

Primary Dealers and the Demand for Government Debt*

Jason Allen^a

Jakub Kastl^b

Milena Wittwer^c

December 30, 2020

Abstract

Leveraging the fact that in many primary debt issuance markets securities of varying maturities are sold simultaneously, we recover participants' full demand systems by generalizing methods for estimating individual demands from bidding data. The estimated preference parameters allow us to partition primary dealers into two main classes. For the first class, which largely (but not fully) coincides with the largest money market players, we find significant complementarities in their demand for Treasury bills in primary markets, while for the second class the patterns in their willingness to pay are mixed and time-varying. We present a dealer-client model that captures the interplay between the primary and secondary market to provide a rationale for our findings. We argue that the complementarity likely arises from the large dealers "making markets", and hence requiring to hold inventory of all securities. Our results are useful both for minimizing the cost of financing of government debt and for optimally implementing financial regulation that is based upon partitioning financial institutions according to their downstream business strategies.

Keywords: Treasury bills, multi-unit auctions, structural estimation, market segmentation, cross-price elasticities

JEL classification: D44, C14, E58, G12

*The presented views are those of the authors, not necessarily of the Bank. All errors are our own. We thank the Treasury, auctions and settlement systems team and the debt-management team at the Bank of Canada for critical insights into debt management and auction operations. We thank Markus Brunnermeier, Jens Christensen, Annette Vissing-Jorgensen, Andreas Uthemann, and Yu Zhu, as well as seminar/conference participants at the Bank of Canada, ECB workshop on money markets and central bank balance sheets, EUI, Princeton and Queen's University. Correspondence to: ^aJason Allen - Bank of Canada, Email: jallen@bankofcanada.ca, ^bJakub Kastl - Princeton University, NBER and CEPR, Email: jkastl@princeton.edu, ^cMilena Wittwer - Stanford University, Email: wittwer@stanford.edu

1 Introduction

The primary objective of Debt Management Offices (DMO) worldwide is to achieve the lowest cost of financing over time. In order to fulfill this objective, a DMO has to decide how to sell government debt: the format of sale, which securities to offer, and how to allocate debt across different maturities.¹ Since primary dealers' valuations of securities of different maturities are not independent (in addition to the expected returns of these securities not being independent, a dealers' business model might generate dependence), the full demand system should be one of the crucial ingredients in these decisions. The main contribution of this paper is to propose a method for quantifying how willingness to pay (WTP) for one maturity depends on (expected) allocations of other maturities. In Treasury auctions for bills with 3-, 6-, and 12-month maturities, we find small own-price elasticities and some degree of complementarities across maturities. This suggests that debt managers can increase revenue by issuing more short-term debt without a large impact on price, and might want to do so in fixed proportions across maturities.

To understand our findings, we focus on how primary dealers link issuers of debt (in our example, government) and the final holders of debt, and how prices in the primary market are influenced by the structure of secondary markets.² Primary dealers buy securities from the DMO and sell them to clients in the secondary market. In the secondary market, diverse clients (asset managers, pension funds, insurance companies, etc.) demand securities with different maturities. Assets with different maturities might not be fully substitutable. This is the classic preferred-habitat justification for the existence of market segmentation (c.f. Culbertson (1957), Modigliani and Sutch (1966), Vayanos and Vila (2009), Guibaud et al. (2013), and Greenwood and Vayanos (2014)). The role of arbitrageurs (primary dealers in our setup) is to intermediate between market participants (clients). Primary dealers, therefore, play an important role in making markets by taking on costly inventory to meet heterogeneous demand. This generates complementarities across maturities in the primary market even if bills are substitutes in the secondary market.

In addition to understanding demand for Treasury bills at auction, our approach can be in-

¹In an ideal frictionless world, the maturity structure of government debt is irrelevant (Wallace (1981)).

²This work therefore complements the burgeoning literature on intermediary asset pricing, e.g. He and Krishnamurthy (2013), Brunnermeier and Sannikov (2014), and He et al. (2017). This is in addition to the literature studying how secondary market structure interacts with Quantitative Easing (c.f. D'Amico and King (2013) and Gorodnichenko and Ray (2017)).

formative about primary dealers’ business models, and their willingness to hold inventory to meet client demands. In particular, the way a primary dealer bids for securities of various maturities at auction is indicative of her relationship with clients in secondary markets. For a typical investor, securities of different (but fairly similar) maturities should behave as imperfect substitutes. However, a primary dealer with many (diverse) clients might view government securities of different maturities as complements. This complementarity arises since such a dealer needs to cover all markets. Regulation that increases the cost of holding inventory reduces the incentive to make markets. For example, the Supplementary Leverage Ratio (SLR) treats Treasuries similarly to risky bonds. Ultimately this impacts dealers willingness to hold Treasuries, and thus demand in the primary market.³ Identifying complementarities for specific dealers therefore not only helps DMOs better allocate issuances across maturities, but can also help regulators better target regulatory intervention and understand how regulation impacts heterogeneous dealers. Dealers, for example, who have a diverse group of clients have to purchase significant inventory across maturities – and this should not be punished by tight aggregate leverage ratios. Our approach allows us to ex-ante identify these dealer-types separately from those dealers who behave more like a typical investor and treat different maturities as substitutes.

Our final contribution belongs to the literature on demand estimation. A central issue that arises when estimating demand systems is unobserved heterogeneity: how to make sure that variation in quantity choices is attributable to variation in prices and not something omitted that is correlated with price, e.g. quality in the case of a typical discrete choice model. In the context of bond markets, it could also be time-varying risk premia. This is usually addressed by employing instruments aimed at isolating such exogenous variation by making the appropriate exogeneity and validity assumptions.⁴ We, instead, utilize a particular institutional feature that is surprisingly common in auctions of government debt such as those run by the US, Japan, Brazil, France, China, and Canada:

³The SLR requires all banks to add up their assets and make sure they have at least 3% of that amount in equity capital, or 5% for global systemically important banks, regardless of the asset type. The rule has been criticized for penalizing the holding of safe assets (Greenwood et al. (2017)). The Federal Reserve, as part of its response to COVID-19, temporarily changed the SLR to exclude deposits and U.S. Treasuries from the SLR calculation. This allowed primary dealers to increase their balance sheet and support household and business lending. In Canada, the banking regulator also exempted central bank reserves and sovereign-issued securities that qualify as High Quality Liquid Assets, including Treasuries.

⁴Koijen and Yogo (2019) adapt the widely used “BLP” approach (Berry et al. (1995)) to estimating demand for differentiated products to financial markets. They introduce the appropriate characteristics space and discuss potential instruments. Allen and Wittwer (2020) estimate the demand of an investor for government bonds. For this, they construct an innovative price instrument from bidding data in the primary auctions.

different securities (Treasury bills and bonds of different maturities) are sold simultaneously in parallel auctions. We extend previous results on identifying WTP from bidding data in auctions (developed by Hortaçsu (2002) and Kastl (2011), who build on the pioneering work of Guerre et al. (2000)) to allow for the willingness to pay to depend not only on the allocation of the underlying security, but also on holdings of securities of other maturities. Since Treasury bill auctions are held simultaneously, valuations that the auction participants attach to the different securities are, among other things, a function of the options available in the parallel auctions. These valuations, therefore, need to be estimated jointly. Overall, the setup enables us to estimate a full demand system in the primary (i.e., auction) market, allowing for flexible substitution patterns across securities, including complementarities.

We use data on all 3-, 6-, and 12-month Canadian Treasury bill auctions from 2002 to 2015 to estimate a model of simultaneous discriminatory price auctions. We find that the average dealer views Treasury bills of different maturity offered in the primary market as weak complements. This result may seem surprising in light of existing literature. To make our argument, we begin by providing evidence of cross-maturity bid-updating. That is, when dealers observe their customer bids in the auction for one maturity, they not only update their bids for the same maturity (as in Hortaçsu and Kastl (2012)), but also make updates to other maturities. Establishing that dealers indeed take into account their bids on different maturities simultaneously, we estimate that the marginal valuation for a 3M bill increases when going from an allocation excluding other maturities to one corresponding to the average observed allocations of 6M and 12M bills (about \$C200 million each) by about 0.14 basis points (bps). This “cross-market” effect is roughly 1/10th the size of the “own-market” effect: the marginal value for the 3M bill drops by about 1.25 bps when going from none to \$C400 million Treasury bills.⁵ The analogous increase in valuations in the 6M and 12M auctions are of similar, relatively small magnitudes, yet statistically significant.

To understand how bills might be complementary in the primary market, we introduce a formal model that captures the motives that drive demand for primary dealers as described above. On the one hand, primary dealers with direct access to the auction might keep some of the bills they win to use as collateral in other financial markets, or to fulfill regulatory requirements.⁶ On the

⁵The average allocation is 400 million bills, or 6% of the supply.

⁶Participants in the Canadian Derivatives Clearing Corporation, for instance, have minimum requirements to post

other hand, a key role of primary dealers is to support the well-functioning of secondary markets; they stand ready to buy and sell (and repo) in the secondary market, thus providing immediacy through their inventory.

Our model predicts that the degree of interdependency in the dealer's demand hinges on the role that dealers play in the secondary market. It suggests that different maturities are complementary for dealers who have a more diverse client base or for whom it is more costly to turn down clients than for dealers who cater to niche clients who might favor specific maturities. Our approach allows us to zoom in on the individual dealer level to gather evidence for this conjecture. We estimate a model of an asymmetric simultaneous discriminatory auction with two latent dealer types: a market-maker type who treats different securities as complements and a niche-clients type whose preferences are more idiosyncratic and hence, for this type, securities of different maturities are often substitutes. As our theory suggests, the estimated dealer types line up market-maker types with banks known to have large fixed-income trading desks and a broad client base.

In line with the predictions of our model, we find that complementarities for dealers in the money-market group are much stronger: the parameters that capture complementarities increase by between 50% and 494% relative to the estimates for the average dealer. Our interpretation of this result is that the market-making effect is stronger for these large dealers. That is, these dealers are less likely to bid at auction to take advantage of small arbitrage opportunities across maturities than they are to bid in fixed proportions in order to support secondary markets across all maturities.

For dealers in the second group, results are less precise. We find that the demand for 3M and 6M and 6M and 12M bills are substitutes or independent, while demand for the 3M and 12M bills are complementary, at least on average. One explanation is that the preferences of these dealers fluctuate more strongly from auction to auction, depending on the current order flows of their clients, than for dealers who trade with many different clients on a regular basis. The market-making effect might therefore be much smaller.

The paper is structured as follows: Section 2 describes the institutional environment and the Treasuries as collateral. See Bartolini et al. (2010) for evidence on the important role of Treasuries as collateral in the repo market.

data set. Section 3 presents evidence for interdependencies across maturities. Section 3.1 begins by documenting some patterns in the raw data that point towards interdependencies; Section 3.2 gives a preview of how we identify interdependencies and summarizes the key identifying assumptions; Sections 3.3 and 3.4 describe the structural model and our estimation strategy, respectively. Estimation findings are presented in Section 3.5. Section 4 concludes. All proofs are in the Appendix.

2 Institutional Environment and Data

2.1 Institutional Environment

In Canada, Treasury bills are issued with three maturities: 3, 6, and 12 months. Since 2002 they are sold every second Tuesday by the Bank of Canada (BoC) in three separate, but parallel, discriminatory price auctions. All three securities have a face value of \$1,000 (Canadian). Due to the large trading sizes, however, throughout the paper the units are in millions. There are two groups of bidders: “dealers” and “customers.” Dealers are either primary dealers or government securities distributors. Customers can only submit bids through primary dealers, but like dealers, they tend to be large financial institutions. They choose not to register as dealers, perhaps to sidestep additional monitoring and dealer-obligations.⁷ One example is Desjardins Securities. As the securities division of one of the largest Canadian financial institutions it is a primary dealer in the bond market, but only a customer in the Treasury market. Similarly, both Casgrain & Company and JP Morgan are not registered as primary dealers and yet are very important players in the Canadian government securities markets (Hortaçsu and Kastl (2012)).

