FTG Working Paper Series

Estimating Demand Systems for Treasuries

by

Milena Wittwer
Jason Allen
Jakub Kastl

Working Paper No. 00111-00

Finance Theory Group

www.financetheory.com

*FTG working papers are circulated for the purpose of stimulating discussions and generating comments. They have not been peer reviewed by the Finance Theory Group, its members, or its board. Any comments about these papers should be sent directly to the author(s).

Estimating Demand Systems For Treasuries*

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

June 9, 2023

Abstract

Leveraging the fact that Treasury auctions of different maturities are often held simultaneously, we propose a method for estimating demand systems for Treasuries, avoiding the usual endogeneity issues in demand estimation. We implement our method using bidding data from Canadian T-bill auctions, and find that different types of bills are only weak substitutes, despite their cash-like nature. We provide a micro-foundation of demand in the primary market to explain this finding, and illustrate how demand elasticities, together with the auction format, determine how to allocate debt on a given day across maturities.

Keywords: Demand estimation, multi-unit auctions, Treasuries

JEL classification: D44, C14, E58, G12

*This paper has been presented under different titles since 2017. An earlier version is part of Milena Wittwer's EUI PhD thesis, published in July 2018. The presented views are those of the authors, not necessarily of the Bank of Canada. 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 also thank Nikhil Agarwal, Rui Albuquerque, Giulia Brancaccio, Markus Brunnermeier, Jens Christensen, Jeremy Fox, Matthew Gentry, Alessandro Gavazza, Pierre-Olivier Pineau, Eric Richert, Marzena Rostek, Paulo Somaini, Andreas Uthemann, Ji Hee Yoon, and Yu Zhu, as well as seminar/conference participants at the Bank of Canada, Boston College, Chicago Booth, Columbia, EUI, Kellogg, Princeton, Penn, Stanford, Queen's University; EARIE 2018/2022, ECB workshop on money markets and central bank balance sheets 2018, EEA 2022, EWM 2021, Joint Bank of Canada-John Deutsch Institute Workshop on Financial Intermediation and Regulation 2018, MIT DSE conference 2022. 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, trillions of dollars' worth of related securities and commodities are sold concurrently. This includes Treasury securities of different maturities, various types of mortgage backed securities, international carbon allowances, and diverse forms of renewable energy. Similarly, exchanges provide a wide array of stocks that investors can trade simultaneously. Technically, all these markets clear through auctions known as multi-unit auctions, wherein bidders demand or supply multiple units of the same good (e.g., Wilson (1979); Kyle (1985); Chen and Duffie (2021); Rostek and Yoon (2021); Wittwer (2021); Budish et al. (2022)).

We leverage institutional features of parallel multi-unit auctions to develop a new framework for identifying demand schedules for multiple assets or commodities. We implement our method using data from Canadian Treasury auctions, and show how to use the demand estimates to change the maturity composition of debt in such a way that lowers the cost of government financing without changing the total amount of debt or the auction format.

The key idea behind our identification of demand systems lies in combining two institutional features. First, parallel auctions take place under the same market rules, with the same set of participants, at the same time, and within the same economic environment. This allows us to control for unobserved heterogeneity. Second, in multi-unit auctions, bidders submit full demand schedules. As a result, there is no need to aggregate data across different time periods and market participants. This approach ensures 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 tackles this issue by employing instrumental variables and making exogeneity and validity assumptions (following Berry et al. (1995) in industrial organization, and Koijen and Yogo (2019) in finance).[1]

Our main data set, detailed in Section 2, includes information on all Canadian T-bill auctions between 2002 until 2015.[2] We observe which security and maturity type is issued

[1]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)).

[2]While Canadian Treasury bonds are sold in different days, the U.S. and other large economies issue both bills and bonds in parallel. Our methodology is easily adapted to these other settings, provided that access to bidding data with unique bidder identifiers is available.

at what amount, all submitted bids, i.e., demand schedules, and where the market clears. We can identify each bidder thanks to unique identifiers, and know whether the bidder is a dealer (typically a large bank) who bids directly to the auctioneer, or a customer (such as a pension or hedge fund) who must bid via a dealer. Additionally, we collect market-level data that allow us to measure the (money-like) premium associated with short bills due to their money-like characteristics (following Greenwood et al. (2015b)), and novel trade-level information from the secondary market on all cash and repo trades.

With these data, we start our analysis by providing descriptive evidence on the interdependence in demand for T-bills in Section 3. For instance, we show that dealers, who observe a customer's bid in one auction, adjust their own bids in all auctions, indicating interdependencies. While one would expect T-bills to be close to perfect substitutes given their cash-like nature, we show that the money-like premium of 3M bills does not change in the same manner when increasing the amount outstanding of 3-month (3M) vs. 6-month (6M) vs. 12M-month (12M) bills, indicating that different bills are imperfect substitutes at best.

The descriptive evidence motivates us to introduce and estimate a structural model of bidding in parallel multi-unit auctions in Section 5. The model allows us to identify the full demand system for T-bills, while we can only get a sense of some of the own-and cross-price demand elasticities from the descriptive evidence. The model overcomes two 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 contingent on prices of multiple securities. This means that, unless demand for different maturities is independent, we only observe parts of the demand schedules.

We find demand for all three maturities of bills to be rather price-insensitive and weakly substitutable. For instance, when the average dealer wins 1% more of the supply of 12M bills, his offer for the 12M bills decreases by 0.24 basis points (bps). Further, if this 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.

To help explain why T-bills are only weakly substitutable, we examine trade-level data from the secondary market. We show that longer maturities are more frequently utilized in repo transactions compared to shorter maturities. Furthermore, distinct investor groups exhibit a tendency to trade specific maturities. These findings indicate that different maturities serve different purposes and imply the presence of a preferred habitat, contributing to the imperfect substitution among T-bills.

Motivated by this evidence, we set up a micro-foundation of demand to better understand what drives substitution patterns in Treasury auctions in Section 4. We assume that bidders—who are mostly primary dealers—buy bonds in the primary market to sell them to preferred habitat clients who have tastes for specific maturities (in the spirit of Vayanos and Vila (2021)). Dealers want to buy enough securities to satisfy client demand because it is costly to turn down a client. Substitutability in dealer demand decreases when it is more costly to turn down a client from a diverse client base.

In the last part of the paper, Section 6, we explain why it is useful for debt managers to know demand elasticities in Treasury auctions. Debt managers aim at issuing debt at the lowest funding cost to the taxpayer.[3] A starting point to achieve this goal is to split the amount of debt to be issued on the same day more strategically across different auctions to achieve a higher revenue. Although such small changes to the maturity composition of debt do not achieve the optimal allocation, it is straightforward to implement; for instance, because it does not require making changes to the calendar of debt issuance, which is determined at the beginning of a fiscal year. Further, it does not require changing established auction formats, which is typically difficult. For example, Klemperer (2010) proposes a combinatorial auction format to sell multiple financial securities. Despite excellent theoretic properties the auction hasn't become popular in practice.[4]

We illustrate how a government can exploit differences in demand elasticities to obtain higher auction revenues. Interestingly, the answer depends on the auction format that is used to allocate debt. For instance, in Canada, where debt is issued via discriminatory price auctions (where bidders pay their bids) it is revenue-increasing to issue more of the price-sensitive bond (longer maturity) and less of the price-insensitive bond (shorter maturity).

[3]For instance, the U.S. office of debt management states its objective here: `https://home.treasury.gov/system/files/276/Debt-Management-Overview.pdf`, accessed on 02/03/2022

[4]Another example is that it took a long academic debate with early contributions by Bikhchandani and Huang (1989) and Back and Zender (1993), as well as 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)).

In the U.S., where debt is allocated via uniform price auctions (where all bidders pay the market clearing price), the opposite is true. Quantitatively, revenue gains for the Canadian government from reshuffling bills are negligible in our sample period (in which demand is price-insensitive and interest rates are low). Revenue gains 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)).

Literature. Our main contribution is to propose and implement a new method for estimating demand systems for multiple assets, and to show how to use these demand systems to obtain higher revenues in Treasury auctions. This adds to different strands of the literature, in addition to those already highlighted.

Our demand analysis complements a literature that analyzes 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)). We consider different types of short-term debt (i.e., bills), more similar to Nagel (2016), Greenwood et al. (2015b), and Krishnamurthy and Li (2022). In contrast to these papers, we zoom in on the primary market, where we can identify demand of individual institutions which act as dealers.

Our methodological 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 security, but also on prices of securities of other maturities. This complements a handful of alternative estimation strategies for identifying product synergies in other auction formats, such as package auctions, sequential auctions, and ascending auctions (e.g., Cantillon and Pesendorfer (2006); Kim et al. (2014); Jofre-Bonnet and Pesendorfer (2003); Kong (2021); Fox and Bajari (2013)).

Of these papers, Gentry et al. (2022) is the closest to ours. They identify preferences in simultaneous single-object first-price auctions with complementarities that depend on observables. We show how to identify interdependencies that depend on unobservables. Moreover, we leverage multi-unit auctions to estimate full demand schedules, establishing a link to standard "BLP" demand estimation following Berry et al. (1995). For this we derive equilibrium conditions for interconnected multi-unit auctions with common empirical features. This stands in contrast to existing theoretical contributions which consider elegant and tractable, yet less realistic settings not suited for estimation (e.g., Wittwer (2020, 2021); Chen and Duffie (2021); Rostek and Yoon (2021)).

Lastly, our counterfactual exercises contribute to two literatures. The first analyzes and proposes changes to (multi-unit) auction rules and formats (see Wilson (2021) for an overview). We reshuffle supply across maturities without changing the total amount of debt or auction format. 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 by including an issuance (or roll-over) cost of debt that absorbs the (mechanical) price difference of bonds with different maturities. Instead, we highlight how a government can reduce its cost of financing by a little bit by making small changes to its current supply split across maturities.

2 Institutional Environment and Data

We use data on Canadian T-bill auctions to illustrate how to leverage institutional features of simultaneous multi-unit auctions to estimate demand systems of interdependent securities, or goods more broadly.

Treasury Auctions. Government bills and bonds of different maturities are issued in separate sealed-bid multi-unit auctions in many countries, including the U.S., Japan, Brazil, France, China, and Canada. In such an auction, bidders submit step-functions 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—used, for instance, in Brazil,

Figure 1: Bids in the Canadian T-Bill Market

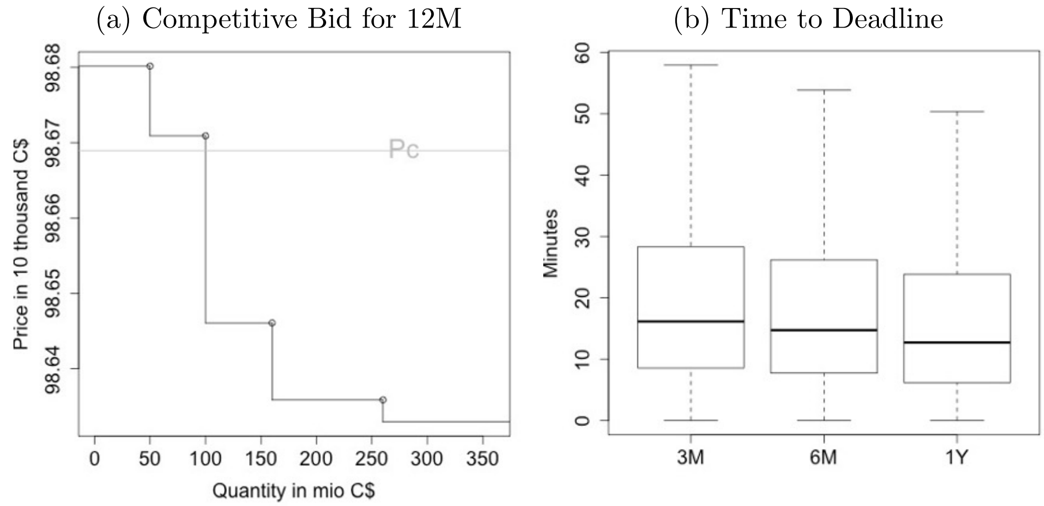

Figure 1a displays an example of a bidding step function. It is the one of the median dealer in a 12M auction, computed as follows: Determine the median number of steps in all competitive bid functions submitted by dealers, and then take the median over all (price, quantity) tuples corresponding to each 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.

Canada, France, Italy, U.K.—, each bidder pays according to what they bid, similar to a first-price auction. In a uniform price auction—used, for instance, in Argentina, Korea, Switzerland, U.S., Norway—, each bidder pays the market clearing price for each unit won.[5]

Canadian T-Bill Auctions. There are three types of T-bills (3M, 6M, 12M) which are sold every second Tuesday by the Bank of Canada in three separate, but parallel, discriminatory price auctions. T-bonds of different maturities are sold on different days. Two groups of bidders, dealers and customers, participate in the auctions. 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, which come in two forms. The first form is a competitive bid. This is a step-function with at most 7 steps. Quantities are denominated in multiples of C$1,000, while values are expressed in terms of yield to maturity. To simplify the analysis, we convert

[5] For detailed descriptions of other empirical settings with multi-unit auctions, see, for instance: Hortaçsu (2002), Février et al. (2004), 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, France, Korea, Czech Republic, U.S., Columbia, Mexico, China; Cramton et al. (2013) for diamonds; Hortaçsu and Puller (2008) and Reguant (2014) for electricity; Ryan (2022) for renewable energy.

