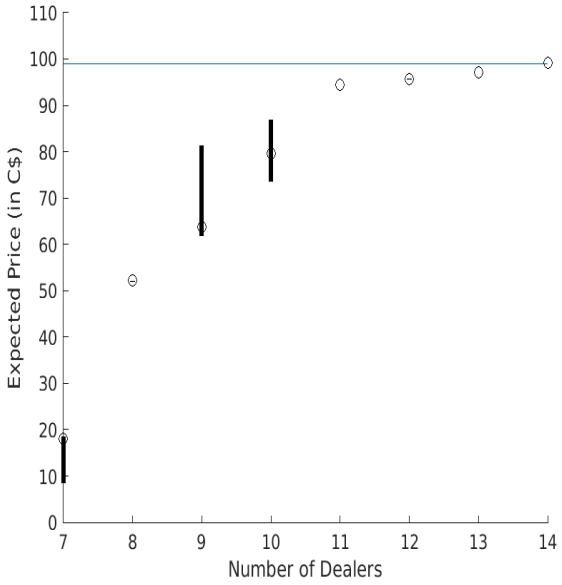
(C) Exp. Revenues (in C$)



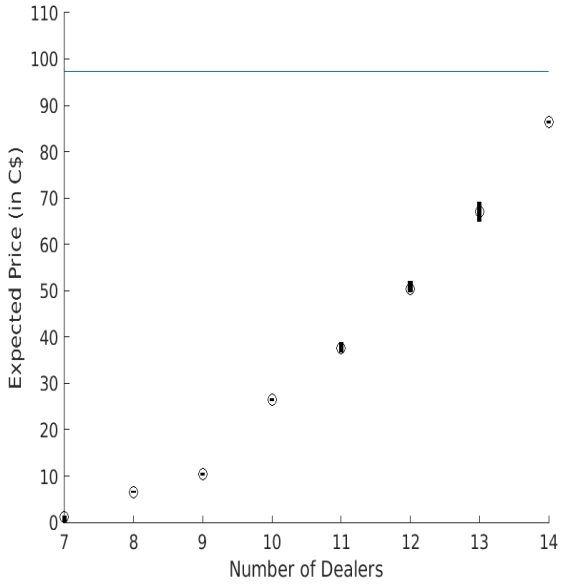
(D) Percentage Change of Expected Price

Appendix Figure A6A shows the probability for a hedge fund (HF) to participate in every 15th auction from 2014 onward 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). Probability is in percentage points. Figure A6B shows the expected number of dealers and hedge funds that participate in each auction in the status quo and the counterfactual. Figure A6C shows the distribution of the expected auction revenues in million C$. Figure A6E a time series of the percentage change in the expected price when going from the status quo to the counterfactual. Prices are in C$ with a face value of 100.

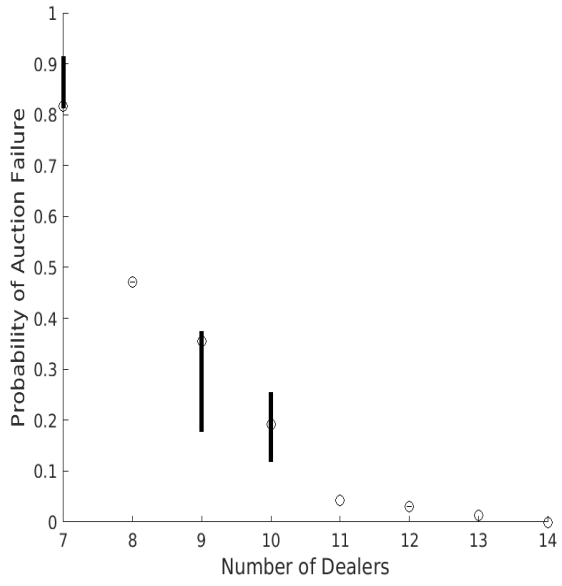
Appendix Figure A7: Expected Price and Number of Dealers for auction p5, and p95



