Interconnected multi-unit auctions: An empirical analysis*

Jason Allen[a] Jakub Kastl[b] Milena Wittwer[c]

August 23, 2022

Abstract

Leveraging an institutional feature that multi-unit auctions for different goods are often held simultaneously, we propose a method for estimating own-and cross-product demand elasticities, avoiding the usual endogeneity issues in demand estimation. We show that these elasticities, together with the auction format, determine how to optimally allocate goods across auctions. To implement our method, we use data from Canadian Treasury auctions. We show that the auctioneer can achieve higher revenues by issuing less of the price-sensitive and more of the price-insensitive security in a discriminatory price auction, and vice versa in a uniform price auction.

Keywords: Multi-unit auctions, structural estimation, market segmentation, government bonds, demand elasticities

JEL classification: D44, C14, E58, G12

*The presented views are those of the authors, not necessarily of the Bank. All errors are our own. We thank the Treasury, auctions and settlement systems team and the debt-management team at the Bank of Canada for critical insights into debt management and auction operations. We thank Markus Brunnermeier, Jens Christensen, Annette Vissing-Jorgensen, Pierre-Olivier Pineau, Andreas Uthemann, and Yu Zhu, as well as seminar/conference participants at the Bank of Canada, Princeton, Penn, ECB workshop on money markets and central bank balance sheets, Econometric Society 2021 winter meetings, EUI, Princeton and Queen's University. Correspondence to: [a]Jason Allen - Bank of Canada, Ottawa, ON K1A0G9, Canada, [b]Jakub Kastl (NBER and CEPR) - Economics, Princeton University, NJ 08540, USA, [c]Milena Wittwer (corresponding author) - Finance, Boston College, MA 02467, USA, Email: wittwer@bc.edu

1 Introduction

Each year, securities and commodities worth trillions of dollars are allocated through multi-unit auctions (Pycia and Woodward (2022)). Surprisingly, many of these auctions are conducted by the same auctioneer, in parallel. For instance, in many countries (including the U.S., Japan, Brazil, France, China, and Canada) government bonds of different maturities are sold in separate, parallel auctions. International carbon allowances regulated by the Emissions Trading Scheme (ETS), diamonds, renewable energy, fish, vegetables, and wine are also sold in parallel auctions in which bidders demand multiple units.[1] Further, trading multiple stocks on an exchange is like participating in parallel multi-unit (double) auctions (Rostek and Yoon (2021a,b, 2022)).

We develop a framework to estimate demand systems for multiple goods by leveraging simultaneous auctions, and show how to use these demand systems to achieve higher revenue without changing the auction format. For illustration, we use data from Canadian Treasury auctions, which is an important market to study in and of itself.

General Framework. In the first part, we introduce a model of the bidding process in simultaneous multi-unit auctions to identify full demand systems, i.e. demand schedules for all goods, and describe how to estimate its parameters.

The idea is to combine two institutional features. First, parallel auctions take place under the same market rules, with the same set of participants, at the same time and in the same economic situation. This allows us to control for unobserved heterogeneity. Second, in multi-unit auctions, bidders submit full demand schedules. This implies that we do not have to pool data across time and market participants to construct demand schedules. Thus, we can ensure that variation in quantities is attributable to variation in prices and not something omitted that is potentially correlated with prices. In contrast, the existing literature addresses this issue by employing instruments aimed at isolating such exogenous variation by making the appropriate exogeneity and validity assumptions (following Berry et al. (1995) in industrial organization, and Koijen and Yogo (2019) in finance).

Our model allows us to overcome two common challenges. First, bidders are strategic and shade their bids which implies that we do not observe their actual demand. Second, by the auction rules, bidders cannot submit multi-dimensional demand schedules that are

[1]For examples of simultaneous auctions of carbon allowances by the EU and UK, click here or here, by Canada and Quebec, click here; of (spot) diamonds, see Cramton et al. (2013); of renewable energy, see Ryan (2022) or click here; of procurement products, see Somaini (2020); of cars, see Larsen and Zhang (2022); of food by Feeding American, click here; of fish, see Carleton (2000), or click here or here; of tomatoes, click here; of wine, click here. All websites to click on were accessed on 07/08/2022.

contingent on prices of multiple securities. This means that, unless demand for different products is independent, we only observe parts of the demand schedules.

Our technical contribution is to extend techniques for identifying demand (or willingness to pay) from bidding data in multi-unit auctions by Guerre et al. (2000), Hortaçsu (2002) and Kastl (2011) to allow demand to depend not only on the allocation of the underlying product, but also on prices of other products. This complements contemporaneous work by Gentry et al. (2022) who identify preferences in simultaneous single-object first-price auctions. Further, we derive equilibrium conditions for interconnected multi-unit auctions with common empirical features. This stands in contrast to existing theoretic contributions which consider elegant and tractable, yet less realistic settings not suited for estimation (e.g., Wittwer (2020, 2021) and Rostek and Yoon (2021a)).

In the second part of the paper, we show that an auctioneer can increase total auction revenue by behaving like a monopolist who price discriminates and exploits differences in demand elasticities. This approach complements a vast literature that analyzes and proposes changes to (multi-unit) auction rules and formats (see Wilson (2021) for an overview). While our approach is by no means the best way to optimize revenue, it is straightforward to implement. This can be useful given that it often difficult to change established auction formats. For instance, Klemperer (2010) proposes a combinatorial auction format to sell multiple financial securities. Despite excellent theoretic properties, however, the auction hasn't become popular in practice.[2]

We begin with a theoretic framework that builds on a simple intuition. Assume the auctioneer seeks to allocate a total amount Q (or production capacity) in the form of two products, S and L.[3] Aggregate demand is price-insensitive for S and price-sensitive for L, meaning that the market price for S decreases less when increasing the supply of S by dQ than the market price for L increases when decreasing the supply of L by the same amount. Then, starting from an equal split of total supply, an auctioneer can increase total revenue by issuing a bit less of L and a bit more of S. Given the difference in price-sensitivity, this increases the revenue in the auction for L by more than it decreases the revenue in the auction for S. However, the lower the supply of L, the lower the revenue gain in the auction for L, even though the price for L increases more strongly. There is a price-quantity trade-off.

[2]Another example is that it took a long academic debate with early contributions by Bikhchandani and Huang (1989) and Back and Zender (1993) and field experiments to change the format of U.S. Treasury auctions from discriminatory price to uniform price. Other countries, such as China, are still debating which auction format to use (e.g., Barbosa et al. (2022)).

[3]In our setting S is short-term Treasuries and L is long-term Treasuries. In the example of fish auctions, S would be one fish species, say halibut, and L another, such as salmon, or in the case of ETS auctions, S would be the carbon allowance for one country and L for another.

We show that this intuition goes through in a uniform price auction in which all winning bidders pay the market clearing price. It would be misleading, however, in a discriminatory price auction in which winning bidders pay the prices they bid, adding a novel aspect to the list of things that distinguish the two auction formats (following Wilson (1979) and Ausubel et al. (2014)). The reason is that unlike in a uniform price auction, the entire aggregate demand curve, and not just the market clearing price, matters in a discriminatory price auction. When supply changes, bidders adjust their bids and the aggregate demand curve adjusts. It becomes an empirical question whether it is revenue-increasing to issue more of S or L, which can be answered using our structural model.

Empirical application. To illustrate the applicability of our method we use data from the Canadian primary market for government debt. Our data covers 15 years from 2002 until 2015. For most of the paper, we focus on bills; bills of different maturities are sold in parallel, while bonds of different maturities are sold on different days.[4] Further, we concentrate on banks that act as primary dealers (dealers) because they buy almost all the debt in an average auction. Non-dealer banks (customers) must place their bids via dealers. We observe which security and maturity type is issued at what amount, in addition to unique anonymized bidder identifiers. Moreover, we observe all submitted bids (i.e., demand schedules) and know at what time and how the bids are submitted, whether directly to the auctioneer or via a dealer. Finally, we observe the identity of the winning bidders, how much they won and at what price.

With these data we first present descriptive evidence to motivate that it is important to consider demand systems that account for interdependencies across different securities. Using the time-stamps on when bids are placed, we show that dealers who observe their customers placing bids for one security before the auction closes, say the 3 month bill (3M), change their own bids not only for the 3M bill (as in Hortaçsu and Kastl (2012)), but also for the 6M and 12M bill. This points towards interdependencies across the different maturities. If instead demand for the 3M bill was entirely independent of the 6M bill, the dealer would bid in the 3M auction as if this auction took place in isolation.

We next estimate our model and find demand for all three maturities of bills to be rather price-insensitive. For instance, when the average dealer wins 1% more of the supply of 12M bills, his price offer for the 12M bills decreases by 0.24 basis points (bps).[5] Perhaps surprisingly, for the average dealer bills are only weak substitutes, despite the claim that

[4]The U.S. and other large economies also issue bonds in parallel. Our methodology is easily portable to these other settings.

[5]Allen and Wittwer (2021) find that the demand elasticity of an average investor in the secondary market is of similar magnitudes, applying a fundamentally different estimation approach.

all bills are cash-like. For instance, if the average dealer wins 1% more of the supply of the 3M bills, the price offered for the 12M bills decreases by 0.06 bps, and of the 6M bills by 0.02 bps.

We use our demand estimates to show that it is revenue-increasing to issue more of the price-sensitive bond (longer maturity) and less of the price-insensitive bond (shorter maturity) in a discriminatory price auction and vice versa in a uniform price auction. In particular, in a discriminatory price auction, assuming that bonds are perfect substitutes would lead us to over-estimate the effects on revenue of maturity-shuffling government debt. On the other hand, assuming independence would lead to an under-estimate of these same revenue effects. Even though the economic magnitudes are small for the Canadian bill market (since demands are overall price-insensitive), the exercise highlights the importance of correctly accounting for interdependencies across goods when calculating revenues. Revenue gains from reshuffling supply are larger when demand is more price-sensitive, as is the case in other Treasury markets, such as the Spanish and Portuguese primary markets (see Bigio et al. (2021); Albuquerque et al. (2022)).

Our empirical application complements two strands of literature in macro-finance. The first analyzes the aggregate demand for government debt with market-level data of the secondary market. Most of this literature focuses on comparing long-term with short-term debt (e.g., Gagnon et al. (2011); D'Amico et al. (2012); Lou et al. (2013); Krishnamurthy and Vissing-Jørgensen (2011, 2012, 2015)).[6] We consider different types of short-term debt (i.e., bills), more similar to Greenwood et al. (2015b) and Krishnamurthy and Li (2022). In contrast to these papers, we zoom in on one market, the primary market, where we can identify demand of individual institutions which act as dealers. Further, we show that the demand of these dealers may differ from the aggregate demand of the market, because dealers function as market makers who connect the primary and secondary markets.[7]

The second literature analyzes whether to issue government debt in the form of long- or short-term bonds (e.g., Missale and Blanchard (1994); Greenwood et al. (2015a,b); Belton et al. (2018); Bhandari et al. (2019); Bigio et al. (2021)). In this literature, the key trade-off for the government is between default and inflation commitment problems and roll-over risks. We set aside these dynamic aspects of the debt allocation problem

[6]A common approach for studying interdependencies across maturities is via term-structure models. To identify the implied correlations of prices (yields) across maturities, various papers rely on changes in the supply of Treasury securities (e.g., Krishnamurthy and Vissing-Jørgensen (2012); D'Amico et al. (2012); Lou et al. (2013)). Other papers provide evidence that even government bonds that are issued by different countries are close substitutes (e.g., Nagel (2016)).

[7]In Appendix A we provide one possible micro-foundation for our demand curve specification in the spirit of Vayanos and Vila (2021). This model highlights how demand and prices in the primary market are affected by the structure of secondary markets.

by including an issuance cost of debt that absorb the (mechanical) price difference of bonds with different maturities. Instead, we highlight how a government can reduce its cost of financing by exploiting the fact that demand for shorter-term bonds tends to be less price-sensitive than demand for longer-term bonds. We view these two approaches as complementary.

General lessons and future research. Our empirical application focuses on the the demand and supply of government debt, yet our method and insights on how to split supply across different goods can be useful in many other settings, including those described above. Unlike standard "BLP" demand estimation following Berry et al. (1995), our method can identify both types of interdependencies, substitutes and complements, which likely arise in many settings for various reasons, for instance, when bidders face budget or capacity constraints. For this, we do not need to impose any correlation structure in the unobserved preferences for different goods which affects the estimates in standard demand models (Gentzkow (2007)).

Our counterfactual exercises highlight that taking these interdependencies seriously can help auctioneers achieve higher auction revenues without having to change the auction format. To illustrate the main insight, let us use another, perhaps more relatable, setting. Consider a fisher who routinely auctions salmon and halibut and charges bidders (here retailers) the market price for each fish. The fisher conjectures that salmon consumers are less price-sensitive than halibut consumers. Then, he should fish more salmon and less halibut to achieve higher revenue given his boat has capacity Q. However, if he charged each bidder the prices that they bid, it might be better to fish more halibut and less salmon.

Future research can expand our framework and estimation in various ways. It might be interesting to quantify how much revenue could be gained when switching to a combinatorial auction format (such as Klemperer (2010)). This could help us understand why parallel auctions of related goods exist in practice, even though we know from the literature on market design that these auctions cannot achieve the first-best revenue or welfare. Further, it might be useful to study alternative objective functions, other than auction revenue. For example, with data on ETS auctions, it might be more natural to reduce total carbon emissions, while we would maximize the total gains from trade with data on limit order books. Finally, our method could be extended to cover sequential auctions which are common in practice.

Outline. The remainder of the paper is structured as follows: Section 2 describes the institutional environment and the data set. Section 3 presents descriptive evidence for interdependencies across goods. Section 4 introduces and estimates the model. Sections 5 discusses the counterfactual exercise. Section 6 concludes. All proofs are in the Appendix. Throughout the paper, random variables are denoted in **bold**.

2 Institutional Environment and Data

Multi-unit auctions. Many divisible goods (such as financial securities, carbon allowances, energy or diamonds) are auctioned via standard (sealed-bid) multi-unit auctions. Other goods (such as fish, cattle or wine) are sold in auctions in which bidders demand multiple units or objects. These auctions can often be approximated by one of the two standard multi-unit auction formats (discriminatory or uniform price), depending on the auction's pricing rule.

In a standard multi-unit auction a bid is a step-function with at most K steps, which specifies how much a bidder offers to pay for specific amounts of the good for sale (as in Figure 1a). When the auction closes, the final bids are aggregated and the market clears where aggregate demand meets total supply. Everyone wins the amount they asked for at the clearing price (subject to pro-rata rationing on-the-margin in case of excess demand at the market clearing price). In a discriminatory price auction, each bidder pays according to what they bid, similar to a first-price auction. In a uniform price auction, each bidder pays the market clearing price for each unit won.[8]

Canadian Treasury auctions. There are three types of Canadian Treasury bills: 3, 6, and 12 month bills. They are sold every second Tuesday by the Bank of Canada in three separate, but parallel, discriminatory price auctions. Two groups of bidders, dealers and customers, participate in the auctions. Dealers are either primary dealers or government securities distributors. Customers can only submit bids through dealers, but like dealers, they tend to be large financial institutions. Our focus lies on dealers, who buy more than 85% of the issued amount in an average auction.

From the time the tender call opens until the auctions close, bidders may submit and update their bids in two forms (see Appendix Figure A1). The first is a competitive bid, which is a step-function with at most 7 steps. The quantity demanded must be stated in multiples of $1,000 and be at least $100,000. The bid shall state the yield to maturity (in

[8] For detailed descriptions of other empirical settings with multi-unit auctions, see, for instance: Hortaçsu (2002), Kang and Puller (2008), Kastl (2011), Hortaçsu et al. (2018), Bonaldi and Ruiz (2021), Cole et al. (2022), Barbosa et al. (2022) for Treasury auctions of Turkey, Korea, Czech Republic, U.S., Columbia, Mexico, China; Cramton et al. (2013) for diamonds; Hortaçsu and Puller (2008) for electricity; Ryan (2022) for renewable energy.

Figure 1: Bids in the Canadian Treasury bill market

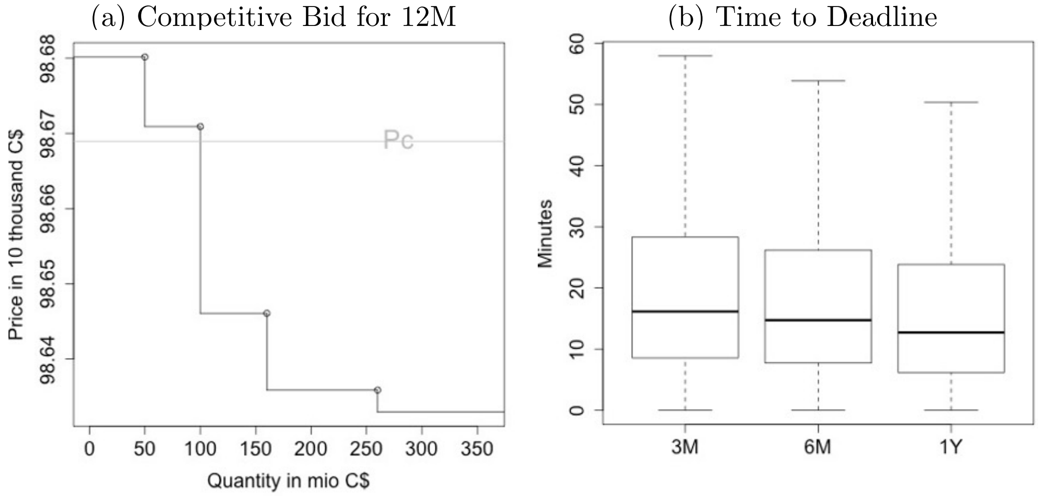

Figure 1a displays an example of a bidding step function. It is the one of the median dealer in a 12M auction, computed as follows: Determine the median number of steps in all competitive bid functions submitted by dealers, and then take the median over all (price, quantity) tuples corresponding to each step by a dealer who submitted the median number of steps. Figure 1b depicts the distribution of the time at which bids arrive prior to the deadline in each of the auctions. Very early outliers and bids that go in after auction closure are excluded.

%) to three decimal places. We convert yields into prices so that demand schedules are decreasing rather than increasing, using a face value of C$ 1 million (as in Figure 1a).[9]

The second form of bidding is a non-competitive bid. This is a quantity order, which the bidder will win for sure, but for which he pays the average price of all accepted competitive bid prices. It is capped at $ 10 million for dealers and $ 5 million for customers, and hence trivial relative to the competitive order sizes—with one exception: the Bank of Canada itself. It utilizes non-competitive bids to reduce the previously announced supply and to purchase Treasuries (assets) to match its issuance of bank notes (liabilities).

Data. To estimate demand systems in multi-unit auctions, the data set must include all winning and losing bids (i.e., full step functions), the issued supply, the market clearing price of each auction, and how much each bidder won. Our data is richer than this.

Our data set consists of all 366 Canadian Treasury bill auctions between 2002 and 2015, in addition to all Treasury bond auctions. Table 1 summarizes the data on bills. On average the Bank of Canada announced issuances of C$6.41 billion for 3M bills and C$2.47 billion for each of the 6M and 12M bills per auction, of which it actually distributed roughly C$5.76 (3M) and C$2.12 billion (6/12M).[10]

[9] yield = (C$1 million - price)/price* 365/days left to maturity.

[10] Online Appendix Figure O1 plots the issuance amounts over the period 2012–2017.

Table 1: Data Summary of 3M/6M/12M Auctions

	Mean			SD			Min			Max		
	3M	6M	12M	3M	6M	12M	3M	6M	12M	3M	6M	12M
Issued amount	5.76	2.12	2.12	1.68	0.52	0.52	3.05	1.22	1.22	10.40	3.80	3.80
Dealers	11.88	11.79	11.03	0.90	0.93	0.83	9	9	9	13	13	12
Global part. (%)	93.67	93.84	98.84	24.34	24.04	10.67	0	0	0	100	100	100
Customers	6.26	5.68	5.35	2.69	2.94	2.54	1	0	0	14	13	15
Global part. (%)	35.66	40.13	39.46	47.90	49.02	48.88	0	0	0	100	100	100
Comp demand as % of announced sup.	16.29	16.91	17.02	7.96	7.61	7.31	0.002	0.019	0.005	25	25	25
Submitted steps	4.83	4.23	4.35	1.86	1.78	1.75	1	1	1	7	7	7
Dealer updates	2.89	2.18	2.48	3.58	2.87	3.18	0	0	0	31	31	42
Customer updates	0.12	0.13	0.19	0.40	0.40	0.58	0	0	0	4	3	9
Non-comp dem. as % of announced sup.	0.05	0.15	0.15	0.03	0.10	0.10	$5/10^5$	$4/10^5$	$2/10^3$	0.24	0.58	0.58

Table 1 displays summary statistics of our sample, which goes from January 2002 until December 2015. There are 366 auctions per maturity. The total number of competitive bids (including updates) in the 3M, 6M, 12M auctions is 66382, 48927, and 56721, respectively. These individual steps make up 18272, 15514, and 17077 different step-functions. The total number of non-competitive bids is 2477, 2378, and 1932. From the raw data we drop competitive bids with missing bid price (133) and competitive or non-competitive bids with missing quantities (69). Global part. is the probability of attending the remaining auctions, conditional on bidding for one maturity. Dollar amounts are in billions of C$.