From the time the tender call opens until the auctions close, bidders may submit and update their bids. There are two types of bids: competitive and non-competitive. A competitive bid is a step-function with at most 7 steps. “These bids must be stated in multiples of \$1,000, subject to the condition that each individual bid be for a minimum of \$100,000. Each bid shall state the yield to maturity to three decimal places” (Bank of Canada (2016)). For the most part we convert

⁷For more details see Sections 10 and 11 in Bank of Canada (2016).

yields into prices:

$$yield = \left(\frac{face\ value - price}{price} \right) \left(\frac{365}{days} \right), \quad (1)$$

with a *face value* of C\$1 million and *days* denoting the days left to maturity. Using prices instead of yields makes bidding as well as demand schedules decreasing rather than increasing. The bid step-function specifies how much a bidder offers to pay for specific amounts of the asset for sale. Figure 1a depicts an example – the choice of the median dealer in a 12M auction.⁸ The dealer offers to pay 98.68 thousand dollars for the first 50 million CAD. For the next 50 he offers to pay less, and so on. In addition to a competitive bid, each bidder may submit one non-competitive bid. This is a quantity order, which the bidder will win for sure, but for which he pays the average price of all accepted competitive bid prices. It is capped at 10 million dollars for dealers and 5 million dollars 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 total amount for sale.⁹ 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) and pays according to what they bid.

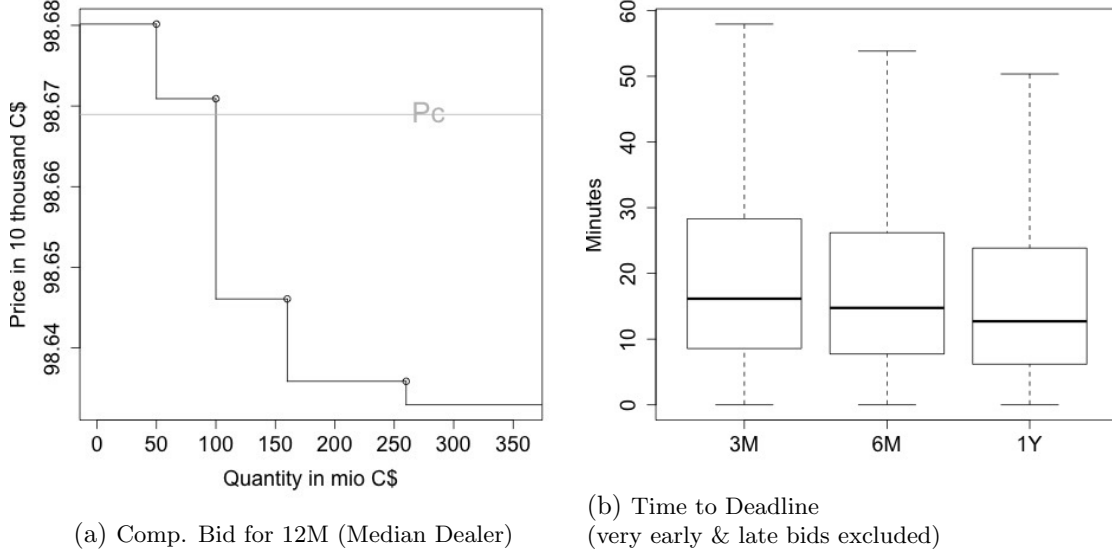
2.2 Data

Our data set consists of all 366 Canadian Treasury bill auctions between 2002 and 2015. Table 1 summarizes the data. On average the BoC announced issuances of C\$6.41 billion for 3M bills and C\$2.47 billion for each of the 6M and 12M bills per auction, of which it actually distributed roughly C\$5.76 (3M) and C\$2.12 billion (6/12M). The total amount issued per year was C\$81 billion for the 3M bills and C\$29 billion of the longer maturities. Figure 2 plots the issuance amounts over the period 2012–2017. Except for the spike in 3M issuance starting with the financial crisis and an

⁸The median step-function is 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 that were submitted by a dealer who submitted the median number of steps.

⁹The amounts purchased are typically divided across maturities as a proportion of what is supplied. The amounts purchased depend on the Bank’s projection of expected future demand for notes and the amount of Treasury bills maturing over the following weeks. See Statement of Policy, 2015.

Figure 1: Bids in the Canadian Treasury bill market



increase in government expenditures, issuances are steady and predictable.

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


The sample starts January 2002 and ends December 2015. There are 366 auctions per maturity. The total number of competitive bids (including updates) in the 3-, 6-, 12-month 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 across maturities is 2477, 2378, and 1932. From the raw data we drop competitive bids with missing bid price (133) and competitive or non-competitive tenders with missing quantities (69). Global participation is the probability of attending the remaining auctions, conditional on bidding for one maturity. Dollar amounts are in billions of C\$.

	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
Updates by dealer	2.89	2.18	2.48	3.58	2.87	3.18	0	0	0	31	31	42
Updates by customer	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

We identify each bidder through a bidder ID, and bidders are classified as a dealer or a customer. In total we observe 21 dealers and 76 customers over the sample period. The average auction has 11 to 12 dealers and 5 to 6 customers. Roughly 71% of participants bid for all three maturities.

Figure 2: Issuance of Canadian 3-, 6-, 12-month Treasury bills

The Bank of Canada follows a predictable issuance strategy. Displayed is 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.



Such “global participation” is even more regular among dealers. To keep their bidder status as government security distributor or primary dealer they have to be active in the primary market.¹⁰ Consequently, almost all who are active in a given auction week go to all three auctions (95%).

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 are within 10 to 20 minutes prior to closing. Figure 1b depicts box plots of the time at which bids arrive prior to the deadline, excluding very early outliers and bids that go in after auction closure. There are very few bids that arrive late (231 out of 57,650); 22 of them win despite being late.¹¹ We therefore keep late bids in our estimation sample. Typically, a dealer updates his bid (competitive or non-competitive) 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).

¹⁰“At every auction, a primary dealer’s bids, and bids from its customers, must total a minimum of 50 per cent of its auction limit and/or 50 per cent of its formula calculation, rounded upward to the nearest percentage point, whichever is less. [...] Each government securities distributor must submit at least one winning competitive or non-competitive bid on its own behalf or on behalf of customers, every six months.” (Bank of Canada (2016), p. 12).

¹¹Bids can show up as late in our data if a bidder manually phones the Bank of Canada to place a bid just before closing and the Bank takes some time to process it.

An average step-function of a competitive bid has 4.5 steps with little difference across maturities. Non-competitive tenders are small in size. On average, bidders only demand 0.1% of the total (announced) supply via non-competitive tender, with a maximal share of 0.58%. Given their size, our structural model abstracts from non-competitive bids, and focuses solely on the decision of placing competitive bids. The BoC, on the other hand, demands substantial amounts via non-competitive bids to reduce the total supply on the day of the auction, which generates uncertainty about the available supply. On average, it takes away 11.13% (3M), 14.35% (6M), 14.26% (12M) with a maximum of 20.45% (3M), 41.66% (6M), 25.00% (12M) of the total previously announced supply. Our model will need to account for unannounced changes in actual supply.

3 Interdependencies

Parallel auctions of different maturities might be interconnected both on the supply and the demand side. On the supply side, the BoC might determine the total amount for sale at each auction jointly, which leads to a non-zero correlation between the sold amounts across maturities.

To understand where interdependencies on the demand side may come from, it is useful to ask what motivates financial institutions' activity in Treasury auctions. For one, they might want to keep some of the bills in their own inventory. Treasury bills serve as collateral in interbank markets and repo transactions and are popular for fulfilling capital and liquidity requirements for safe assets. Second, most bidders (primary dealers) act as market makers in the secondary market. Therefore they buy securities of different maturities in order to sell (or repo) them to clients. To avoid having to turn down clients with demand for different maturities in the days that follow the auction, dealers want to buy bundles of maturities. How much each bidder values the particular securities offered at auctions depends on the bank's own balance sheets and other factors that are internal to the institution. It is the presence of such private information that makes it complicated to measure interdependencies on the demand side. Bidders with private information (that might be correlated across maturities) have incentives to shade their bids so as to minimize the prices they will have to pay for each unit they win. To estimate how strong the interdependencies are, we first have to back out how much bidders are truly willing to pay – a problem that is at the heart

of virtually all empirical analyses of auction markets.

3.1 Preliminary Empirical Evidence of Interdependencies

Table 2 displays correlations on the supply (2a) and demand side (2b) of Canadian Treasuries. The supply that the BoC announces exhibits perfect positive correlation across maturities. In fact, over our long sample the BoC always announces the exact same issuance size for the 6M and 12M bills. The amount it actually distributes on the auction day is also almost perfectly correlated.¹² We observe a similar pattern on the demand side. The total amount financial institutions demand (via competitive or non-competitive tender) when the auction closes is highly positively correlated across maturities, about 0.91–0.92. This pattern is suggestive of banks having preferences for buying assets in some fixed proportion, pointing towards complementarities. Since the correlation between quantities actually won drops to 0.54–0.57 (for all maturities), it seems that primary dealers do not always succeed in achieving this goal.

Table 2: Cross-Market Correlations

\bar{Q}_m is the announced issuance amount, Q_m the distributed supply for $m = 3, 6, 12M$. $q_{m,i}^D$ is bidder i 's demand, $q_{m,i}^*$ the amount won for $m = 3, 6, 12M$.

(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								
	$q_{3M,i}^D$	$q_{6M,i}^D$	$q_{12M,i}^D$		$q_{3M,i}^*$	$q_{6M,i}^*$	$q_{12M,i}^*$	
$q_{3M,i}^D$	1.00			$q_{3M,i}^*$	1.00			
$q_{6M,i}^D$	0.92	1.00		$q_{6M,i}^*$	0.57	1.00		
$q_{12M,i}^D$	0.91	0.91	1.00	$q_{12M,i}^*$	0.54	0.57	1.00	

¹²Canadian policy-makers perform stochastic simulations to determine a debt strategy that is desirable over a long horizon, e.g. 10 years. The model (publicly available at <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 3/6/12-month bills and $1/16^{th}$ of each of the 2/5/10/30-year bonds (see Bolder (2003)). Final issuance decisions are taken based on model simulations and judgment. “The typical practice [of the Bank of Canada] is to split the total amount purchased by the Bank, so that the Bank’s purchases approximate the same proportions of issuance by the government across the three maturity tranches” (Bank of Canada (2015) p. 5).

A further piece of evidence suggesting dependencies across auctions concerns updating behavior by dealers. Observing their customer orders, dealers may update their own bids. This can be because the customer bids provide information about competition or also about the fundamental security value (Hortaçsu and Kastl (2012)). The demand for bills across auctions is likely interconnected if dealers, upon observing a customer order flow (which may be concentrated only in one maturity), update their own bids across all maturities. To be more concrete, say a dealer observes a customer bid in the 3M auction. This triggers the dealer to update his own bid for the 3M bill. If his demand for 3-, 6-, and 12-month bills are interrelated, this should then also lead to an update of bids for the other maturities. To get a preliminary look at this pattern, we run the following Probit regression on competitive bids placed by dealers:

$$update_{i,m} = \alpha_i + \sum_m I_m (\beta_m customer_m + \delta_{m,-m} customer_{-m}) + \varepsilon_{i,m}. \quad (2)$$

To avoid double counting, each step-function (as in Figure 1a) is treated as one observation. The dependent variable *update* takes value 1 if the dealer updated his bid in an auction, and 0 otherwise. I_m is an indicator variable equal to 1 if the update occurs in the auction for maturity m . The independent variables $customer_l$ (for $l = m$ or $-m$) are also indicator variables. They are 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. The second specification builds on this benchmark but takes a longer sequence of events into account. It acknowledges that it takes time to calculate bids, enter them manually (which until 2019 is the rule rather than exception), and transfer them electronically. Table 3 provides an example of such a sequence. It shows the last 10 minutes of events of a dealer before auction closure on 10 February 2015. Having observed a customer in the 3M auction, he takes action himself and places several bids in a row. Specification (1) assigns value 0 to $customer_{3M}$ in the 6M auction because the dealer has not received an order for the 3M maturity immediately before bidding on his own behalf for the 6M bills (second-to-last column). He first bids for the 12M bills. The second specification assigns a value of 1 (last column). Here $customer_l$ is 1 for all bids the dealer places in a sequence (each with a time difference of 20 seconds) if he has received an order for maturity l within one minute before he places his own bid in auction m , or the latest

order the dealer achieved is for maturity l .