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$.

the yields into prices by employing a face value of C$1 million. This conversion ensures that the demand schedules are decreasing, rather than increasing (as shown in Figure 1a).

The second form of bidding is a non-competitive bid. This is a quantity order that the bidder will win for sure, at 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).

Auction Data. To estimate demand systems in multi-unit auctions, the data set must include all winning and losing bids, the issued supply, the market clearing price of each auction, and how much each bidder wins. Our data set is richer than this. It consists of all 366 Canadian T-bill auctions between 2002 and 2015, summarized in Table 1.

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, and distributes roughly C$5.76 in the form of 3M and C$2.12 billion in the form of 6M and 12M bills (see

Appendix Figure A2). Dealers tend to participate in all parallel auctions to keep their dealer status (see Bank of Canada (2016), p. 12, for details).

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).

The 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 (other than the Bank of Canada) 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 of placing competitive bids.

Market Data. In order to further support our understanding of how demand for different T-bills are related, we complement our auction data with market-level information.

We obtain the daily amount outstanding of all securities, and the daily average market yield for all maturity categories (3M, 6M, 12M) between 2002 and 2015 from the Bank of Canada. Further, we download the time series of daily yields curves for zero-coupon bonds.[6] Of these daily yields, we use the shortest maturity (3M). It tells us the yield at which the 3M bills should trade based on an extrapolation from the rest of the yield curve, and is useful to capture the money-like premium that the market is willing to pay for the shortest bills. In addition, we collect the time series of nominal Canadian GDP, which is reported quarterly.

Finally, we acquire trade-level information from the Debt Securities Transaction Reporting System, MTRS2.0, collected by IIROC since November 2015. The sample contains all bond trades with registered brokers or dealers in the cash market from 2016 to 2019 (as in Allen and Wittwer (2023)). In addition, we add the repo market from 2016 until 2022. For each transaction, we observe the time, price, and size of trade, the security traded, and the type of trade (buy/sell/repo/reverse-repo). We also know who trades, which is useful, for instance, to study whether different investor types have preferences for different securities.

[6] The data is available at: `https://www.bankofcanada.ca/rates/interest-rates/bond-yield-curves/`, accessed on 03/17/2023.

Most comparable data sets, collected in other countries, such as TRACE for the U.S., do not provide this information.

3 Motivating Interdependencies in Demands

Before estimating demand for different securities, it is useful to present evidence suggesting that studying auctions for individual securities in isolation provides an incomplete picture of demand. Throughout the paper, we use the same coefficients in different regressions; $\delta \neq 0$ always suggests that different bills are not independent.

Interdependencies. Different securities 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 securities. In our setting, supply is perfectly correlated across auctions (see Table 2). In fact, over our long sample the Bank of Canada always announces the exact same issuance size for the 6M and 12M bills.[7]

On the demand side, dealers might want to buy bundles of securities to satisfy the demand of consumers downstream post-auction. In our data, the total amount dealers demand 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 dealers do not always achieve this goal. Overall, these cross-maturity correlations suggest that dealers don't value different bills as independent.

To provide evidence for this conjecture, we exploit the institutional feature of Canadian auctions that customers must place their bids with a dealer, who can observe these bids before passing them to the auctioneer. We know that a dealer updates his own bid, say, for the 3M bills, upon observing a customer bid in the 3M auction (Hortaçsu and Kastl (2012)).

[7] 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}
$\bar{Q}_{3M}$	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.

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:

$$update_{m,i} = \alpha + \sum_m I_m \left(\lambda_m customer_m + \delta_{m,-m} customer_{-m}\right) + \varepsilon_{m,i}. \tag{1}$$

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 dealers 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.

Our findings, reported in Table 3, confirm the conjecture that demand for different bills are interdependent, given that the $\delta_{m,-m}$ coefficients are positive and statistically significant.

Table 3: Probability of Dealer Updating Bids

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

Table 3 shows the results of the Probit regression (1). 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.

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.

Type of Interdependencies. The evidence presented so far suggests that demand for different maturities are interdependent, but it does not tell us anything about the type of interdependency. Intuitively, we expect bills with different maturities to be substitutes, given that they are cash-like. Thus, we conjecture that a bidder is willing to pay more for one maturity when he expects to win less of the other maturity.

With bidding data, we can test whether the bidder offers more or less for one security when he wins more of the other securities. For this, we regress bid k of dealer i on day t for security 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:

$$b_{t,m,i,k} = \alpha_m + \lambda_m q_{t,m,i,k} + \sum_{l \neq m} \delta_{m,l} won_{t,l,i} + \zeta_t + \zeta_i + \epsilon_{t,m,i,k}. \tag{2}$$

Counterintuitively, we find that all δ parameters are statistically significant and positive, suggesting that bills might be complements (see Table 4). One explanation for this odd finding is that the estimates are biased because the bid and the amount won are determined

Table 4: 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 4 shows the estimation results of regression (2) 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.

simultaneously in equilibrium. Our structural model will allow us to solve the endogeneity problem.

Using data from outside of the auctions, we provide evidence that does not suffer from the same endogeneity problem as in equation (2). Inspired by Greenwood et al. (2015b) who provide evidence that U.S. T-bills and longer-term bonds are imperfect substitutes, we analyze how the money-like premium of the 3M responds to changes in the total amount outstanding of the 3M, 6M, and 12M bills relative to GDP. If 3M bills and 6M bills were perfect substitutes, the 3M premium would increase by the same amount independently of whether there is more of the 3M or of the 6M bills. In Appendix A we show that this is not the case, which indicates that different bills are at best imperfect substitutes.

One reason for imperfect substitutability could be that different types of bills serve different purposes. For example, data from the repo market suggest that 12M bills are more heavily used in repo agreements than 3M bills (see Figure 2). Another reason could be that different investors prefer different maturities, as highlighted by Figure 3. Public entities, including governments, are more likely to trade longer maturities, perhaps to reduce the frequency by which they need to roll-over debt, whereas asset managers and banks are more likely to trade shorter maturities, perhaps to fulfill liquidity requirements.

Take Away. Taken together, the evidence suggests that bills are interdependent, but not perfectly substitutable. This motivates the need for a methodology that can identify full demand systems of interdependent securities. Before doing so, however, it is useful to take a step back and ask what drives the demand of dealers.

Figure 2: Volume Traded in the Repo Market

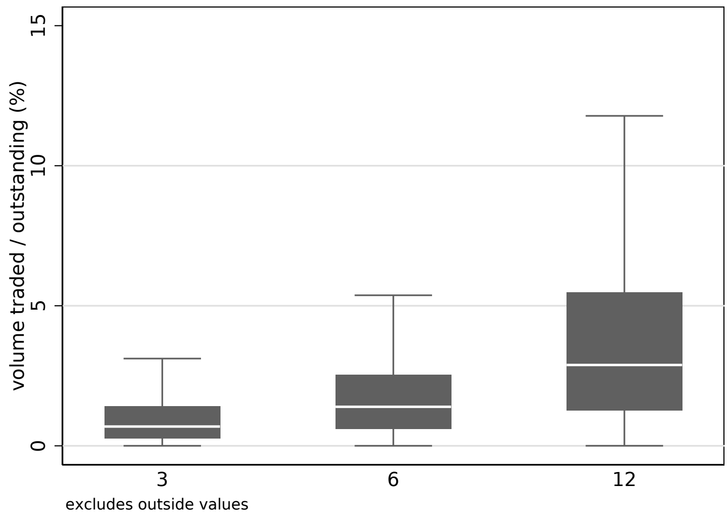

Figure 2 shows the daily amount exchanged in repurchase transactions and buy/sell-backs (which are both types of repo) that involve 3M, 6M, and 12M bills, as a percentage of total amount outstanding of bills in the respective categories. This information is constructed from MTRS.2 data, using transactions between 01/2016 until 01/2022.

Figure 3: Preferred Habitat for Short (3M) vs. Long (6M/12M) Bills

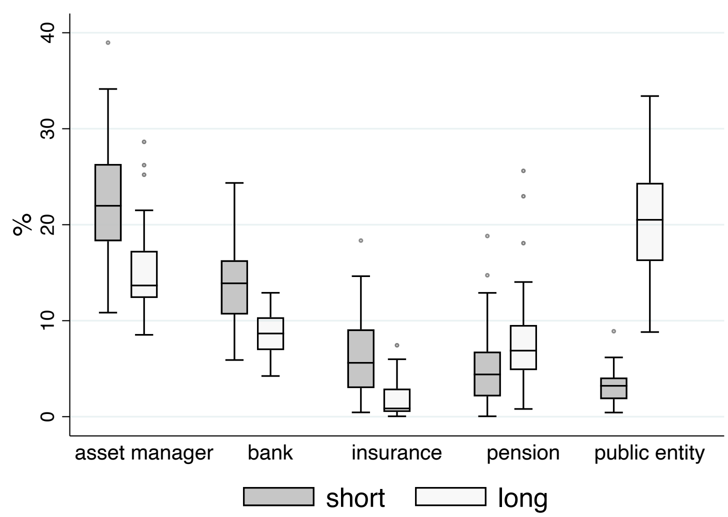

Figure 3 shows two box plots for each of the largest types of investors: asset managers, banks, insurance companies, pension funds and public entities. A point in one of the gray box plots tell us the percentage of total volume of short (3M) bills traded on a day that is executed by the respective investor type, for instance asset managers in case of the first box plot from the left. A point in one of the white box plots shows the analogue for longer bills (6M & 12M). In all box plots, the distribution is taken over days. The data is taken from Allen and Wittwer (2023) who have access to trade-level information of essentially all cash-trades in the Canadian government bond market from 01/2016 until 12/2019.

4 Micro-Foundation of Dealer Demand

We present a model that helps explain what might drive demand in the primary market. Our micro-foundation is by no means exclusive. There are many other reasons that might drive interdependencies and heterogeneity in dealer demand, including regulatory or budget constraints, as well as different repo needs.

Motivated by the evidence presented in Figure 3, our model features market segmentation in the spirit of Vayanos and Vila (2021). Dealers buy securities with different maturities in the primary market to sell them to investors who have preferences for specific maturities (preferred habitat) in the secondary market. We restrict the number of maturities to $M = 2$, and drop the superscript g and the subscripts i, τ for simplicity. Generalizing the model to more than two maturities is straightforward but mathematically cumbersome and brings little additional insight. All proofs are in Appendix F.

Each dealer has a type $\boldsymbol{s}$, which decomposes into ν (known by all dealers) and $\boldsymbol{t}$ (private information): $\boldsymbol{s} = (\boldsymbol{t}, \nu)$ with $\boldsymbol{t} = (\boldsymbol{t_1}, \boldsymbol{t_2})$ and $\nu = (a, b, e, \gamma, \kappa_1, \kappa_2, \rho)$. Here $\boldsymbol{s}$ and $\boldsymbol{t}$ are in **bold** because they are random variables—a convention we adopt throughout the paper.

Rather than assuming that dealers are risk-averse, we assume that they face a cost of not meeting investor demand.[8] A dealer who draws type s obtains the following gross benefit from "consuming" amounts $(1 - \kappa_1)q_1$ and $(1 - \kappa_2)q_2$:

$$U(q_1, q_2, s) = t_1(1 - \kappa_1)q_1 + t_2(1 - \kappa_2)q_2. \tag{3}$$

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 investors with random demand $\{\boldsymbol{x_1}, \boldsymbol{x_2}\}$ arrives to the dealer. Equivalently, you may imagine that there are two types of investors, 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. Specifically,

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

$\{\boldsymbol{x_1}, \boldsymbol{x_2}\}$ 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 investors 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 investors' 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, e.g., for maturity 1:

$$p_{i,1}(x_1, x_2 | q_1, q_2) = \begin{cases} a - bx_1 - ex_2 & \text{for } x_1 \leq \kappa_1 q_1 \text{ and } x_2 \leq \kappa_2 q_2 \\ a - bx_1 & \text{for } x_1 \leq \kappa_1 q_1 \text{ and } x_2 > \kappa_2 q_2 \\ 0 & \text{for } x_1 > \kappa_1 q_1 \text{ and } x_2 > \kappa_2 q_2. \end{cases} \tag{4}$$

There are three cases. In the first, investors 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 investors 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 investor demand, he cannot sell.