We identify each bidder through a bidder ID, and know whether the bidder is a dealer or a customer. The average auction has 11 to 12 dealers and 5 to 6 customers. Dealers tend to participate in all parallel auctions to keep their dealer status.[11]

We observe all bids submitted from the opening of the tender call until the auction closes. The updating period lasts one week, although most bids arrive within 10 to 20 minutes prior to closing (see Figure 1b). Typically, a dealer updates his (competitive or non-competitive) bid once or twice. The median number of updates is one. The higher average (2.26) is driven by outliers. Customers are less likely to update, with an average number of 0.1 (and a median of no updates).

An average step-function of a competitive bid has 4.5 steps with little difference across maturities. Non-competitive bids are small in size. On average, bidders only demand 0.1% of the total (announced) supply via non-competitive bid. Given their size, our structural model abstracts from non-competitive bids, and focuses solely on the decision

[11] "At every auction, a primary dealer's bids, and bids from its customers, must total a minimum of 50 per cent of its auction limit and/or 50 per cent of its formula calculation, rounded upward to the nearest percentage point, whichever is less. [...] Each government securities distributor must submit at least one winning competitive or non-competitive bid on its own behalf or on behalf of customers, every six months." (Bank of Canada (2016), p. 12).

of placing competitive bids. The Bank of Canada, on the other hand, demands substantial amounts via non-competitive bids to reduce the total supply on the day of the auction. On average, it takes away 11.13% (3M), 14.35% (6M), 14.26% (12M) with a maximum of 20.45% (3M), 41.66% (6M), 25.00% (12M) of the total previously announced supply. Our empirical model will need to account for unannounced changes in actual supply.

3 Motivating Interdependencies in Demands

Before estimating demand for different goods, it is useful to present evidence suggesting that studying auctions for individual goods in isolation provides an incomplete picture of demand. Different goods might be interconnected both on the supply and the demand side. On the supply side, the seller might determine the total amount for sale at each auction jointly, which leads to a non-zero correlation between the sold amounts across goods. On the demand side, bidders might want to buy bundles of goods to satisfy the demand of consumers downstream after the auction.

Cross-auctions correlations. A natural starting point to look for dependencies across auctions is to analyze correlations on the supply and demand side (see Table 2). In our data, the supply that the Bank of Canada announces exhibits perfect positive correlation across maturities. In fact, over our long sample the Bank of Canada always announces the exact same issuance size for the 6M and 12M bills. The amount it actually distributes on the auction day is also almost perfectly correlated.[12]

We observe a similar pattern on the demand side. The total amount bidders demand (via competitive or non-competitive bid) when the auction closes is highly positively correlated across maturities, about 0.91–0.92. The correlation between quantities actually won drops to 0.54–0.57, suggesting that bidders do not always achieve this goal.

These correlations suggest that bidders don't value bills as independent. Bills could be complements or substitutes. This depends on how much bidders are willing to pay for one maturity when winning more of the other maturities, and not on the correlation patterns of supply and demand quantities alone. Concretely, if the bidder is willing to pay more

[12]Policy-makers perform stochastic simulations to determine a debt strategy that is desirable over a long horizon, e.g. 10 years. The model (https://github.com/bankofcanada/CDSM) trades off risks and costs of different ways to decompose debt over the full spectrum of government securities. Part of the simulation routine is to specify ratios between maturities, for instance $1/4^{th}$ of each of the 3/6/12M bills and $1/16^{th}$ of each of the 2/5/10/30-year bonds. Final issuance decisions are taken based on model simulations and judgment. "The typical practice is to split the total amount purchased by the Bank [of Canada], so that the Bank's purchases approximate the same proportions of issuance by the government across the three maturity tranches" (Bank of Canada (2015)).

Table 2: Cross-Market Correlations

(a) Supply Side

	$\bar{Q}_{3M}$	$\bar{Q}_{6M}$	$\bar{Q}_{12M}$		Q_{3M}	Q_{6M}	Q_{12M}
Q_{3M}	1.00			Q_{3M}	1.00		
$\bar{Q}_{6M}$	1.00	1.00		Q_{6M}	0.99	1.00	
$\bar{Q}_{12M}$	1.00	1.00	1.00	Q_{12M}	0.99	1.00	1.00

(b) Demand Side

	$q^D_{3M,i}$	$q^D_{6M,i}$	$q^D_{12M,i}$		$q^*_{3M,i}$	$q^*_{6M,i}$	$q^*_{12M,i}$
$q^D_{3M,i}$	1.00			$q^*_{3M,i}$	1.00		
$q^D_{6M,i}$	0.92	1.00		$q^*_{6M,i}$	0.57	1.00	
$q^D_{12M,i}$	0.91	0.91	1.00	$q^*_{12M,i}$	0.54	0.57	1.00

Table 2a displays the correlation between the announced issuance amount, $\bar{Q}_m$, and the distributed supply, Q_m, for the three maturities, $m = 3, 6, 12M$. Table 2b correlates bidder i's demand $q^D_{m,i}$ and the amount he won $q^*_{m,i}$ across the different maturities.

for one good when he wins more of the other good, the goods would be complements, and substitutes vice versa.

Bid regression. We cannot observe how much a bidder is willing to pay, but we can test whether the bidder offers more or less for one good when he wins more of the other goods. For this, we regress bid k of dealer i on day t for good m on the quantity demanded at that step, $q_{t,m,i,k}$ and the amount the dealer won in the other two auctions, $won_{t,l,i}$, plus a day and bidder fixed effect, ζ_t, ζ_i:

$$b_{t,m,i,k} = \alpha_m + \lambda_m q_{t,m,i,k} + \sum_{l \neq m} \delta_{m,l} won_{t,l,i} + \zeta_t + \zeta_i + \epsilon_{t,m,i,k}. \tag{1}$$

We find that all δ parameters are statistically significant and positive, suggesting that bills might be complements (see Table 3). Given that bills are cash-like, this seems counterintuitive. In fact, the estimates from this regression are unbiased only if bidders do not shade their bids and have a sufficiently precise idea of how much they will win in the other auctions while bidding. Both are strong assumptions which may fail to hold, but can be tested and avoided using our model.

Bid updating. In our setting, in which dealers observe customer bids, we can provide another piece of evidence in favor of interdependencies. We know that a dealer updates his own bid, say, for the 3M bills, upon observing a customer bid in the 3M auction,

Table 3: Bid regression

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	-4.593	(0.519)	λ_{6M}	-6.338	(0.667)	λ_{12M}	-11.86	(1.130)
$\delta_{3M,6M}$	$+0.627$	(0.114)	$\delta_{6M,3M}$	$+1.289$	(0.220)	$\delta_{12M,3M}$	$+1.633$	(0.641)
$\delta_{3M,12M}$	$+0.475$	(0.180)	$\delta_{6M,12M}$	$+2.422$	(0.396)	$\delta_{12M,6M}$	$+4.735$	(0.786)
$\hat{\alpha}_{3M}$	995189.4	(4.847)	$\hat{\alpha}_{6M}$	990957.6	(6.081)	$\hat{\alpha}_{12M}$	980312.6	(7.568)
N	18307			15641			16302	

Table 3 shows the estimation results of regression (1) for each maturity, using final bids by dealers. Bids are in C$ and quantities in % of auction supply. Standard errors are in parentheses, clustered at the bidder level. Results are robust to including customer bids and bid updates.

for instance, because this provides information about competition or the fundamental security value (Hortaçsu and Kastl (2012)). If the three maturities are interdependent, the dealer should also update his bids for the other bills. To test this, we run the following Probit regression on competitive bids placed by dealers, treating each step function as one observation:

$$update_{m,i} = \alpha + \sum_m I_m \left(\beta_m customer_m + \delta_{m,-m} customer_{-m}\right) + \varepsilon_{m,i}. \tag{2}$$

The dependent variable $update_{m,i}$ takes value 1 if dealer i updated his bid in an auction for m, and 0 otherwise. I_m is an indicator variable equal to 1 if the update occurs in the auction for maturity m. $customer_l$ (for $l = m$ or $-m$) is also an indicator variable, which is created in two different ways. In the more conservative specification (1) $customer_l$ takes value 1 only if the dealer received a competitive order by his customer for maturity l immediately before taking action in auction m himself. Specification (2) builds on this benchmark but takes a longer sequence of events, which are less than 20 seconds apart, into account (e.g., as in Appendix Table A1). This acknowledges that the auction interface (shown in Appendix Figure A1) does not allow bidders to submit bids for different maturities at the same time. Further, it takes time to calculate bids, enter them manually—which until 2019 is the rule rather than exception—and transfer them electronically.

Table 4 displays the estimated coefficients for specifications (1) and (2), in columns (1) and (2), respectively. The significant positive $\hat{\beta}_m$ coefficients support existing evidence by Hortaçsu and Kastl (2012) on dealer updating. The significantly positive $\hat{\delta}_{m,-m}$ suggest that dealers also update their bids across maturities. As expected, the level of significance increases when taking into account the fact that in practice dealers' bids are hardly ever simultaneous, but instead placed in close sequence.

Table 4: Probability of Dealer Updating Bids

Coefficient	Verbal description	(1)		(2)	
$\hat{\beta}_{3M}$	update in $3M$ after order for $3M$	0.533	(0.056)	0.711	(0.053)
$\hat{\delta}_{3M,6M}$	update in $3M$ after order for $6M$	0.405	(0.064)	0.531	(0.061)
$\hat{\delta}_{3M,12M}$	update in $3M$ after order for $12M$	0.303	(0.057)	0.446	(0.054)
$\hat{\delta}_{6M,3M}$	update in $6M$ after order for $3M$	0.086	(0.063)	0.248	(0.059)
$\hat{\beta}_{6M}$	update in $6M$ after order in $6M$	0.848	(0.076)	0.929	(0.070)
$\hat{\delta}_{6M,12M}$	update in $6M$ after order in $12M$	0.729	(0.080)	0.762	(0.074)
$\hat{\delta}_{12M,3M}$	update in $12M$ after order for $3M$	0.556	(0.070)	0.664	(0.066)
$\hat{\delta}_{12M,6M}$	update in $12M$ after order for $6M$	0.120	(0.059)	0.244	(0.056)
$\hat{\beta}_{12M}$	update in $12M$ after order for $12M$	0.828	(0.061)	0.934	(0.059)
$\hat{\alpha}$	constant	0.476	(0.007)	0.448	(0.007)

Table 4 shows the results of the Probit regression (2). In column (1) $customer_l$ is an indicator variable equal to 1 if the dealer received a competitive order from a customer for maturity l immediately before taking action in auction m himself. In column (2) $customer_l$ is an indicator variable equal to 1 if the dealer received an order for maturity l within one minute before placing his own bid in auction m, or if the dealer's bid is part of a sequence of bids which are each less then 20 seconds apart, starting less than one minute after the customer's order. The total number of observations is 39,271. Standard errors are in parentheses.

Take away. Taken together, the evidence suggests that securities are interdependent. This motivates the need for a methodology that can identify full demand systems of interdependent goods.

4 Identifying Demand Systems

It is challenging to consistently estimate the full demand system for two main reasons. First, bidders have private information about how much they value each good. This generates incentives to misrepresent the true demand. As in a first-price auction, bidders shade their bids to reduce the total payments they must make to win. Thus, we cannot infer their true demands by looking at bids. Second, even if bidders wanted to report their true demands, by the rules of the auction, they can, in auction m, only submit a one-dimensional bidding step-function that depends on amounts of good m, not on goods $-m$ (such as in Figure 1a). This implies that we only observe parts of the demand system. To solve these challenges, we model the auction process.

4.1 Model of Simultaneous Multi-Unit Auctions

We describe an auction model that fits our empirical application, but highlight how to adjust assumptions for other applications.

M perfectly divisible goods, indexed m, are auctioned in M separate multi-unit auctions that run in parallel. The auction format may be uniform or discriminatory price. We focus on the latter in the main text, but present the equilibrium conditions for uniform price auctions in Appendix B.

There are $g \geq 1$ commonly known groups of bidders. In our application, there are dealers (d) and customers (c). In addition, all or some bidders could carry a latent type, as illustrated in a model extension in Online Appendix A. The total number of bidders in each group, N_g, is commonly known.

Over the course of the auction, bidder i of group g draws a private signal $\boldsymbol{s^g_{i,\tau}} \equiv (\boldsymbol{s^g_{1,i,\tau}} \dots \boldsymbol{s^g_{M,i,\tau}})$ at time τ. The signal may be multi-dimensional. To account for differences between bidder groups, it may be drawn from different distributions for customers and dealers.

Assumption 1. *Dealers' and customers' private signals $\boldsymbol{s^d_{i,\tau}}$ and $\boldsymbol{s^c_{i,\tau}}$ are for all bidders i independently drawn from common atomless distribution functions F^d and F^c with support $[0,1]^M$ and strictly positive densities f^d and f^c.*

Notably, a bidder's signal can be persistent since we do not pool bids from auctions held at different points in time. The signal must only be independent from all other signals conditional on anything that everyone knows at the time of the auction. In our empirical application, this includes a reference price-range provided by the auctioneer (see Appendix Figure A1), in addition to all public information that is available in the active forward (when-issued) market. The presence of this market implies that most, if not all, information relevant for price-discovery is aggregated prior to the auction and that any private information about future resale value can be arbitraged away. Thus the heterogeneity of information at the time of the auction is likely driven mostly by idiosyncratic factors such as the structure of the balance sheet, investment opportunities or repo needs—which do not depend on private information of other dealers. Consistent with this, Hortaçsu and Kastl (2012) fail to reject that dealers only learn about competition from observing customers' bids, which provides support for assuming that valuations are (conditionally) private.[13]

[13] In other settings, the independent signal assumption might be too strong. For example, Boyarchenko et al. (2021) provide evidence of information sharing in U.S. Treasury auctions. Estimating bidder valuations in such settings without having to make strong functional form assumptions remains an open

The bidder's signal affects his true (inverse) demand or marginal willingness to pay.

Assumption 2. *The marginal willingness to pay (or valuation) of a bidder with signal $s^g_{m,i,\tau}$ for amount q_m conditional on purchasing q_{-m} of the other securities $-m$ is*

$$v_m(q_m, q_{-m}, s^g_{m,i,\tau}) = f_m(s^g_{m,i,\tau}) + \lambda_m q_m + \delta_m \cdot q_{-m}, \quad (3)$$

where $f_m(\cdot)$ maps any realization of $\boldsymbol{s^g_{m,i,\tau}}$ into $\mathbb{R}^+$ for all m, and $\lambda_m < 0$, $|\delta_m| < \lambda_m$, α_m are sufficiently high such that the marginal willingness to pay does not drop below 0 for any amount that might be for sale.

Generally, $f_m(\cdot)$, λ_m and δ_m could be bidder-specific, e.g., $\delta_{m,i}$. Importantly, δ_m and q_{-m} are vectors when there are more than two maturities—a simplified notation we adopt throughout the paper. The vector of δ_m parameters measures interdependencies across maturities. Take the example of the $m = 3M$ auction, where $q_{-m} \equiv (q_{6M}\ q_{12M})'$ and $\delta_m \equiv (\delta_{3M,6M}\ \delta_{3M,12M})$. If $\delta_{3M,6M} < 0$, bidders are willing to pay less for any amount of the 3M maturity the more they purchase of the 6M bills, hence the bills are substitutes. When $\delta_{3M,6M} > 0$ they are complementary, and independent if $\delta_{3M,6M} = 0$.

Relative to existing papers that ignore interdependencies across auctions, we impose more structure on the bidder's marginal willingness to pay. To justify our functional form assumption, we provide one possible micro-foundation in Appendix A that helps us understand what drives the demand of a bidder in a multi-unit auction who seeks to sell some of the amount won to clients that arrive after the auction. We show that demand schedules can be approximated by linear functions (Proposition 2), and that different goods are less substitutable (even complementary) for bidders for whom it is more costly to turn down clients after the auction. For large bidders—whom we identify in our application as market makers—who can more easily satisfy client demand, bills are stronger substitutes.

Knowing their own true demands, each bidder chooses how to bid. A bid in auction m consists of a set of quantities in combination with prices. It is a step-function which characterizes the price the bidder would like to pay for each amount.

Assumption 3. *In auction m each bidder has the following action set each time he places a bid:*

question in the literature. Bonaldi and Ruiz (2021) take a first step in this direction for uniform price auctions.

$$A_m = \begin{cases} (b_m, q_m, K_m) : dim\,(b_m) = dim(q_m) = K_m \in \{1, ..., \overline{K}_m\} \\ b_{m,k} \in [0, \infty) \;\; and \; q_{m,k} \in [0, 1] \\ b_{m,k} > b_{m,k+1} \;\; and \; q_{m,k} > q_{m,k+1} \forall k < K_m. \end{cases}$$

To compare bids in auctions with different sizes of supply, $q_{m,k} \in [0, 1]$, representing the share of total supply. A bid of 0 denotes non-participation.

In Canadian Treasury auctions bidders may update their bids over the course of the auction. To capture this process we follow Hortaçsu and Kastl (2012) and assume that new information may arrive at a discrete number of time slots $\tau = 0, ..., \Gamma$, giving the bidder incentives to update the bid.[14] For settings with no updating set $\Gamma = 0$.

The rules of the Canadian auctions do not allow for customers to submit their own bids. At each time τ at which a customer seeks to place a bid, he is matched to a dealer who observes the customer's bid. This provides the dealer with additional private information, which we denote by $\boldsymbol{Z_{m,i,\tau}}$. Thus, if the dealer observed a customer's bid in all three auctions, his information set is $\boldsymbol{\theta^g_{i,\tau}} = (\boldsymbol{s^g_{i,\tau}}, \boldsymbol{Z_{1,i,\tau}}, \boldsymbol{Z_{2,i,\tau}}, \boldsymbol{Z_{3,i,\tau}})$. If he only has a customer in one auction, say for maturity 1, $\boldsymbol{\theta^g_{i,\tau}} = (\boldsymbol{s^g_{i,\tau}}, \boldsymbol{Z_{1,i,\tau}})$, and so on. By Assumption 1, $(\boldsymbol{s^g_{i,\tau}}, \boldsymbol{Z_{i,\tau}})$ are independent across dealers and time. However, $\boldsymbol{s^g_{i,\tau}}$ and $\boldsymbol{Z_{i,\tau}}$ can be correlated within a dealer across τ. For auctions in which all bidders bid directly to the auctioneer: $\boldsymbol{\theta^g_{m,\tau}} = \boldsymbol{s^g_{m,\tau}}$.

Definition 1. *A pure-strategy is a mapping from the bidder's set of types at each time τ to the action space of all three auctions: $\Theta^g_{i,\tau} \to A_1 \times A_2 \times A_3$.*

A choice in auction m by a bidder who draws type $\theta^g_{i,\tau}$ may be summarized as bidding function $b^g_{m,i,\tau}(\cdot, \theta^g_{i,\tau})$. When auction m closes at $\tau = \Gamma$, the auctioneer aggregates the bidders' final bids, and the market clears at the lowest price P^c_m at which aggregate demand satisfies aggregate supply.

In our application, auction supply is the announced amount for sale net of what the Bank of Canada demands in the form of non-competitive bids during the auction plus all other competitive bids by bidder i's competitors. In other settings, auction supply is fixed at Q_m. In that case, set $\underline{Q}_m = \overline{Q}_m = Q_m$.

[14] At $\tau = 0$, a bidder draws an iid random variable $\boldsymbol{\Psi_i} \in [0, 1]$ which is one dimension of the bidder's private signal and thus unobservable to competitors. It corresponds to the mean of an iid Bernoulli random variable, $\boldsymbol{\Omega_i}$, which determines whether the bidder's later bids will make it in time to be accepted by the auctioneer. Thus, for $\tau > 0$, the bidder's information set includes the realizations $\omega_i \in \{0, 1\}$ of $\boldsymbol{\Omega_i}$, where $\omega_i = 1$ means that the bid of time τ will make it in time. This gives an incentive to bid at each arrival of new information because there might not be an opportunity to successfully bid in the future.

Assumption 4. *Supply* $\{\boldsymbol{Q_1}, \boldsymbol{Q_2}, \boldsymbol{Q_3}\}$ *is a random variable distributed on* $[\underline{Q}_1, \overline{Q}_1] \times [\underline{Q}_2, \overline{Q}_2] \times [\underline{Q}_3, \overline{Q}_3]$ *with strictly positive marginal density conditional on* $s^g_{i,\tau}$ $\forall i, g = c, d$ *and* τ.

If aggregate demand equals total supply exactly there is a unique market clearing price P^c_m. Each bidder wins their demand at the market clearing price and pays for all units according to their individual price offers. When there are several prices at which total supply equals aggregate demand by all bidders, the auctioneer chooses the highest one. Finally, in the event of excess demand at the market clearing price, bidders are rationed pro-rata on-the-margin.[15]

Denoting the amounts bidder i gets allocated by $q^c_i = (q^c_{1,i}\ q^c_{2,i}\ q^c_{3,i})$ when submitting $b^g_{i,\tau}(\cdot, \theta^g_{i,\tau}) \equiv (b^g_{1,i,\tau}(\cdot, \theta^g_{i,\tau})\ b^g_{2,i,\tau}(\cdot, \theta^g_{i,\tau})\ b^g_{3,i,\tau}(\cdot, \theta^g_{i,\tau}))$ his total surplus is

$$TS(b^g_{i,\tau}(\cdot, \theta^g_{i,\tau}), s^g_{i,\tau}) = V(q^c_i, s^g_{i,\tau}) - \sum_{m=1}^{3} \int_0^{q^c_{m,i}} b^g_{m,i,\tau}(x, \theta^g_{i,\tau}) dx \tag{4}$$

in the event in which τ is the time of his final bid, with $V(q^c_i, s^g_{i,\tau})$ given by $\frac{\partial V(q_m, q_{-m}, s^g_{i,\tau})}{\partial q_m} = v_m(q_m, q_{-m}, s^g_{m,i,\tau})$ in (3). It is the total utility he achieves from obtaining the amounts he wins minus the total payments he must make. Ex ante, when placing a bid, the bidder knows neither how much he will win nor at which price the market will clear. His optimal choice maximizes the expected total surplus.