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

Table 3: Sequence of Events of a Dealer on 02/10/2015 in last 10 Min Before Auction Closure

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

3.2 A Preview of Our Identification Strategy

Our goal is to consistently estimate a parameter that measures by how much a bidder's marginal willingness to pay (MWTP) for some quantity of bills with maturity m changes the more he expects to win of the other maturities $-m$. As a first step we must understand what drives the MWTP. We introduce a formal model that captures the key motives for purchasing bills in the primary market. Within our model the true MWTP can be approximated by a linear function. To be precise, let bidder i of type $s_{m,i,\tau}^g$ in bidder group $g \in \{d = \text{dealer}, c = \text{customer}\}$ at time τ during the auction week have the following WTP for amount q_m in auction m conditional on winning q_{-m} of the other two maturities and keeping a share $(1 - \kappa_m)$ on its own balance sheet:

$$v_m(q_m, q_{-m}, s_{m,i,\tau}^g) = \alpha + (1 - \kappa_m)s_{m,i,\tau}^g + \lambda_m q_m + \delta_m \cdot q_{-m}. \quad (3)$$

Table 4: Probability of Dealer Updating Bids

The results of this table are based on the Probit regression, (2). In column (1) $customer_l$ is an indicator variable equal to 1 if the dealer received a competitive order from a customer for maturity l immediately before taking action in auction m himself. In column (2) $customer_l$ is an indicator variable equal to 1 if the dealer received an order for maturity l within one minute before placing his own bid in auction m , or the latest order received for maturity l . The total number of observations is 39,271. *p<0.1; **p<0.05; ***p<0.01.

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

The vector of δ_m parameters measures interdependencies across maturities. Take the example of the $m = 3M$ auction, where $q_{-m} \equiv (q_{6M} \quad q_{12M})'$ and $\delta_m \equiv (\delta_{3M,6M} \quad \delta_{3M,12M})$. If $\delta_{3M,6M} < 0$, bidders are willing to pay less for any amount of the 3M maturity the more they purchase of the 6M bills, hence the bills are substitutes. When $\delta_{3M,6M} > 0$ they are complementary, and independent if $\delta_{3M,6M} = 0$.

Estimating the parameters of interest consistently is challenging for two main reasons. First, the bank has private information about how much it values the securities.¹³ In our model, $s_{m,i,\tau}^g$

¹³Treasury bills have very active forward markets (“when-issued”). The presence of this market implies that a lot of information relevant for price-discovery is aggregated prior to the auction. It is therefore not unreasonable to assume that the heterogeneity in valuations at the time of the auction itself is driven mostly by idiosyncratic

is the bank’s private signal (or an index aggregate of a multidimensional signal). This generates incentives to misinterpret the true MWTP. As in the well-known first-price auction, bidders shade their bids to reduce the total payments they must make to win. By looking at the bids we are thus unable to differentiate between bidders reducing their bids for strategic reasons or because they are purchasing an interdependent good at the same time (Problem 1: Bid-shading).¹⁴ Second, even if a bidder wanted to report their true MWTP, $v_m(q_m, q_{-m}, s_{m,i,\tau}^g)$, the disconnected auction design does not allow it. By the rules of the auction, a bidder can, in auction m , only submit a one-dimensional bidding step-function (such as in Figure 1a) that depends on amounts of security m , not on securities $-m$ (Problem 2: Disconnected market design). Summarizing both challenges: we observe bidding functions that specify a price for amounts of one maturity only, q_m , not the true MWTP, which is a function of all maturities $v_m(q_m, q_{-m}, s_{m,i,\tau}^g)$ without knowing $s_{m,i,\tau}^g$.

Our two-stage estimation procedure solves both of these problems. First, we estimate the joint distribution of market clearing prices and recover how much each bank would bid if it were bidding truthfully. This solves the problem of strategic bid-shading. Here we extend the structural estimation techniques developed by Hortaçsu (2002), Kastl (2011), and Hortaçsu and Kastl (2012) to the case of simultaneous auctions of potentially related goods. Our estimates are consistent under the identifying assumptions that (i) private information about all maturities $s_{i,\tau}^g \equiv \begin{pmatrix} s_{3M,i,\tau}^g & s_{6M,i,\tau}^g & s_{12M,i,\tau}^g \end{pmatrix}$ of each bidder i at each bid-update-time τ is *iid* across bidders conditional on observed auction and date characteristics, and (ii) that all bidders are ex ante symmetric within their bidder group (dealer or customer) and play a (type-) symmetric Bayesian Nash Equilibrium (BNE) each time new bills are issued. In an extension we relax the latter assumption.

Given the disconnected market design, the schedule a bidder would submit if it were truthful, call it $\tilde{v}_m(q_m, s_{m,i,\tau}^g)$, is not his true MWTP, $v_m(q_m, q_{-m}, s_{m,i,\tau}^g)$. This is because his actual marginal benefit from winning amount q_m depends on how much he will win of the other assets, $q_{-m,i}^*$. Since

factors such as the structure of the balance sheet, investment opportunities or repo needs – which do not depend on private information of other dealers. Any private information about future resale value can be arbitrated away in the when-issued market. Hence, after conditioning on the public information, one can view the information structure as corresponding to private values. This is consistent with the results in Hortaçsu and Kastl (2012), who fail to reject the null hypothesis of no learning about fundamentals from clients’ bids – which is one of the necessary conditions for private values.

¹⁴Generally, so-called “demand-reduction” can be a severe problem in multi-unit auctions in which bidders have demand for more than one unit (e.g. Ausubel et al. (2014)). Bid-shading should play a minor role for Treasury bills since they are highly liquid in secondary market trading. Conditional on observables, such as the when-issued price of these bills, or the spot price in the secondary market, bidders can infer one another’s preferences fairly accurately.

auctions take place in parallel, this is unknown. In equilibrium, these random quantities $\mathbf{q}_{-m,i}^*$ need to be integrated out:

$$\tilde{v}_m(q_m, s_{m,i,\tau}^g) = \mathbb{E}[v_m(q_m, \mathbf{q}_{-m,i}^*, s_{m,i,\tau}^g) | \text{win } q_m].$$

In the first stage of our estimation procedure we estimate $\tilde{v}_m(q_m, s_{m,i,\tau}^g)$. In addition, we estimate the joint distribution of market clearing prices which allows us to estimate the conditional expectation $\mathbb{E}[\mathbf{q}_{-m,i}^* | \text{win } q_m]$. Assuming the true MWTP is linear, as in (3), we can estimate the parameters of interest, δ_m , in a linear regression with bidder-auction-time fixed effects that control for $\alpha + (1 - \kappa_m)s_{m,i,\tau}^g$.

We now proceed to describe the model and estimation strategy, before presenting our estimation results. Throughout, random variables are denoted in **bold**.

3.3 A Dealer-Client Model

M perfectly divisible goods, indexed m , are auctioned in M separate discriminatory price auctions, run in parallel. In each auction, there are two groups (g) of bidders: dealers (d) and customers (c). We assume that the total number of potential dealers N_d and customers N_c is commonly known, and denote the total number of bidders by $N = N_c + N_d$. Over the course of the auction, new information may arrive at a discrete number of time slots $\tau = 0, \dots, \Gamma$. How much each bidder bids each τ depends on their WTP. Before modelling the auction process, we introduce a stylized model that captures the key driving factors of individual demand in the primary market.

Financial institutions participate in Treasury auctions for different purposes. They have private information about how much they need the bills supplied at auction. Formally, we let a bidder i of group g draw a private signal at time τ he places his bid: $\mathbf{s}_{i,\tau}^g \equiv \left(s_{1,i,\tau}^g \dots s_{M,i,\tau}^g \right)$. This type is multi-dimensional. To account for differences between bidder groups, it may be drawn from different distributions for customers and dealers.

Assumption 1. *Dealers' and customers' private signals $s_{i,\tau}^d$ and $s_{i,\tau}^c$ 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 .*

Notably, we do not need to impose any restrictions on time dependence. The reason is that we will not pool bids from auctions that took place at different points in time. A bidder's type can therefore be persistent across time.

3.3.1 Micro-Foundation of Individual Demand

In the spirit of Vayanos and Vila (2009), our model features market segmentation, where investor/clients may have preferences for specific maturities and dealers function across maturities by participating in the primary market and making markets in secondary trading. Rather than assume dealers are risk-averse, we assume that dealers face a cost of not meeting client demand.¹⁵

For simplicity, in this section we restrict the number of maturities to $M = 2$.¹⁶ For notational convenience we drop the superscript g and the subscripts i, τ for the remainder of the section. Further, we label one part of the private type by ν , and the other t :

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

A bidder of type s obtains the following gross benefit from “consuming” amounts $(1 - \kappa_1)q_1$ and $(1 - \kappa_2)q_2$:

$$U(q_1, q_2, s) = t_1(1 - \kappa_1)q_1 + t_2(1 - \kappa_2)q_2. \quad (4)$$

The private type determines how much a bidder 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. Bidders, in particular dealers, function as market makers in the secondary market where they distribute the rest of the bills $\{\kappa_1 q_1, \kappa_2 q_2\}$ among investors who are yet to arrive. To incorporate future resale opportunities we let there be a second stage following the primary auction. In the secondary market a (mass of) client(s) with random demand $\{\mathbf{x}_1, \mathbf{x}_2\}$ arrives to the bidder.¹⁷ Equivalently,

¹⁵A practical reason for why we do not model dealers as risk neutral is that it is much harder to estimate auction models with risk-averse bidders than having a cost of not meeting demand.

¹⁶Generalizing to more than two maturities is straightforward but mathematically cumbersome and brings no major additional insights.

¹⁷The terms “client” and “customer” denote different players. Customers participate in the auction by placing bids with dealers, while clients buy in the secondary market.

you may imagine that there are two types of clients, each with a random demand for one of the two maturities. For simplicity we assume that each of $\{\mathbf{x}_1, \mathbf{x}_2\}$ is on-the-margin uniformly distributed on $[0, 1]$ but allow both amounts to be correlated. More specifically, $\{\mathbf{x}_1, \mathbf{x}_2\}$ 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 $\rho \in [-\frac{1}{3}, +\frac{1}{3}]$.

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

$$p_{i,1}(x_1, x_2 | q_1, q_2) = \begin{cases} a - bx_1 - ex_2 & \text{for } x_1 \leq \kappa_1 q_1 \text{ and } x_2 \leq \kappa_2 q_2 \\ a - bx_1 & \text{for } x_1 \leq \kappa_1 q_1 \text{ and } x_2 > \kappa_2 q_2 \\ 0 & \text{for } x_1 > \kappa_1 q_1 \text{ and } x_2 > \kappa_2 q_2. \end{cases} \quad (5)$$

The price function for maturity 2 is analogous. It splits into three cases. In the first, clients for both bills arrive and the bidder has enough of both in their inventory for resale. The bidder 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 bidder can only sell maturity 1. This might be because only clients with demand for this maturity arrive or because the bidder does not have enough of the other maturity in inventory for resale, $x_2 > \kappa_2 q_2$. The price the bidder 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 bidder does not hold enough of either bill to satisfy the demand of client(s) he cannot sell. Notice that the magnitudes of the resale prices are characterized by three parameters $\{a, b, e\}$. A higher intercept $a > 0$ increases the bidder's bargaining power, and with it the price he can charge for each unit sold. Parameter $b > 0$ governs the price-sensitivity of clients. Large clients (who demand more) have more negotiating power and can drive down the price. When $e > 0$ bills are substitutes in the secondary market, and vice versa for complements.