Overall, 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 investors. Large investors (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:

$$revenue(x_1, x_2|q_1, q_2) = p_1(x_1, x_2|q_1, q_2)x_1 + p_2(x_1, x_2|q_1, q_2)x_2. \tag{5}$$

Turning down investors is costly for the dealer. An unhappy investor 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 investor demands in the secondary market so as to avoid losing his investor in the longer run. This is costly for the dealer because it is expensive to borrow or buy additional T-bills on the secondary market when demand is high. In our model, dealers face the following cost function:

$$cost(x_1, x_2|q_1, q_2) = \begin{cases} 0 & \text{if } x_1 \leq \kappa_1 q_1 \text{ and } x_2 \leq \kappa_2 q_2 \\ \gamma x_1 & \text{if } x_1 > \kappa_1 q_1 \text{ and } x_2 \leq \kappa_2 q_2 \\ \gamma x_2 & \text{if } x_1 \leq \kappa_1 q_1 \text{ and } x_2 > \kappa_2 q_2 \\ \gamma x_1 x_2 & \text{if } x_1 > \kappa_1 q_1 \text{ and } x_2 > \kappa_2 q_2. \end{cases} \tag{6}$$

This function captures the idea that it is more costly to turn down larger investors, 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. This means that the marginal cost from turning down a investor 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 in the primary market:

$$V(q_1, q_2, s) = U(q_1, q_2, s) + \mathbb{E}\left[revenue(\boldsymbol{x_1}, \boldsymbol{x_2}|q_1, q_2) - cost(\boldsymbol{x_1}, \boldsymbol{x_2}|q_1, q_2)\right], \tag{7}$$

which determines how much he is willing to pay on-the-margin: $v_m(q_m, q_{-m}, s) = \frac{\partial V(q_m, q_{-m}, s)}{\partial q_m}$.

Proposition 1. *The marginal willingness to pay for amount q_m of a dealer with signal $s = (s_m, s_{-m})$, conditional on winning q_{-m} in the other auction $-m$ can be approximated by*

$$v_m(q_m, q_{-m}, s) = f_m(s_m) + \lambda_m q_m + \delta_m q_{-m} \tag{16}$$

for $m = 1, 2 - m \neq m$, where $f_m(s_m) = \alpha_m + (1 - \kappa_m)t_m$ and $\alpha_m, \lambda_m, \delta_m$ are polynomials of parameters $\{\kappa_1, \kappa_2, \gamma, \rho, a, b, e\}$.

To derive an intuition for the willingness to pay and to see why bills might be substitutable or complementary, let us contrast the extreme cases when the dealer keeps all of maturity 1 ($\kappa_1 = 0$), when he keeps all of maturity 2 ($\kappa_2 = 0$), or sells all of both ($\kappa_1 = \kappa_1 = 1$) under independent investor demand ($\rho = 0$). The willingness to pay for maturity 1 is:

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

When buying only for its own account 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 at least some of maturity 1, his demand in auction 1 decreases in q_1 as long as his investors are sufficiently price-elastic, i.e., b is sufficiently high. If he sells all of both maturities, the demand is independent of his private type t_1, but depends on the amount he wins of the other maturity. Bills can be substitutes or complements depending on the relative magnitudes of γ and e.

Expanding beyond our specific example, we derive the following result that sheds light on the factors influencing the degree of substitutability.

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

The corollary yields two noteworthy implications. First, it highlights that bills might be strongly substitutable in the secondary market ($e > 0$), but weakly substitutable in the primary market. Second, the corollary tells us that for dealers, for whom it is less costly to turn down investors (low γ)—for instance, thanks to their central position in the trading network—bills are stronger substitutes than for other dealers. This motivates a model extension, presented in Appendix C, where dealers are differentiated based on their latent business type as market makers or non-market makers.

Take Away. Our micro-foundation provides insights into the factors that can drive demand in Treasury auctions, revealing that this demand may differ from aggregate investor demand in the secondary market. These insights will be instrumental in the estimation and interpretation of demand systems.

5 Identifying and Estimating Demand Systems

It is challenging to consistently estimate the full demand system using bidding data of Treasury auctions for two reasons. First, banks have private information about how much they value securities. This generates incentives to misrepresent their true demand. As in a first-price auction, bidders shade their bids to reduce the total payments they must make to win, which implies that 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 security m, not on securities $-m$ (such as in Figure 1a). This implies that we only observe parts of the demand system. To solve these challenges, we introduce and estimate a structural model.

5.1 Model of Simultaneous Multi-Unit Auctions

M perfectly divisible goods (here bills), indexed m (for maturity), are auctioned in M parallel multi-unit auctions to $g \geq 1$ groups of bidders who participate in all auctions.[9] In our application, there are dealers (d) and customers (c), and the auction format is discriminatory price. However, our approach generalizes to uniform price auctions. The number of bidders in each group, N_d, N_c, is commonly known. In addition, all or some bidders could carry a latent type, as illustrated in a model extension in Appendix C.

Each bidder i of group g draws a private, potentially multi-dimensional signal $\boldsymbol{s_i^g} = (\boldsymbol{s_{1,i}^g}, \ldots, \boldsymbol{s_{M,i}^g})$. A bidder's signal must be independent from the other bidders' signals conditional on anything that everyone knows at the time of the auction. 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 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

[9]Given our data, auction participation is exogenous (recall Table 1). In other settings, it might be natural to identify groups of bidders who participate in sets of auctions based on observable characteristics. Such a model extension would be straightforward. A more complicated extension would be to endogenize the bidder's decision to participate in each auction. Allen et al. (2023) take first steps in this direction.

idiosyncratic factors, such as the structure of the balance sheet, client order flows, investment opportunities, or repo needs—which do not depend on private information of other dealers. Consistent with this, Hortaçsu and Kastl (2012) and Allen et al. (2023) conduct statistical tests to provide support for the assumption that values are conditionally independent and private.[10] To capture the fact that balance sheet factors and client order flows are likely correlated across different maturities, we let signals for different maturities m of the same bidder i be affiliated.

Assumption 1. *(i) Dealers' and customers' private signals* $\boldsymbol{s_i^d}$ *and* $\boldsymbol{s_i^c}$ *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*. (ii) Signals for different maturities* m *of the same bidder* i *are affiliated.*

The bidder's signal affects his true marginal willingness to pay, or (inverse) demand. Formally, the marginal willingness to pay or "value" of a dealer with signal $s^g_{m,i}$ for amount q_m conditional on purchasing vector q_{-m} of the other securities $-m$ is $v^g_m(q_m, q_{-m}, s^g_i)$.

Assumption 2. $v^g_m(q_m, q_{-m}, s^g_i)$*, is non-negative, bounded, strictly increasing in* s^g_i *for all* q_m*, and non-increasing and continuous in* q_m *for all* s^g_i *and* m*.*

We say that two maturities, for instance, 3M and 6M, are substitutable if the dealer is willing to pay less for 3M bills when purchasing more of the 6M bills: $\frac{\partial v^g_{3M}(q_{3M}, q_{-3M}, s^g_i)}{\partial q_{6M}} < 0$ for all q_{3M}, q_{-3M}. We call maturities complementary if this partial derivative is positive at all quantities and independent if it is zero at all quantities. If the sign is not consistent across quantities we use the term mixed.

Over the course of the auction, each bidder may place (competitive) bids at a discrete number of time slots $\tau = 0, ..., \Gamma$, where $\Gamma \geq 0$. Each time a bidder places a bid, he thinks that this might be his last bid, and therefore the one that counts for market clearing (as in Hortaçsu and Kastl (2012)).[11] Each time a customer seeks to bid, he is matched to a

[10] 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 values in such settings without having to make strong functional form assumptions remains an open question in the literature. Bonaldi and Ruiz (2021) take a first step in this direction for uniform price auctions.

[11] Formally, 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

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 M auctions, his information set is $\boldsymbol{\theta_i^g} = (\boldsymbol{s_i^g}, \boldsymbol{Z_{1,i,\tau}}, ..., \boldsymbol{Z_{M,i,\tau}})$. If he only has a customer in one auction, say for maturity 1, $\boldsymbol{\theta_{i,\tau}^g} = (\boldsymbol{s_i^g}, \boldsymbol{Z_{1,i,\tau}})$, and so on.

A bid in auction m consists of a set of quantities in combination with prices—a step-function which characterizes the price the bidder would like to pay for each amount. To compare bids in auctions with different sizes of supply, we normalize quantities to be between 0 and 1, representing the share of total supply. A zero bid denotes non-participation.

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

$$A_m = \begin{cases} (b_m, q_m, K_m) : dim\,(b_m) = dim(q_m) = K_m \in \{1, ..., \overline{K}_m\} \\ b_{m,k} \in [0, \infty) \;\; and \; q_{m,k} \in [0,1] \\ b_{m,k} > b_{m,k+1} \;\; and \; q_{m,k} > q_{m,k+1} \forall k < K_m. \end{cases} \tag{8}$$

When the auctions close, the auctioneer aggregates the final bids within each auction. Each auction clears at the lowest price P_m^c at which aggregate demand satisfies aggregate supply. The supply (to competitive bidders) can be random, and thus unknown to bidders when placing their bids, for instance, due to non-competitive bids.

Assumption 4. *Supply $(\boldsymbol{Q_1}, ..., \boldsymbol{Q_M})$ is drawn from a distribution with strictly positive marginal density on support $[\underline{Q}_1, \overline{Q}_1] \times ... \times [\underline{Q}_M \overline{Q}_M]$ where $\underline{Q}_m \leq \overline{Q}_m$ for all m.*

If aggregate demand equals total supply exactly there is a unique market clearing price P_m^c. 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 (see Kastl (2011) for details).

We focus on group-symmetric Bayesian Nash equilibrium (BNE), which is natural given that bidders from the same group are ex ante identical. To define such BNE, note that a

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.

pure-strategy is a mapping from the bidder's set of information sets to the action space (8) of all auctions, and that a choice in auction m by a bidder who draws information $\theta^g_{i,\tau}$ may be summarized as bidding function $b^g_{m,i}(\cdot, \theta^g_{i,\tau})$. Further, let $q^c_i = (q^c_{1,i} \ldots q^c_{M,i})$ denote the amount bidder i gets allocated when submitting $b^g(\cdot, s^g_i)$, and $V(q^c_i, s_i)$ be the gross payoff from obtaining q^c_i, given by $v^g_m(q_m, q_{-m}, s^g_i) = \frac{\partial V^g(q_m, q_{-m}, s^g_i)}{\partial q_m}$.

Definition 1. *A BNE is a collection of functions $b^g_{i,\tau}(\cdot, \theta^g_{i,\tau}) \equiv (b^g_{1,i,\tau}(\cdot, \theta^g_{i,\tau}) \ldots b^g_{M,i,\tau}(\cdot, \theta^g_{i,\tau}))$ that for each bidder i and almost every signal $\theta^g_{i,\tau}$ maximizes his expected total surplus, that is, his gross payoff minus the total payment:*

$$\mathbb{E}\left[V(q^c_i, s^g_i) - \sum_{m=1}^{M} \int_0^{q^c_{m,i}} b^g_{m,i,\tau}(x, \theta^g_{i,\tau}) dx \middle| \theta^g_{i,\tau}\right]. \tag{9}$$

The BNE is group-symmetric if $b^g_{i,\tau}(\cdot, \theta^g_{i,\tau}) = b^g(\cdot, \theta^g_{i,\tau})$ for all g, i, τ.

Given our focus on dealers, who buy most of the issued amount in our sample, we characterize their equilibrium. Kastl (2011) proves that there is an equilibrium in distributional strategies in a discriminatory price auction with k-step functions. His proof generalizes to the case of simultaneous auctions.

For ease of notation, we drop the d-subscript for the remainder of the paper.

Proposition 2. *Consider a dealer i with information $\theta_{i,\tau}$ who submits $\hat{K}_{m,i,\tau}$ steps in auction m. In any group-symmetric BNE every step k in his bid function $b_m(\cdot, \theta_{i,\tau})$ has to satisfy*

$$\tilde{v}_{m,}(q_{m,k}, s_i|\theta_{i,\tau}) = b_{m,k} + \frac{\Pr\left(b_{m,k+1} \geq \boldsymbol{P^c_m}|\theta_{i,\tau}\right)}{\Pr\left(b_{m,k} > \boldsymbol{P^c_m} > b_{m,k+1}|\theta_{i,\tau}\right)}(b_{m,k} - b_{m,k+1}) \; \forall k < \hat{K}_{m,i,\tau} \tag{10}$$

$$\text{with } \tilde{v}_m(q_{m,k}, s_i|\theta_{i,\tau}) \equiv \mathbb{E}\left[v_m\left(q_{m,k}, \boldsymbol{q^*_{-m,i}}, s_i\right)\middle| b_{m,k} \geq \boldsymbol{P^c_m} > b_{m,k+1}, \theta_{i,\tau}\right] \tag{11}$$

for all m with $-m \neq m$, and $b_{m,k} = \tilde{v}_m(\bar{q}_m(\theta_{i,\tau}), s_i|\theta_{i,\tau})$ at $k = \hat{K}_{m,i,\tau})$ where $\bar{q}_m(\theta_{i,\tau})$ is the maximal amount the bidder may be allocated in an equilibrium.

To build an intuition for the equilibrium condition, it helps to review what we know about bidding in isolated discriminatory price auctions. Here, a dealer 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).