Definition 2. *A Bayesian Nash equilibrium (BNE) is a collection of functions* $b^g_{i,\tau}(\cdot, \theta^g_{i,\tau})$ *that for each bidder* i *and almost every type* $\theta^g_{i,\tau}$ *at each time* τ *maximizes the expected total surplus,* $\mathbb{E}[TS(b^g_{i,\tau}(\cdot, \theta^g_{i,\tau}), s^g_{i,\tau})|\theta^g_{i,\tau}]$.

We focus on type-symmetric BNE of the auction game, in which bidders who are ex ante identical follow the same strategies. Dealers who draw the same signal employ the same function, and similarly for customers: $b^d_{i,\tau}(\cdot, \theta_{i,\tau}) = b^d(\cdot, \theta_{i,\tau})$ and $b^c_{i,\tau}(\cdot, \theta_{i,\tau}) = b^c(\cdot, \theta_{i,\tau})$ $\forall i, \tau$. Across bidder groups strategies might be asymmetric.

[15] "Under this rule, all bids above the market clearing price are given priority, and only after all such bids are satisfied, the remaining marginal demands at exactly price $P^c = p$ are reduced proportionally by the rationing coefficient so that their sum exactly equals the remaining supply. An alternative rationing rule would, for example, not give bids at higher prices priority." (Kastl (2011)). The rationing coefficient satisfies $R_m(P^c_m) = \frac{Q_m - TD^+_m(P^m_c)}{TD_m(P^c_m) - TD^+_m + (P^c_m)}$ where $TD_m(P^c_m)$ denotes the total demand at price P^c_m, and $TD^+_m(P^c_m) = \lim_{p_m \downarrow P^c_m} TD_m(p_m)$.

4.2 Estimation Strategy

To identify how much bidders are willing to pay, we first characterize an equilibrium by combining insights from Kastl (2011) and Wittwer (2020). We then assume that bidders in our data play this equilibrium and estimate the joint distribution of market clearing prices by extending resampling techniques developed by Hortaçsu (2002), Kastl (2011) and Hortaçsu and Kastl (2012) for isolated auctions. This allows us to back out each bidder's true willingness to pay from the equilibrium condition.

Equilibrium Condition. Bidding incentives in simultaneous discriminatory price auctions are similar to those in an isolated auction. In an isolated auction ($\delta = 0$), a bidder chooses his bids to maximize total surplus subject to market clearing. Similar to a first-price auction, he trades off the expected surplus on the marginal infinitesimal unit versus the probability of winning it (see Kastl (2017), p. 237 for more details).

The key difference when auctioned goods are interdependent ($\delta \neq 0$), is that the bidder takes this interconnection across auctions into account. His demand for one good, say q_1, now depends on how much of the other goods the bidder gets allocated, $v_1(q_1, q_{-1}, s^g_{1,i,\tau})$. Ideally, the bidder would want to condition his bid $b_{1,k}$ for amount $q_{1,k}$ on how much he will purchase of the other securities in equilibrium, $q^*_{-1,i} \equiv (q^*_{2,i}\ q^*_{3,i})'$. However, he cannot do this by the rules of the auction. Thus, he takes an expectation conditional on winning $q_{1,k}$—which happens when $b_{1,k} \geq \boldsymbol{P^c_1} > b_{1,k+1}$—and equates the expected marginal benefit, $\mathbb{E}\left[v_1\left(q_{1,k}, \boldsymbol{q^*_{-1,i}}, s^g_{1,i,\tau}\right) \middle| b_{1,k} \geq \boldsymbol{P^c_1} > b_{1,k+1}, \theta^g_{i,\tau}\right]$, with the marginal cost of changing the bid. Since the auctions clear separately, the cost is identical to the cost in an isolated auction, only that the market clearing prices follow a joint distribution.

Proposition 1 (Discriminatory price auctions)**.** *Consider a dealer i with private information $\theta^g_{i,\tau}$ who submits $\hat{K}_m(\theta^g_{i,\tau})$ steps in auction m at time τ. Under Assumptions 1-4 in any type-symmetric BNE every step k in his bid function $b^g_m(\cdot, \theta^g_{i,\tau})$ has to satisfy*

$$\tilde{v}_m(q_{m,k}, s^g_{m,i,\tau}|\theta^g_{i,\tau}) = b_{m,k} + \frac{\Pr\left(b_{m,k+1} \geq \boldsymbol{P^c_m}|\theta^g_{i,\tau}\right)}{\Pr\left(b_{m,k} > \boldsymbol{P^c_m} > b_{m,k+1}|\theta^g_{i,\tau}\right)}(b_{m,k} - b_{m,k+1})\ \forall k < \hat{K}_m(\theta^g_{i,\tau})$$

with $\tilde{v}_m(q_{m,k}, s^g_{m,i,\tau}|\theta^g_{i,\tau}) \equiv \mathbb{E}\left[v_m\left(q_{m,k}, \boldsymbol{q^*_{-m,i}}, s^g_{m,i,\tau}\right) \middle| b_{m,k} \geq \boldsymbol{P^c_m} > b_{m,k+1}, \theta^g_{i,\tau}\right]$ for all m with $-m \neq m$, and $b_{m,k} = \tilde{v}_m(\bar{q}_m(\theta^g_{i,\tau}), s^g_{m,i,\tau}|\theta^g_{i,\tau})$ at $k = \hat{K}_m(\theta^g_{i,\tau})$ where $\bar{q}_m(\theta^g_{i,\tau})$ is the maximal amount the bidder may be allocated in an equilibrium.

Resampling Procedure. To back out the bidders' valuations from the equilibrium conditions, we estimate the distribution of market clearing prices $\boldsymbol{P^c_{t,m}}$ and, equally important, the corresponding amount $\boldsymbol{q^*_{t,-m,i}}$ of each maturity bidder i would win at the market clearing price, by resampling. Unlike to the case of isolated auctions, the resampling must be done jointly for all—in our case three—maturities.

A natural starting point is to extend Hortaçsu (2002)'s resampling procedure: fix a triplet of bids simultaneously submitted by a bidder i and draw a random subsample of $N-1$ bid-vector triplets with replacement from the sample of bids.[16] From this, construct the bidder's realized residual supply for all maturities were others to submit these bids to determine the realized clearing prices $P^c = (P^c_{3M}\ P^c_{6M}\ P^c_{12M})$ and the amount $q^*_i = (q^*_{3M,i}\ q^*_{6M,i}\ q^*_{12M,i})$ this bidder would have won for all q^*_i, P^c. Repeating this procedure a large number of times provides an estimate of the joint distribution of market clearing prices and, equally important, the corresponding amount of each security i would win.

Our resampling procedure is more complex due to the fact that customers place bids via dealers (as in Hortaçsu and Kastl (2012)). With simultaneous auctions, there are two complications. A customer might bid via different dealers for different maturities, and two bids for the same maturity but by different customers might go through the same dealer. Neither of these cases happens more than a handful of times. Therefore, we assume that the information set of dealers who observe the same customer is independent across maturities, conditional on his own signal. In addition, we restrict the number of possible observed customer bids to two given that most customers only submit one bid and that there are many more dealers than customers in a typical auction.

Our procedure goes as follows: draw N_c customer bids from the empirical distribution of customer bids at date t. If a customer did not participate in one auction, replace his bid by 0. For each customer, find the dealer(s) who observed this customer's bid(s). If the customer submitted only one bid, we take the dealer who observed it. If the customer submitted more than one bid, draw uniformly over dealer-bids having observed this customer. Finally, if the total number of dealers drawn is at this point lower than the total number of potential dealers, draw the remaining bids from the pool of uninformed dealers, i.e., those who do not observe a customer bid in any of the three auctions. Note that—while theory allows for many updates—we restrict the number of possible observed customer bids to two in order to simplify our resampling algorithm. This includes most

[16] In practice, bids are not submitted at the exact same time given electronic or human delays (see the example in Appendix Table A1). We define bids to be "simultaneous" if they are the closest bids of all bids a bidder places within 200 seconds, or they are the last bids made before the auction deadline, i.e. final bids. Setting an upper bound of 200 seconds seems sensible when looking at the number of seconds between bids across maturities which we know were determined "simultaneously". Those are cases where the bidder does not update his bids over the course of the auctions. On average 551(383) seconds pass between such bids for different maturities by dealers (customers). Excluding outliers reduces the time (see Online Appendix Figure O2).

cases as most bidders only update once or twice.

Our resampling gives consistent estimates under two scenarios: in the benchmark, all bidders (customers and dealers) are ex ante symmetric. Dealers do not know whether their rivals have complementary, substitutable, or independent preferences for different maturities. This is plausible if we believe that these preferences are mostly driven by fluctuating factors in the secondary market. In the extended model presented in the online appendix, there are two latent dealer types. They consistently display different preferences, for example, because they follow different business models. Each dealer is aware of how many dealers are of each type but they do not know dealer identities.

Obtaining demand coefficients. With the estimated joint distributions, we can estimate how much each bidder expects to win of the other maturities $-m$ if he were to win a given quantity in maturity m:

$$\hat{\mathbb{E}}\left[\boldsymbol{q^*_{t,-m,i}}|...\right] = \mathbb{E}\left[\boldsymbol{q^*_{t,-m,i}}|b_{t,m,i,\tau,k} \geq \boldsymbol{P^c_{t,m}} > b_{t,m,i,\tau,k+1}, \theta^g_{t,i,\tau}\right] + \varepsilon^q_{t,m,i,\tau,k}. \quad (5)$$

Moreover, using Proposition 1, we can back out each bidder's valuations given the observed bids at all steps:

$$\hat{v}_{t,m,i,\tau,k} = \mathbb{E}\left[v_m\left(q_{t,m,i,\tau,k}, \boldsymbol{q^*_{t,-m,i}}, s^g_{m,i,\tau}\right)\middle| b_{t,m,i,\tau,k} \geq \boldsymbol{P^c_{t,m}} > b_{t,m,i,\tau,k+1}, \theta^g_{t,i,\tau}\right] + \varepsilon^v_{t,m,i,\tau,k}. \quad (6)$$

Finally, with (5), (6) and Assumption 2, we can estimate all demand coefficients by running the following regressions with auction-bidder-time fixed effects, $u_{t,m,i,\tau} = f_m(s^g_{t,m,i,\tau})$:

$$\hat{v}_{t,m,i,\tau,k} = u_{t,m,i,\tau} + \lambda_m q_{t,m,i,\tau,k} + \delta_m \cdot \hat{\mathbb{E}}\left[\boldsymbol{q^*_{t,-m,i}}|...\right] + \varepsilon_{t,m,i,\tau,k} \quad \forall m, m \neq -m \quad (7)$$

on a subsample with competitive bids with at least two steps. Appendix Figure A2 shows that it is the case for virtually all dealer bids: almost all submit more than one step.[17]

Given that the expected amounts, $\hat{\mathbb{E}}\left[\boldsymbol{q^*_{t,-m,i}}|...\right]$, are estimated and thus might be measured with error, the δ estimates or regression (7) might be downward biased. We illustrate how to grasp the size of the measurement bias in Online Appendix C. In our empirical application, this bias seems small.

[17]In Online Appendix Table O1 we show that our findings are robust when focusing on bids with at least 3 steps.

4.3 Estimation Findings: Demand Coefficients

We restrict attention to dealers, and impose valuations $\tilde{v}_m(\cdot, s^d_{m,i,\tau}|\theta^d_{i,\tau})$ to be weakly decreasing. Furthermore, to correct for outliers that occasionally occur due to small values of the denominator in the estimated (marginal) hazard rate of the market clearing price, $\hat{\Pr}\left(b_{m,k} > \boldsymbol{P^c_m} > b_{m,k+1}|\theta^g_{i,\tau}\right)$, we trim our estimated valuations. Specifically, we restrict each to be lower than the bidder's maximal bid plus a markup of about 5 bps (C\$500 for 12M, C\$250 for 6M, C\$125 for 3M).[18] This approach is conservative in light of the distribution of how bidders shade the untrimmed estimated valuations per step (see Online Appendix Figures O3). The less we trim, the larger, in absolute value, all coefficients (see Online Appendix Table O2).

Regressions with bids. As a starting point, we estimate all regressions (7) using observed bids rather than estimated valuations (see Table 5 (a) and Online Appendix Table O3). All δ coefficients are positive if statistically significant. This is similar to the reduced-form regression (1), reported in Table 3, which uses the amounts bidders actually win (ex-post) rather than the amounts they expect to win when bidding (ex ante). This means that the average dealer is willing to pay a higher price when winning more of the other maturities, implying that bills are complements. This is a somewhat surprising result—there is a long literature which classify securities of similar term and risk as substitutes.

Regressions with valuations. To determine if bid-shading leads to biased estimates, we re-estimate the regressions using the estimated valuations. We do this both with valuations expressed as prices (in C\$) and yields (in bps), but only report results for prices (see Table 5 (b)).[19] In contrast to the case of using bids, all δ coefficients are now negative, implying that bills are substitutes.[20] This highlights how important it is to eliminate bid-shading and use the true valuations to identify interdependencies. This is

[18] Note that 1 bp of a 12M T-bill with a face value of 1 mil corresponds to 1 mil/10,000=C\$100. Hence, 1 bp for a 3M T-bill corresponds approximately to C\$25 and for 6M T-bill to C\$50.

[19] Comparing the δ coefficients across auctions, we may notice that the estimates are not symmetric. For example, $\hat{\delta}_{3M,6M} \neq \hat{\delta}_{6M,3M}$. The main reason for this asymmetry is that the price of a bill mechanically increases as it approaches maturity. If we estimate the demand coefficients using yields we obtain δ estimates that are more symmetric across auctions up to some estimation error. We prefer to work with prices since it is more natural to think of demand schedules as downward sloping, especially moving to the counterfactual exercise and for other empirical applications. Alternatively, we could impose symmetry in the estimation.

[20] In previous versions of the paper, we falsely reported that bills are complements. This was because there was a typo in the estimation code due to which the estimated values were trimmed so much that they resembled the submitted bids. Therefore, the biased estimation results were similar to Table 5 (a).

Table 5: Demand coefficients

(a) With bids as independent variables

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	-5.033	(0.025)	λ_{6M}	-7.990	(0.046)	λ_{1Y}	-15.87	(0.084)
$\delta_{3M,6M}$	$+0.167$	(0.055)	$\delta_{6M,3M}$	$+0.435$	(0.101)	$\delta_{1Y,3M}$	-0.014	(0.212)
$\delta_{3M,1Y}$	$+0.411$	(0.059)	$\delta_{6M,1Y}$	$+0.737$	(0.110)	$\delta_{1Y,6M}$	$+1.557$	(0.214)
N	58542			42282			50408	

(b) With valuations as independent variables

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	-6.726	(0.033)	λ_{6M}	-11.53	(0.066)	λ_{1Y}	-24.03	(0.135)
$\delta_{3M,6M}$	-0.921	(0.073)	$\delta_{6M,3M}$	-2.343	(0.146)	$\delta_{1Y,3M}$	-6.317	(0.339)
$\delta_{3M,1Y}$	-0.140	(0.079)	$\delta_{6M,1Y}$	-0.514	(0.159)	$\delta_{1Y,6M}$	-2.561	(0.342)
N	58542			42282			50408	

Table 5 (a) reports the coefficients for equation (7), but with the observed competitive bids by dealers with more than one step as independent variables rather than the estimated true valuations. Table 5 (b) reports the coefficients with valuations. Bids and valuations are in C$ and quantities in % of auction supply. Alternatively, we could express bids in yields (bps) and quantities in units (C$). The findings are qualitatively the same. The first three columns show the estimates for the 3M Bill auction, the next three for the 6M Bill auction and the last three for the 12M Bill auction. The point estimates are in the second, fifth and eight column. Standard errors are in parentheses.

true even though we estimate small shading factors in absolute terms.

The magnitudes of all coefficients are relatively small in absolute terms, which is not surprising given that the bidding curves in bill auctions are very flat. For instance, when the dealer wins 1% more of the supply of the 6M bills, his price for the 6M bills decreases by λ_{6M} =C$11.53. Instead, if he wins 1% more of the supply of the 3M bills the price for the 6M bills decreases by $\delta_{6M,3M}$ =C$2.343, and of the 12M bills by $\delta_{6M,1Y}$ =C$0.514.

The δ estimates imply that the dealer's valuation for the 6M bill decreases by about C$2.343*7.3+C$0.514*6.9 $\approx$ C$20.65 when obtaining the average amount of the 3M (7.3% of supply) and 12M bills (6.9% of supply), rather than nothing. In the 3M auction the analogous price decrease is about C$0.921*6.7 + C$0.140*6.9$\approx$C$7.14 and in the 12M auctions about C$63.27. These price drops are not negligible in comparison to the difference between the maximal and minimal bid in the average bidding function, which is C$142.

Take away. Taken together, our analysis highlights that bills are at best imperfect substitutes, despite being cash-like. Using our extended model, presented in the online appendix, we find that larger dealers (market makers) view bills as substitutes, but that smaller dealers (non-market makers), may view bills are complements.

5 Policy Takeaway

Now we demonstrate how to use demand estimates to analyze how to split total supply across different goods strategically so as to maximize auction revenues in a discriminatory or uniform price auction.

We make two simplifications. First, we consider the supply split of two goods pairwise, and ask under what conditions total revenue increases when issuing a little bit more of one good, and a little bit less of the other, holding the total amount of supply (or production capacity) constant. Second, we abstract from non-competitive bids because they are not present at every multi-unit auction and because in our application the large majority of these are allocated to the auctioneer (Bank of Canada). Therefore, most of the revenue that the auctioneer collects from non-competitive bids comes from its own pocket, and can thus be viewed as revenue-neutral in-house transfer. It is straightforward to include non-competitive bids in the revenue calculations with our framework.

5.1 Theory: How to Allocate Goods?

Assume that there are only two goods, S for short and L for long, which are auctioned in two separate auctions. Three factors determine how to split total supply (or production capacity), across both goods: (i) price levels, (ii) price sensitivities, (iii) the auction format.

Price levels. When S predictably sells at a higher price but is not more costly to issue than L, it is revenue maximizing to only issue S. This case is trivial. We focus on the more interesting case where the predictable price difference is due to a difference in issuance cost: $c_S = P_S - P_L$. For instance, the price for salmon might always be higher than the price for halibut because it requires more expensive machinery to catch salmon. Then, even though the auction for salmon will clear at a higher price, the fisher does not necessarily make more money if he only fishes salmon.

In our empirical application, shorter bonds are typically sold at higher prices than longer bond, $P_S > P_L$.[21] Thus, if the government's objective were to maximize the auction revenue on a single day, it should issue only the short bond. In practice, however, governments do not take this strategy. They seek, instead, to maximize revenues over

[21] In normal times, the yield curve of government bonds is upward sloping, implying that bond prices decline in term to maturity.

a planning horizon—typically one year. Doing so, they take into account that the short bond must be rolled-over more frequently than the long bond to maintain the same level of expenditures, all else equal. Rolling over debt is costly not only because it involves running more auctions, but also because it is risky. For instance, if the level of interest rates in the economy suddenly increases, future auction revenue is lower than expected. To a first-order approximation, the issuance cost must be such that the government does not want to issue more of the short bond and less of the long bond only to cash in the higher price of the short bond: $c_S = P_S - P_L$.

Price sensitivity. The price sensitivity tells us by how much the market price changes in response to a 1% change in supply. It is the inverse of the price elasticity of aggregate demand. Formally, to fix ideas, let $P_m(Q_m)$ sum all bids for good $m \in \{S, L\}$ per unit of supply. Assume, for illustration, that the market clears at $\{P_m, Q_m\}$ and that $P_m(\cdot)$ is differentiable (which is not the case in the data). Then the price sensitivity is $\frac{\partial P(Q_m)}{\partial Q_m}\frac{Q_m}{P_m}$.

The important feature is that this market price sensitivity not only depends on the bidders' price sensitivity when winning more in the auction—the own good effect (the λ's)—but also on how this sensitivity changes when winning more of the other goods—the cross-good effect (the δ's).

In our case, the aggregate demand for the long bond is typically more price-sensitive than for the short bond, which means that the price for the long bond responds more strongly to a change in auction supply than the price for the short maturity. This is true when bonds are independent and when they are substitutes. It may not hold when they are highly complementary—a case we exclude from our discussion since it seems not to be empirically relevant.

Auction format. In a uniform price auction, the difference in price sensitivities implies that the auctioneer can increase total auction revenue by issuing less of the price-sensitive good (here the long bond) and more of the price insensitive good (here the short bond), without changing total supply (see Figure 2 (a)-(b)). The reason is that everyone pays the market prices, and the market price of the long bond increases more strongly than the market price of the short bond decreases.