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

$$revenue(x_1, x_2|q_1, q_2) = p_1(x_1, x_2|q_1, q_2)x_1 + p_2(x_1, x_2|q_1, q_2)x_2. \quad (6)$$

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

$$cost(x_1, x_2|q_1, q_2) = \begin{cases} 0 & \text{if } x_1 \leq \kappa_1 q_1 \text{ and } x_2 \leq \kappa_2 q_2 \\ \gamma x_1 & \text{if } x_1 > \kappa_1 q_1 \text{ and } x_2 \leq \kappa_2 q_2 \\ \gamma x_2 & \text{if } x_1 \leq \kappa_1 q_1 \text{ and } x_2 > \kappa_2 q_2 \\ \gamma x_1 x_2 & \text{if } x_1 > \kappa_1 q_1 \text{ and } x_2 > \kappa_2 q_2. \end{cases} \quad (7)$$

This function captures the idea that it is more costly to turn down larger clients, i.e. those with larger demand. The important feature for our results is that it is supermodular in x_1, x_2 , i.e. has increasing differences.¹⁸ This means that the marginal cost from turning down a client who demands one maturity is higher the larger the order for the other maturity.

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

$$V(q_1, q_2, s) = U(q_1, q_2, s) + \mathbb{E}[revenue(\mathbf{x}_1, \mathbf{x}_2|q_1, q_2) - cost(\mathbf{x}_1, \mathbf{x}_2|q_1, q_2)]. \quad (8)$$

The gross payoff determines how much a bidder is willing to pay on-the-margin. Consider auction 1. At time τ the bidder is willing to pay $v_1(q_1, q_2, s) = \frac{\partial V(q_1, q_2, s)}{\partial q_1}$ for amount q_1 conditional on winning q_2 of the other maturity. The appendix shows that $v_1(\cdot, \cdot, s)$ is a third-order polynomial for any s . It can be approximated by a linear function. Taking the first-order Taylor expansion around $(\mathbb{E}[\mathbf{x}_1], \mathbb{E}[\mathbf{x}_2]) = (1/2, 1/2)$ we obtain the following result.

¹⁸Supermodularity is for functions that map from $\mathbb{R}^n \rightarrow \mathbb{R}$ equivalent to increasing differences: $cost(x'_1, x'_2|q_1, q_2) - cost(x_1, x'_2|q_1, q_2) \geq cost(x'_1, x_2|q_1, q_2) - cost(x_1, x_2|q_1, q_2)$ for $x'_1 \geq x_1$ and $x'_2 \geq x_2$.

Proposition 1. *The marginal willingness to pay of a bidder with type $s_{m,i,\tau}^g$ for amount q_m conditional on winning q_{-m} in the other auction can be approximated by*

$$v_m(q_m, q_{-m}, s_{m,i,\tau}^g) = \alpha_{m,i,\tau}^g + (1 - \kappa_{m,i,\tau}^g)t_{m,i,\tau}^g + \lambda_{m,i,\tau}^g q_m + \delta_{m,i,\tau}^g q_{-m} \text{ for } m = 1, 2 - m \neq m, \quad (9)$$

where $\alpha_{m,i,\tau}^g, \lambda_{m,i,\tau}^g, \delta_{m,i,\tau}^g$ are polynomials of the exogenous parameters $\{\kappa_{1,i,\tau}^g, \kappa_{2,i,\tau}^g, \gamma_{i,\tau}^g, \rho_{i,\tau}^g, a_{i,\tau}^g, b_{i,\tau}^g, e_{i,\tau}^g\}$.

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

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

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

When buying only for its own account ($\kappa_1 = 0$) a bidder 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 MWTP in auction 1 decreases in q_1 as long as his clients are sufficiently price-elastic (i.e. b is sufficiently high). If he sells all of both maturities ($\kappa_1 = \kappa_2 = 1$) the MWTP is independent of his private type t_1 . How much he is willing to pay for one maturity now hinges on the amount he wins of the other maturity. Whether bills are substitutes or complements in the primary market depends on how large γ is relative to e . More generally one can derive the following corollary which will be useful when interpreting our estimation results. It holds for the general case where clients' demand might be correlated ($\rho \neq 0$) and the bidder keeps any amount of bills ($\kappa_1, \kappa_2 \in [0, 1]$).

Corollary 1. *Securities in the primary market become more complementary for bidder of type $s_{i,\tau}$ when*

- (i) *they are weaker substitutes in the secondary market ($e_{i,\tau}^g \downarrow$),*
- (ii) *it is more costly to turn down clients ($\gamma_{i,\tau}^g \uparrow$), or*
- (iii) *it is more likely that clients with demand for different maturities arrive ($\rho_{i,\tau}^g \uparrow$).*

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

Second, the corollary tells us that it is possible that some bidders view bills as substitutes and others as complements, as long as the joint distribution of parameters $\mathbf{v}_{i,\tau}^g$ is not degenerate. There could, for example, be a group of bidders for which it is more costly to turn down clients (high γ_i) or whose clients are more likely to demand variety (high ρ_i) than for other bidders. We might expect financial institutions whose primary business is to trade Treasury bills to be part of this group, while it might be relatively less costly to turn down clients (low γ_i) or less likely to serve clients who seek to buy different maturities (low ρ_i) for financial institutions whose primary business lies outside the money market. This observation motivates our model of asymmetric simultaneous discriminatory auctions with two latent primary dealer types that we present next.

3.3.2 The Auctions

In modeling the auction process we build on Hortag su and Kastl (2012)’s model of a stand-alone auction.¹⁹ Motivated by the previous section, we assume that the MWTP is linear, but different across bidder types, g .

Assumption 2. *The marginal willingness to pay of a bidder with type $s_{m,i,\tau}^g$ for amount q_m con-*

¹⁹See Kastl (2017) for a review of advances in the application of Industrial Organization tools in finance.

ditional on purchasing q_{-m} of the other two securities $-m$ is

$$v_m(q_m, q_{-m}, s_{m,i,\tau}^g) = \alpha_{i,m}^g + (1 - \kappa_{i,m}^g) t_{m,i,\tau}^g + \lambda_{i,m}^g q_m + \delta_{i,m}^g \cdot q_{-m}, \quad (9)$$

with $\lambda_{i,m}^g < 0$, $|\delta_{i,m}^g| < \lambda_{i,m}^g$ and α_i^g sufficiently high such that the marginal willingness to pay does not drop below 0 for any amount that might be for sale.

Furthermore, in line with the discussion in the previous section we assume that there are two types of primary dealers: a market-maker type, mm , and the niche-customer type, nc .

Assumption 3. Dealers can be partitioned into two types: $\mathbb{N}_d = \mathbb{N}_{d,mm} \cup \mathbb{N}_{d,nc}$, such that $\forall m \in \mathbb{N}_{mm} : \delta_{i,m}^{d,mm} \geq 0$.

Note that Assumption 1 should be simply augmented to have private signals drawn independently from $F^{g,mm}$, $F^{g,nc}$ and F^c .

A bid in auction for maturity m consists of a set of quantities in combination with prices. It is a step-function which characterizes the price the bidder would like to pay for each amount.

Assumption 4. In auction m each bidder has the following action set each time an offer is placed:

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

Notice that $q_{m,k} \in [0, 1]$. It represents the share of total supply. This allows us to compare bids in auctions with different sizes of supply. A bid of 0 denotes non-participation.

To capture the updating process of bids prior to auction closure, we assume that new information may arrive at each time slot τ . At $\tau = 0$, a bidder draws an *iid* random variable $\Psi_i \in [0, 1]$. It is one dimension of the bidder's private signal and thus unobservable to competitors. It corresponds to the mean of an *iid* Bernoulli random variable, Ω_i , which determines whether the bidder's later

bids will make it in time to be accepted by the auctioneer. More specifically, for $\tau > 0$, the bidder's information set includes the realizations $\omega_i \in \{0, 1\}$ of Ω_i , where $\omega_i = 1$ means that the bid of time τ will make it in time. This gives an incentive to bid at each arrival of new information because there might not be an opportunity to successfully bid in the future.

Given that the rules of the auction do not allow for customers to submit their own bids, at each time τ all customers who want to place an order are matched to a dealer. The dealer can observe his customer's bid. This provides him with additional information at time τ – one that is unavailable to other dealers or customers. A dealer might have the same customer in all three auctions. Denoting the information obtained from observing a customer's bids at time τ in auction m by $Z_{m,i,\tau}$, dealer i 's information set or, equivalently, his “type” is $\theta_{i,\tau}^g = (s_{i,\tau}^g, Z_{1,i,\tau}, Z_{2,i,\tau}, Z_{3,i,\tau})$. If he only has a customer in one auction, say for maturity 1, $\theta_{i,\tau}^g = (s_{i,\tau}^g, Z_{1,i,\tau})$, and so on. Notice that by Assumption 1, $(s_{i,\tau}^g, Z_{i,\tau})$ are independent across dealers and time. However, $s_{i,\tau}^g$ and $Z_{i,\tau}$ can be correlated within a dealer across τ .

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

A choice in auction m by a bidder with information $\theta_{i,\tau}^g$ may be summarized as bidding function $b_{m,i,\tau}^g(\cdot, \theta_{i,\tau}^g)$ or equivalently as a demand function $y_{m,i,\tau}^g(\cdot, \theta_{i,\tau}^g)$. The latter specifies how much an agent demands at each admissible price. When auction m closes at $\tau = \Gamma$, the auctioneer aggregates individual demands of the bidders' final bids. The market clears at the lowest price P_m^c at which aggregate demand, denoted $\sum_{i=1}^{N_c} y_{m,i,\Gamma}^c(p_m, \theta_{i,\Gamma}^c) + \sum_{i=1}^{N_d} y_{m,i,\Gamma}^d(p_m, \theta_{i,\Gamma}^d)$, satisfies aggregate supply. The latter is the announced amount for sale net of what the BoC demands in the form of non-competitive bids during the auction plus all other competitive bids by bidder i 's competitors.

Assumption 5. *Supply $\{Q_1, Q_2, Q_3\}$ is a random variable distributed on $[\underline{Q}_1, \overline{Q}_1] \times [\underline{Q}_2, \overline{Q}_2] \times [\underline{Q}_3, \overline{Q}_3]$ with strictly positive marginal density conditional on $s_{i,\tau}^g \forall i, g = c, d$ and τ .*

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.²⁰

Denoting the amounts bidder i gets allocated by $q_i^c = \begin{pmatrix} q_{1,i}^c & q_{2,i}^c & q_{3,i}^c \end{pmatrix}$ when submitting $b_{i,\tau}^g(\cdot, \theta_{i,\tau}^g) \equiv \begin{pmatrix} b_{1,i,\tau}^g(\cdot, \theta_{i,\tau}^g) & b_{2,i,\tau}^g(\cdot, \theta_{i,\tau}^g) & b_{3,i,\tau}^g(\cdot, \theta_{i,\tau}^g) \end{pmatrix}$ his total surplus is

$$TS(b_{i,\tau}^g(\cdot, \theta_{i,\tau}^g), s_{i,\tau}^g) = V(q_i^c, s_{i,\tau}^g) - \sum_{m=1}^3 \int_0^{q_{m,i}^c} b_{m,i,\tau}^g(x, \theta_{i,\tau}^g) dx \quad (10)$$

in the event in which τ is the time of his final bid, with $V(q_i^c, s_{i,\tau}^g)$ given by (9). It is the total utility he achieves from obtaining the amounts he wins minus the total payments he must make. Ex ante, when placing a bid, the bidder knows neither how much he will win nor at which price the market will clear. His optimal choice maximizes the expected total surplus.