In simultaneous auctions of related securities, the dealer's demand for one security, say q_1, depends on how much of the other securities he gets allocated, $v_1(q_1, q_{-1}, s_{1,i})$. Ideally, the dealer 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, $q^*_{-1,i} \equiv (q^*_{2,i} \dots q^*_{M,i})'$. 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, $\mathbb{E}\left[v_1\left(q_{1,k}, \boldsymbol{q^*_{-1,i}}, s_i\right) \middle| b_{1,k} \geq \boldsymbol{P^c_1} > b_{1,k+1}, \theta_{i,\tau}\right]$, 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.

5.2 Identification

The goal is to learn about each dealer's demand system, given by $v_m(q_m, q_{-m}, s_i)$, at all q_m, q_{-m} for all $m, -m$ using bidding data. This can be done in various ways, depending on details of the auction setting that generates the data. In all cases, we first infer "pseudo-values", $\tilde{v}_m(q_{m,k}, s_i|\theta_{i,\tau})$, at each submitted step, from how a dealer bids, under the assumption that all bidders play the equilibrium strategies of Proposition 2. Generalizing Kastl (2012), we know that these pseudo-values are point-identified. Then, we back out the actual values, $v_m(q_{m,k}, q_{-m}, s_i)$, that are implied by the pseudo-values given (11) and the observed distribution of winning quantities, $\boldsymbol{q^*_{m,i}}$ at step k. This can be done non-parametrically.

Here, we provide two non-parametric identification approaches to back out actual values from pseudo-values before discussing the parametric approach that we implement with our data. For illustration, we consider the case with 3 maturities ($M = 3$), and assume that the distribution of a dealers' winning quantities has finite support. The insights from this example generalize to settings with more maturities and continuous winning quantities.

To implement the first identification approach, we need to observe sufficiently many dealer updates between the auction opening and closing. Under Assumption 1, the demand-relevant private information, s_i, does not change across updates τ. Therefore, with sufficiently many updates, we can use variation in updated bids to identify values, and with that the shape of the dealer's demand. To see this, note that we have $\Gamma \times K$ equations having estimated pseudo-values, $\tilde{v}_m(q_{m,k}, s_i|\theta_{i,\tau})$, at each step k and update τ:

$$\tilde{v}_{3M}(q_{3M,k}, s_i|\theta_{i,\tau}) = \sum_{h=1}^{H}\sum_{r=1}^{R} p_{3M,i,\tau}^{k,h,r} v_{3M}\left(q_{3M,k}, (q_{6M,h}, q_{12M,r}), s_i\right) \text{ for } \tau = 1, ..., \Gamma, \tag{12}$$

$$\text{where } p_{3M,i,\tau}^{k,h,r} = \Pr(\boldsymbol{q^*_{6M,i}} = q_{6M,h} \cap \boldsymbol{q^*_{12M,i}} = q_{12M,r}|b_{3M,i,k} \geq \boldsymbol{P^c_{3M}} > b_{3M,i,k+1}, \theta_{i,\tau}) \tag{13}$$

is the probability that $q_{6M,h}$ and $q_{12M,r}$ will be the quantities awarded in the 6M and 12M auction, respectively, conditional on winning the k'th step in the 3M auction. This is a system of $\Gamma \times K$ linear equations in $H \times R \times K$ unknowns. If we have $H \times R$ updates at the same $q_{3M,k}$ for each step k, there is a unique solution given an appropriate rank condition.

If we don't observe bid updates, we can still identify demand non-parametrically if we are willing to assume that private signals are drawn iid across auctions and dealers from the same signal distribution, F^g, given in Assumption 1. In this setting, we use observed variation in the auction environment that affects the dealer's information, but not his true willingness-to-pay for a given signal s_i, such as exogenous variation in the available supply.

To see this, assume, for illustration, that there are L auctions that are ex-ante identical but differ in their supply distribution (specified in Assumption 4). This means that true demand for a given signal, $v_m(q_{m,k}, q_{-m}, s_i)$, is the same across auctions within the same group, but pseudo-values, $\tilde{v}_m(q_{m,k}, s_i|\theta_{i,l})$, differ since information, $\theta_{i,l}$, differs across auctions l. Given that $\tilde{v}_m(q_{m,k}, s_i|\theta_{i,l})$ is strictly increasing in s_i for any $l, q_{m,k}$ (by Lemma 1 in Appendix F), any quantile p of the distribution of pseudo-values at some step k, corresponds to the quantile of the signal distribution: $Quantile_p(\tilde{v}_m(q_{m,k}, s_i|\theta_{i,l})) = Quantile_p(s_i)$. Using variation across auctions l, by holding the quantile of the pseudo-values at the same quantity point $q_{m,k}$ fixed, we can follow the same arguments as above to identify the underlying demand functions just by replacing τ (the update index) with l (the exogenous shifter index) in equation (12).

Once we know the shape of $v_m(\cdot, \cdot, s_i)$, it is straightforward to back out the degree of interdependency. For instance, the dependency between 3M and 6M bills when winning $(q_{3M}, q_{6M}, q_{12M})$ is

$$\delta_{3M,6M}(q_{3M}, q_{6M}, q_{12M}, s_i) = \frac{\partial v_{3M}(q_{3M}, (q_{6M}\ \ q_{12M}), s_i)}{\partial q_{6M}}. \tag{14}$$

With finite data, we can point-identify finite points on $v_m(\cdot, \cdot, s_i)$ and approximate (14) by

$$\delta^{k,h,r}_{3M,6M}(s_i) = \frac{v_{3M}(q_{3M,k}, (q_{6M,h+1}\ q_{12M,r}), s_i) - v_{3M}(q_{3M,k}, (q_{6M,h}\ q_{12M,r}), s_i)}{q_{6M,h+1} - q_{6M,h}}\ \forall k, h, r. \quad (15)$$

Proposition 3. *Under Assumptions 1-4, $\delta^{k,h,r}_{3M,6M}(s_i)$ is identified non-parametrically. Moreover, when there are sufficiently many updates within an auction, $\delta^{k,h,r}_{3M,6M}(s_i)$ is also identified without requiring Assumption 4, or single-dimensionality of signals, and allowing signal distributions to differ across auction days. The analogous result holds for all other maturity combinations.*

In our empirical application we choose to implement a middle ground. We would like to avoid pooling data across auctions because the distributions of signals may easily shift over time, especially after important macroeconomic events, or monetary policy announcements. However, we do not observe sufficiently many updates to identify demand non-parametrically by relying on variation across updates. Therefore, as an approximation, we rely on our microfoundation (Proposition 1), and assume that the dealer's willingness to pay is linear:

$$v_m(q_m, q_{-m}, s_{m,i}) = f_m(s_{m,i}) + \lambda_m q_m + \delta_m q_{-m}, \quad (16)$$

with $\lambda_m < 0$, $|\delta_m| < \lambda_m$, and sufficiently large $f_m(s_{m,i}) > 0$ for all $s_{m,i}$ such that the marginal willingness to pay does not drop below 0 for any amount that might be for sale. Importantly, δ_m and q_{-m} are vectors when there are more than two maturities—a simplified notation we adopt throughout the paper. The δ_m vector captures cross-maturity demand coefficients. For example, in the $3M$ auction, where $q_{-m} \equiv (q_{6M}\ q_{12M})'$ and $\delta_m \equiv (\delta_{3M,6M}\ \delta_{3M,12M})$, 3M and 6M bills are substitutes when $\delta_{3M,6M} < 0$, complements when $\delta_{3M,6M} > 0$, and independent otherwise.

With linear demand, we can identify the demand coefficients (λ_m, δ_{-m}) using variation across steps k of pseudo-values of a bidder with some unknown signal s_i and information $\theta_{i,\tau}$.

5.3 Estimation

To estimate the dealer's pseudo-values on auction day t, we need to estimate the probabilities that the market clears at different prices, or equivalently the distribution of the winning quantities of each bidder. For this we extend resampling techniques of the auction literature to take dependencies across auctions into account. Here we describe the simplest resampling procedure with simultaneous auctions with N ex ante identical bidders and no updates; in Appendix B, we discuss the more complex procedure we implement which takes customer bidding and dealer updating into account.

We fix a triplet of bids simultaneously submitted by a bidder i. We then draw a random subsample of $N-1$ bid-vector triplets with replacement from the sample of bids to construct M realizations of residual supply curves that the bidder faces in the auctions. Next we determine at what points the residual supply curves intersect with the fixed bidding function. This gives one realization of the clearing prices $P_t^c = (P_{t,3M}^c\ P_{t,6M}^c\ P_{t,12M}^c)$ and the amount $q_{t,i}^* = (q_{t,3M,i}^*\ q_{t,6M,i}^*\ q_{t,12M,i}^*)$ this bidder would have won for all $q_{t,i}^*, P_t^c$. Repeating this procedure many times provides a consistent estimate of the joint distribution of market clearing prices and the corresponding amount of each security i would win.[12]

Using the joint distributions, we derive the pseudo-values, abbreviated by $\hat{v}_{t,m,i,k}$, implied by the observed bids and the necessary equilibrium condition stated in Proposition 2. Additionally, we estimate the expected amount each dealer anticipates winning from the other maturities $-m$. In settings without bid updating, neither the pseudo-values nor the expectations are dependent on the time at which the bid was placed. For instance, we have $\hat{\mathbb{E}}q_{t,-m,i,k}^* \equiv \mathbb{E}\left[\boldsymbol{q_{t,-m,i}^*} | b_{t,m,i,k} \geq \boldsymbol{P_{t,m}^c} > b_{t,m,i,k+1}, \theta_{t,i}\right]$. In settings with bid updates, both the pseudo-values and the expectations become time-dependent.

Finally, with these expectations and the estimated pseudo-values, we estimate all demand coefficients in the dealer's demand (16) in a linear regression with dealer-auction-time fixed effects $u_{t,m,i} = f_{t,m}(s_{t,m,i})$, and with $\varepsilon_{t,m,i,k}$ representing noise coming from the resampling procedure:

$$\hat{v}_{t,m,i,k} = u_{t,m,i} + \lambda_m q_{t,m,i,k} + \delta_m \hat{\mathbb{E}} q_{t,-m,i,k}^* + \varepsilon_{t,m,i,k} \quad \forall m, m \neq -m. \tag{17}$$

[12] Cassola et al. (2013) establish consistency of the estimator for a single discriminatory price auction. Their proof generalizes to simultaneous discriminatory price auctions.

We use a subsample of competitive dealer bids with at least two steps, which includes virtually all dealer bids (see Appendix Figure A4). Our findings remain robust even when we narrow our analysis to bids with at least three steps, as demonstrated in Appendix Table A2.[13]

5.4 Estimated Demand Coefficients

As a starting point, we estimate equation (17) using observed bids instead of estimated values. The results are presented in Table 5 (a). Notably, all the estimated δ coefficients are positive and statistically significant. This finding is consistent with regression (2) reported in Table 4, which utilizes the amounts bidders actually win rather than the amounts they expect to win when bidding. The positive coefficients indicate that dealers are willing to pay higher prices when they win more of the other maturities. This result suggests that bills are complements, which contrasts with the conventional notion in the literature that securities of similar term and risk are substitutes.

In the next step, we re-estimate the regressions using the estimated values to examine whether bid-shading leads to biased estimates. In contrast to the case of using bids, all δ coefficients, reported in Table 5 (b), are now negative, indicating that bills are substitutes.[14] This finding underscores the importance of eliminating bid-shading and utilizing the true values to accurately identify interdependencies, even when shading factors appear small.

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

[13] Since the expected amounts, $\hat{\mathbb{E}}q^*_{t,-m,i,k}$, are estimated, they might be subject to measurement with error, which implies that the estimates of regression (17) might be biased. While it is challenging to grasp the size of such bias, we provide in Appendix D some suggestive evidence that in our empirical application the bias might be small.

[14] Comparing the δ coefficients across auctions, one may notice that the estimates are not symmetric. For example, $\delta_{3M,6M} \neq \delta_{6M,3M}$. 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.

In previous versions of the paper, we mistakenly reported that bills are complements. This error occurred due to a typo in the estimation code, which led to the estimated values being heavily trimmed, resulting in their resemblance to the submitted bids. Consequently, the biased estimation results were similar to those presented in 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 Values 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 (17), 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.

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 estimated δ coefficients imply that the dealer's value for the 6M bill decreases by approximately C$20.65 when obtaining the average amount of the 3M (7.3% of supply) and 12M bills (6.9% of supply), compared to receiving nothing. In the 3M auction the corresponding price decrease is about C$7.14 and in the 12M auctions about C$63.27. These price drops are not negligible when compared to the difference between the maximum and minimum bid in the average bidding function, which is C$142.

Take Away. Our analysis emphasizes that bills are, at best, imperfect substitutes, despite their cash-like nature. In our extended model presented in Appendix C, we uncover that larger dealers perceive bills as substitutes, while smaller dealers may perceive bills as complements, consistent with Corollary 1. This finding sheds light on the nuanced perspectives of different market participants regarding the substitutability of bills.