However, there is a price-quantity trade-off. Starting from an equal supply split across goods, when the auctioneer moves one dollar from the long into the short bonds, the price of the short bond drops less than the price of the long bond increases. Thus, while the revenue of the short bond auction decreases, the revenue of the long bond auction increases by more. Total revenue increases. Yet, the more debt is issued as short rather than long, the lower the revenue gain in the long bond auction given that the higher price

Figure 2: Simplified example of revenue under two auction formats

(a) Uniform price auction: S

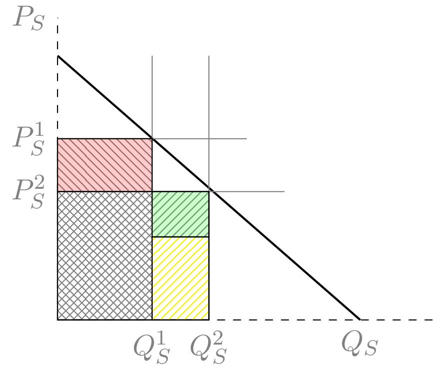

(b) Uniform price auction: L

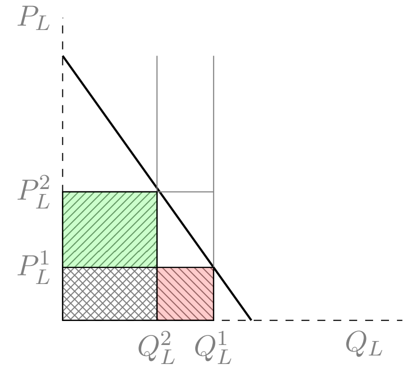

(c) Pay-as-bid auction: S

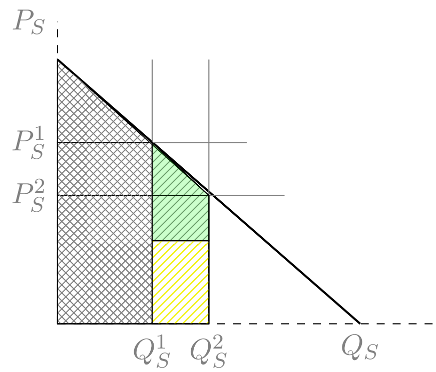

(d) Pay-as-bid auction: L

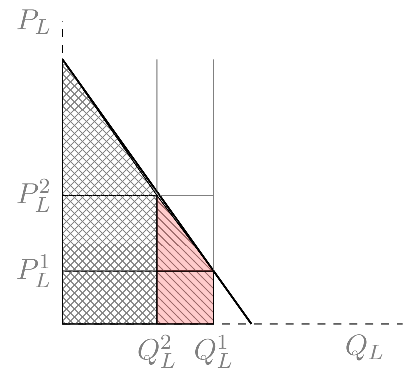

The figures on the LHS show the revenue gain (in green) and loss (in red) when issuing dQ more of S depending on the auction format, in addition to the change in total issuance cost (in yellow). The figures on the RHS show the analogous changes in revenue when issuing dQ less of L. In all cases, we assume that $\frac{\partial P_m(Q_m)}{\partial Q_m} = -\mu_m$ for $m \in \{S, L\}$ does not change. Formally, before the change in supply, $Q_S^1 = Q_L^1 = Q$, $P_S^1 > P_L^1$, $c_S^1 = P_S^1 - P_L^1$. After issuing dQ more for S and dQ less of L, $Q_S^2 = Q + dQ$, $Q_L^2 = Q - dQ$ $P_S^2 = P_S^1 - \mu_S dQ$, $P_L^2 = P_L^1 + \mu_L dQ$. In the uniform price auction, the total change in revenue is $[-\mu_S(Q+dQ)+\mu_L(Q-dQ)]dQ > 0$ when dQ is small, $dQ > 0, \mu_L > \mu_S$. In the discriminatory price auction it is $[-\mu_S/2dQ - \mu_L/2dQ]dQ < 0$ when $dQ > 0, \mu_L > 0, \mu_S > 0$.

for the long bond is multiplied by a smaller and smaller amount. In the extreme, when the auctioneer goes from a mixed supply split to issuing only short bonds, no one pays the hypothetical high price for the long bond that would clear the market, and therefore total revenue decreases.

A similar price-quantity trade-off can arise in the discriminatory price auction (see Figure 2 (c)-(d)). There are two differences. First, shifting supply towards the short bond may decrease total revenue. Second, the revenue of one auction is determined by the area underneath the aggregate demand curve. The key is that this area shrinks in the long bond auction by more than it increases in the short bond auction when decreasing the supply of the long bond—unless bidders adjust the price offers for small amounts of the bond.

If aggregate demand curves were linear as in Figure 2 and no bidder adjusted bids given the new supply quantities, we could formalize the price-quantity trade-off for both auction formats, and determine the revenue-maximizing supply split. In reality, the optimal supply split cannot be determined as easily.

From theory to practice. In practice, there are two complications. First, bidders respond to changes in supply. Therefore, the aggregate demand curves change. This is especially important when the auction is discriminatory price since the auction revenue is determined by the shape of the entire aggregate demand curve, and not only the market clearing price. Take Figure 2 (c)-(d), as an example. Due to the change in the aggregate demand curve, it is actually not true that the gray area is the same before and after the change in supply. Generally, it is an empirical question as to whether total revenue increases or decreases because the theoretical effect is ambiguous.

Second, bidders submit step functions based on their individual willingness to pay and shade their bids. This implies that it is not straightforward to compute the steepness of the aggregate demand curves. These curves have steps and cannot be constructed based on any single parameter (such as the λ's) that we can estimate.

5.2 Empirical Implementation

Now we illustrate how to compute revenue gains or losses as proof-of-concept to qualitatively test the insights from our theoretical model. For this, we use demand schedules with prices and quantities, both expressed in C\$.[22] We focus on reshuffling the 6M and 12M bills and keep the 3M bills at the observed amount, but illustrate how sensitive these revenue gains are when the aggregate demand is more price sensitive—as is the case for bonds with longer maturities than 12 months. Reshuffling supply from the 12M to 3M bills leads to slightly higher revenue impact since demand for the 3M bills is less price sensitive than demand for the 6M bills.

Bids and aggregate submitted demand. It is still an open question in the literature on how to characterize equilibrium strategies in discriminatory price auctions that are sufficiently complex to capture real world markets. Our idea is to extrapolate from the

[22] Alternatively, we could use demand schedules with yields. The key insights are independent of the units of measurements we use.

observed shading factors to the counterfactual ones, given that there are by now a fair number of papers that find shading factors of similar magnitudes for different settings (e.g., Kang and Puller (2008); Kastl (2011); Hortaçsu et al. (2018)).[23]

We assume that the shading factor changes sufficiently little when going from the status quo to the counterfactual and approximate the counterfactual (final) bid for amount q_m of a bidder i for maturity m on day t by his demand minus the fixed shading factor:

$$b^{cf}_{t,m,i}(q_m) = \hat{v}^{cf}_{t,m,i}(q_m) - \hat{s}hading_{t,m,i}(q_m) \tag{8}$$

$$\text{with } \hat{v}^{cf}_{t,m,i}(q_m) = \hat{u}_{t,m,i} + \hat{\lambda}_m q_m + \hat{\delta}_m \cdot \hat{\mathbb{E}}[\boldsymbol{q^{cf*}_{t,-m,i}}|q_m] \tag{9}$$

$$\text{and } \hat{s}hading_{t,m,i}(q_m) = \hat{v}_{t,m,i}(q_m) - b_{t,m,i}(q_m)\ \forall m, i. \tag{10}$$

The counterfactual aggregate (submitted) demand sums across these bids: $P_{t,m}(Q_m) = \sum_i b^{cf}_{t,m,i}(q_m)$. Here, $\hat{v}_{t,m,i}(q_m)$ is estimated demand for amount q_m, and $b_{t,m,i}(q_m)$ is the observed (final) bid. $\hat{v}^{cf}_{t,m,i}(q_m)$ and $b^{cf}_{t,m,i}(q_m)$ are the counterfactual demand and (final) bid. Both depend on the slope parameters, $\hat{\lambda}_m$ and $\hat{\delta}_m$, the estimated fixed effect, $\hat{u}_{t,m,i}$, and on the amount the bidder expects to win in the counterfactual, $\hat{\mathbb{E}}[\boldsymbol{q^{cf*}_{t,-m,i}}|q_m]$.

$\hat{\mathbb{E}}[\boldsymbol{q^{cf*}_{t,-m,i}}|q_m]$ depends on how everyone bids in an auction, and thus can be found by solving a fixed point problem. Solving this problem is computationally intensive since it involves finding a fixed point for each bidder and each auction. To reduce the computational complexity, we show by means of examples that $\hat{\mathbb{E}}[\boldsymbol{q^{cf*}_{t,-m,i}}|q_m]$ is typically very close to the amount one obtains when rescaling the original expectation by the factor by which total supply of m is changed in the counterfactual, and approximate the fixed points by rescaling (see Appendix D for details).

To highlight how important the market price sensitivity is when determining the supply split, we illustrate how our findings change as price sensitivities become larger by rescaling all λ and δ parameters by factors, for instance by 10. In our application, we expect sensitivities to be higher for longer-term bonds because the observed bidding curves become steeper in maturity length (see Table 6).[24] For instance, we use a λ-factor of 10 to approximate the bids for the 5Y and 10Y bonds, combined with varying δ-factors from 0 (independent) up to high enough to make bonds perfect substitutes (i.e., set the δ's equal to the λ's, e.g., $\lambda_{6M} = \delta_{6M,3M} = \delta_{6M,12M}$). Given that all λ estimates for bills are under-estimated when using bids rather than valuations (recall Table 5), we conjecture

[23]This assumption is stronger when we switch the auction format or scale the demand coefficients. As a robustness check we verify that our qualitative findings go through when we abstract from bid-shading and assume that bidders submit their true demands as is the case in a perfectly competitive auction.

[24]We cannot estimate the demand systems for bonds of different maturities because these are sold on different days, which implies that we cannot directly implement the method developed in this paper.

Table 6: Slopes of bidding curves per maturity

	3M	6M	12M	2Y	3Y	5Y	10Y	30Y
q_m	-0.0710	-0.297	-0.534	-1.319	-2.562	-4.444	-11.34	-46.03
	(0.000342)	(0.00171)	(0.00280)	(0.0120)	(0.0316)	(0.0363)	(0.0839)	(0.438)
N	68045	42224	56880	29804	11413	25421	23246	13195

Table 6 shows the estimated slope of an average bidding curve for each maturity length that Canada issues, using data of all final bids submitted in all Canadian Treasury Bill and Bond auctions from 2002 until 2015. Each column displays the β estimate when regressing the submitted price bids of a bidder in an auction for maturity m on the quantity that this bidder asked for at that price and a bidder-auction fixed effect: $b_{t,i,m,k} = u_{t,i,m} + \beta q_{t,i,m,k} + \epsilon_{t,i,m,k}$. Bids and quantities are in million C$ to facilitate revenue calculations. Note that the units are different in Table 5 so that the numbers are not directly comparable. Standard errors are in parentheses.

that using factors based on bidding data provides lower bounds of all effects.

Issuance costs. To identify the indirect revenue effect that comes from different market price sensitivities, we eliminate the mechanical price effect that makes short bonds achieve higher revenues than long bonds by including an issuance cost.[25] We compute the prices at which an auction clears when bidders bid as in (8) and the government issues the observed supply, normalize the cost of the long maturity to zero, $c_{t,L} = 0$, and define the cost of the short maturity relative to the long maturity of an auction as $c_{t,S} = P^c_{t,S} - P^c_{t,L}$.[26] When we make out-of-sample statements about bonds, we recompute the cost as the difference between the market clearing prices that arise if the λ's and δ's in the bidders' willingness to pay was scaled up by a specific set of factors.

Revenue gains. To quantify how much revenue can be gained when moving slightly away from the observed supply split, we compute by how much the revenue of one auction day ($Revenue_t$) changes when issuing 1% more of total debt in form of the short maturity and 1% less of the long maturity, and vice versa.

$$Revenue_t = \begin{cases} \sum_{m\in\{S,L\}} \sum_{i=1}^{N_{t,m}} \int_0^{q^*_{t,m,i}} (b^{cf}_{t,m,i}(q_m) - c_{t,m}) dq_m & \text{if discriminatory price} \\ \sum_{m\in\{S,L\}} (P^c_{t,m} - c_{t,m}) Q_{t,m} & \text{if uniform price,} \end{cases}$$

[25] It is an open question in the literature as to how to estimate issuance costs for government debt. Providing a precise answer to this question is beyond the scope of this paper.

[26] Alternatively, we could compute the costs that rationalize the supply split that we observe in the data, assuming that the Bank of Canada chooses the supply split that maximizes the revenue of an auction day, or on average in a year. These cost-estimates are similar to the ones we pick. We prefer our more transparent approach to eliminate the mechanical price effect.

Figure 3: Example on an auction

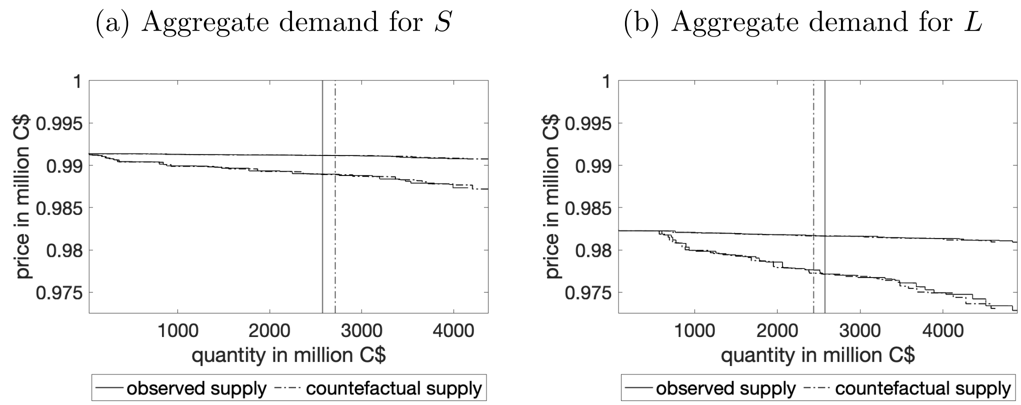

Figure 3 shows aggregate demand curves for two auctions that took place on the same day at some point in our sample. Each graph plots four curves. Two of the curves look rather flat. In 3a, the flat curves correspond to the aggregate demand for 6M bills when the Bank of Canada issues the supply as we observe it, and when we increase the supply of the 6M bills by 1% of the total debt issued on that auction day. In 3b, we see the same curves but for the 12M bills. The steeper curves correspond to the aggregate demand curves when scaling the λ parameters by a factor of 10 and making bills perfect substitutes. Here we can see how the aggregate demand curve changes in responds to the change in supply.

where $N_{t,m}$ is the observed number of bidders who participate in the auction for maturity m on day t, $b_{t,m,i}(q_m)$ is a bid for amount q_m of a bidder i, $q^*_{t,m,i}$ is the amount this bidder wins at market clearing, $P^c_{t,m}$ is the market clearing price, $c_{t,m}$ is the maturity's issuance cost, and $Q_{t,m}$ is the supply issued to competitive bidders.

We measure the gain (or loss) in revenue in bps of the revenue before reshuffling supply locally. For example, a revenue gain of 1 bps means that the government earned 0.01% more money in a single auction.

5.3 Counterfactual Findings

Example. We start with an example shifting supply between the $S = 6M$ and $L = 12M$ bills in one auction in our sample (see Figure 3). Given the observed supply, the 6M auction clears at $P^1_S = C\$991,162$ and $P^1_L = C\$981,627$ and $Q^S = Q^L = 2.575$ billion. Shuffling 1% of total debt between the 6M and 12M auction, we get the following prices: $P^2_S = P^1_S - 5, P^2_L = P^1_L + 25$. The revenue gain is small: +0.09 bps in a uniform price auction and −0.04 bps in a discriminatory price auction, similar to Figure 2.

To illustrate the impact of reshuffling longer-dated debt, for example, 5Y and 10Y bonds, let us scale up all λ's by a factor of 10. Now, a uniform price auction achieves a larger revenue gain—both when bonds are independent (+0.32 bps) and when bonds are

Table 7: Average gain (in bps) per auction when reshuffling 1% of debt

Demand coefficients		$S\uparrow L\downarrow$ Uniform	$S\uparrow L\downarrow$ Discrim	$S\downarrow L\uparrow$ Uniform	$S\downarrow L\uparrow$ Discrim
Independent:	factor$_\lambda$=1, factor$_\delta$=0	+0.020	+0.007	−0.023	−0.010
Weak substitutes:	factor$_\lambda$=1, factor$_\delta$=1	+0.016	−0.002	−0.024	+0.001
Perfect substitutes:	factor$_\lambda$=1, $\delta=\lambda$	+0.011	−0.052	−0.016	+0.048
Independent:	factor$_\lambda$=10, factor$_\delta$=0	+0.234	−0.028	−0.297	+0.007
Weak substitutes:	factor$_\lambda$=10, factor$_\delta$=1	+0.225	−0.036	−0.292	+0.016
Perfect substitutes:	factor$_\lambda$=10, $\delta=\lambda$	+0.119	−0.609	−0.189	+0.590
Independent:	factor$_\lambda$=100, factor$_\delta$=0	+2.344	−0.446	−2.9757	+0.191
Weak substitutes:	factor$_\lambda$=100, factor$_\delta$=1	+2.341	−0.455	−2.970	+0.200
Perfect substitutes:	factor$_\lambda$=100, $\delta=\lambda$	+1.313	−6.720	−1.956	+6.624

Table 7 shows the revenue gains when issuing 1% of debt more for the short maturity and 1% less of the long maturity in the second and third column ($S\uparrow L\downarrow$) and vice versa in the fourth and fifth column $S\downarrow L\uparrow$ when the auction format is uniform price (Uniform) and when it is discriminatory price (Discrim). The first three rows (factor$_\lambda$=1) correspond to the demand estimates of the 6M and 12M bills assuming different degrees of substitution. The fourth-sixth row and seventh-ninth row correspond to hypothetical auctions in which the λ parameters in the bidder's demand are scaled by a factor of 10, and 100, respectively. The revenue gain is in bps of the original revenue. Online Appendix Table O6 shows the analogous revenue gains in the extended model with heterogeneous dealers.

perfect substitutes (+0.10 bps)—because the aggregate demand curve is steeper than it is at the estimated λ. The revenue loss in a discriminatory price auction is also larger than before. If bonds are independent, the loss is −0.25 bps; when bonds are perfect substitutes, the loss is −1.55 bps. The non-negligible difference between these predictions highlights the importance of taking substitution patterns into account when comparing revenues across auction formats.

Average revenue gains. On average, it is revenue-increasing to issue more of the more price-sensitive bond (here the long maturity) and less of the more price-insensitive bond (here the short maturity) in a discriminatory price auction, and vice versa in a uniform price auction (see Table 7). The revenue effect increases when scaling up the λ's since the aggregate demand curves become more price-sensitive. This suggests that it might pay off to reshuffle supply across longer bonds.

It is important to have a good understanding of the degree of substitutability of different maturities before changing supply. In particular, when the auction format is discriminatory price, we over-estimate the revenue effect when assuming that different maturities are perfect substitutes, and under-estimate the effect when we assume they are independent (see Appendix Figure A3). This is because the revenue effect is determined by the shape of the entire aggregate demand curve and not just the point at which the market clears.

Price-quantity trade-off. So far, we have considered relatively moderate changes in supply. Next, we present graphically the price-quantity trade-off described above, which pins down the (two-dimensional) revenue-maximizing supply split. For illustration, we consider one auction day in our sample. The qualitative findings of other auction days are similar. Further, when scaling up the λ or δ parameters, the price-quantity trade-off is more pronounced.

When the auction is uniform price (as in Figure 4a), revenue increases when going from issuing no short bonds to issuing some short bonds until 61% of debt is issued as short and 39% as long. Until that point, the positive price effect dominates the negative quantity effect in the auction for the long bonds. When further increasing the supply of the short bond and decreasing the supply of the long bond, the negative quantity effect dominates and total revenue decreases.

In the discriminatory price auction, we see a similar pattern (see in Figure 4b). The difference is that the highest revenue gain is achieved when issuing less of the short (39%) and more of the long bond (61%).

Back-of-the-envelope calculation. We conclude the discussion with a back-of-the-envelope calculation to get a rough sense of how much the Canadian government could save if it changed its current supply split only marginally. For illustration, we consider the issuance in 2021. In 2021 the Canadian government issued C$416 billion in form of bills and C$277 billion in form of bonds. Taking Table 7 at face-value, issuing slightly more of the longer maturities would have brought a revenue gain of +0.001 bps per bill-auction and roughly +0.02 bps per-bond auction, assuming bonds are weak substitutes. This sums to moderate savings of C$595,600.