Definition 2. A BNE is a collection of functions $b_{i,\tau}^g(\cdot, \theta_{i,\tau}^g)$ that for each bidder i and almost every type $\theta_{i,\tau}^g$ at each time τ maximizes the expected total surplus, $\mathbb{E}[TS(b_{i,\tau}^g(\cdot, \theta_{i,\tau}^g), s_{i,\tau}^g)]$.

We focus on type-symmetric BNE of the auction game, in which bidders who are ex ante identical follow the same strategies. Dealers of the same latent type who draw the same signal employ the same function, and similarly for customers. Across bidder groups strategies might be asymmetric.

$$\forall \omega \in \{mm, nc\} : b_{i,\tau}^{d,\omega}(\cdot, \theta_{i,\tau}) = b^{d,\omega}(\cdot, \theta_{i,\tau}) \text{ and } b_{i,\tau}^c(\cdot, \theta_{i,\tau}) = b^c(\cdot, \theta_{i,\tau}) \quad \forall i, \tau.$$

²⁰ “Under this rule, all bids above the market clearing price are given priority, and only after all such bids are satisfied, the remaining marginal demands at exactly price $P^c = p$ are reduced proportionally by the rationing coefficient so that their sum exactly equals the remaining supply. An alternative rationing rule would, for example, not give bids at higher prices priority.” (Kastl (2011)). The rationing coefficient satisfies $R_m(P_m^c) = \frac{Q_m - TD_m^+(P_m^c)}{TD_m(P_m^c) - TD_m^+(P_m^c)}$ where $TD_m(P_m^c)$ denotes the total demand at price P_m^c , and $TD_m^+(P_m^c) = \lim_{p_m \downarrow P_m^c} TD_m(p_m)$.

3.4 Estimation Strategy

3.4.1 First Stage of the Estimation Strategy

To solve the problem of strategic bid-shading we recover what the bidder would bid if he were truthful by extending Hortaçsu (2002), Kastl (2011), Kastl (2012), and Hortaçsu and Kastl (2012) to the case of simultaneous auctions of potentially related goods. To determine which marginal valuations rationalize the observed bids we must first characterize the optimality conditions for the type-symmetric BNE of the game. Here we extend Wittwer (2020), who characterizes the equilibrium in simultaneous discriminatory price auctions under more stringent assumptions than we impose in this paper.

Bidding incentives in simultaneous discriminatory price auctions are similar to those in an isolated auction. To fix ideas, we begin with the benchmark case of auctions of independent goods. Securities in our model are unrelated if all δ parameters are equal to 0. In this case gross utility is additively separable across maturities and the WTP for one maturity $v_m(q_m, s_{m,i,\tau}^g)$ is independent of the amount allocated to this bidder in auctions of other maturities. In addition, all markets clear separately. A bid offered for good 1 will not affect the payment the agent has to make for good 2 because the agent's demand for good 1 can, by the rules of a standard discriminatory price auction, only depend on the price for good 1. Since neither utility nor payments are interrelated, strategic incentives are identical to those in an isolated auction. In determining his best reply to all others, the bidder can, therefore, focus on each auction in isolation. If the bidder knew the residual supply curve when choosing his bids, he would just pick a point on this curve that maximizes his total surplus. Yet, when making his choices, he does not know this curve as it depends on the random total supply and the private information of his competitors. He thus has to integrate out the uncertainty about the market clearing price and evaluate marginal benefits and costs of changing a bid. The marginal cost is losing the surplus on the last infinitesimal unit demanded, which happens exactly when the price is between bids, defined by the k^{th} and $k+1^{st}$ step. The marginal benefit is saving the difference between these bids whenever the market clearing price ends up being actually weakly lower than b_{k+1} .

Proposition 2 (Unrelated Goods). *Consider a bidder i of group g with private information $\theta_{i,\tau}^g$ who submits $\hat{K}_m(\theta_{i,\tau}^g)$ steps in auction m at time τ . Under Assumptions 1-5 in any type-symmetric BNE every step k in his bid function $b_m^g(\cdot, \theta_{i,\tau}^g)$ has to satisfy*

$$v_m(q_{m,k}, s_{m,i,\tau}^g) = b_{m,k} + \frac{\Pr(b_{m,k+1} \geq \mathbf{P}_m^c | \theta_{i,\tau}^g)}{\Pr(b_{m,k} > \mathbf{P}_m^c > b_{m,k+1} | \theta_{i,\tau}^g)} (b_{m,k} - b_{m,k+1}) \quad \forall k < \hat{K}_m(\theta_{i,\tau}^g)$$

and $b_{m,k} = v_m(\bar{q}_m(\theta_{i,\tau}^g), s_{m,i,\tau}^g)$ at $k = \hat{K}_m(\theta_{i,\tau}^g)$ where $\bar{q}_m(\theta_{i,\tau}^g)$ is the maximal amount the bidder may be allocated in the equilibrium.

This equation allows us to estimate the marginal valuations that rationalize the observed bids at all steps $b_{m,k}$ of all bidders at all times and auctions. Following Hortaçsu (2002) the approach is to estimate the distribution of the market clearing price using a resampling procedure. It relies on the assumption that private information is not interdependent across bidders, so that “each bidder i cares about others’ bidding strategies only insofar as they affect the distribution of bidder i ’s residual supply” (Hortaçsu and McAdams (2010)). The choice of bid by bidder i transforms the distribution of the residual supply into the distribution of the market clearing price. In a standard multi-unit auction with no bid-updating in which N bidders draw independent private information and play the symmetric BNE in T auctions with identical covariates, the resampling procedure works as follows: Fix bidder i . For all bidders that did not bid in an auction, augment the data with their bids being 0. Draw a random sample $N - 1$ bid-vector with replacement from the sample of NT bids. Construct bidder i ’s realized residual supply were others to submit these bids. Repeating this routine many times gives a consistent estimate of the distribution of the market clearing price. Our setup is, even if there are no interdependencies across auctions, more complicated. First, we have two bidder groups (dealers and customers) which may be ex ante asymmetric. Second, bidders may update their bids within an auction with dealers observing their customers’ bids. Hortaçsu and Kastl (2012) have extended the resampling procedure for this more complex environment.²¹ We extend this approach for stand-alone auctions to the cases of parallel

²¹The procedure is as follows: Drawing N_c customer bids from the empirical distribution of customer bids. If a customer did not participate, replace his bid by a 0. For each customer bid vector, draw a corresponding dealer bid. If a zero customer bid is drawn, draw from the pool of uninformed dealers (those who did not observe any customer bids). If a nonzero customer bid is drawn, draw from the pool of dealers’ bids, which have been submitted having observed a “similar” customer bid with equal probabilities. Those are customer bids whose quantity-weighted

auctions below.

It is highly unlikely that demands for Treasury bills of different maturities are independent, in particular when the maturities are similar. Bidders take this interconnection across auctions into account when determining their optimal bidding strategies. Consider an auction for maturity $m = 1$. When preferences are no longer separable across maturities, the agent's MWTP for amount q_1 depends on how much of the other goods he gets allocated, $v_1(q_1, q_{-1}, s_{1,i,\tau}^g)$. Ideally, he would want to condition his price $b_{1,k}$ for amount $q_{1,k}$ on how much he will purchase of the other securities in equilibrium, $q_{-1,i}^* \equiv (q_{2,i}^* \quad q_{3,i}^*)'$. Since the rules of the auction do not allow participants to express their preferences in this way, they have to integrate out the uncertainty. Conditional on winning $q_{1,k}$, which happens when $b_{1,k} \geq \mathbf{P}_1^c > b_{1,k+1}$, a bidder expects a marginal benefit of $\mathbb{E} \left[v_1 \left(q_{1,k}, \mathbf{q}_{-1,i}^*, s_{1,i,\tau}^g \right) \middle| b_{1,k} \geq \mathbf{P}_1^c > b_{1,k+1}, \theta_{i,\tau}^g \right]$. Analogous to the decision process in an isolated auction, the agent equates the benefit of winning the bid with its marginal cost. Since auctions clear separately the cost is identical to the cost in an isolated auction with one important difference. With stochastic dependence across auctions, market clearing prices are connected. With M maturities, they are drawn from a joint M -dimensional distribution.

Proposition 3 (Related goods). *Consider a bidder i of group g with private information $\theta_{i,\tau}^g$ who submits $\hat{K}_m(\theta_{i,\tau}^g)$ steps in auction m at time τ . Under Assumptions 1-5 in any type-symmetric BNE every step k in his bid function $b_m^g(\cdot, \theta_{i,\tau}^g)$ has to satisfy*

$$\tilde{v}_m(q_{m,k}, s_{m,i,\tau}^g | \theta_{i,\tau}^g) = b_{m,k} + \frac{\Pr \left(b_{m,k+1} \geq \mathbf{P}_m^c | \theta_{i,\tau}^g \right)}{\Pr \left(b_{m,k} > \mathbf{P}_m^c > b_{m,k+1} | \theta_{i,\tau}^g \right)} (b_{m,k} - b_{m,k+1}) \quad \forall k < \hat{K}_m(\theta_{i,\tau}^g)$$

with

$$\tilde{v}_m(q_{m,k}, s_{m,i,\tau}^g | \theta_{i,\tau}^g) \equiv \mathbb{E} \left[v_m \left(q_{m,k}, \mathbf{q}_{-m,i}^*, s_{m,i,\tau}^g \right) \middle| b_{m,k} \geq \mathbf{P}_m^c > b_{m,k+1} | \theta_{i,\tau}^g \right]$$

for $m = 1 \dots M$ with $-m \neq m$, and $b_{m,k} = \tilde{v}_m(\bar{q}_m(\theta_{i,\tau}^g), s_{m,i,\tau}^g | \theta_{i,\tau}^g)$ at $k = \hat{K}_m(\theta_{i,\tau}^g)$ where $\bar{q}_m(\theta_{i,\tau}^g)$ is the maximal amount the bidder may be allocated in an equilibrium.

Analogously to a stand-alone auction, we can estimate the marginal valuations by estimating the bid price are sufficiently close (according to a pre-defined bandwidth). The resulting estimate of the distribution of clearing prices is consistent (Hortaçsu and Kastl (2012)).

distribution of residual supply curves, now jointly for all maturities. With $M = 3$ parallel auctions, the benchmark resampling procedure of Hortaçsu (2002) must be changed in that a choice of a bidder is now a triplet of bidding functions submitted on a given auction day. Fixing such a triplet of bids submitted by a bidder, one then draws a random subsample of $N - 1$ bid-vector triplets with replacement from the sample of NT bids, and constructs bidder i 's realized residual supply $\forall m$ were others to submit these bids to determine the realized clearing prices $P^c = \begin{pmatrix} P_{3M}^c & P_{6M}^c & P_{12M}^c \end{pmatrix}$, and the amount i would have won $q_i^* = \begin{pmatrix} q_{3M,i}^*, q_{6M,i}^*, q_{12M,i}^* \end{pmatrix}$ for all q_i^*, P^c . Repeating this procedure a large number of times provides an estimate of the joint distribution of market clearing prices and, equally important, the corresponding amount of each security i would win.


There are two complications when auctions are not considered separately. First, bids in different auctions are not submitted at the exact same time given electronic or human delays (see the example in Table 3). In our procedure, we define bids to be “simultaneous” if they are the closest bids of all bids a bidder places within 200 seconds, or they are the last bids made before the auction deadline, i.e. final bids. Setting an upper bound of 200 seconds seems sensible when looking at the number of seconds between bids across maturities which we know were determined “simultaneously”. Those are cases where the bidder does not update his bids over the course of the auctions. On average 551(383) seconds pass between such bids for different maturities by dealers(customers). Excluding outliers reduces the time (see Figure 3).