6 Policy Takeaway

Knowing demand elasticities is useful for debt managers who try to minimize the cost of debt funding. One practical approach is to more strategically allocate the total amount of debt to be issued on a given day across different maturities, with the goal of maximizing auction revenues. To illustrate this approach, we consider the supply split of two securities (S and L) pairwise, and ask under what conditions total revenue increases when issuing a little bit more of one security, and a little bit less of the other, holding the total amount of supply constant.[15] This depends on three factors: price levels, price sensitivities, and the auction format.

Price Levels. Shorter maturities are typically sold at higher prices than longer maturities given that the Treasury yield curve is upward sloping in normal times: $P_S > P_L$. One explanation for this is that investors attach a money-like premium to shorter bills. Additionally, the higher price of shorter bills can reflect the costs and risks associated with more frequent rollovers compared to longer bills. Rolling over debt involves additional auctions and exposes the government to the risk of interest rate fluctuations.

The price difference between short and long bills suggests that, if the government aimed to maximize auction revenue on a single day, it would issue only short bills. In practice, however, governments don't take this strategy, because they seek to maximize revenues over a long horizon. For instance, they consider the fact that short bills need to be rolled over more frequently than long bills to sustain the same expenditure level. This dynamic maximization problem is complex and beyond the scope of this paper.

Our goal is to emphasize the indirect revenue effect resulting from different market price sensitivities. To isolate this effect, we introduce a price wedge ($c_S = P_S - P_L$) that eliminates the mechanical price effect causing short bills to generate higher revenues than long bills.[16]

[15] It is straightforward to include non-competitive bids in the revenue calculations. We exclude them because most of these bids are allocated to the auctioneer, and therefore represent a revenue-neutral in-house transfer.

[16] There are different ways to micro-found the price-wedge. For instance, Greenwood et al. (2015a)'s model "generates a simple trade-off between the monetary services provided by issuing more short-term debt, and the increased rollover risk that comes as a result" (p. 1700). At issuance, the price of the long bond is $P_L = 1$, while the price of the short bond is $P_S = 1 + c_S$, with c_S reflecting the money premium. Alternatively, Bigio et al. (2021) solve for the equilibrium price of a perfectly competitive auction that is followed by over-the-counter trading. In their model, c_S is shaped by different liquidity costs associated with

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. To fix ideas, let $P_m(Q_m)$ sum all bids for security $m \in \{S, L\}$ given supply Q_m, and assume, for illustration, that $P_m(\cdot)$ is differentiable (which is not the case in the data), and that the market clears at $\{Q_m, P_m\}$. 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-security effect (the λ's)—but also on how this sensitivity changes when winning more in the other auction—the cross-security effect (the δ's).

In our case, the aggregate demand for the long maturity is typically more price-sensitive than for the short maturity, which means that the price for the long bills responds more strongly to a change in auction supply than the price for the short bill. This is true when bills 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 security (here the long bill) and more of the price insensitive security (here the short bill), without changing total supply (see Figure 4 (a)-(b)). The reason is that everyone pays the market prices, and the market price of the long bill increases more strongly than the market price of the short bill decreases.

However, there is a price-quantity trade-off. Starting from an equal supply split across securities, when the auctioneer moves one dollar from the long into the short bill, the price of the short bill drops less than the price of the long bill increases. Thus, while the revenue of the short bill auction decreases, the revenue of the long bill 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 bill auction given that the higher price for the long bill is multiplied by a smaller and smaller amount. In the extreme, when the auctioneer goes from a mixed supply split to issuing only short bills, no one pays the hypothetical high price for the long

trading the short versus long bond.

Figure 4: Simplified Example of Revenue Under two Auction Formats

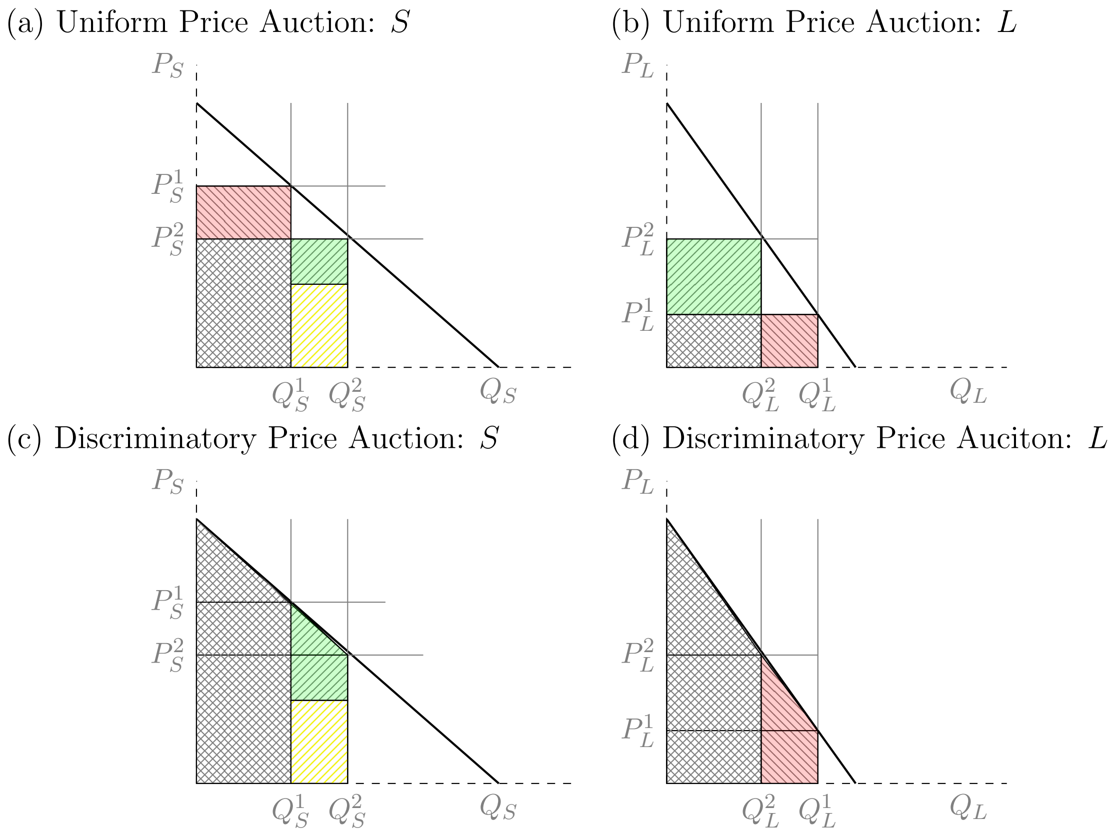

The figures on the LHS show the revenue gain (in green) and loss (in red) when issuing dQ more of S, as well as the change coming from the price-wedge (in yellow). The figures on the RHS show the analogous changes in revenue when issuing dQ less of L. In all cases, we assume that $\frac{\partial P_m(Q_m)}{\partial Q_m} = -\mu_m$ 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 $dQ > 0, \mu_L > 0, \mu_S > 0$.

bill 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 4 (c)-(d)). There are two differences. First, shifting supply towards the short bill may decrease total revenue. Second, the revenue of one auction is determined by the area underneath the aggregate demand curve.

If aggregate demand curves were linear, as in Figure 4, and no bidder adjusted their 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 this easily.

From Theory to 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 4 (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.

Second, bidders submit step functions 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.

Given these complications it becomes an empirical question as to whether total revenue increases or decreases when reshuffling supply—a question we can answer with our empirical framework (see Appendix E for details). As proof-of-concept, we use demand schedules with prices and quantities, both expressed in C$, and reshuffle the 6M and 12M bills, keeping the 3M bills at the observed amount. Reshuffling from 12M to 3M bills, fixing the 6M bills, would lead to slightly higher revenue impact since demand for the 3M bills is less price sensitive than demand for the 6M bills.

Example. For illustration, we consider one auction day in our sample in Figure 5, and shift 1% of total debt from the 12M to the 6M auction. This decreases the market clearing price of 6M bills by C$ 5, and increases the one of 12M bills by C $25. Given that demand curves are relatively flat, the associated revenue gain is small: +0.09 bps in a uniform price auction and −0.04 bps in a discriminatory price auction. In other markets, for instance, the Portuguese government bond market, demand curves are steeper (as shown by Albuquerque et al. (2022)). If we make demand curves steeper by scaling up all λ's by a factor of 10, the revenue effects become larger in both auction formats. For instance, the revenue loss in the discriminatory price auction is −0.25 bps when bills are independent ($\delta = 0$), and −1.55 bps when they are perfect substitutes ($10\lambda = 10\delta$). The non-negligible difference between these predictions highlights the importance of taking substitution patterns into account when comparing revenues across auction formats.

Figure 5: Example of Reshuffling Supply

(a) Aggregate demand for S (b) Aggregate demand for L

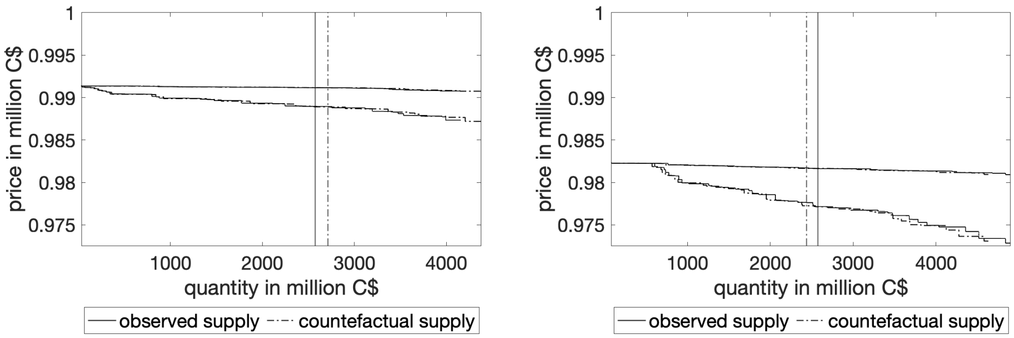

Figure 5 shows aggregate demand curves for two auctions that took place on the same day at some point in our sample. Each graph plots four curves. Two of the curves look rather flat. In 5a, the flat curves correspond to the aggregate demand for 6M bills when the Bank of Canada issues the supply as we observe it, and when we increase the supply of the 6M bills by 1% of the total debt issued on that auction day. In 5b, 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.

Average Revenue Gains. The example's qualitative insights generalize to the full sample. On average it is revenue-decreasing to issue more of the short bill and less of the long bill in discriminatory price auctions, and vice versa in uniform price auctions (see Appendix Table A5). Further, 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; it is less clear for uniform price auctions (see Appendix Figure A6).

Figure 6 visualizes the price-quantity trade-off for both auction formats. When the format is uniform price, revenue increases when going from issuing no short bills to issuing some short bills 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 bills. In a discriminatory price auction, we see a similar pattern, but the highest revenue gain is achieved when issuing less of the short (39%) and more of the long bill (61%).

Back-of-the-Envelope Calculation. Due to the low price-sensitivity in the demand for bills, cost-savings from changing the current maturity-split are very small. Issuing 1% of the 12M bills and 1% less of the 6M bills gives an average revenue gain of 0.001 bps per

Figure 6: Illustration of the Price-Quantity Trade-off

(a) Uniform Price Auction (b) Discriminatory Price Auction

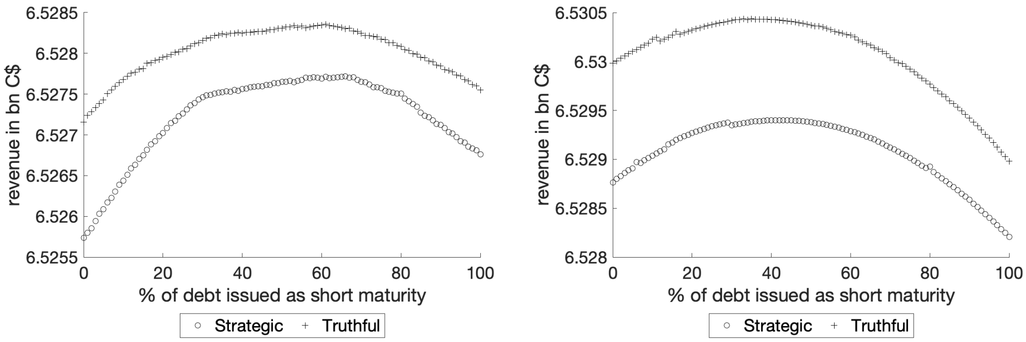

Figures 6 depict the price-quantity trade-off when the auction is uniform price (a), and discriminatory price (b) using the estimated λ and δ parameters. We consider two specifications. First, we allow bidders to shade their bids (Strategic). This involves solving a complicated fixed-point problem, as explained in Appendix E. Second, we assume that they submit their actual willingness to pay (Truthful) as a sanity check. 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%.

auction. For 2021, when the government issued C$416 billion in form of bills, this implies annual savings of C$41,000.

In other markets, in which demand is more price-sensitive, savings would be larger.[17] 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 price sensitivity could lead to sizable annual cost savings for taxpayers.