In other markets, in which demand is more price-sensitive, savings would be larger. For example, Albuquerque et al. (2022) estimate an average price elasticity of demand of -379 in Portuguese bond auctions between 2014 and 2019, implying a price sensitivity of 1/-379 $= -0.0026$. In comparison, with Canadian data, we find an average price sensitivity of -0.0012 if we, like Albuquerque et al. (2022), assume that securities are independent. The cross-country difference suggests that cost savings could be larger in the Portuguese market. Further, savings increase when considering larger changes in supply. For instance, Bigio et al. (2021) study a one percentage point increase in monthly issuances over annual GDP in the Spanish primary market. They find that this reduces auction prices between 8 bps for the 3 year bonds and 56 bps for the 30 year bonds. Exploiting this difference in

Figure 4: Illustration of the price-quantity trade-off

(a) Uniform price auction (b) Discriminatory price auction

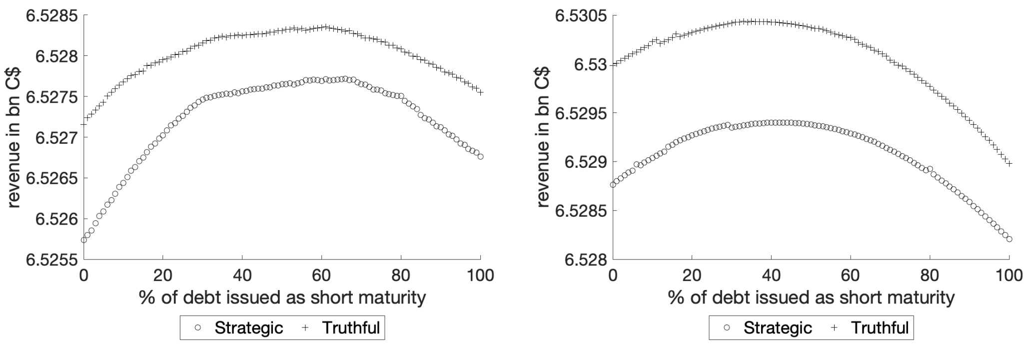

Figures 4 depict the price-quantity trade-off when the auction is uniform price (a), and discriminatory price (b) using the estimated λ and δ parameters. On the y-axis is the total revenue earn from issuing both maturities (in billion C$) when issuing x% of the short maturity and (1-x)% of the long maturity. The x-axis scales up x from 0% to 100%. Online Appendix Figure O5 shows the analogous trade-off in the extended model with heterogeneous dealers.

price sensitivity could lead to sizable annual cost savings for taxpayers.

Take away. We introduce a simple framework to guide auctioneers in their decision on how to split goods across auctions without changing the auction format since this can be difficult in practice. The key idea is that the auctioneers should behave like a monopolist who price discriminates subject to the auction rules. We show that it is generally revenue-increasing to issue more of the relatively price-insensitive good and less of the price-sensitive good when the auction is uniform price, and vice versa when it is discriminatory price.

6 Conclusion

We leverage the institutional feature that many multi-unit auctions of related goods are held in parallel with an overlapping set of bidders to develop a new methodology on how to estimate demand systems for multiple goods that can account for any degree of substitution or complementarity between goods. We illustrate how to use these demand systems to better target the auctioneers' objective. In our empirical application (Canadian Treasury auctions), the objective is to maximize total auction revenue, which implies lower financing costs of debt for tax payers. In other settings, the objective might be different, yet it would still depend on the estimated demand systems.

References

Albuquerque, R., Cardoso-Costa, J., and Faias, J. A. (2022). Price elasticity of demand and risk-bearing capacity in sovereign bond auctions. Working paper.

Allen, J. and Wittwer, M. (2021). Centralizing over-the-counter markets? Working paper.

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

Back, K. and Zender, J. F. (1993). Auctions of divisible goods: On the rationale for Treasury experiment. *The Review of Financial Studies*, 6(4):733–764.

Bank of Canada (2015). Statement of policy governing the acquisition and management of financial assets for the Bank of Canada's balance sheet.

Bank of Canada (2016). Standard terms for auctions of Government of Canada securities.

Barbosa, K., Silva, D. G. D., Yang, L., and Yoshimoto, H. (2022). Auction mechanisms and Treasury revenue: Evidence from the Chinese experiment. *American Economic Journal: Microeconomics*, forthcoming.

Belton, T., Dawsey, K., Greenlaw, D., Li, H., Ramaswamy, S., and Sack, B. (2018). Optimizing the maturity strucutre of U.S. Treasury debt: A model-based framework. Working paper.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.

Bhandari, A., Evans, D., Golosov, M., and Sargent, T. J. (2019). The optimal maturity of government debt. Working paper.

Bigio, S., Nuño, G., and Passadore, J. (2021). Debt-maturity management with liquidity costs. Working paper.

Bikhchandani, S. and Huang, C.-f. (1989). Auctions with Resale Markets: An Exploratory Model of Treasury Bill Markets. *The Review of Financial Studies*, 2(3):311–339.

Bonaldi, P. and Ruiz, A. (2021). Quantifying private information about fundamentals in Treasury auctions. Working paper.

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

Carleton, C. (2000). Electronic fish auctions. Strategies for securing and maintaining comparative advantages in the seafood trade. Nautilus Consultants, Report.

Cassola, N., Hortaçsu, A., and Kastl, J. (2013). The 2007 subprime market crisis in the EURO area through the lens of ECB repo auctions. *Econometrica*, 81(4):pp. 1309–1345.

Cole, H., Neuhann, D., and Ordonez, G. (2022). Asymmetric information and sovereign debt: Theory meets Mexian data. *Journal of Political Economy*, page forthcoming.

Cramton, P., Dinkin, S., and Wilson, R. (2013). A competitive sales process for BHB Billitons Ekati diamonds. In: The Handbook of Market Design, Z. Neeman, A. Roth and N. Vulkan (eds.), Oxford University Press.

D'Amico, S., English, W., Lopes-Salido, D., and Nelson, E. (2012). The Federal Reserve's large-scale asset purchase programs: Rationale and effects. *Economic Journal*, 122(564):F415.

Gagnon, J. E., Raskin, M., Remache, J., and Sack, B. P. (2011). The financial market effects of the federal reserve's large-scale asset purchases. *International Journal of Central Banking*, 7(3-43).

Gentry, M., Komarova, T., and Schiraldi, P. (2022). Preferences and performance in simultaneous first-price auctions: A structural analysis. *The Review of Economic Studies*, 0:1–27.

Gentzkow, M. (2007). Valuing new goods in a model with complementarity: Online newspapers. *American Economic Review*, 97(3):713–744.

Greenwood, R., Hanson, S. G., Rudolph, J. S., and Summers, L. (2015a). The optimal maturity of government debt. In: The $13 trillion question: How America manages its debt, edited by David Wessel, 1-41. Brookings Institution Press.

Greenwood, R., Hanson, S. G., and Stein, J. C. (2015b). A comparative-advantage approach to government debt maturity. *The Journal of Finance*, 70(4):1683–1722.

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

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

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

Hortaçsu, A., Kastl, J., and Zhang, A. (2018). Bid shading and bidder surplus in the U.S. Treasury auction. *American Economic Review*, 108(1):147–69.

Hortaçsu, A. and Puller, S. L. (2008). Understanding strategic bidding in multi-unit auctions: a case study of the Texas electricity spot market. *The RAND Journal of Economics*, 1:86–114.

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

Kastl, J. (2011). Discrete bids and empirical inference in divisible good auctions. *The 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.

Klemperer, P. (2010). The product-mix auction: a new auction design for differentiated goods. *Journal of the European Economic Association*, 8:526–536,.

Koijen, R. S. J. and Yogo, M. (2019). A demand system approach to asset pricing. *Journal of Political Economy*, 127(4):1475–1515.

Krishnamurthy, A. and Li, W. (2022). The demand for money, near-money, and Treasury bonds. Working paper.

Krishnamurthy, A. and Vissing-Jørgensen, A. (2011). The effects of quantitative easing on interest rates: Channels and implications for policy. *Brookings Papers on Economic Activity*.

Krishnamurthy, A. and Vissing-Jørgensen, A. (2012). The aggregate demand for Treasury debt. *Journal of Political Economy*, 120(2):233–267.

Krishnamurthy, A. and Vissing-Jørgensen, A. (2015). The impact of Treasury supply on financial sector lending and stability. *Journal of Financial Economics*, 118:571–600.

Larsen, B. and Zhang, A. L. (2022). Quantifying bargaining power under incomplete information: A supply-side analysis of the used-car industry. Working paper.

Lou, D., Yan, H., and Zhang, J. (2013). Anticipated and repeated shocks in liquid markets. *The Review of Financial Studies*, 26(8):1891–1912.

Missale, A. and Blanchard, O. J. (1994). The debt burden and debt maturity. *American Economic Review*, 84(1):309–319.

Nagel, S. (2016). The liquidity premium of near-money assets. *The Quarterly Journal of Economics*, 131(4):1927–1971.

Pycia, M. and Woodward, K. (2022). Auctions of homogeneous goods: A case for pay-as-bid. Working paper.

Rostek, M. and Yoon, J. (2021a). Exchange design and efficiency. *Econometrica*, 89:2887–2928.

Rostek, M. and Yoon, J. H. (2021b). Design of synthetic financial products in decentralized markets. Woring paper.

Rostek, M. and Yoon, Y. H. (2022). Innovation in decentralized markets: Trading technology vs. synthetic products. Working paper.

Ryan, N. (2022). Holding up green energy. Working paper.

Somaini, P. (2020). Identification in auction models with inderdependent costs. *Journal of Political Economy*, 128(1):3820–3871.

Vayanos, D. and Vila, J.-L. (2021). A preferred-habitat model of the term structure of interest rates. *Econometrica*, 89(1):77–112.

Wilson, R. (1979). Aucitons of shares. *The Quarterly Journal of Economics*, 93(4):675–689.

Wilson, R. (2021). Strategic analysis of auctions. *Econometrica*, 89(2):555–561.

Wittwer, M. (2020). Interconnected pay-as-bid auctions. *Games and Economic Behavior*, 121:506–530.

Wittwer, M. (2021). Connecting disconnected financial markets? *American Economic Journal: Microeconomics*, 13(1):252–82.

Appendix

A Micro-Foundation of Demand

We present a model that helps explain what might drive the demand of a bidder who plans to sell some of its auction purchases to clients after the auction. In line with our empirical application, we call the bidder a dealer and think of his clients as investors. However, the model could be adjusted to fit other resale settings.

We view our model as one possible micro-foundation that is by no means exclusive. In our application, there are other reasons that might drive interdependencies and heterogeneities in dealer demand across maturities. For instance, dealers face regulatory constraints that could affect their needs across different maturities (or their budget constraints for these auctions).

Our model features market segmentation in the spirit of Vayanos and Vila (2021). Investors/clients may have preferences for specific maturities and dealers function across maturities by participating in the primary market and making markets in secondary trading. For simplicity, we restrict the number of maturities to $M = 2$, and drop the superscript g and the subscripts i, τ for the remainder of the section with exception of the formal statements.[27]

Each dealer has a type $\boldsymbol{s}$, which decomposes into ν (known by all bidders) and $\boldsymbol{t}$ (private information):

$$\boldsymbol{s} = (\boldsymbol{t}, \nu) \text{ with } \boldsymbol{t} = (\boldsymbol{t_1}, \boldsymbol{t_2}) \text{ and } \nu = (a, b, e, \gamma, \kappa_1, \kappa_2, \rho).$$

Rather than assuming that dealers are risk-averse, we assume that dealers face a cost of not meeting client demand.[28] A dealer who draws type s obtains the following gross benefit from "consuming" amounts $(1 - \kappa_1)q_1$ and $(1 - \kappa_2)q_2$:

$$U(q_1, q_2, s) = t_1(1 - \kappa_1)q_1 + t_2(1 - \kappa_2)q_2. \tag{11}$$

The private type determines how much a dealer benefits from keeping a share $(1 - \kappa_m) \in [0, 1)$ of the purchased bill m in his own inventory or to fulfill existing customer orders. Dealers function as market makers in the secondary market where they distribute the rest of the bills $\{\kappa_1 q_1, \kappa_2 q_2\}$ among investors who are yet to arrive. To incorporate future resale opportunities we let there be a second stage following the primary auction.

In the secondary market a (mass of) client(s) with random demand $\{\boldsymbol{x_1}, \boldsymbol{x_2}\}$ arrives to the dealer.[29] Equivalently, you may imagine that there are two types of clients, each with a random demand for one of the two maturities. We assume that each of $\{\boldsymbol{x_1}, \boldsymbol{x_2}\}$ is on-the-margin uniformly distributed on $[0, 1]$ but allow both amounts to be correlated. More specifically, $\{\boldsymbol{x_1}, \boldsymbol{x_2}\}$ assumes the following (Farie-Gumbel-Morgenstern cupola) density $f(x_1, x_2) = 1 + 3\rho(1 - 2F_1(x_1))(1 - 2F_2(x_2))$ with marginal distributions $F_m(x_m) = x_m$ and correlation parameter $\rho \in \left[-\frac{1}{3}, +\frac{1}{3}\right]$.

The dealer sells to clients who arrive as long as there is enough for resale: $x_m \leq \kappa_m q_m$. Selling x_m brings a payment of $p_m x_m$. The prices depend on the clients' willingness to pay, or the aggregate demand in the secondary market more generally. For simplicity we assume that it is linear and symmetric across maturities. The inverse demand schedule for maturity 1 in the secondary market takes the following form:

[27] Generalizing to more than two maturities is straightforward but mathematically cumbersome and brings no major additional insights.

[28] A practical reason for why we model dealers as risk neutral is that it is much harder to estimate auction models with risk-averse bidders than having a cost of not meeting demand.

[29] The terms "client" and "customer" denote different players. Customers participate in the auction by placing bids with dealers, while clients buy in the secondary market.

$$p_{i,1}(x_1, x_2|q_1, q_2) = \begin{cases} a - bx_1 - ex_2 & \text{for } x_1 \leq \kappa_1 q_1 \text{ and } x_2 \leq \kappa_2 q_2 \\ a - bx_1 & \text{for } x_1 \leq \kappa_1 q_1 \text{ and } x_2 > \kappa_2 q_2 \\ 0 & \text{for } x_1 > \kappa_1 q_1 \text{ and } x_2 > \kappa_2 q_2. \end{cases} \tag{12}$$

The price function for maturity 2 is analogous. It splits into three cases. In the first, clients for both bills arrive and the dealer has enough of both in their inventory. The dealer charges a bundle price of $\{p_1(x_1, x_2|q_1, q_2), p_2(x_1, x_2|q_1, q_2)\}$ for selling $\{x_1, x_2\}$. In the second case the dealer can only sell maturity 1. This might be because only clients with demand for this maturity arrive or because the dealer does not have enough of the other maturity in inventory for resale, $x_2 > \kappa_2 q_2$. The price the dealer charges is independent of the maturity he does not sell, $p_1(x_1, x_2|q_1, q_2) = a - bx_1$. Finally, if the dealer does not hold enough of either bill to satisfy the demand of client(s) he cannot sell.

Notice that the magnitudes of the resale prices are characterized by three parameters $\{a, b, e\}$. A higher intercept $a > 0$ increases the dealer's bargaining power, and with it the price he can charge for each unit sold. Parameter $b > 0$ governs the price-sensitivity of clients. Large clients (who demand more) have more negotiating power and can drive down the price. When $e > 0$ bills are substitutes in the secondary market, and vice versa for complements.

Selling $\{x_1, x_2\}$ generates a resale revenue of:

$$revenue(x_1, x_2|q_1, q_2) = p_1(x_1, x_2|q_1, q_2)x_1 + p_2(x_1, x_2|q_1, q_2)x_2. \tag{13}$$

Turning down clients is costly for the dealer. An unhappy client is, for instance, less likely to contact the dealer again in the future. In reality, a dealer might even want to source the security a client demands in the secondary market so as to avoid losing his client in the longer run. This is costly for the dealer because it is expensive to borrow or buy additional Treasury bills on the secondary market when demand is high. In our model, dealers face the following cost function:

$$cost(x_1, x_2|q_1, q_2) = \begin{cases} 0 & \text{if } x_1 \leq \kappa_1 q_1 \text{ and } x_2 \leq \kappa_2 q_2 \\ \gamma x_1 & \text{if } x_1 > \kappa_1 q_1 \text{ and } x_2 \leq \kappa_2 q_2 \\ \gamma x_2 & \text{if } x_1 \leq \kappa_1 q_1 \text{ and } x_2 > \kappa_2 q_2 \\ \gamma x_1 x_2 & \text{if } x_1 > \kappa_1 q_1 \text{ and } x_2 > \kappa_2 q_2. \end{cases} \tag{14}$$

This function captures the idea that it is more costly to turn down larger clients, i.e. those with larger demand. The important feature for our results is that it is supermodular in

x_1, x_2, i.e. has increasing differences.[30] This means that the marginal cost from turning down a client who demands one maturity is higher the larger the order for the other maturity.

Taken together, a dealer expects to derive the following payoff from winning q_1, q_2 at time τ in the primary market:

$$V(q_1, q_2, s) = U(q_1, q_2, s) + \mathbb{E}\left[revenue(\boldsymbol{x_1}, \boldsymbol{x_2}|q_1, q_2) - cost(\boldsymbol{x_1}, \boldsymbol{x_2}|q_1, q_2)\right]. \tag{15}$$

The gross payoff determines how much a dealer is willing to pay on-the-margin. Consider auction 1. At time τ the dealer is willing to pay $v_1(q_1, q_2, s) = \frac{\partial V(q_1,q_2,s)}{\partial q_1}$ for amount q_1 conditional on winning q_2 of the other maturity. The appendix shows that $v_1(\cdot, \cdot, s)$ is a third-order polynomial for any s. It can be approximated by a linear function. Taking the first-order Taylor expansion around $(\mathbb{E}[\boldsymbol{x_1}], \mathbb{E}[\boldsymbol{x_2}]) = (1/2, 1/2)$ we obtain the following result.

Proposition 2. *The marginal willingness to pay of a dealer with type $s^g_{m,i,\tau}$ for amount q_m conditional on winning q_{-m} in the other auction can be approximated by*

$$v_m(q_m, q_{-m}, s^g_{m,i,\tau}) = f_{m,i}(s^g_{m,i,\tau}) + \lambda_{m,i} q_m + \delta_{m,i} q_{-m} \tag{3}$$

for $m = 1, 2$ $-m \neq m$, where $f_{m,i}(s^g_{m,i,\tau}) = \alpha_{m,i} + (1 - \kappa_{m,i}) t^g_{m,i,\tau}$ and $\alpha_{m,i}, \lambda_{m,i}, \delta_{m,i}$ are polynomials of parameters $\{\kappa_{1,i}, \kappa_{2,i}, \gamma_i, \rho_i, a_i, b_i, e_i\}$.

The higher the private marginal benefit t_1 from keeping a share $(1 - \kappa_1)$ of the bill for personal usage, the more the dealer is willing to pay. Bills might be substitutable or complementary depending on the underlying exogenous parameters.

To understand this result, let us contrast the extreme cases where the dealer keeps all of maturity 1 ($\kappa_1 = 0$), keeps all of maturity 2 ($\kappa_2 = 0$), or sells all of both ($\kappa_1 = \kappa_1 = 1$) and the demand of clients is stochastically independent ($\rho = 0$).

$$v_1(q_1, q_2, s_1) = \begin{cases} t_1 & \text{if } \kappa_1 = 0 \\ \frac{1}{4}\kappa_1(b\kappa_1^2 - 2\gamma) + (1 - \kappa_1)t_1 + \kappa_1^2((a - b\kappa_1) + \frac{1}{2}\gamma)q_1 & \text{if } \kappa_2 = 0 \\ \frac{1}{8}(2(b + e) - 6\gamma) + ((a - b) - \frac{1}{4}e + \frac{7}{8}\gamma)q_1 + \frac{1}{4}(3\gamma - 2e)q_2 & \text{if } \kappa_1 = \kappa_2 = 1 \end{cases}$$

When buying only for its own account ($\kappa_1 = 0$) a dealer is willing to pay the marginal value that the bill brings to his own institution, t_1. When he anticipates that he will sell

[30] Supermodularity is for functions that map from $\mathbb{R}^n \to \mathbb{R}$ equivalent to increasing differences: $cost(x_1', x_2'|q_1, q_2) - cost(x_1, x_2'|q_1, q_2) \geq cost(x_1', x_2|q_1, q_2) - cost(x_1, x_2|q_1, q_2)$ for $x_1' \geq x_1$ and $x_2' \geq x_2$.

at least some of maturity 1, his demand in auction 1 decreases in q_1 as long as his clients are sufficiently price-elastic (i.e. b is sufficiently high). If he sells all of both maturities ($\kappa_1 = \kappa_2 = 1$) the demand is independent of his private type t_1. How much he is willing to pay for one maturity now hinges on the amount he wins of the other maturity. Whether bills are substitutes or complements in the primary market depends on how large γ is relative to e.

More generally one can derive the following corollary which will be useful when interpreting our estimation results. It holds for the general case where clients' demand might be correlated ($\rho \neq 0$) and the dealer keeps any amount of bills ($\kappa_1, \kappa_2 \in [0, 1]$).

Corollary 1. *Securities in the primary market become less substitutable for a dealer when (i) they are weaker substitutes in the secondary market ($e_i \downarrow$),*
(ii) it is more costly to turn down clients ($\gamma_i \uparrow$), or
(iii) it is more likely that clients with demand for different maturities arrive ($\rho_i \uparrow$).

The corollary has two interesting implications. First, it highlights that bills might be substitutable for clients, or more generally for traders in the secondary market ($e_i > 0$), but complementary for dealers who purchase in the primary auctions to sell in the secondary market. Through the lens of our model, the existing literature using market-level data to estimate the degree of substitutability between government securities (e.g., Koijen and Yogo (2019)) estimates the mean of parameter e_i. We, instead, focus on the preferences of dealers in the primary market.

Second, the corollary tells us that it is possible that some dealers view bills as substitutes and others as complements, depending on ν_i. For some dealers it could be more costly to turn down clients (high γ_i), for instance, because they are not at the core of the market's trade network, such as the key market makers. For these dealers bills are less substitutable—potentially even complementary—than for the market makers. This insight motivates us to allow dealers to have a latent business type (market makers versus non-market makers) in an extension of our structural model.