Second, a customer might place his order via different dealers in an auction week. He might, for instance, go via one dealer in the 3M auction and via another in the 6M auction. Furthermore, two bids for the same maturity but by different customers might go through the same dealer. Neither of these two cases happens more than a handful of times. Therefore, we assume that the information set of dealers who observe the same customer is independent across maturities, conditional on his own signal. In addition, we restrict the number of possible observed customer bids to two. Given that most customers only submit one bid and that there are many more dealers than customers in a typical auction, this simplifying restriction is reasonable.

With these simplifications our procedure is 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

Figure 3: Time Between Bids of Those Who Do Not Update

This figure plots the distribution of times between bids for both dealers and customers who do not update their bids. Time is in seconds.



customer submitted only one bid, 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 update once or twice.

The above described estimation procedure gives consistent estimates under two scenarios: In the benchmark, all bidders (customers and dealers) are ex ante symmetric. In particular, dealers do not know whether their rivals have complementary, substitutable, or independent preferences for different maturities. This is plausible if we believe that these preferences are mostly driven by fluctuating factors in the secondary market. In the second scenario, there are two groups of dealers. They consistently display different preferences, for example, because they follow different business models. Each dealer is aware of how many dealers are in each group but they do not know dealer identities.²²

²²With this specification, our estimates are consistent if the number of bidders is large enough.

3.4.2 Estimation with Asymmetric Dealers

Alternatively, we assume that dealers' partition between *mm*- and *nc*-types is common knowledge among bidders, and further extend the resampling procedure as in Cassola et al. (2013) to account for asymmetric types. Since the econometrician does not observe dealers' (ex-ante) types, we introduce a latent variable capturing type. The estimation proceeds in three steps: (i) Partition dealers into the two groups, (ii) Estimate a model, where resampling is conditional on that assignment,²³ (iii) Use the estimated demands to classify dealers into types.²⁴ 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.

In our main specification we impose marginal valuations $\tilde{v}_m(\cdot, s_{m,i,\tau}^g | \theta_{i,\tau}^g)$ to be weakly decreasing. Furthermore, to correct for outliers that occasionally occur due to small values of the denominator in the estimated (marginal) hazard rate of the market clearing price, $\tilde{\text{Pr}}(b_{m,k} > P_m^c > b_{m,k+1} | \theta_{i,\tau}^g)$, we trim our estimated marginal values. Specifically, we restrict each to be lower than the bidder's maximal bid plus a markup of about 0.4 bps (C\$40 for 12M, C\$20 for 6M, C\$10 for 3M). The size of the markup is motivated by the distribution of how bidders shade the untrimmed estimated marginal values per step, i.e., $\hat{v}_{t,m,i,\tau,k} - b_{t,m,i,\tau,k}$. Figure 4 shows that the vast majority of the (untrimmed) shading factor lie below 0.4 bps and the median bid-shading in all cases is zero.²⁵

3.4.3 Second Stage of the Estimation Strategy

Our resampling procedure delivers a consistent estimate of the joint distribution of market clearing prices and the amount bidder i wins in equilibrium conditional on the information he has at time τ . This allows us to estimate how much he expects to win of the other maturities if he were to win a given quantity in maturity m :

$$\hat{\mathbb{E}}[q_{t,-m,i}^* | \dots] = \mathbb{E}[q_{t,-m,i}^* | b_{t,m,i,\tau,k} \geq P_{t,m}^c > b_{t,m,i,\tau,k+1}, \theta_{t,i,\tau}] + \varepsilon_{t,m,i,\tau,k}^q. \quad (11)$$


²³ A *mm*-type needs to integrate over bids of $N_{mm} - 1$ other *mm*-types and N_{nc} niche-client types and vice versa.

²⁴ Since we have 3 maturities, we have 6 coefficients in the demand system given by (9) governing the substitution patterns. We assign a dealer to *mm*-type if at most 2 of those are negative.

²⁵ Results are robust to reasonable alternative trimming adjustments.

Figure 4: Distribution of the untrimmed shading factor

This figure plots the distribution of the untrimmed (raw) shading factor for the three maturities, 3, 6, and 12 months. The shading factor is measured in Canadian dollars.



Moreover, using Proposition 3, we can use the marginals of this joint distribution to obtain an estimate of how much bidder i would be willing to pay at step k at time τ in auction m of week t given the observed bid:

$$\hat{v}_{t,m,i,\tau,k} = \mathbb{E} \left[v_m \left(q_{t,m,i,\tau,k}, \mathbf{q}_{t,-m,i}^*, s_{m,i,\tau}^g \right) \middle| b_{t,m,i,\tau,k} \geq \mathbf{P}_{t,m}^c > b_{t,m,i,\tau,k+1}, \theta_{t,i,\tau} \right] + \varepsilon_{t,m,i,\tau,k}^v. \quad (12)$$

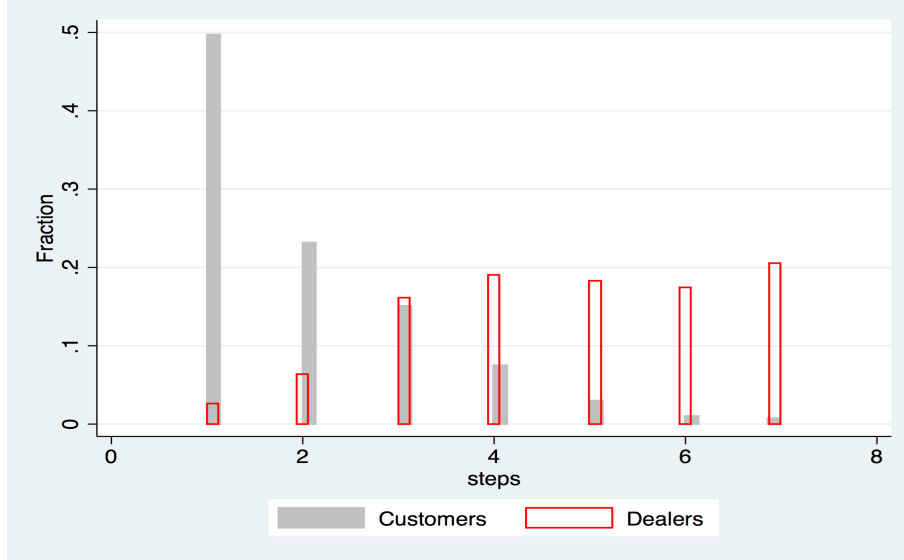
Assuming linearity of the MWTP (Assumption 2), we can now estimate the following linear regression with auction-bidder-time fixed effects $u_{t,m,i,\tau} \equiv \alpha + (1 - \kappa_m)t_{t,m,i,\tau}^g$:

$$\hat{v}_{t,m,i,\tau,k} = u_{t,m,i,\tau} + \lambda_m q_{t,m,i,\tau,k} + \delta_m \cdot \hat{\mathbb{E}} \left[\mathbf{q}_{t,-m,i}^* | \dots \right] + \varepsilon_{t,m,i,\tau,k}, \quad (13)$$

for $m = 1 \dots M, m \neq m$ on a subsample with competitive bids of more than one step to identify the parameters of interest. Figure 5 shows that it is the case for virtually all dealer bids: almost all submit more than one step. In Appendix 6 we show that our findings are robust to excluding the first step in addition.

Figure 5: Steps by Bidder Groups

This figure plots the frequency of bid-steps taken by the customers (gray) and dealers (red).



3.5 Estimation Results

We restrict attention to dealers. First, we present the estimation results for “an average dealer” (Table 5) when estimating a model assuming ex-ante symmetric dealers. We then present results from estimating the asymmetric version of the model with two dealer groups to allow for heterogeneity in preferences (Tables 6 and 7). Our model predicts that dealers with a broad client base whose main business it is to trade in the money market are more likely to have complementary preferences. To test whether our data supports this conjecture, we split the dealers into two groups. The first group consists of the five dealers with large money market desks. They are the most active players with a broad client base. The second group includes dealers whose primary business lies outside money markets. Huyn et al. (2017) use a similar breakdown to highlight the increase in agent-based trading relative to principal-based for dealers in the “other” group.

In addition, we run all regressions (13) using observed bids ($b_{t,m,i,\tau,k}$) to complement our results using $\hat{v}_{t,m,i,\tau,k}$ (see Appendix 6). This gives an idea of whether bid-shading ($b_{t,m,i,\tau,k} - \hat{v}_{t,m,i,\tau,k}$) is sensitive to interdependencies across auctions. If we use bids rather than valuations, the λ parameters are biased downwards. This would suggest flatter marginal valuations than estimated. The δ parameters, on the other hand, are biased upwards. Using bids overestimates the degree of

complementarities compared to the case where we allow dealers to shade their bids strategically.

In line with the model, we report results with marginal valuations and bids expressed in C\$ (prices), not yields. Whenever we quote a number in bps it is the estimated value from the corresponding regressions performed with yields.²⁶ All quantities are expressed as a percentage of the total amount issued in the auction, thereby facilitating the comparison of auctions with different supply.

Table 5: Preferences of the average dealer

Using bids for all dealers, this table reports estimates for equation (13). Estimates are interdependencies given by the δ parameters. Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	-5.229***	(0.0267)	λ_{6M}	-8.450***	(0.0485)	λ_{1Y}	-16.96***	(0.0860)
$\delta_{3M,6M}$	0.178**	(0.0625)	$\delta_{6M,3M}$	0.626***	(0.106)	$\delta_{1Y,3M}$	1.087***	(0.207)
$\delta_{3M,1Y}$	0.241***	(0.0669)	$\delta_{6M,1Y}$	0.437***	(0.114)	$\delta_{1Y,6M}$	0.418	(0.215)
const	995661***	(0.367)	const	991657***	(0.721)	const	981633***	(1.258)
N	58542			42282			50410	

Table 6: Preferences of dealer group 1

Using bids for dealer group 1 (main dealers), this table reports estimates for equation (13). Estimates are interdependencies given by the δ parameters. Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	-6.213***	(0.0487)	λ_{6M}	-9.499***	(0.0848)	λ_{1Y}	-19.82***	(0.152)
$\delta_{3M,6M}$	1.054***	(0.111)	$\delta_{6M,3M}$	1.217***	(0.177)	$\delta_{1Y,3M}$	0.887**	(0.342)
$\delta_{3M,1Y}$	0.363**	(0.123)	$\delta_{6M,1Y}$	0.940***	(0.200)	$\delta_{1Y,6M}$	1.412***	(0.388)
const	995671***	(0.543)	const	991420***	(1.058)	const	981251***	(1.863)
N	28592			21406			25134	

Table 7: Preferences of dealer group 2

Using bids for dealer group 2 (other dealers), this table reports estimates for equation (13). Estimates are interdependencies given by the δ parameters. Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

	3M Bill Auction			6M Bill Auction			12M Bill Auction		
λ_{3M}	-4.708***	(0.0288)	λ_{6M}	-7.829***	(0.0538)	λ_{1Y}	-14.98***	(0.0913)	
$\delta_{3M,6M}$	-0.306***	(0.0683)	$\delta_{6M,3M}$	0.196	(0.122)	$\delta_{1Y,3M}$	1.484***	(0.238)	
$\delta_{3M,1Y}$	0.275***	(0.0716)	$\delta_{6M,1Y}$	0.139	(0.125)	$\delta_{1Y,6M}$	-0.210	(0.224)	
const	995662***	(0.465)	const	991917***	(0.928)	const	982107***	(1.572)	
N	29863		20818		25097				

²⁶ As a rule of thumb, 100/50/25 C\$ of a 12M/6M/3M bill are approximately 1 basis point of a yield.