Take Away. Summarizing, we introduce a simple framework to guide auctioneers in their decision on how to split securities across auctions without changing the auction format, which can be challenging in practice. The main idea is that auctioneers should behave like monopolists who price discriminate subject to auction rules. We show that it is generally revenue-increasing to issue more of the price-insensitive security and less of the price-sensitive security when the auction is uniform price, and vice versa when it is discriminatory price.

[17] 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 $1/\text{-}379 = -0.0026$. With Canadian data, we find an average price sensitivity of -0.0012 if we, like Albuquerque et al. (2022), assume that securities are independent.

7 Conclusion

Leveraging institutional features of parallel multi-unit auctions, we develop a methodology to identify demand systems capable of accommodating 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., Hortaçsu, A., Richert, E., and Wittwer, M. (2023). Entry and exit in Treasury auctions. Working Paper.

Allen, J. and Wittwer, M. (2023). Centralizing over-the-counter markets? *Journal of Political Economy*, forthcoming.

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.

Budish, E., Cramton, P., Kyle, A. S., Lee, J., and Malec, D. (2022). Flow trading. Working paper.

Cantillon, E. and Pesendorfer, M. (2006). Combination bidding in multi-unit auctions. Working paper.

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.

Chen, D. and Duffie, D. (2021). Market fragmentation. *American Economic Review*, 111(7):247–2274.

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.

Février, P., Préget, R., and Visser, M. (2004). Econometrics of share auctions. Working Paper, University of Chicago and INRA.

Fox, J. T. and Bajari, P. (2013). Measuring the efficiency of an FCC spectrum auction. *American Economic Journal: Microeconomics*, 5(1):100–146.

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.

Jofre-Bonnet, M. and Pesendorfer, M. (2003). Estimation of a dynamic auction game. *Econometrica*, 71:1442–1489.

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.

Kim, S., Olivares, M., and Weintraub, G. (2014). Measuring the performance of large-scale combinatorial auctions: A structural estimation approach. *Management Science*, 60:1180–1201.

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.

Kong, Y. (2021). Sequential auctions with synergy and affiliation across auctions. *Journal of Political Economy*, 129(1):148–181.

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.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335.

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.

Reguant, M. (2014). Complementary bidding mechanisms and startup costs in electricity markets. *The Review of Economic Studies*, 81(4):1708–1742.

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

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

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.

ONLINE APPENDIX

Estimating Demand Systems for Treasuries

by Jason Allen, Jakub Kastl, and Milena Wittwer

The appendix has five sections. Section A provides additional descriptive evidence that bills of different maturities are not perfectly substitutable. Section B describes our resampling procedure. Section C presents a model extension that allows dealers to carry latent business types. Section D analyzes the potential bias stemming from measurement error in the estimated expected winning quantities. Section E provides details on how we solve the fixed point problem when determining the counterfactual bids. Section F presents all proofs.

A Additional Descriptive Evidence

We follow Greenwood et al. (2015b) (GHS) and construct daily "z-spreads" of the shortest available maturity (3M) to analyze how these spreads correlate with the total amount outstanding of the 3M, 6M, and 12M bills. The 3M z-spread on any given day represents the extent to which the yield of 3M T-bills differs from what one would expect based on an extrapolation of the rest of the yield curve on that day. Intuitively, z-spreads "reflect a money-like premium on short-term T-bills, above and beyond the liquidity and safety premia embedded in longer term Treasury yields" (GHS, p. 1687).

GHS show that the z-spread of the 10-week U.S. T-bill increases more strongly when the total amount outstanding of U.S. T-bill increases than when the total amount outstanding of U.S. bonds increases. This suggests that bills and bonds are imperfect substitutes. If they were perfect substitutes, the z-spread should change by the same amount independently of whether there are more bills or bonds.

Extending their logic, we estimate the following regression:[18]

$$z\_spread_t = \alpha + \sum_m \delta_m outstanding_{m,t} + \gamma t + \epsilon_t, \tag{18}$$

[18] Note that regression (18) is equivalent to regressing the daily 3M yield on the daily amount outstanding of each of the maturities, a control variable that captures unobservable factors that affect the term structure (the 3M yield that is predicted based on the rest of the yield curve on that day), and a time trend.

where z_spread_t is the 3M z-spread on day t, and $outstanding_{m,t}$ is the amount outstanding on day t of bills in maturity class, normalized by GDP. Following GHS we include a linear time trend, t, and standard errors are assumed to follow an AR(1) process.

As a sanity check, we replicate GHS's results and also regress the z-spread on the total amount of bills versus bonds outstanding, rather than the amounts outstanding of different types of bills. As GHS point out, the OLS regression with bills and bonds might be biased because yields and supply are simultaneously determined in equilibrium. In particular, we are worried about unobservable shocks that drive up the demand for bills and lead the government to issue more bills than bonds. We are less concerned about this possibility when estimating the analogues regression with different types of bills, because the Bank of Canada has followed a strict rule during our sample period. They always issue the exact same amounts of 6M and 12M bills, and 1.2-1.3 times this amount of 3M bills (see Appendix Figure A2).

We find that the 3M z-spread increases by about 0.2 bps when there is 1% more of 3M bills; by about 0.4 bps when there is 1% more of the 6M bills and by about 1 bps when there is 1% more of the 12M bills (see Appendix Table A7 (a)). The fact that the coefficients increase in maturity could suggest that 3M bills are scarce relative to the longer bills. Importantly, all three coefficients are statistically significantly different from one another, suggesting that bills of different maturities are at best imperfect substitutes.

In line with GHS's findings, the z-spread increases by much less (0.05 bps) when there is 1% more of bonds relative to an increase in bills (see Appendix Table A7 (b)). Our effects are smaller than GHS's, perhaps because the average interest rate level was much lower in our sample (2002–2015) compared to theirs (1983–2009).

B Resampling Procedure

The resampling procedure that we adopt is more complex than the one described in the main text due to the fact that customers place bids via dealers and dealers may update their bids (as in Hortaçsu and Kastl (2012)).

Three complications arise in this setting. First, bids may not be submitted at the exact same time given electronic or human delays (recall the example in Appendix Table A1).

We define bids to be "simultaneous" if they are the closest bids placed by a bidder within a 200-second window or if they are the final bids made before the auction deadline. The choice of a 200-second upper bound seems reasonable based on the observed time differences between bids across different maturities, where bidders did not update their bids during the auctions (see Appendix Figure A3). Second, a customer might bid via different dealers for different maturities, and third, two bids for the same maturity but by different customers might go through the same dealer. Neither of these cases happens often. 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.

With this, 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 cases as most bidders only update once or twice.

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, we trim our estimated pseudo-values. 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). This approach is conservative in light of the distribution of how bidders shade the untrimmed estimated pseudo-values per step (see Appendix Figures A5). The less we trim, the larger, in absolute value, are all demand coefficients (see Appendix Table A3).

C Model Extension: Latent Business Type

We consider a model extension in which dealers have a latent business type χ. Each dealer is either a market-maker type ($\chi = mm$), for whom bills are substitutes, or a niche-customer type ($\chi = nc$).[19] The econometrician does not know which dealer is of which type, but the bidders do. Private signals are now drawn independently from three distributions $F^{d,mm}$, $F^{d,nc}$ and F^{c}, and the marginal willingness to pay may be bidder specific: $v_{m,i}(q_m, q_{-m}, s^g_{m,i}) = f_{m,i}(s^g_{m,i}) + \lambda_{m,i} q_m + \delta_{m,i} q_{-m}$. All other assumptions remain unchanged.

To estimate demand coefficients, we adjust Proposition 2 and extend the resampling procedure as in Cassola et al. (2013) to account for asymmetric latent 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. (*iii*) Use the estimated demands to classify dealers into types.[20] Repeat until (*iii*) yields the same assignment as we started 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 (17) 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 Appendix Table A6). 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 (recall 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 an investor wants. For dealers in the second group, who tend to win less at auction, this might not always be true.

[19] 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.

[20] Since we have 3 maturities, we have 6 coefficients in the demand system given by (16) governing the substitution patterns. We assign a dealer to *mm*-type if at most 2 of those are negative.

D Measurement Error

Assessing the potential bias in our demand coefficients due to measurement errors in the expected winning quantities presents a significant challenge. In an attempt to gain a rough understanding of this bias, we compare the estimates derived from regression (2) with estimates obtained from a similar regression that utilizes expected rather than actual winning amounts. Here, we use bids instead of estimated values to explicitly attribute any measurement errors to the expectations themselves. Importantly, these regressions do not account for bid shading, and therefore, this exercise only serves as a means to gauge the magnitude of the measurement bias.

In addition, we estimate this regression using observable variables that approximate how much bidders expect to win. For instance, we use the amount a 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.

Appendix Table A8 shows the estimation findings for the 3M auction; the other auctions show similar patterns. We find that the δ coefficients range between roughly 0.2 and 0.8 when using the observable variables, while they are around 0.3 when using the expected values. All of the δ estimates are small compared to the λ estimate, which is -3.7 across specifications. Therefore, it seems as if the measurement error in the expected quantities is not causing us to vastly understate the degree of interdependencies.

E Details Regarding the Counterfactual

To quantify how much revenue can be gained when moving slightly away from the observed supply split, we must compute counterfactual bids. This is challenging because it is impossible to solve for an equilibrium analytically. We can only characterize necessary equilibrium conditions.

Counterfactual Bids. Our idea is to extrapolate from the 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)).[21] Concretely, 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:

$$b^{cf}_{t,m,i}(q_m) = \hat{v}^{cf}_{t,m,i}(q_m) - \hat{s}hading_{t,m,i}(q_m) \tag{19}$$

$$\text{with } \hat{v}^{cf}_{t,m,i}(q_m) = \hat{u}_{t,m,i} + \hat{\lambda}_m q_m + \hat{\delta}_m \hat{\mathbb{E}}[\boldsymbol{q^{cf*}_{t,-m,i}}|q_m] \tag{20}$$

$$\text{and } \hat{s}hading_{t,m,i}(q_m) = \hat{v}_{t,m,i}(q_m) - b_{t,m,i}(q_m) \;\forall m, i. \tag{21}$$

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 bid. $\hat{v}^{cf}_{t,m,i}(q_m)$ and $b^{cf}_{t,m,i}(q_m)$ are the counterfactual demand and 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, conditional on obtaining q_m in auction m, $\hat{\mathbb{E}}[\boldsymbol{q^{cf*}_{t,-m,i}}|q_m]$.

The expected winning quantities depend on where the counterfactual auction clears and therefore on how everybody bids. As a result, finding $\hat{\mathbb{E}}[\boldsymbol{q^{cf*}_{t,-m,i}}|q_m]$ of all bidders and all auctions constitutes a complicated fixed point problem.

To illustrate how to find a fixed point, consider one auction t and assume we change the supply from $Q_{t,m}$ to $Q^{cf}_{t,m}$ for all m. We start by rescaling all amounts demanded and expectations:

$$q^{cf}_{t,m,i,k} = \frac{Q^{cf}_{t,m}}{Q_{t,m}} q_{t,m,i,k} \tag{22}$$

$$\hat{\mathbb{E}}[\boldsymbol{q^{cf*}_{t,-m,i}}|q_{t,m,i,k}]^{old} = \frac{Q^{cf}_{t,m}}{Q_{t,m}} \hat{\mathbb{E}}[\boldsymbol{q^{*}_{t,-m,i}}|q_{t,m,i,k}] \text{ for all } m, -m, i, k, \tag{23}$$

and computing counterfactual bids according to (19). Ideally, we then estimate how much each bidder expects to win in the other auctions by simulating market clearance for each bidder and maturity many times, and update all expectations, $\hat{\mathbb{E}}[\boldsymbol{q^{cf*}_{t,-m,i}}|q_{t,m,i,k}]^{new}$. With the updated expectations, we update all bids, and repeat until updating expectations results in no change.

In practice, this algorithm is computationally infeasible because it involves conducting numerous auction simulations. As a solution, we suggest a statistical routine that can find

[21] 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.

an approximation of the fixed point while introducing some level of estimation noise. Rather than estimating expectations by simulating market clearance, we regress

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

to determine whether we need to shrink or expand the expectations. We then update all expectations. $\hat{\mathbb{E}}[\boldsymbol{q^{cf*}_{t,-m,i}}|q_{t,m,i,k}]^{new}$ become $\hat{\mathbb{E}}[\boldsymbol{q^{cf*}_{t,-m,i}}|q_{t,m,i,k}]^{old}$ and the new $\hat{\mathbb{E}}[\boldsymbol{q^{cf*}_{t,-m,i}}|q_{t,m,i,k}]^{new}$ is given by $\hat{\beta}_m \hat{\mathbb{E}}[\boldsymbol{q^{cf*}_{t,-m,i}}|q_{t,m,i,k}]^{old}$ for all m, i, k. We 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 illustrate that this method works reasonably well (see Appendix Figure A7a). Second, we show that the fixed point is sufficiently close to the rescaled expectations (23) with which we start the fixed point routine (see Appendix Figure A7b). This motivates us to use the rescaled expectations to conduct the counterfactual exercises on a broader scale.