B Equilibrium Condition for Uniform Price Auctions

As in the discriminatory price auction, the equilibrium condition for simultaneous uniform price auctions is like the condition for an independent uniform price auction (see Kastl (2011), Proposition 1). The only difference is that the bidder must take the expectation over how much he expects to benefit from winning a specific amount in one auction.

Proposition 1 (Uniform price auctions)**.** *Consider a dealer i with private information $\theta^g_{i,\tau}$ who submits $\hat{K}_m(\theta^g_{i,\tau})$ steps in uniform price auction m at time τ. Under Assumptions 1-4 in any type-symmetric BNE every step k in his bid function $b^g_m(\cdot,\theta^g_{i,\tau})$ has to satisfy*

$$\tilde{v}_m(q_{m,k}, s^g_{m,i,\tau}|\theta^g_{i,\tau}) = \mathbb{E}[\boldsymbol{P^c_m}|b_{m,k} > \boldsymbol{P^c_m} > b_{m,k+1}, \theta^g_{i,\tau}] + \frac{\frac{\partial \mathbb{E}[\boldsymbol{P^c_m}|b_{m,k}\geq \boldsymbol{P^c_m} \geq b_{m,k+1}, \theta^g_{i,\tau}]}{\partial q_{m,k}}}{\frac{\Pr(b_{m,k} > \boldsymbol{P^c_m} > b_{m,k+1}|\theta^g_{i,\tau})}{q_{m,k}}} \ \forall k < \hat{K}_m(\theta^g_{i,\tau})$$

with $\tilde{v}_m(q_{m,k}, s^g_{m,i,\tau}|\theta^g_{i,\tau}) \equiv \mathbb{E}\left[v_m\left(q_{m,k}, \boldsymbol{q^*_{-m,i}}, s^g_{m,i,\tau}\right)\middle| b_{m,k} \geq \boldsymbol{P^c_m} > b_{m,k+1}, \theta^g_{i,\tau}\right]$ *for all m with $-m \neq m$, and $b_{m,k} = \tilde{v}_m(\bar{q}_m(\theta^g_{i,\tau}), s^g_{m,i,\tau}|\theta^g_{i,\tau})$ at $k = \hat{K}_m(\theta^g_{i,\tau})$ where $\bar{q}_m(\theta^g_{i,\tau})$ is the maximal amount the bidder may be allocated in an equilibrium.*

C Proofs

Proposition 1. Consider discriminatory price auctions. Take the perspective of dealer i. Fix his type, a time slot τ, as well as one of his information sets $\theta^g_{i,\tau}$, and let all other agents $j \neq i$ play a type-symmetric equilibrium. In this equilibrium it must be optimal for the bidder to choose the same set of functions $\{b^g_1(\cdot,\theta^g_{i,\tau}), ...b^g_M(\cdot,\theta^g_{i,\tau})\}$ as all other bidders in his bidder group with information $\theta^g_{i,\tau}$. These M functions must *jointly* maximize the bidder's expected total surplus. It must therefore be the case that each of the functions $b^g_m(\cdot,\theta^g_{i,\tau})$ maximizes his expected total surplus *separately* when fixing all the other bidding functions $-m$ at the optimum. To determine necessary conditions of the type-symmetric equilibrium we can consequently fix the agent's strategy in all but one auction at the equilibrium.

The remainder of the proof extends Kastl (2012)'s proof for a K-step equilibrium of a discriminatory price auction that takes place in isolation (on pp. 347–348). There are two main differences to the original proof. First, our framework allows bidders to update their bids due to arrival of new information. Such information arrives at discrete time slots $\tau = 1...\Gamma$. Bidding functions do not (only) depend on the bidder i's type $s^g_{i,\tau}$ drawn at time τ but on the (entire) information set at that time $\theta^g_{i,\tau}$. It includes the type, $s^g_{i,\tau} \subseteq \theta^g_{i,\tau}$. Since only final bids count, bidders bid as if it was their last bid each time they place a bid. We can just keep some τ fixed throughout the proof. Second, following Hortaçsu and Kastl (2012) we allow for asymmetries in bidding behavior between dealers and customers. They draw types from (potentially) different distributions and may have different information available. The original proof extends to this setup. For details see the Online Appendix.

The proof for the uniform price follows the same steps, and is therefore omitted. □

Proposition 2. For notational convenience we drop the superscript g and the subscript i of all parameters $\{\kappa^g_{1,i}, \kappa^g_{2,i}, \gamma^g_i, \rho^g_i, a^g_i, b^g_i, e^g_i\}$. Given the aggregate inverse demand of the dealer's clients (12), the dealer expects the following payoff (15) from owning q_1, q_2:

$$
\begin{aligned}
V(q_1,q_2,s) = U(q_1,q_2,s) &+ \int_0^{\kappa_1 q_1}\int_0^{\kappa_2 q_2} [p_1(x_1,x_2)x_1 + p_2(x_2,x_1)x_2] f(x_1,x_2)dx_1dx_2 \\
&+ \int_0^{\kappa_1 q_1}\int_{\kappa_2 q_2}^1 [p_1(x_1)x_1 - \gamma x_2] f(x_1,x_2)dx_1dx_2 + \int_{\kappa_1 q_1}^1\int_0^{\kappa_2 q_2} [p_2(x_2)x_2 - \gamma x_1] f(x_1,x_2)dx_1dx_2 \\
&- \int_{\kappa_1 q_1}^1\int_{\kappa_2 q_2}^1 [\gamma x_1 x_2] f(x_1,x_2)dx_1dx_2.
\end{aligned}
$$

Insert the functional forms (11), (12), and $f(x_1,x_2) = 1 + 3\rho(1-2F_1(x_1))(1-2F_2(x_2))$, integrate and take the partial derivative w.r.t. q_1. Then do a Taylor expansion around $\left(\frac{1}{2},\frac{1}{2}\right)$ to obtain

$$
v_1(q_1,q_2,s_1) = (1-\kappa_1)t_1 + h_0(\kappa_1,\kappa_2,\gamma,\rho) + h_1(\kappa_1,\kappa_2,\gamma,a,b,e,\rho)q_1 + h_2(\kappa_1,\kappa_2,e,\rho)q_2
$$

with

$$
\begin{aligned}
h_0(\kappa_1,\kappa_2,\gamma,\rho) =& \frac{1}{16}(4b\kappa_1^3 + 2e\kappa_1^2\kappa_2^2(2 + (6 - 9\kappa_1 - 6\kappa_2 + 8\kappa_1\kappa_2)\rho)) \\
&+ \frac{1}{16}(\gamma\kappa_1(8(-1+\rho) + \kappa_1^2(-2+\kappa_2)(2+\kappa_2(-11+8\kappa_2))\rho)) \\
&+ \frac{1}{16}(\gamma\kappa_1(+2\kappa_2^2(-1-3\rho+4\kappa_2\rho) + 2\kappa_1\kappa_2(-2+\kappa_2-3(-1+\kappa_2)(-2+3\kappa_2)\rho))) \\
h_1(\kappa_1,\kappa_2,\gamma,a,b,e,\rho) =& \frac{1}{8}\kappa_1^2(8a - 8b\kappa_1 - 2e\kappa_2^2(1 + (-1+2\kappa_1)(-3+2\kappa_2)\rho)) \\
&+ \frac{1}{8}\kappa_1^2(\gamma(4 + 4\kappa_2 - \kappa_2^2 - (-2+\kappa_2)(-6+3\kappa_2-6\kappa_2^2+2\kappa_1(-2+\kappa_2)(-1+2\kappa_2))\rho)) \\
h_2(\kappa_1,\kappa_2,\gamma,e,\rho) =& -\frac{1}{4}\kappa_1\kappa_2(-2\gamma\kappa_1 + \gamma(-2+\kappa_1)\kappa_2 + 2e\kappa_1\kappa_2)(1+3(-1+\kappa_1)(-1+\kappa_2)\rho) \quad \square
\end{aligned}
$$

Proof of Corollary 1. Securities become less substitutable when $h_2(\kappa_1,\kappa_2,\gamma,e,\rho)$ increases. Therefore, for any $\kappa_m \in [0,1]$ and any $\rho \in [-1/3, 1/3]$: $\frac{\partial h_2(\kappa_1,\kappa_2,\gamma,e,\rho)}{\partial e} \le 0$, $\frac{\partial h_2(\kappa_1,\kappa_2,\gamma,e,\rho)}{\partial \gamma} \ge 0$, $\frac{\partial h_2(\kappa_1,\kappa_2,\gamma,e,\rho)}{\partial \rho} \ge 0$. $\square$

Appendix Figure A1: Auction interface–what bidders see when bidding

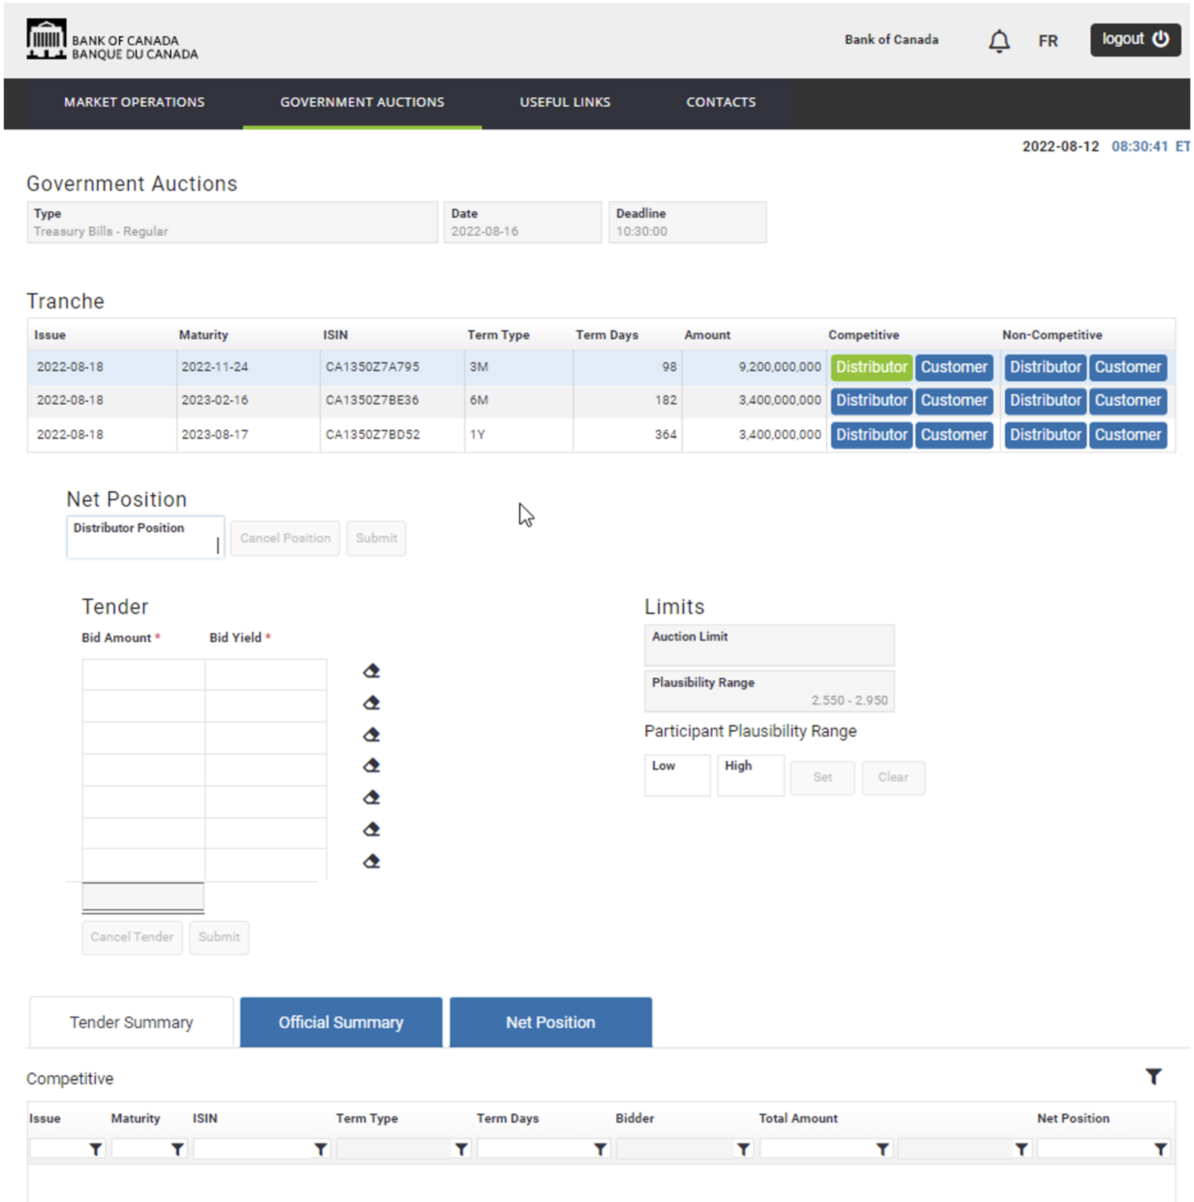

Appendix Figure A1 shows the a screen shot of what a bidder sees when bidding. He sees three "Tranches", listing the three different securities for sale. If he is a dealer, placing a competitive bid for his own account, he clicks on "Distributor" (in green). He fills in what "Net Position" he currently holds of the security, and his bid "Tender" of maximally 7 steps. On the RHS he sees a non-binding "Plausibility Range" suggesting yields at which the auction might clear. When the dealer is content with his bid, he clicks "Submit". Then he can choose a different "Tranche" to bid in a different auction. At the bottom of the page, he sees an overview of all his submitted bids for all auctions.

Appendix Figure A2: Steps by Bidder Groups

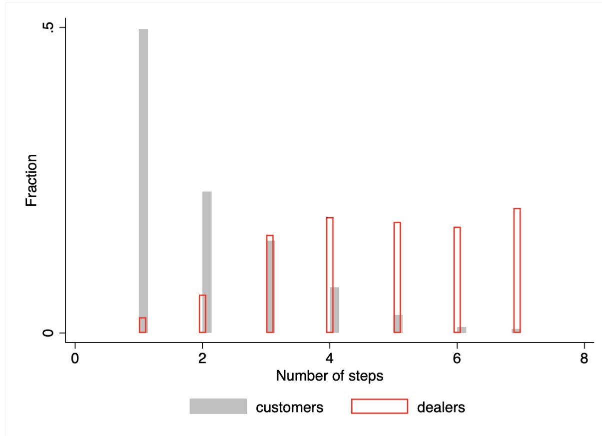

Appendix Figure A2 shows a histogram of the number of steps customers (in gray) and dealers (in red) submit. The fraction is measured in percentage points.

Appendix Table A1: Bid Updating

			Update in 12M for 3M order		Update in 6M for 3M order	
Bid by	Time	Maturity	(1)	(2)	(1)	(2)
Customer	10:19:52	3M	.	.	.	.
Dealer	10:21:59	12M	1	1	0	0
Dealer	10:22:17	6M	0	0	0	1
Dealer	10:22:34	3M	0	0	0	0
Dealer	10:26:52	12M	0	0	0	0
Dealer	10:27:16	12M	0	0	0	0
Customer	10:28:34	3M	.	.	.	.
Dealer	10:28:44	3M	0	0	0	0

Appendix Table A1 illustrates the sequence of events from a random dealer and their customer for the last 10 minutes before the auction closes on 02/10/2015. Having observed a customer in the 3M auction (visible in the first row), the dealer takes action himself and places several bids in a row (as shown in the second until sixth row). He first bids in the 12M auction. Therefore $customer_{3M}$ assume value 1 in specification (1) and (2) shown in the fourth and sixth column. Then the dealer bids in the 6M auction. Now, the $customer_{3M}$ variable switches to 1 only in specification (2) in the seventh column, but not in specification (1) in the sixth column. This is because the dealer has observed a customer in the 3M auction one minute before placing a bid in the 6M auction but not immediately before that.

Appendix Figure A3: Time series

(a) Proxy of market price sensitivities

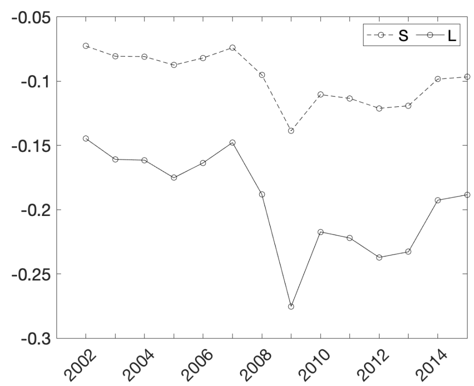

(b) Revenue gains in uniform price auction

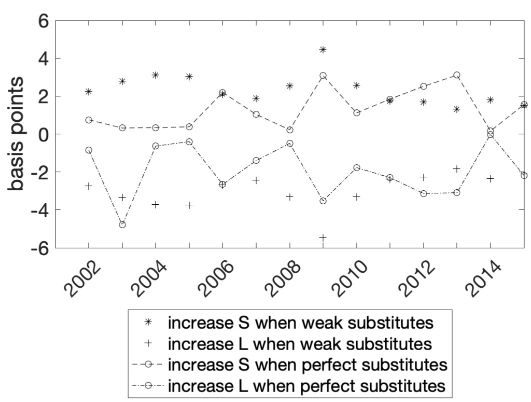

(c) Revenue gains in discrim. price auction

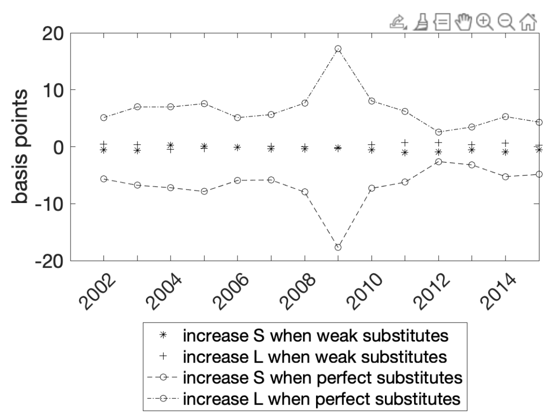

Appendix Figure A3 shows two time series. An observation in A3a shows the yearly average market price sensitivity when scaling the λ parameters by a factor of 100 and setting all δ's to zero: $\frac{1}{T_y}\sum_{t=1}^{T_y}(-100)\hat{\lambda}_m\frac{Q_{t,m}}{P^c_{tm}}$, where T_y is the total number of auction in year y. An observation in A3b and A3c is the gain in total revenue of the two maturities on a day when issuing 1% of total debt more of the short and less of the long bill, or vice versa, averaged across all auction days in a year. The revenue gain is computed for a discriminatory price auctions and is measured in bps of the revenue earned when issuing the observed supply. We scale up the λ and δ parameters to make the time trends visible.

Online Appendix

The Online Appendix has four sections. Section A presents our model extension in which dealers have a latent type. Section B walks through a more detailed proof of Proposition 1. Section C analyzes the potential bias stemming from measurement error in the estimated expected winning quantities. Section D provides details on how we solve the fixed point problem when determining the counterfactual bids. Online Appendix Figures and Tables are presented at the end.

A Model Extension

We consider a model extension in which dealers have a latent business type χ. Each dealer is either a market-maker type ($\chi = mm$) or a niche-customer type ($\chi = nc$). In theory one could allow for more than two types. In the estimation, this is feasible only if there are sufficiently many bidders that participate in an auction, which is not the case in our setting. Otherwise, one would need to pool auctions that take place on different dates and lose the ability to control for unobservable auction characteristics.

Assumptions 1 and 2 adjust in that the private signals draw independently from three distributions $F^{d,mm}$, $F^{d,nc}$ and F^c. The marginal willingness to pay may now be bidder specific: $v_{m,i}(q_m, q_{-m}, s^g_{m,i,\tau}) = f_{m,i}(s^g_{m,i,\tau}) + \lambda_{m,i} q_m + \delta_{m,i} \cdot q_{-m}$. All other assumptions remain unchanged, but there is an additional assumption.

Assumption 5 (Model extension)**.** *Dealers can be partitioned into two types:* $\mathbb{N}_d = \mathbb{N}_{d,mm} \cup \mathbb{N}_{d,nc}$*, such that* $\forall\, m \in \mathbb{N}_{mm} : \delta^{d,mm}_{m,i} \leq 0$*.*

The econometrician does not know which dealer is of which type. The bidders know who is of which type. For them these latent types just mean that there are more than two bidder groups. Therefore, there are more than two strategies in the type-symmetric equilibrium: $\forall \chi : b^{d,\chi}_{i,\tau}(\cdot, \theta_{i,\tau}) = b^{d,\chi}(\cdot, \theta_{i,\tau})$ and $b^c_{i,\tau}(\cdot, \theta_{i,\tau}) = b^c(\cdot, \theta_{i,\tau})$ $\forall i, \tau$.

To recover the valuations, $\hat{v}_{t,m,i,\tau,k}$, we adjust Proposition 1 and extend the resampling procedure as in Cassola et al. (2013) to account for asymmetric types. The resampling proceeds in three steps: (i) Partition dealers into the two groups. (ii) Estimate a model, where resampling is conditional on that assignment.[31] (iii) Use the estimated demands to classify dealers into types.[32] Repeat until (iii) yields the same assignment as we started

[31] An $mm-$type needs to integrate over bids of $N_{mm} - 1$ other mm-types and N_{nc} niche-client types and vice versa.

[32] Since we have 3 maturities, we have 6 coefficients in the demand system given by (3) governing the substitution patterns. We assign a dealer to mm-type if at most 2 of those are negative.

with in (i). While there is no formal argument that this procedure will converge, in practice it converges within 2 or 3 steps. Finally, we estimate regression (7) for each dealer group separately, identifying the group-specific average δ^χ and λ^χ parameters.