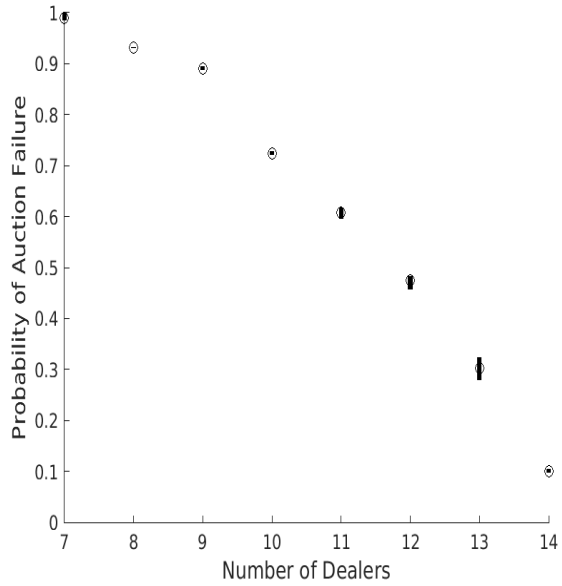
(A) Expected Price



(B) Expected Price



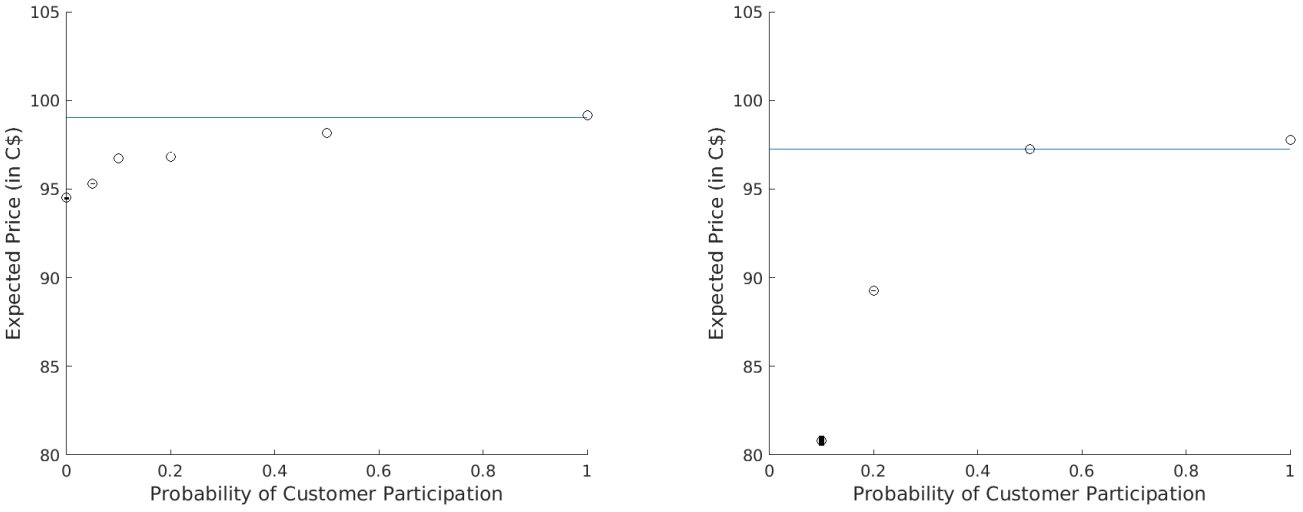
(C) Auction Failure



(D) Auction Failure

Appendix Figure A7 is analogues to Figure 8. It shows how the range of expected price (in C$) at which an auction clears varies as the number of dealers increases from 7 to 14 in two representative auctions, which issue the average supply of all auctions with 1 participating customer in (A), and with 4 participating customers in (B). Here, 1 and 4 are the 5th and 95th percentile of the observed distribution of the number of participating customers. In (C) and (D) we show the auction failure for these two auctions, respectively. In theory, there is one counterfactual equilibrium for each fixed number of dealers. In practice, we determine a range of prices (marked in black) 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) for details). The blue horizontal line shows the average observed bid which is close to the observed market clearing price.

Appendix Figure A8: Typical Auction: Varying Customer Participation Probability



(A) (B)

Appendix Figure A8 shows how the range of expected price (in C$) at which an auction clears varies as the customer participation probability varies between 0 and 1 in an auction in two representative auctions, which issue the average supply of all auctions with 1 participating customer in (A), and with 4 participating customers in (B). Here, 1 and 4 are the 5th and 95th percentile of the observed distribution of the number of participating customers. In theory, there is one expected price for each counterfactual. In practice, we determine a range of prices (marked in black) 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) for details). The blue horizontal line shows the average observed bid which is close to the observed market clearing price.