An average dealer. The estimation results when pooling across all dealers are displayed in Table 5. Recall that the face value of a bill is in millions of Canadian dollars. Before discussing the estimated degree of interdependency (the δ parameters), consider first by how much a dealer’s MWTP for maturity m changes in q_m (the λ parameters of the first column in all tables). This helps provide a sense of magnitudes. As expected, marginal utility is strictly decreasing (all λ ’s are significantly negative). They are not large in magnitude, however, indicating that valuations are fairly flat with respect to quantity. Increasing the amount of the 3M bills by 1% of total supply, for example, decreases a dealer’s marginal benefit from owning the 3M bill by C\$5.229 or about 0.21 bps. Given the large average allocated amounts (about 8.12% of the issued supply or C\$544 million of 3M bills per dealer), there is a more sizable difference between the value of the first and last Treasury bill: about 1.7 bps. The other two maturities exhibit similar patterns: In the 6M/12M bill auction, the bidder’s marginal valuation drops by C\$8.45/16.96, or about 0.17 bps. This translates to a difference in values for the first and last won Treasury bill (on average 8.12% of the issued supply or C\$204 million per dealer) of about 1.37 bps for both, the 6M and 12M bills. Given the small price variation in the submitted bid step-functions (recall Figure 1a) this finding should not be that surprising. Intuitively, financial institutions have a rather precise idea of the price at which the primary market will clear, since all securities we consider are highly liquid. To avoid paying unnecessarily high prices, they submit bids that vary closely around the clearing price they expect, conditional on publicly available information (such as prices in the when-issued market).

With these “own” effects in mind, we can turn to the discussion of our main parameters of interest. All δ parameters are positive but relatively small. This suggests that bills are weak complements for the average dealer. In the 3M auction, the estimates imply that the bidders’ valuation increases by about $(0.178 * 8.12\% + 0.241 * 8.12\%)/25 \approx 0.14$ bps when obtaining the average amount of the 6M and 12M bills (8.12% of the supply of each), rather than nothing. For auctions of the longer maturities the complementarities are similar. The valuation for bidders in the 6M auction increases by about $(0.626 * 8.12\% + 0.437 * 8.12\%)/50 \approx 0.1 + 0.07 = 0.17$ bps and the valuation of bidders in the 12M auction by about $(1.087 * 8.12\% + 0.418 * 8.12\%)/100 \approx 0.09 + 0.03 = 0.12$ bps when going from owning nothing of the other two maturities to obtaining the average amounts won.

Heterogeneities in preferences. Our model of demand predicts that dealers with different business models have different preferences. According to Corollary 1, bills are more complementary for dealers for whom it is more costly to turn down clients or whose clients are more likely to demand different maturities. Both is more likely to be true for dealers whose primary business lies in the money market. They have a broader base of clients with demand (or a preferred habitat) for different bills and might worry more about sustaining this base.

To test for preference heterogeneity we split the dealers into two groups. Our first group contains dealers with large money market desks and many clients for whom we suspect to find stronger complementarities (called “main dealers”). The second includes “other dealers”.

Our results, displayed in Tables 6 and 7, confirm our conjecture. Compared to an average dealer (Table 5), bills are more complementary for a dealer with a money market focus. All of the δ parameters, with the exception of $\delta_{1Y,3M}$, increase relative to our baseline estimates. The percentage increase is relatively large, ranging from 50% up to 494%.²⁷ All estimates of complementarities are statistically significant. Yet they are relatively small: The dealer’s valuation in the 3/6/12M auction increases by about $0.32 + 0.1 = 0.42/0.19 + 0.14 = 0.33/0.07 + 0.11 = 0.18$ bps when going from owning nothing of the other two maturities to obtaining the average amounts won. This compares to an “own-maturity” drop in valuation for 3/6/12M bills of 1.99/1.42/1.43 bps when going from winning nothing in the auction for 3/6/12M bills to winning the average amounts. Therefore, complementarities are relatively weak.

The preferences of dealers who have smaller money market trading desks differ from those whose primary business lies in the money market. The estimation results of Table 7 are less convincing than the results presented thus far. Only three out of the six point-estimates are statistically significant. One reason for this could be that preferences of dealers in this second group are more volatile over time. Client demand might fluctuate more from week to week for smaller dealers relative to those with a broad client base. Taken together, it seems as if 3M and 6M bills, as well as 6M and 12M bills, might be substitutes, while 12M and 3M bills are complementary.

To study heterogeneity in preferences, we classified dealers into two groups ex ante and show

²⁷The percentage increase is calculated by taking the difference between the parameters of Tables 5 and 6 over the baseline parameter of Table 5.

differences in bidding behavior. There is a group of dealers for which bills are complementary – large money market firms; and a group of smaller dealers whose demand is more idiosyncratic. Our methodology, however, could be used to uncover groups, especially in settings with many players. These groups might in fact be relevant for regulation in markets where regulation restricts trading depending on dealer-client relationships.

4 Conclusion

In this paper we study interdependencies in the demand for securities of different maturities. Using data from Canadian Treasury bill auctions over a 15-year period, we find that 3-, 6-, and 12-month bills are oftentimes weakly complementary in the primary market. To explain our findings we present a model that captures the interplay between the primary and secondary market. We argue that the typical bidder of a primary auction buys bills not only for his own balance sheets but also (or even primarily) to distribute them in the secondary market where different clients demand different maturities. A bidder anticipates that it will be costly to turn down clients in case he did not buy sufficiently many bills at auction, or to satisfy their demand by purchasing the bills from other financial institutions at higher prices. This generates complementarities in the primary market even if Treasury bills are substitutes in other financial markets. Primary dealers, therefore, are supporting liquidity across maturities in the Canadian money market.

References

- Allen, J. and M. Wittwer (2020). Centralizing over-the-counter markets? Working paper.
- Ausubel, L. M., P. Cramton, M. Pycia, M. Rostek, and M. Weretka (2014). Demand reduction and inefficiency in multi-unit auctions. *The Review of Economic Studies* 81(4), 1366–1400.
- 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.
- Bartolini, L., S. Hilton, S. Sundaresan, and C. Tonetti (2010). Collateral values by asset class: Evidence from primary securities dealers. *The Review of Financial Studies* 24(1), 248–278.
- Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica* 63(4), 841–890.
- Bolder, D. J. (2003). A stochastic simulation framework for the Government of Canada’s debt strategy. Working Paper, Bank of Canada.
- Brunnermeier, M. K. and Y. Sannikov (2014). A macroeconomic model with a financial sector. *American Economic Review* 104(2), 379–421.
- Cassola, N., A. Hortaçsu, and J. Kastl (2013). The 2007 subprime market crisis in the EURO area through the lens of ECB repo auctions. *Econometrica* 81(4), pp. 1309–1345.
- Culbertson, J. M. (1957). The term structure of interest rates. *The Quarterly Journal of Economics* 71(4), 485–517.
- D’Amico, S. and T. B. King (2013). Flow and stock effects of large-scale Treasury purchases: Evidence on the importance of local supply. *Journal of Financial Economics* 108(2), 425–448.
- Gorodnichenko, Y. and W. Ray (2017). The effects of quantitative easing: Taking a cue from Treasury auctions. Working Paper 24122, National Bureau of Economic Research.
- Greenwood, R., S. Hanson, J. Stein, and A. Sunderam (2017). Strengthening and streamlining bank capital regulation. Brookings Paper.
- Greenwood, R. and D. Vayanos (2014). Bond supply and excess bond returns. *The Review of Financial Studies* 27(3), 663–712.
- Guerre, E., I. Perrigne, and Q. Vuong (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica* 68(3), 525–574.
- Guibaud, S., Y. Nosbusch, and D. Vayanos (2013). Bond market clienteles, the yield curve, and the optimal maturity structure of government debt. NBER Working Paper No. 18922.
- He, Z., B. Kelly, and A. Manela (2017). Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics* 126, 1–35.
- He, Z. and A. Krishnamurthy (2013). Intermediary asset pricing. *American Economic Review* 103(2), 732–770.
- Hortaçsu, A. (2002). Bidding behavior in divisible good auctions: Theory and evidence from the Turkish Treasury auction market. Working Paper.
- Hortaçsu, A. and J. Kastl (2012). Valuing dealers’ informational advantage: A study of Canadian Treasury auctions. *Econometrica* 80(6), 2511–2542.

- Hortaçsu, A. and D. McAdams (2010). Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the Turkish Treasury auction market. *Journal of Political Economy* 118(5), 833–865.
- Huyn, D., J. Johal, and C. Garriott (2017). Do Canadian broker-dealers act as agents or principals in bond trading? Bank of Canada SAN No. 2017-11.
- Kastl, J. (2011). Discrete bids and empirical inference in divisible good auctions. *Review of Economic Studies* 78, 974–1014.
- Kastl, J. (2012). On the properties of equilibria in private value divisible good auctions with constrained bidding. *Journal of Mathematical Economics* 48, 339–352.
- Kastl, J. (2017). Recent advances in empirical analysis of financial markets: Industrial organization meets finance. In B. Honoré, A. Pakes, M. Piazzesi, and L. Samuelson (Eds.), *Advances in Economics and Econometrics: Eleventh World Congress (Econometric Society Monographs)*, pp. 231–270. Cambridge: Cambridge University Press.
- Koijen, R. S. J. and M. Yogo (2019). A demand system approach to asset pricing. *Journal of Political Economy* 127(4), 1475–1515.
- Modigliani, F. and R. Sutch (1966). Innovations in interest rate policy. *The American Economic Review* 56(1/2), 178–197.
- Vayanos, D. and J.-L. Vila (2009). A preferred-habitat model of the term structure of interest rates. NBER Working Paper No. 15487.
- Wallace, N. (1981). A Modigliani-Miller theorem for open-market operations. *The American Economic Review* 71(3), 267–274.
- Wittwer, M. (2020). Interconnected pay-as-bid auctions. *Games and Economic Behavior* 121, 506–530.

Appendix

5 Proofs

5.1 Proof of Proposition 1 and Corollary 1

For notational convenience we drop the superscript g and the subscript i, τ of all parameters $\{\kappa_{1,i,\tau}^g, \kappa_{2,i,\tau}^g, \gamma_{i,\tau}^g, \rho_{i,\tau}^g, a_{i,\tau}^g, b_{i,\tau}^g, e_{i,\tau}^g\}$.