We find that there are two dealer groups with different preferences (see Online Appendix Table O4). For the 11 dealers in group one, bills are (in most cases) more substitutable than for the average dealer in our benchmark model. For the 4 dealers in the second group, preferences are mixed.

Our micro-foundation is able to rationalize these findings (see Corollary 1). Dealers in the first group win on average larger amounts in the auctions than dealers in the second group. They are the bigger players in the market who are not concerned about turning down clients, either because they hold large inventory positions or because they can rely on their trading network to quickly and cheaply find the security a clients wants. For dealers in the second group, who tend to win less at auction, this might not always be true.

B Detailed Proof of Proposition 1

We present the proof for discriminatory auctions. The proof for uniform price auctions follows the same steps, and is therefore omitted.

Take the perspective of dealer i. Fix his type, a time slot τ, as well as one of his information sets $\theta^g_{i,\tau}$, and let all other agents $j \neq i$ play a type-symmetric equilibrium. In this equilibrium it must be optimal for the bidder to choose the same set of functions $\{b^g_1(\cdot, \theta^g_{i,\tau}), ...b^g_M(\cdot, \theta^g_{i,\tau})\}$ as all other bidders in his bidder group with information $\theta^g_{i,\tau}$. These M functions must *jointly* maximize the bidder's expected total surplus. It must therefore be the case that each of the functions $b^g_m(\cdot, \theta^g_{i,\tau})$ maximizes his expected total surplus *separately* when fixing all the other bidding functions $-m$ at the optimum. To determine necessary conditions of the type-symmetric equilibrium we can consequently fix the agent's strategy in all but one auction at the equilibrium. Without loss take this auction to be for security 1, and denote the inverse of bid function $b^g_1(\cdot, \theta^g_{i,\tau})$ by $y^g_1(\cdot, \theta^g_{i,\tau})$.

The remainder of the proof follows Kastl (2012)'s original proof. To facilitate the comparison with the original proof (on pp. 347–348 of Kastl (2012)) we copy it as closely as possible but adopt our notation.

We drop subscripts τ, i as well as superscript g. We refer to the amount a bidder with information θ wins at market clearing in auction m (for a given set of strategies in the event that τ is the time of the bidder's final bid) by $\boldsymbol{q^c_1}$, and the amount he wins in equilibrium by $\boldsymbol{q^*_1}$. Notice that both, $\boldsymbol{q^c_1}$ and $\boldsymbol{q^*_1}$ are (for given strategies of all agents) functions of the total supply $\boldsymbol{Q_1}$ and the information of all agents $\{\boldsymbol{\theta_i}\}_{i=1}^N$. They are implicitly defined by market clearing.

The proof of the proposition relies on three lemmas. The second and third are taken from Kastl (2012).

Lemma 1. *Fix a bidder with information θ. Denote his marginal willingness to pay in auction m at step k when submitting some function $b_1'(\cdot,\theta)$ with $\{(b'_{1,k}, q'_{1,k-1}), (b'_{1,k+1}, q'_{1,k})\}$ by $\tilde{v}_1(q_1, \theta | b'_{1,k}, b'_{1,k+1}) \equiv \mathbb{E}\left[v_1\left(q_1, \boldsymbol{q^*_{-1}}, s_1\right) \middle| b'_{1,k} \geq \boldsymbol{P_1^c} > b'_{1,k+1}, \theta\right]$ for $q_1 \in (q'_{1,k-1}, q'_{1,k}]$.*

- *(i) $\tilde{v}_1(q_1, \theta | b'_{1,k}, b'_{1,k+1})$ is bounded.*
- *(ii) In equilibrium, where the bidder submits function $b_1(\cdot,\theta)$ with $\{(b_{1,k}, q_{1,k-1}), (b_{1,k+1}, q_{1,k})\}$, $\tilde{v}_1(q_1, \theta | b_{1,k}, b_{1,k+1})$ is decreasing in q_1 and right-continuous in $b_{1,k}$.*

Proof of Lemma 1. (i) By Assumption 2

$$\tilde{v}_1(q_1, \theta | b'_{1,k}, b'_{1,k+1}) \overset{(3)}{=} f_1(s_1) + \lambda_1 q_1 + \delta_1 \cdot \mathbb{E}\left[\boldsymbol{q^*_{-1}} | b'_{1,k} \geq \boldsymbol{P_1^c} > b'_{1,k+1}, \theta\right]$$

for $q_1 \in (q'_{1,k-1}, q'_{1,k}]$. Since types and total supply are drawn from distributions with bounded support by Assumptions 1 and 4, $\mathbb{E}\left[\boldsymbol{q^*_{-1}} | b'_{1,k} \geq \boldsymbol{P_1^c} > b'_{1,k+1}, \theta\right]$ and with it $\tilde{v}_1(q_1, \theta | b'_{1,k}, b'_{1,k+1})$ is bounded.

(ii) In equilibrium $\tilde{v}_1(q_1, \theta | b_{1,k}, b_{1,k+1})$ must be decreasing in q_1 or it could not give rise to a decreasing bidding function that fulfills the necessary conditions of Proposition 1.

To see why $\tilde{v}_1(q_1, \theta | b_{1,k}, b_{1,k+1})$ is right-continuous in $b_{1,k}$ note first that it can only jump discontinuously if changing $b_{1,k}$ breaks a tie between this bidder and at least one other bidder. Since there can be only countably many prices on which a tie might occur, however, there must exist a neighborhood at any $b_{1,k}$ for which for any price in that neighborhood there are no ties. Therefore, when perturbing b_k, there cannot be any discontinuous shift in the conditional probability measure and thus in the object of interest. □

Lemma 2. *Fix a bidder with information θ. If at some step k in auction* 1*,* $\Pr(\boldsymbol{q_1^c} \geq q_{1,k} | \theta) > 0$*, then $b_{1,k} \leq \tilde{v}_1(q_1, \theta | b_{1,k}, b_{1,k+1})$.*

Proof of Lemma 2. The proof is analogous to Kastl (2012)'s proof of Lemma 2. It suffices to replace $v(q,s)$ by $\tilde{v}_1(q_1, \theta | b_{1,k}, b_{1,k+1})$ and rely on Lemma 1. □

Lemma 3. (i) *Ties occur with zero probability for a.e. θ in any K-step equilibrium of simultaneous discriminatory price auctions except possibly at the last step ($k_1 = K_1$).*

(ii) *If a tie occurs with positive probability at the last step, a bidder with information θ must be indifferent between winning or losing all units between the lowest share he gets* *allocated after rationing in the event of a tie $\underline{q}_1^{RAT}$ and the last infinitesimal unit he may be allocated in equilibrium, $\bar{q}_1$:*

$$b_{1,K_1} = \tilde{v}_1(\bar{q}_1, \theta | b_{1,K_1}) \text{ where } \bar{q}_1 = \sup_{\{Q_1, \theta_{-i}\}} y_1(b_{1,K_1}, \theta | Q_1, \theta_{-i}) \; \forall q_1 \in [\underline{q}_1^{RAT}, \bar{q}_1].$$

Proof of Lemma 3. The proof is analogous to the proof of Lemma 1 in Kastl (2012). In essence, it suffices to replace the bidder's true valuation $v(q,s)$ in Kastl (2012) by $\tilde{v}_1(\cdot, \theta | b_k, b_{k+1})$ in equilibrium and $\tilde{v}_1(\cdot, \theta | b'_k, b'_{k+1})$ for deviations and rely on Lemma 1.

To facilitate this conversion, we demonstrate the beginning of the proof: Suppose that there exists an equilibrium, in which for a bidder i with information set θ a tie between at least two bidders can occur with positive probability $\pi_1 > 0$ in auction 1. Since there can be only finitely many prices that can clear the market with positive probability, in order for a tie to be a positive probability event, it has to be the case that there exists a positive measure subset of information sets $\hat{\Theta}_{-i} \in [0,1]^{N-1}$ such that for some bidder j, and all profiles of information sets $\theta_{-i} \in \hat{\Theta}'_{-i} \subset \hat{\Theta}_{-i}$ (another positive measure subset) and some step k and l we have $b_{1,k}(\theta_i) = b_{1,l}(\theta_j) = P_1^c$. Without loss, suppose that this event occurs at the bid $(b_{1,k}, q_{1,k})$, and that the maximum quantity allocated to i after rationing is $\bar{q}_1^{RAT} < q_{1,k}$. Let $\bar{S}_{1\pi}^R$ denote the maximal level of the residual supply at $b_{1,k}$ in the states leading to rationing at $b_{1,k}$.

Consider a deviation to a step $b'_{1,k} = b_{1,k} + \varepsilon$ and $q'_{1,k} = q_{1,k}$ where ε is sufficiently small. This deviation increases the probability of winning $q_{1,k} - q_{1,k-1}$ units. Most importantly in the states that led to rationing under the original bid, the bidder with information θ will now obtain $q_1^u > \bar{q}_1^{RAT}$ where $q_1^u \geq \min\{q_{1,k}, \bar{S}_{1\pi}^R\}$. Notice that since we hypothesized a positive probability of a tie at $b_{1,k}$, we need to have $q_{1,k-1} < \bar{q}_1^{RAT} < q_{1,k}$ due to rationing pro-rata on-the-margin. Therefore, the lower bound on the increase in θ's expected gross surplus from such a deviation is

$$ED_\varepsilon = \pi_1 \left(\tilde{V}_\varepsilon(q_1^u, \theta) - \tilde{V}(\bar{q}_1^{RAT}, \theta) \right) \qquad (ED_\varepsilon)$$

where

$$\tilde{V}_\varepsilon(q_1^u, \theta) \equiv \int_0^{\bar{q}_1^{RAT}} \tilde{v}_1(q_1, \theta | b_1(q_1|\theta)) + \int_{\bar{q}_1^{RAT}}^{q_1^u} \tilde{v}_1(q_1, \theta | b'_{1,k}, b'_{1,k+1}) dq_1$$

and

$$\tilde{V}(\bar{q}_1^{RAT}, \theta) \equiv \int_0^{\bar{q}_1^{RAT}} \tilde{v}_1(q_1, \theta | b_1(q_1|\theta)) dq_1$$

with $\tilde{v}_1(q_1, \theta | b_1(q_1|\theta))$ denoting the true valuation when submitting $b_1(q_1|\theta)$ not just at step k, as $\tilde{v}_1(q_1, \theta | b_{1,k}, b_{1,k+1})$, but including all previous steps (if any).

To continue, let us first focus on steps other than the last one, $k < K_1$, and suppose that $\tilde{v}_1(\cdot, \theta|b_{1,k}, b_{1,k+1})$ is strictly decreasing. The increased bid $b_{1,k} + \varepsilon$ also results in an increase in the payment for the share requested at this step. This increase, however, is bounded by $(q_{1,k} - q_{1,k-1})\varepsilon$. Comparing the upper bound on the change in expected payment with the lower bound on the change in expected gross utility, in order for this deviation to be strictly profitable we need to obtain

$$(q_{1,k} - q_{1,k-1})\varepsilon < \pi_1 ED_\varepsilon. \tag{16}$$

As $b_{1,k} \leq \tilde{v}_1(q_{1,k}, \theta|b_{1,k}, b_{1,k+1})$ by Lemma 2 and $\tilde{v}_1(q_{1,k}, \theta|b_{1,k}, b_{1,k+1}) < \tilde{v}_1(q_1^u, \theta|b_{1,k}, b_{1,k+1})$, the LHS of (16) goes to 0 and the RHS to a strictly positive number as $\varepsilon \to 0$. Since $\tilde{v}_1(q_1, \theta|b_{1,k}, b_{1,k+1})$ is for any $q_1 \in [\bar{q}_1^{RAT}, q_{1,k}]$ right-continuous in $b_{1,k}$, the proposed deviation would indeed be strictly profitable for the bidder with information θ. Moreover, there can be only countable many θ's with a profitable deviation, otherwise bidder i could implement this deviation jointly and thus for *a.e.* information sets θ ties have zero probability in equilibrium for all bidders i.

Relying on Lemma 1, the remainder of the proof is analogous to the original proof. It suffices to replace $v(q, s)$ by $\tilde{v}_1(\cdot, \theta|b_k, b_{k+1})$ in equilibrium and $\tilde{v}_1(\cdot, \theta|b'_k, b'_{k+1})$ when deviating, as well as $V(q^*, s) - V(\bar{q}_i^{RAT}, s)$ by ED_ε. In our environment with updating, a tie may occur with positive probability only at the last step and the bidder with information θ (at the previously fixed time τ) must not prefer winning any units in $\left[\underline{q}_1^{RAT}, \overline{q}_1\right]$ where $\overline{q}_1 = \sup_{\{Q_1, \theta_{-i}\}} y_1(b_{1,K_1}, \theta|Q_1, \theta_{-i})$ is the maximal quantity the bidder may be allocated in an equilibrium (in the event that τ is the time of his final bid). $\square$

At step $k = K_1$ Lemma 2 specifies the optimal bid-choice. At steps $k < K_1$ Lemma 3 can be applied. Kastl (2012) perturbs the kth step to $q'_1 = q_{1,k} - \epsilon$ and takes the limit as $q'_1 \to q_{1,k}$. The original proof goes through without complications. It suffices to replace the type s by the information set θ, $\mathbb{E}\left[V(Q_i^c(Q, \boldsymbol{S}, \boldsymbol{y}(\cdot|S)), s_i)|\,\text{states}\right]$ by $\mathbb{E}\left[V(\boldsymbol{q_1^*}, \boldsymbol{q_{-1}^*}, s)\middle|\,\theta, \text{states}\right]$ with all states as specified in the original proof, and similarly $\mathbb{E}\left[V(Q_i^c(Q, \boldsymbol{S}, \boldsymbol{y}'(\cdot|S)), s_i)|\,\text{states}\right]$ by $\mathbb{E}\left[V(\boldsymbol{q_1^c}, \boldsymbol{q_{-1}^*}, s)\middle|\,\theta, \text{states}\right]$ where $\boldsymbol{q_1^c}$ denotes the amount the bidder wins at market clearing under the deviation in our simplified notation. $\square$

C Measurement Error

It is generally difficult to get a sense of whether and by how much our demand coefficients might be downward biased due to measurement error in the expected winning quantities:

$$\hat{\mathbb{E}}\left[\boldsymbol{q_{t,-m,i}^*}|\ldots\right] = \mathbb{E}\left[\boldsymbol{q_{t,-m,i}^*}|b_{t,m,i,\tau,k} \geq \boldsymbol{P_{t,m}^c} > b_{t,m,i,\tau,k+1}, \theta_{t,i,\tau}^g\right] + \varepsilon_{t,m,i,\tau,k}^q. \tag{5}$$

This is because we can construct proxy variables of how much a bidder might expect to win in total $\hat{\mathbb{E}}\left[\boldsymbol{q^*_{t,-m,i}}\right]$, but we cannot approximate how much the bidder expected to win at a particular bid without estimating our model.

To get at least a rough sense of the size of the bias coming, we return to regression (1), which regresses bids on winning quantities. The idea is that this regression only relies on observable variables. This means that we can rule out that variables other than $\hat{\mathbb{E}}\left[\boldsymbol{q^*_{t,-m,i}}|...\right]$ are measured with errors that come from simulating auction clearance in the first stage of our estimation. Since the coefficients from this regression are not the coefficients that we obtain when we take into account that bidders shade their bids, this exercise is useful only to the extend that it helps us grasp the size of the measurement bias. The coefficient estimates or their signs alone, rather than in comparison, bring no additional insights.

We compare the estimates of regression (1) with the estimates of the analogous regression where we use the amounts a bidder expected to win $\hat{\mathbb{E}}\left[\boldsymbol{q^*_{t,-m,i}}|...\right]$ instead of the amount the bidder wins. In addition, we use other observable variables that help bound how much a bidder might expect to win in an auction but are observable. Here, we display two of those: the amount the bidder demanded at the highest step that he ever wins of a maturity in a year, and the total amount the bidder demands at auction if this amount is less than the 1% highest amount the bidder ever wins of the maturity during a year.

Online Appendix Table O5 shows the estimation findings for the 3M auction for illustration. The other auctions show similar patterns. We find that the δ coefficients range between roughly 0.1 and 0.7 when using the observable variables, while they are around 0.3 when using the expected values. Most importantly, all of the δ estimates are small compared to the λ estimate which is 3.7. Therefore, it seems as if the measurement error in the expected values is not causing us to vastly understand the interdependencies.

D Fixed Point Problem and Approximation

In order to compute how bidders bid when we change supply, we must determine how much each bidder expects to win in the other auctions, $\hat{\mathbb{E}}[\boldsymbol{q^{cf*}_{t,-m,i}}|q_m]$. This depends on how all bidders bid in all auctions. Therefore, finding $\hat{\mathbb{E}}[\boldsymbol{q^{cf*}_{t,-m,i}}|q_m]$ of all bidders and all auctions is a complicated fixed point problem. Below we fix one auction date and omit the day subscript.

Exact fixed point routine. Assume we change the supply from Q_m to Q_m^{cf} for all m.

Step 1. Rescale all amounts demanded and expectations:

$$q_{m,i,k}^{cf} = \frac{Q_m^{cf}}{Q_m} q_{m,i,k} \tag{17}$$

$$\hat{\mathbb{E}}[\boldsymbol{q_{-m,i}^{cf*}} | q_{m,i,k}]^{old} = \frac{Q_m^{cf}}{Q_m} \hat{\mathbb{E}}[\boldsymbol{q_{-m,i}^{*}} | q_{m,i,k}] \text{ for all } m, -m, i, k. \tag{18}$$

Then compute the counterfactual bids for each step k, bidder i and maturity m according to (8):

$$b_{m,i,k}^{cf} = \hat{u}_{m,i} + \hat{\lambda}_m q_{m,i,k}^{cf} + \hat{\delta}_m \cdot \hat{\mathbb{E}}[\boldsymbol{q_{-m,i}^{cf*}} | q_{m,i,k}]^{old} - \hat{s}hading_{m,i,k}. \tag{8}$$

Step 2. Given the counterfactual bids, estimate how much each bidder expects to win in the other auctions by simulating market clearance for each bidder and maturity many times (e.g., 5,000 times). Update all expectations, $\hat{\mathbb{E}}[\boldsymbol{q_{-m,i}^{cf*}} | q_{m,i,k}]^{new}$.

Step 3. With the updated expectations, update all bids. Repeat steps 2-3 until none of the expectations change when updated.

Statistical fixed point routine. It is computational infeasible to implement the exact fixed point routine. Therefore, we propose a routine that finds the fixed point with some estimation noise.

Steps 1-2 are as before. Step 3. Find out whether the expectations are too large or too small, by regressing:

$$\hat{\mathbb{E}}[\boldsymbol{q_{-m,i}^{cf*}} | q_{m,i,k}]^{new} = \alpha_m + \beta_m * \hat{\mathbb{E}}[\boldsymbol{q_{-m,i}^{cf*}} | q_{m,i,k}]^{old} + \epsilon_{m,i,k} \text{ for all } m.$$

Update all expectations: $\hat{\mathbb{E}}[\boldsymbol{q_{-m,i}^{cf*}} | q_{m,i,k}]^{new}$ become $\hat{\mathbb{E}}[\boldsymbol{q_{-m,i}^{cf*}} | q_{m,i,k}]^{old}$ and the new $\hat{\mathbb{E}}[\boldsymbol{q_{-m,i}^{cf*}} | q_{m,i,k}]^{new}$ $= \hat{\beta}_m * \hat{\mathbb{E}}[\boldsymbol{q_{-m,i}^{cf*}} | q_{m,i,k}]^{old}$ for all m, i, k. Repeat this step until all $\hat{\beta}_m$ estimates are within the 95% confidence interval around 1.

We determine fixed points using our statistical routine for a couple of randomly selected auction days. We do this for two reasons. First, we want to illustrate that this method works reasonably well (see Online Appendix Figure O4a). Second, we want to show that the fixed point is sufficiently close to the rescaled expectations (18) with which we start the fixed point routine (see Online Appendix Figure O4b). This motivates us to use the rescaled expectations in our counterfactual exercises.

Online Appendix Figure O1: Issuance of Canadian 3M, 6M, 12M Treasury bills

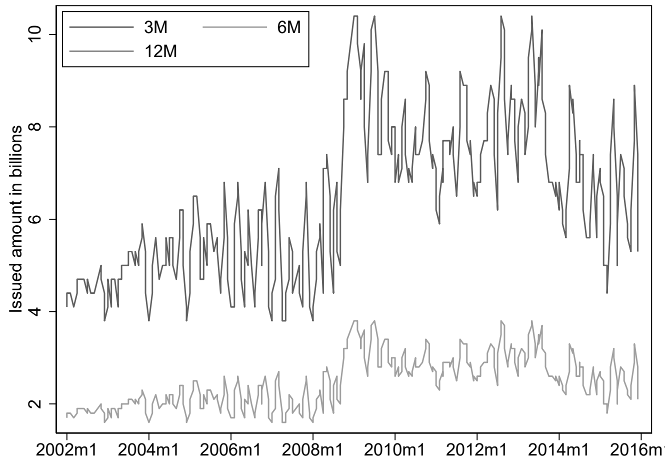

Online Appendix Figure O1 displays a time series of the issued supply of the 3M, 6M, and 12M bills, where the 6M issuance do not appear in the graph because they are identical to 12M issuance. The Bank of Canada always issues as many 6M bills as 12M bills. Over time, the amounts issued of the different maturities are perfectly correlated.