Counterfactual Revenue. With the bids, we compute 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.

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

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, and $Q_{t,m}$ is the supply. We 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 $c_{t,S} = P^c_{t,S} - P^c_{t,L}$.[22]

[22] 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.

F Proofs

Proof of Proposition 1. Given the aggregate inverse demand of investors (4), the dealer expects the following payoff (7) 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 (3), (4), 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 for illustration. 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, and given our parameter restrictions, $\frac{\partial h_2(\kappa_1,\kappa_2,\gamma,e,\rho)}{\partial e} \leq 0$, $\frac{\partial h_2(\kappa_1,\kappa_2,\gamma,e,\rho)}{\partial \gamma} \geq 0$, $\frac{\partial h_2(\kappa_1,\kappa_2,\gamma,e,\rho)}{\partial \rho} \geq 0$ for any $\kappa_m \in [0,1]$ and any $\rho \in [-1/3, 1/3]$. $\square$

Proof of Proposition 2. Take the perspective of dealer with information θ, and let all other agents $j \neq i$ play a type-symmetric equilibrium. In this equilibrium it must be optimal for the dealer to choose the same set of functions $\{b_1(\cdot,\theta),...,b_M(\cdot,\theta)\}$ as all other dealers with information θ. These M functions must jointly maximize the dealer's expected total surplus. It must therefore be the case that each of the functions $b_m(\cdot,\theta)$ 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 dealer's strategy in all but one auction at the equilibrium.

With this insight, it is not surprising that the remainder of the proof follows the same steps as Kastl (2012)'s proof for a K-step equilibrium of a discriminatory price auction that takes place in isolation (on pp. 347–348). We know from Hortaçsu and Kastl (2012) that this proof extends when allowing for asymmetries in bidding behavior between dealers and customers and when dealers can update their bids. The only thing that we need to show for their proofs to apply, is that the pseudo-value functions, $\tilde{v}_m(q_m, s|\theta)$, behave like the value functions, $v(q, s)$, in Kastl (2012). This is shown in Lemma 1.

Lemma 1. *Fix some dealer i in auction m with signal s and information θ. $\tilde{v}_m(q_m, s|\theta)$ is (i) non-negative, measurable, and bounded, (ii) weakly decreasing in q_m for all s, (iii) strictly increasing in s_m for all q_m, and (iv) right-continuous in $b_{m,k}$.*

Proof of Lemma 1. Statement (i) follows directly from Assumptions 1-4. (ii) In equilibrium $\tilde{v}_m(q_m, s|\theta)$ must be decreasing in q_m or it could not give rise to a decreasing bidding function that fulfills the necessary conditions of Proposition 2. (iii) By Assumption 2, $v_m(q_m, q_{-m}, s)$ is strictly increasing in s_m for any q_{-m} and q_m, and by the Envelope Theorem, $\boldsymbol{q_m^*}$ must be increasing in s_m for all m. Further, given that s_m and s_{-m} are affiliated by Assumption 1, a bidder expects to win more of the other maturity $-m$ conditional on winning more of maturity m when drawing a high s_m. Therefore $\tilde{v}_m(q_m, s|\theta)$ must strictly increase in s_m at all q_m. (iv) To see why $\tilde{v}_m(q_m, s|\theta)$ is right-continuous in $b_{m,k}$ note first that it can only jump discontinuously if changing $b_{m,k}$ breaks a tie between this dealer 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 $b_{m,k}$ for which for any price in that neighborhood there are no ties. Therefore, when perturbing $b_{m,k}$, there cannot be any discontinuous shift in the conditional probability measure and thus in the object of interest. □

Appendix Figure A1: Auction Interface—What Bidders See When Bidding

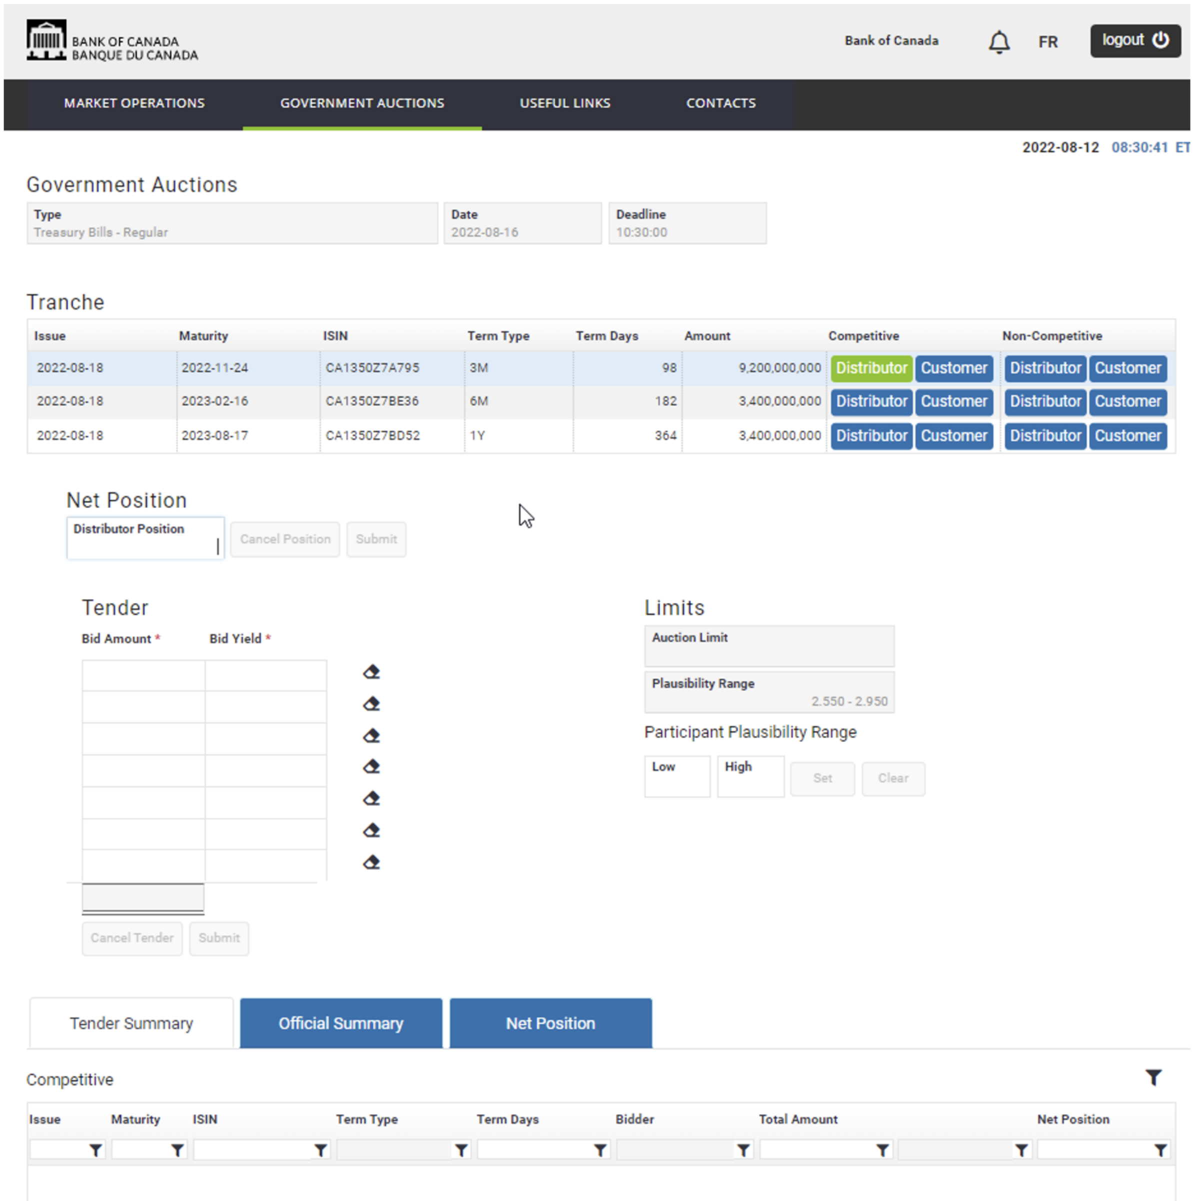

Appendix Figure A1 shows the a screen shot of what a bidder sees when bidding. He sees three "Tranches", listing the three different securities for sale. If he is a dealer, placing a competitive bid for his own account, he clicks on "Distributor" (in green). He fills in what "Net Position" he currently 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: Issuance of Canadian 3M, 6M, 12M T-Bills

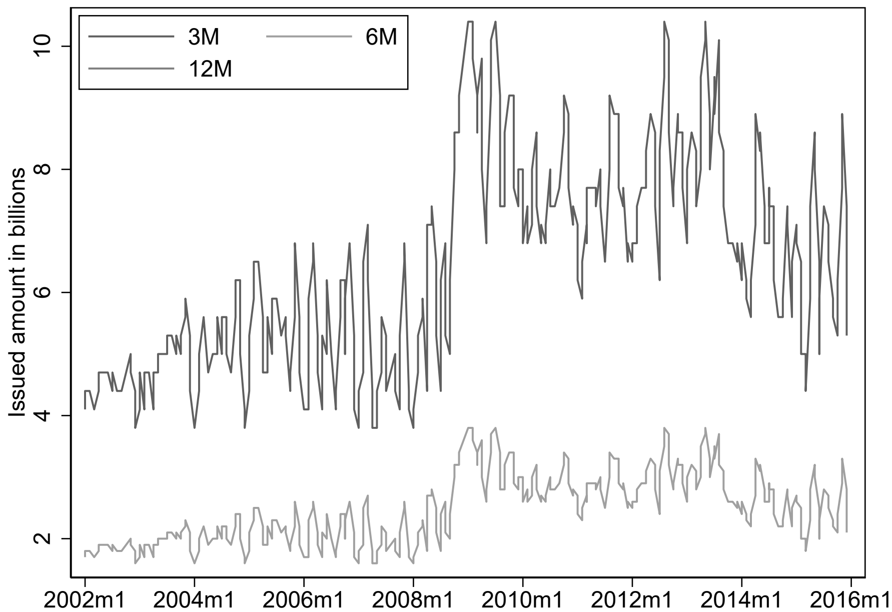

Appendix Figure A2 displays a time series of the issued supply of the 3M, 6M, and 12M bills, where the 6M issuance do not appear in the graph because they are identical to 12M issuance. The Bank of Canada always issues as many 6M bills as 12M bills. Over time, the amounts issued of the different maturities are perfectly correlated.

Appendix Figure A3: Time Between Bids of Those who do not Update

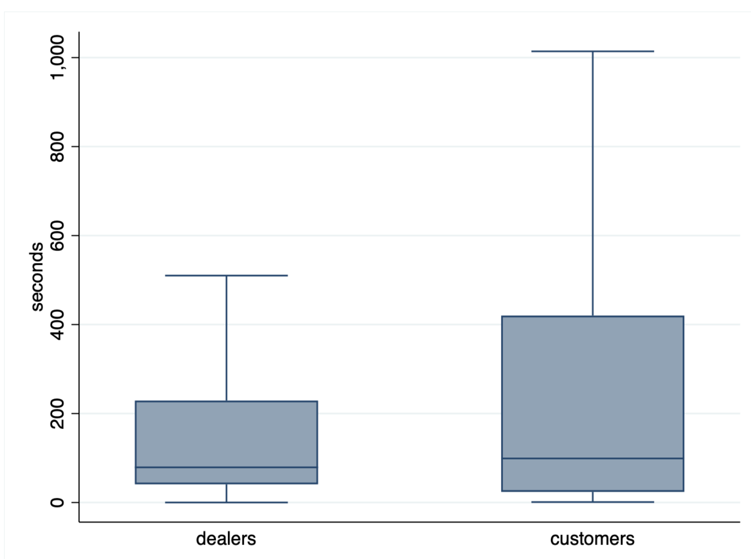

Appendix Figure A3 shows the distribution of the time difference (measured in seconds) between the bids that a dealer and a customer who does not update the bids places in different auctions. Outliers are excluded.

Appendix Figure A4: Steps by Bidder Groups

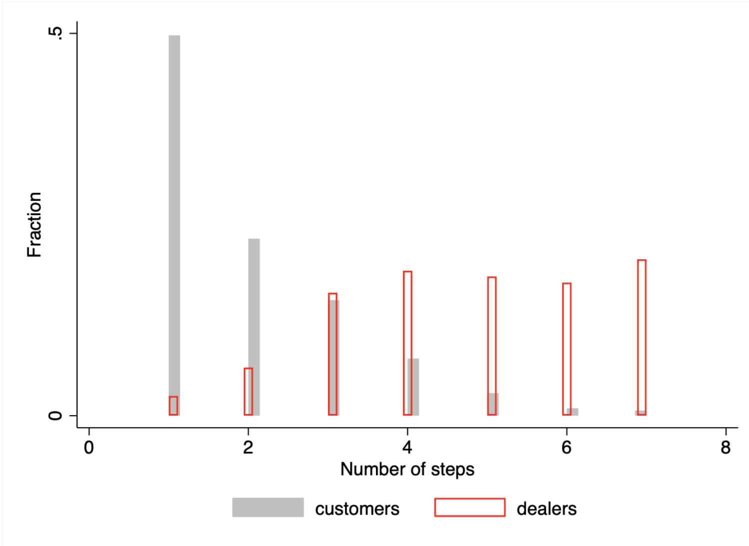

Appendix Figure A4 shows a histogram of the number of steps customers (in gray) and dealers (in red) submit. The fraction is measured in percentage points.