Online Appendix Figure O2: Time between bids of those who do not update

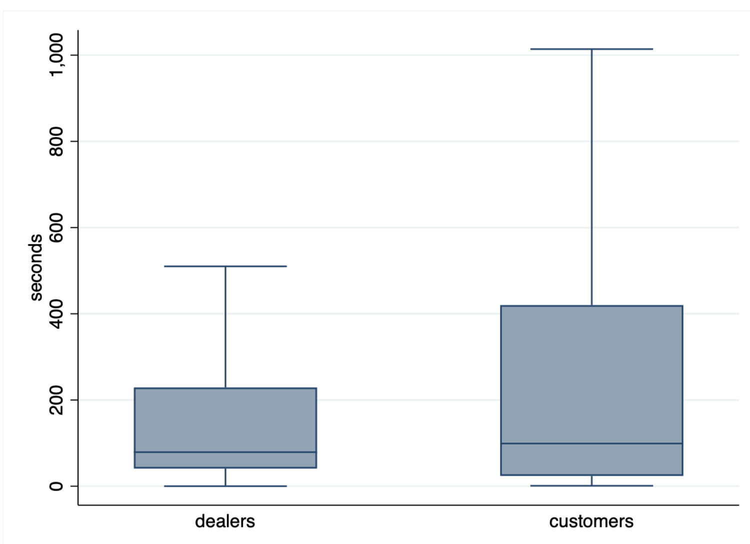

Online Appendix Figure O2 shows the distribution of the time difference (measured in seconds) between the bids that a dealer and a customer who does not update the bids places in different auctions.

Online Appendix Figure O3: Distribution of the untrimmed shading factor

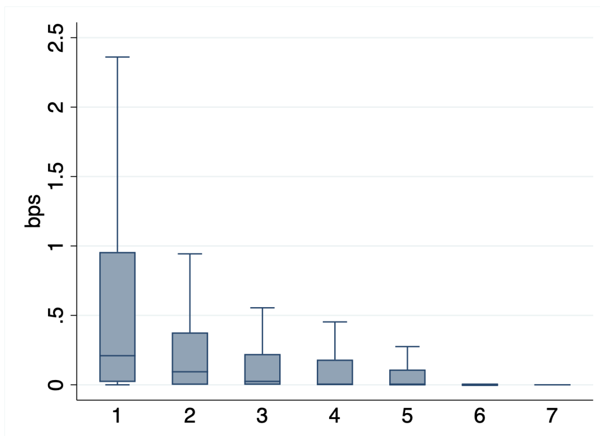

Online Appendix Figure O3 shows box plots of the untrimmed shading factor, $\hat{v}_{t,m,i,,\tau,k} - b_{t,m,i,,\tau,k}$, per step $\in \{1, 2, 3, 3, 5, 6, 7\}$ in a bidding function. For each step, the distribution is taken over dealers i, days t and time τ and maturities m. The shading factor is in bps.

Online Appendix Figure O4: Expectations on 3 auction days

(a) Did we find a fixed point? (b) Fixed point vs. rescaled expectations

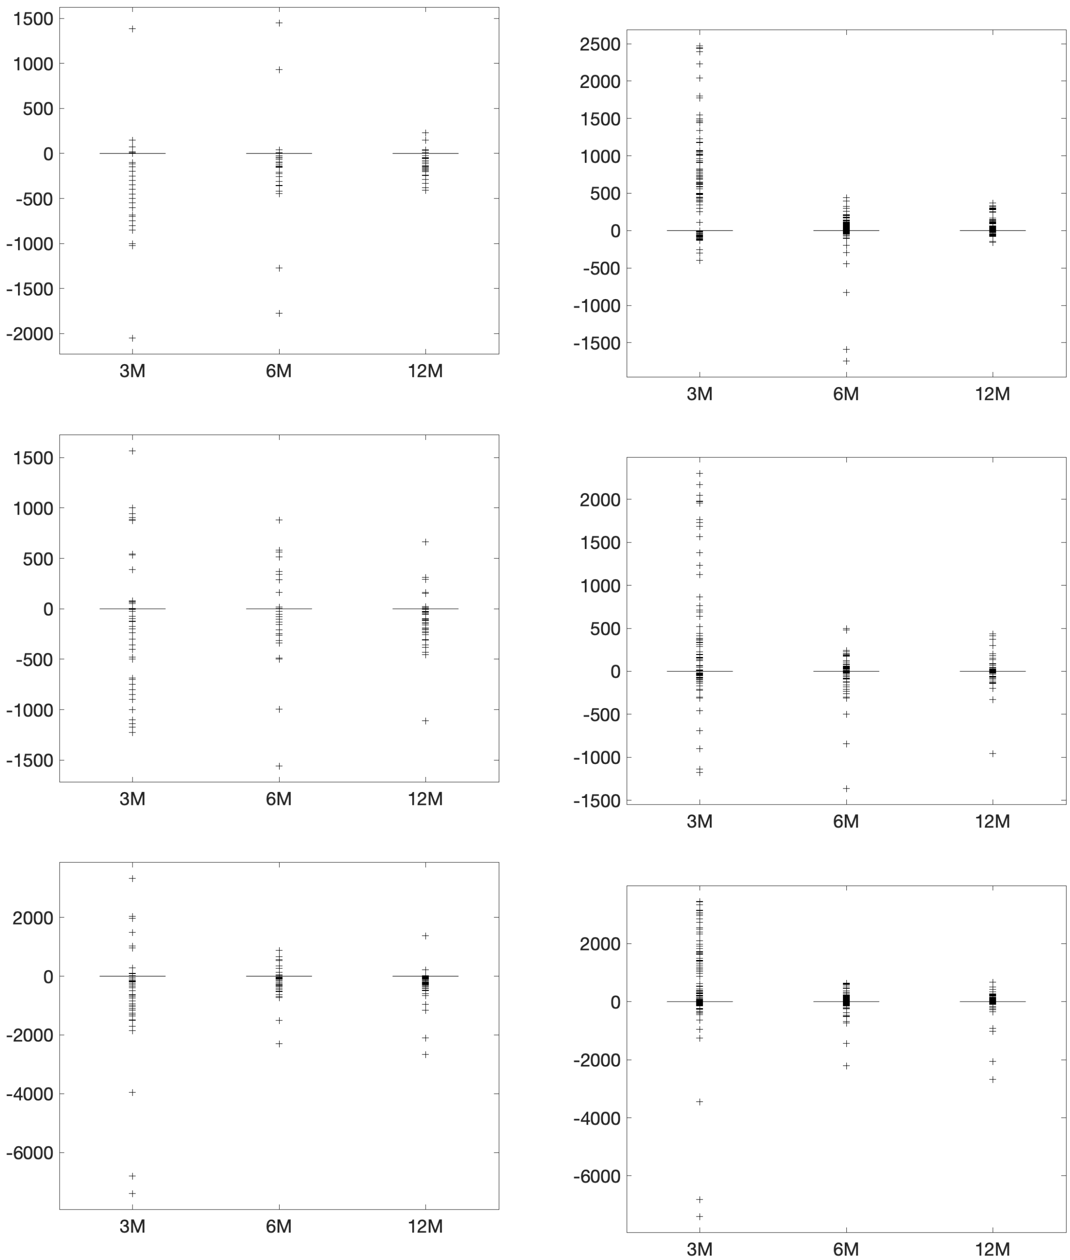

Online Appendix Figures O4a shows the distributions of the difference (in million C$) between the last two iterations of updating expectations in our statistical fixed point routine for all three maturities on three different auction days. We claim to find a fixed point (up to measurement noise) if the median difference is zero and there are only occasional outliers. Figure O4b shows the difference (in million C$) between the rescaled expectations (18) and the expectations that we find using our statistical fixed point routine. The median difference is again zero.

Online Appendix Figure O5: Illustration of the price-quantity trade-off (extended model with heterogeneous dealers)

(a) Uniform price auction

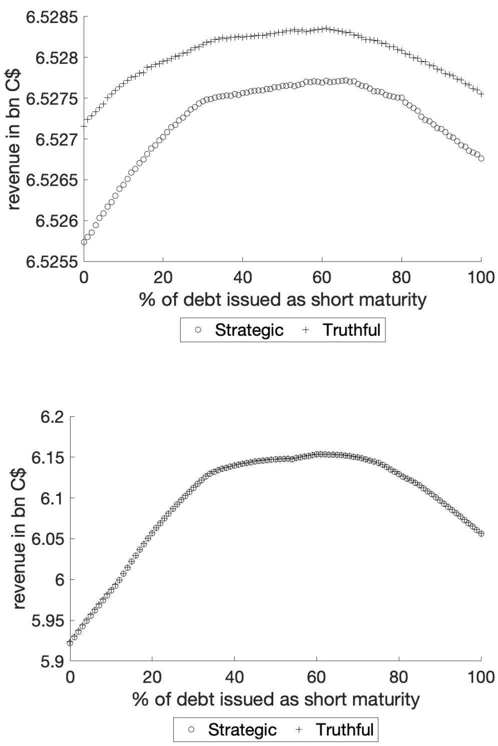

(b) Discriminatory price auction

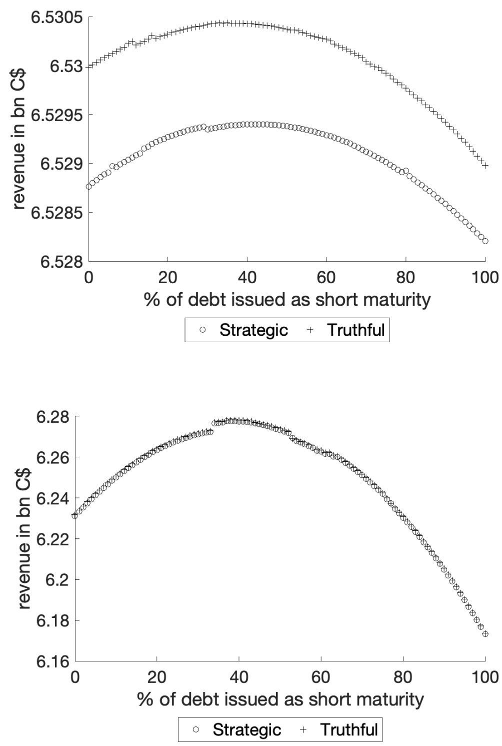

Online Appendix Figures O5 is the analogue to Figure 4 but using the extended model with heterogeneous dealers. depict the price-quantity trade-off when the auction is uniform price (a), and discriminatory price (b) using the estimated λ and δ parameters in the upper graphs and scaling the parameters by 100 in the lower graphs. On the y-axis is the total revenue earn from issuing both maturities (in billion C$) when issuing x% of the short maturity and (1-x)% of the long maturity. The x-axis scales up x from 0% to 100%.

Online Appendix Table O1: Demand coefficients with valuations with more than 3 steps

(a) Average dealer

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	-6.777	(0.034)	λ_{6M}	-11.81	(0.069)	λ_{1Y}	-24.46	(0.138)
$\delta_{3M,6M}$	-0.931	(0.074)	$\delta_{6M,3M}$	-2.396	(0.149)	$\delta_{1Y,3M}$	-6.336	(0.345)
$\delta_{3M,1Y}$	-0.171	(0.080)	$\delta_{6M,1Y}$	-0.552	(0.163)	$\delta_{1Y,6M}$	-2.647	(0.348)
N	55822			38856			46778	

(b) Dealer group 1

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	-6.165	(0.034)	λ_{6M}	-11.07	(0.069)	λ_{1Y}	-23.09	(0.140)
$\delta_{3M,6M}$	-1.158	(0.074)	$\delta_{6M,3M}$	-2.290	(0.146)	$\delta_{1Y,3M}$	-5.498	(0.344)
$\delta_{3M,1Y}$	-0.281	(0.080)	$\delta_{6M,1Y}$	-1.105	(0.163)	$\delta_{1Y,6M}$	-4.281	(0.352)
N	42937			30456			37820	

(c) Dealer group 2

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	-11.13	(0.106)	λ_{6M}	-17.29	(0.224)	λ_{1Y}	-35.04	(0.469)
$\delta_{3M,6M}$	$+0.236$	(0.237)	$\delta_{6M,3M}$	-1.608	(0.639)	$\delta_{1Y,3M}$	-7.224	(1.463)
$\delta_{3M,1Y}$	$+1.243$	(0.246)	$\delta_{6M,1Y}$	$+3.524$	(0.537)	$\delta_{1Y,6M}$	$+7.319$	(1.189)
N	12885			8400			8958	

Online Appendix Table O1 (a)-(c) are analogous to Tables 5 (a) and O4. They report the coefficients for equation (7), but estimated on a subsample of valuations estimated from bidding functions with strictly more than two steps, instead of one step. Valuations are in C$ and quantities in % of the auction supply. The first three columns show the estimates for the 3M bill auction, the next three for the 6M bill auction and the last three for the 12M bill auction. The point estimates are in the second, fifth and eight column. Standard errors are next to them in parentheses.

Online Appendix Table O2: Demand coefficients for the average dealer with trimmed valuations

(a) 3M Bill auction

markup	4 bps		10 bps		20 bps		40 bps	
λ_{3M}	−6.496	(0.031)	−7.767	(0.046)	−9.609	(0.075)	−12.89	(0.135)
$\delta_{3M,6M}$	−0.752	(0.069)	−1.692	(0.101)	−3.040	(0.163)	−5.499	(0.293)
$\delta_{3M,1Y}$	−0.040	(0.074)	−0.605	(0.108)	−1.449	(0.175)	−2.806	(0.314)
N	58542		58542		58542		58542	

(b) 6M Bill auction

markup	4 bps		10 bps		20 bps		40 bps	
λ_{6M}	−11.05	(0.061)	−13.62	(0.096)	−17.25	(0.162)	−23.75	(0.296)
$\delta_{6M,3M}$	−1.892	(0.134)	−4.350	(0.209)	−7.910	(0.351)	−14.02	(0.644)
$\delta_{6M,1Y}$	−0.308	(0.147)	−1.446	(0.228)	−2.994	(0.383)	−5.763	(0.701)
N	42282		42282		42282		42282	

(c) 1Y Bill auction

markup	4 bps		10 bps		20 bps		40 bps	
λ_{1Y}	−22.89	(0.123)	−29.14	(0.202)	−38.03	(0.345)	−54.03	(0.637)
$\delta_{1Y,3M}$	−5.102	(0.309)	−12.25	(0.507)	−23.42	(0.869)	−44.35	(1.603)
$\delta_{1Y,6M}$	−1.895	(0.312)	−5.630	(0.512)	−11.27	(0.877)	−21.93	(1.618)
N	50408		50408		50408		50408	

Appendix Table O2 (a)-(c) report the coefficients for equation (7), estimated using competitive bids of more than one step that were placed by dealers for different valuations of the markup (4 bps, 10 bps, 20 bps, 40 bps). The estimates for a markup of 5 bps, our favorite specification, are in the main text. Valuations are in C$, quantities % of auction supply. Standard errors are in parentheses next to the point estimates.

Online Appendix Table O3: Demand coefficients per dealer group with bids as independent variables

(a) Dealer group 1

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	−4.498	(0.023)	λ_{6M}	−7.266	(0.040)	λ_{1Y}	−14.59	(0.077)
$\delta_{3M,6M}$	−0.081	(0.051)	$\delta_{6M,3M}$	+0.538	(0.086)	$\delta_{1Y,3M}$	+0.710	(0.191)
$\delta_{3M,1Y}$	+0.305	(0.055)	$\delta_{6M,1Y}$	+0.145	(0.096)	$\delta_{1Y,6M}$	−0.070	(0.196)
N	45405			33464			40956	

(b) Dealer group 2

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	−8.879	(0.086)	λ_{6M}	−13.43	(0.183)	λ_{1Y}	−25.88	(0.340)
$\delta_{3M,6M}$	+1.613	(0.193)	$\delta_{6M,3M}$	+1.156	(0.526)	$\delta_{1Y,3M}$	+0.993	(1.072)
$\delta_{3M,1Y}$	+1.760	(0.201)	$\delta_{6M,1Y}$	+5.234	(0.442)	$\delta_{1Y,6M}$	+12.16	(0.875)
N	13137			8818			9452	

Online Appendix Tables O3 (a) and (b) are analogous to Table 5 (a). They report the coefficients for equation (7), but with the observed competitive bids by dealers with more than one step as independent variables rather than the estimated true valuations. Bids are in C$ and quantities in % of auction supply. The first three columns show the estimates for the 3M Bill auction, the next three for the 6M Bill auction and the last three for the 12M Bill auction. The point estimates are in the second, fifth and eight column. Standard errors are next to them in parentheses.

Online Appendix Table O4: Demand coefficients per dealer group with valuations as independent variables

(a) Dealer group 1

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	−6.107	(0.033)	λ_{6M}	−10.75	(0.066)	λ_{1Y}	−22.53	(0.135)
$\delta_{3M,6M}$	−1.158	(0.073)	$\delta_{6M,3M}$	−2.249	(0.142)	$\delta_{1Y,3M}$	−5.478	(0.336)
$\delta_{3M,1Y}$	−0.243	(0.078)	$\delta_{6M,1Y}$	−1.080	(0.158)	$\delta_{1Y,6M}$	−4.258	(0.344)
N	45405			33464			40956	

(a) Dealer group 2

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	−11.19	(0.106)	λ_{6M}	−17.42	(0.221)	λ_{1Y}	−35.75	(0.462)
$\delta_{3M,6M}$	+0.285	(0.237)	$\delta_{6M,3M}$	−1.666	(0.636)	$\delta_{1Y,3M}$	−6.957	(1.459)
$\delta_{3M,1Y}$	+1.216	(0.247)	$\delta_{6M,1Y}$	+3.748	(0.536)	$\delta_{1Y,6M}$	+7.607	(1.190)
N	13137			8818			9452	

Online Appendix Tables O4 (a) and (b) are analogous to Table 5 (b). They report the coefficients for equation (7). Valuations are in C$ and quantities in % of auction supply. The first three columns show the estimates for the 3M bill auction, the next three for the 6M bill auction and the last three for the 12M bill auction. The point estimates are in the second, fifth and eight column. Standard errors are next to them in parentheses.

Online Appendix Table O5: Bidding regression in 3M auction

	(1)	(2)	(3)	(4)
λ_{3M}	-3.750	-3.726	-3.713	-3.722
	(0.0320)	(0.0316)	(0.0318)	(0.0318)
$\delta_{3M,6M}$	0.316	0.678	0.289	0.551
	(0.0428)	(0.0351)	0.0441)	(0.0531)
$\delta_{3M,12M}$	0.348	0.359	0.170	0.0857
	(0.0438)	(0.0347)	(0.0423)	(0.0567)
N	59718	59583	59583	59583

Online Appendix Table O5 shows the estimation results of regression (1) but using different explanatory variables: the estimated expected winning quantities in column (1), the actual winning quantities in column (2), the the amount the bidder demanded at the highest step that he ever wins of a maturity in a year in column (3) and the total amount the bidder demands at auction if this amount is less than the 1% highest amount the bidder ever wins of the maturity during a year in column (4). We use all bids and not just final bids of dealers similar to Table 5. Therefore, the estimates in column (2) are not identical to the estimates of Table 3. Bids and valuations are in C$ and quantities in % of auction supply. Standard errors are in parentheses, clustered at the bidder level.

Online Appendix Table O6: Average gain (in bps) per auction when reshuffling 1% of debt in the extended model with heterogeneous dealers

		$S \uparrow L \downarrow$ Uniform	$S \uparrow L \downarrow$ Discrim	$S \downarrow L \uparrow$ Uniform	$S \downarrow L \uparrow$ Discrim
Independent:	factor$_\lambda$=1, factor$_\delta$=0	+0.021	+0.006	−0.024	−0.008
Weak substitutes:	factor$_\lambda$=1, factor$_\delta$=1	+0.012	−0.004	−0.020	+0.001
Perfect substitutes:	factor$_\lambda$=1, $\delta = \lambda$	+0.011	−0.056	−0.015	+0.047
Independent:	factor$_\lambda$=1, factor$_\delta$=0	+0.234	−0.029	−0.295	−0.002
Weak substitutes:	factor$_\lambda$=1, factor$_\delta$=1	+0.227	−0.036	−0.290	+0.007
Perfect substitutes:	factor$_\lambda$=1, $\delta = \lambda$	+0.089	−0.586	−0.208	+0.589
Independent:	factor$_\lambda$=1, factor$_\delta$=0	+2.365	−0.448	−2.996	+0.185
Weak substitutes:	factor$_\lambda$=1, factor$_\delta$=1	+2.361	−0.445	−2.992	+0.144
Perfect substitutes:	factor$_\lambda$=1, $\delta = \lambda$	+1.009	−6.591	−2.113	+6.523

Online Appendix Table O6 is analogous to Table 7 but builds on the extended model with two dealer groups (market makers and non-market makers). It shows the revenue gains when issuing 1% of debt more for the short maturity and 1% less of the long maturity in the second and third column ($S \uparrow L \downarrow$) and vice versa in the fourth and fifth column $S \downarrow L \uparrow$ when the auction format is uniform price (Uniform) and when it is discriminatory price (Discrim). The first three rows (factor$_\lambda$=1) correspond to the demand estimates of the 6M and 12M bills assuming different degrees of substitution. The fourth-sixth row and seventh-ninth row correspond to hypothetical auctions in which the λ^χ parameters in the bidder's demand are scaled by a factor of 10, and 100, respectively. The revenue gain is in bps of the original revenue.