Appendix Figure A5: Distribution of the Untrimmed Shading Factor

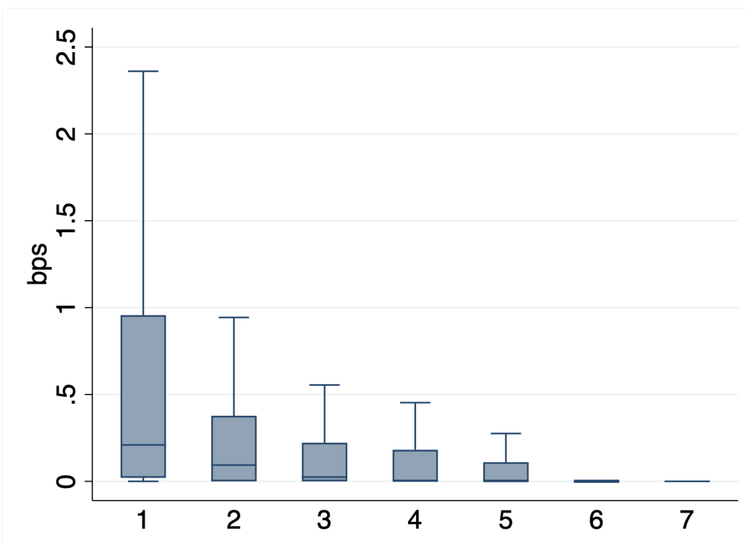

Appendix Figure A5 shows box plots of the untrimmed shading factor, $\hat{v}_{t,m,i,,\tau,k} - b_{t,m,i,,\tau,k}$, per step $\in \{1, 2, 3, 3, 5, 6, 7\}$ in a bidding function. For each step, the distribution is taken over dealers i, days t and time τ and maturities m. The shading factor is in bps.

Appendix Figure A6: Time series

(a) Proxy of Market Price Sensitivities

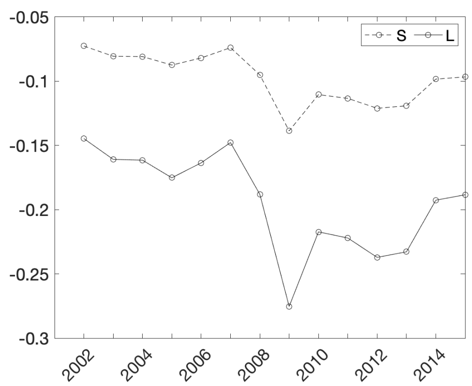

(b) Revenue Gains in Uniform Price Auction

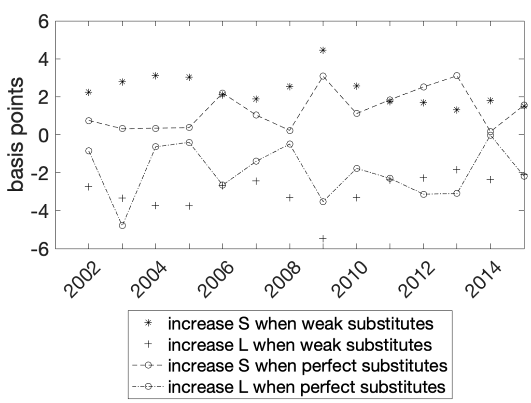

(c) Revenue Gains in Discrim. Price Auction

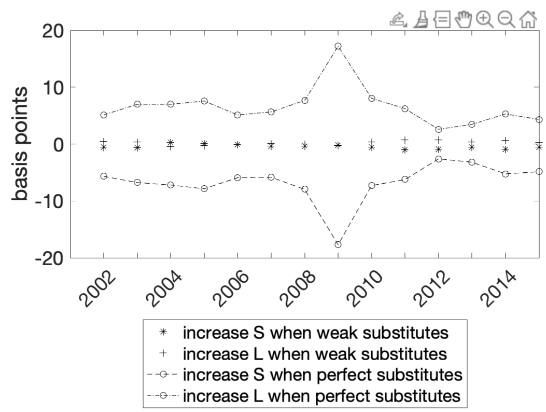

Appendix Figure A6 shows two time series. An observation in A6a shows the yearly average market price sensitivity when scaling the λ parameters by a factor of 100 and setting all δ's to zero: $\frac{1}{T_y}\sum_{t=1}^{T_y}(-100)\hat{\lambda}_m\frac{Q_{t,m}}{P^c_{tm}}$, where T_y is the total number of auction in year y. An observation in A6b and A6c 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.

Appendix Figure A7: Expectations on 3 Auction Days

(a) Did we Find a Fixed Point? (b) Fixed Point vs. Rescaled Expectations

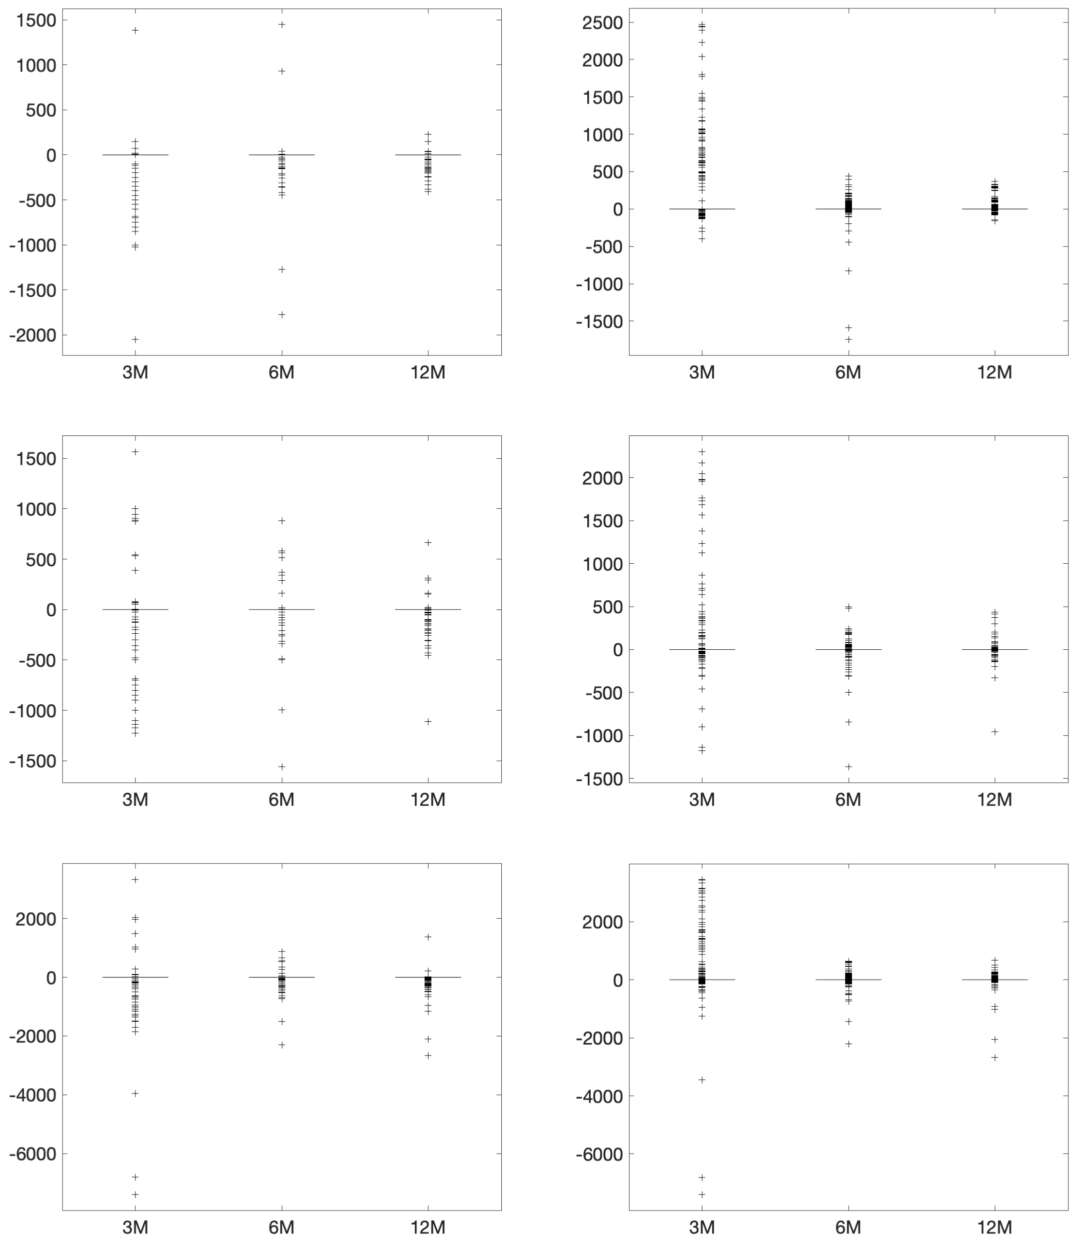

Appendix Figures A7a shows the distributions of the difference (in million C$) between the last two iterations of updating expectations in our statistical fixed point routine for all three maturities on three different auction days. We claim to find a fixed point (up to measurement noise) if the median difference is zero and there are only occasional outliers. Figure A7b shows the difference (in million C$) between the rescaled expectations (23) and the expectations that we find using our statistical fixed point routine. The median difference is again zero.

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 Table A2: Demand Coefficients with Values 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	

Appendix Table A2 (a)-(c) are analogous to Tables 5 (a) and A6. They report the coefficients for equation (17), 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.

Appendix Table A3: Demand Coefficients for the Average Dealer with Trimmed Values

(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 A3 (a)-(c) report the coefficients for equation (17), 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.

Appendix Table A4: 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	

Appendix Tables A4 (a) and (b) are analogous to Table 5 (a). They report the coefficients for equation (17), 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.

Appendix Table A5: 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
(1)	Independent	$+0.020$	$+0.007$	-0.023	-0.010
	Weak substitutes	$+0.016$	-0.002	-0.024	$+0.001$
	Perfect substitutes	$+0.011$	-0.052	-0.016	$+0.048$
(2)	Independent	$+2.344$	-0.446	-2.976	$+0.191$
	Weak substitutes	$+2.341$	-0.455	-2.970	$+0.200$
	Perfect substitutes	$+1.313$	-6.720	-1.956	$+6.624$

Appendix Table A5 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, marked by (1), correspond to the demand estimates of the 6M and 12M bills assuming different degrees of substitution. (2) corresponds to hypothetical auctions in which the λ parameters in the bidder's demands are scaled by a factor of 100. Scaling by a factor of 10 results in similar outcomes than those in (2)/10. All revenue gains are in bps of the original revenue. We do not include standard errors because it is computationally intense to conduct each counterfactual.

Appendix Table A6: Demand Coefficients per Dealer Group with Values 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	

Appendix Tables A6 (a) and (b) are analogous to Table 5 (b). They report the coefficients for equation (17). 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.

Appendix Table A7: Evidence for Imperfect Substitutes

	(a)			(b)	
3M bills	0.222	(0.0529)	bills	0.349	(0.0273)
6M bills	0.360	(0.0643)	bonds	0.058	(0.0153)
12M bills	1.042	(0.129)			
N	3,481			3,481	

Appendix Table A7 provides evidence that different types of bills are not perfectly substitutable, and neither are bills compared to bonds, using data from 2002 until 2015. Column (a) shows the estimation results of regression (18), which regresses the daily z-spread of 3M bills (in bps) on the daily total amount outstanding of 3M, 6M, and 12M bills relative to the quarterly GPD (in %), and a linear time trend. We classify bills that have more than 30 and less than 91 days left to maturity as 3M bills, those that have more than 91 and less than 182 days left as 6M bills and bills with more than 182 days as 12M bills. Alternatively, we could count bills with less than 30 days to maturity—1M bills, which are issued as separate category in cash-management auctions—as part of the 3M bills. The qualitatively findings would not change. Column (b) replicates regression (1) of GHS. Here, we regress the 3M z-spread on the total amount outstanding of all bills, and of all bonds, plus a linear time trend. In both cases, standard errors are in parentheses. They are assumed to follow an AR (1) process.

Appendix Table A8: 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

Appendix Table A8 shows the estimation results of regression (2) 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 4. Bids and valuations are in C$ and quantities in % of auction supply. Standard errors are in parentheses, clustered at the bidder level.