Proposition 1. Recall that the dealer expects the following payoff from owning q_1, q_2 :

$$V(q_1, q_2, s) = U(q_1, q_2, s) + \mathbb{E}[\text{revenue}(\mathbf{x}_1, \mathbf{x}_2|q_1, q_2) - \text{cost}(\mathbf{x}_1, \mathbf{x}_2|q_1, q_2),] \quad (8)$$

with $\text{revenue}(x_1, x_2|q_1, q_2) = \sum_{m=1}^2 p_m(x_1, x_2|q_1, q_2)x_m$. Given the aggregate inverse demand of the dealer's clients (5):

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

Inserting the assumed functional forms (4), (5), and $f(x_1, x_2) = 1 + 3\rho(1 - 2F_1(x_1))(1 - 2F_2(x_2))$, integrating and taking the partial derivative w.r.t. q_1 we obtain:

$$\begin{aligned} v_1(q_1, q_2, s_1) &= 1/2\gamma\kappa_1(-1 + \rho) - 2\gamma\kappa_1\kappa_2^3q_2^3\rho + 1/2\gamma\kappa_1\kappa_2^2q_2^2(1 + 3\rho) \\ &+ q_1^2(-6\gamma\kappa_1^3\kappa_2q_2\rho + 3(3\gamma + 2e)\kappa_1^3\kappa_2^2q_2^2\rho - 4(\gamma + 2e)\kappa_1^3\kappa_2^3q_2^3\rho + \kappa_1^3(-b + \gamma\rho)) \\ &+ q_1(2(3\gamma + 2e)\kappa_1^2\kappa_2^3q_2^3\rho + c\kappa_1^2\kappa_2q_2(1 + 3\rho) + 1/2\kappa_1^2(2a + \gamma - 3\gamma\rho) - 1/2\kappa_1^2\kappa_2^2q_2^2(\gamma + 2e + 15\gamma\rho + 6e\rho)) \\ &+ (1 - \kappa_1)t_1. \end{aligned}$$

A Taylor expansion around $(\frac{1}{2}, \frac{1}{2})$ gives

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

with

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

Corollary 1. Securities become more complementary when $h_2(\kappa_1, \kappa_2, \gamma, e, \rho)$ increases. For any $\kappa_m \in [0, 1]$ and any ρ that is within the allowed range of correlation parameters of the Farlie-Gumbel-Morgenstern Distributions with uniform marginal distributions, $[-1/3, 1/3]$:

$$\begin{aligned}\frac{\partial h_2(\kappa_1, \kappa_2, \gamma, e, \rho)}{\partial e} &= -(1/2) \underbrace{\kappa_1^2 \kappa_2^2 (1 + 3(-1 + \kappa_1)(-1 + \kappa_2)\rho)}_{\geq 0} \leq 0 \\ \frac{\partial h_2(\kappa_1, \kappa_2, \gamma, e, \rho)}{\partial \gamma} &= -(1/4) \underbrace{\kappa_1(\kappa_1(-2 + \kappa_2) - 2\kappa_2)}_{\leq 0} \underbrace{\kappa_2(1 + 3(-1 + \kappa_1)(-1 + \kappa_2)\rho)}_{\geq 0} \geq 0 \\ \frac{\partial h_2(\kappa_1, \kappa_2, \gamma, e, \rho)}{\partial \rho} &= -(1/4) \underbrace{\kappa_1(\kappa_1(-2 + \kappa_2) - 2\kappa_2)}_{\leq 0} \underbrace{\kappa_2(1 + 3(-1 + \kappa_1)(-1 + \kappa_2)\rho)}_{\geq 0} \geq 0 \quad \square\end{aligned}$$

5.2 Proof of Proposition 2

The proposition follows from Proposition 3 when all δ parameters are 0. \square

5.3 Proof of Proposition 3

Take the perspective of bidder i who belongs to group $g \in \{c, d\}$. Fix his type, a time slot τ , as well as one of his information sets $\theta_{i,\tau}^g$, and let all other agents $j \neq i$ play a type-symmetric equilibrium. In this equilibrium it must be optimal for the bidder to choose the same set of functions $\{b_1^g(\cdot, \theta_{i,\tau}^g), \dots, b_M^g(\cdot, \theta_{i,\tau}^g)\}$ as all other bidders in his bidder group with information $\theta_{i,\tau}^g$. These M functions must *jointly* maximize the bidder's expected total surplus. It must therefore be the case that each of the functions $b_m^g(\cdot, \theta_{i,\tau}^g)$ maximizes his expected total surplus *separately* when fixing all the other bidding functions $-m$ at the optimum. To determine necessary conditions of the type-symmetric equilibrium we can consequently fix the agent's strategy in all but one auction at the equilibrium. Without loss take this auction to be the one for security 1.

The remainder of the proof extends Kastl (2012)'s proof for a K-step equilibrium of a discriminatory price auction that takes place in isolation without difficulties. To facilitate the comparison with the original proof (on pp. 347–348 of Kastl (2012)) we copy it as closely as possible but adopt the notation used in this article.

There are two main differences to the original proof. First, our framework allows bidders to update their bids due to arrival of new information. Such information arrives at discrete time slots $\tau = 1 \dots \Gamma$. Bidding functions do not (only) depend on the bidder i 's type $s_{i,\tau}^g$ drawn at time τ but on the (entire) information set at that time $\theta_{i,\tau}^g$. It includes the type, $s_{i,\tau}^g \subseteq \theta_{i,\tau}^g$. Since only final bids count, bidders bid as if it was their last bid each time they place a bid. We can just keep some τ fixed throughout the proof.

Second, following Hortagsu and Kastl (2012) we allow for asymmetries in bidding behavior between dealers and customers. They draw types from (potentially) different distributions and may have different information available. The original proof extends to this setup.

Simplified Notation. We drop subscripts τ, i as well as superscript g . We refer to the amount a bidder with information θ wins at market clearing in auction m (for a given set of strategies in the

event that τ is the time of the bidder's final bid) by \mathbf{q}_1^c , and the amount he wins in equilibrium by \mathbf{q}_1^* .

Notice that both, \mathbf{q}_1^c and \mathbf{q}_1^* are (for given strategies of all agents) functions of the total supply \mathbf{Q}_1 and the information of all agents $\{\theta_i\}_{i=1}^N$. They are implicitly defined by market clearing.

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

Lemma 1. *Fix a bidder with information θ .*

Denote his marginal willingness to pay in auction m at step k when submitting some function $b'_1(\cdot, \theta)$ with $\{(b'_{1,k}, q'_{1,k-1}), (b'_{1,k+1}, q'_{1,k})\}$ by

$$\tilde{v}_1(q_1, \theta | b'_{1,k}, b'_{1,k+1}) \equiv \mathbb{E} [v_1(q_1, \mathbf{q}_{-1}^*, s_1) | b'_{1,k} \geq \mathbf{P}_1^c > b'_{1,k+1}, \theta] \text{ for } q_1 \in (q'_{1,k-1}, q'_{1,k}]. \quad (14)$$

(i) $\tilde{v}_1(q_1, \theta | b'_{1,k}, b'_{1,k+1})$ is bounded.

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

Proof of Lemma 1. (i) By Assumption 2

$$\tilde{v}_1(q_1, \theta | b'_{1,k}, b'_{1,k+1}) \stackrel{(9)}{=} \alpha + (1 - \kappa_1)s_1 + \lambda_1 q_1 + \delta_1 \cdot \mathbb{E} [\mathbf{q}_{-1}^* | b'_{1,k} \geq \mathbf{P}_1^c > b'_{1,k+1}, \theta]$$

for $q_1 \in (q'_{1,k-1}, q'_{1,k}]$. Since types and total supply are drawn from distributions with bounded support by Assumptions 1 and 5, $\mathbb{E} [\mathbf{q}_{-1}^* | b'_{1,k} \geq \mathbf{P}_1^c > b'_{1,k+1}, \theta]$ and with it $\tilde{v}_1(q_1, \theta | b'_{1,k}, b'_{1,k+1})$ is bounded.

(ii) In equilibrium $\tilde{v}_1(q_1, \theta | b_{1,k}, b_{1,k+1})$ must be decreasing in q_1 or it could not give rise to a decreasing bidding function that fulfills the necessary conditions of Proposition 3.

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

Lemma 2. *Fix a bidder with information θ .*

If at some step k in auction 1, $\Pr(\mathbf{q}_1^c \geq q_{1,k} | \theta) > 0$, then $b_{1,k} \leq \tilde{v}_1(q_1, \theta | b_{1,k}, b_{1,k+1})$.

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

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

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

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

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

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

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

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

where

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

and

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

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

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

obtain

$$(q_{1,k} - q_{1,k-1})\varepsilon < \pi_1 ED_\varepsilon. \quad (15)$$

As $b_{1,k} \leq \tilde{v}_1(q_{1,k}, \theta | b_{1,k}, b_{1,k+1})$ by Lemma 2 and $\tilde{v}_1(q_{1,k}, \theta | b_{1,k}, b_{1,k+1}) < \tilde{v}_1(q_1^u, \theta | b_{1,k}, b_{1,k+1})$, the LHS of (15) goes to 0 and the RHS to a strictly positive number as $\varepsilon \rightarrow 0$. Since $\tilde{v}_1(q_1, \theta | b_{1,k}, b_{1,k+1})$ is for any $q_1 \in [\bar{q}_1^{RAT}, q_{1,k}]$ right-continuous in $b_{1,k}$, the proposed deviation would indeed be strictly profitable for the bidder with information θ . Moreover, there can be only countable many θ 's with a profitable deviation, otherwise bidder i could implement this deviation jointly and thus for *a.e.* information sets θ ties have zero probability in equilibrium for all bidders i .

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

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

6 Robustness

6.1 Regressions with bids rather than marginal valuations

In this section we report results measuring interdependence across maturities using observed bids rather than using the marginal valuations estimated using the auction model. Across all specifications, when using bids the λ parameters are downward biased relative to using marginal valuations while the δ parameters are upward biased. The former would lead us to underestimate the steepness of valuations while the latter would lead us to overestimate the degree of complementarities across maturities.

Table 8: The average dealer (with bids as independent variables)

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	-4.924***	(0.0256)	λ_{6M}	-7.789***	(0.0465)	λ_{1Y}	-15.54***	(0.0815)
$\delta_{3M,6M}$	0.384***	(0.0599)	$\delta_{6M,3M}$	1.034***	(0.102)	$\delta_{1Y,3M}$	1.606***	(0.196)
$\delta_{3M,1Y}$	0.367***	(0.0642)	$\delta_{6M,1Y}$	0.642***	(0.109)	$\delta_{1Y,6M}$	1.112***	(0.204)
const	995651***	(0.351)	const	991639***	(0.692)	const	981593***	(1.193)
N	58542		42282				50410	

Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 9: Dealer group 1 (with bids as independent variables)

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	-5.871***	(0.0471)	λ_{6M}	-8.738***	(0.0826)	λ_{1Y}	-18.23***	(0.146)
$\delta_{3M,6M}$	1.330***	(0.107)	$\delta_{6M,3M}$	1.541***	(0.172)	$\delta_{1Y,3M}$	0.957**	(0.327)
$\delta_{3M,1Y}$	0.435***	(0.119)	$\delta_{6M,1Y}$	1.131***	(0.195)	$\delta_{1Y,6M}$	2.403***	(0.372)
const	995661***	(0.524)	const	991402***	(1.031)	const	981210***	(1.782)
N	28592		21406				25134	

Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 10: Dealer group 2 (with bids as independent variables)

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	-4.425***	(0.0272)	λ_{6M}	-7.228***	(0.0503)	λ_{1Y}	-13.66***	(0.0849)
$\delta_{3M,6M}$	-0.146*	(0.0646)	$\delta_{6M,3M}$	0.665***	(0.115)	$\delta_{1Y,3M}$	2.424***	(0.221)
$\delta_{3M,1Y}$	0.429***	(0.0677)	$\delta_{6M,1Y}$	0.350**	(0.117)	$\delta_{1Y,6M}$	0.303	(0.208)
const	995652***	(0.440)	const	991899***	(0.868)	const	982067***	(1.463)
N	29863		20818				25097	

Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

6.2 Robustness with respect to the number of steps included in the regression

Here we display the estimation results when using bidding functions with more than 2 rather than more than 1 step as in the main specification. On the left it uses the marginal values as independent variable, on the right the bids. Qualitatively, all findings are as in our main specification. Results for the 6M and 12M auctions are available upon request.

Table 11: All dealer - 3M Bill Auction (using bidding functions with > 2 steps)

	v_k		b_k	
λ_{3M}	-5.285***	(0.0271)	-4.979***	(0.0260)
$\delta_{3M,6M}$	0.169**	(0.0641)	0.378***	(0.0614)
$\delta_{3M,1Y}$	0.240***	(0.0684)	0.366***	(0.0655)
Constant	995696.4***	(0.367)	995686.6***	(0.351)
Observations	55822		55822	

Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 12: Dealer group 1 - 3M Bill Auction (using bidding functions with > 2 steps)

	v_k		b_k	
λ_{3M}	-6.214***	(0.0492)	-5.870***	(0.0474)
$\delta_{3M,6M}$	1.057***	(0.112)	1.333***	(0.108)
$\delta_{3M,1Y}$	0.372**	(0.125)	0.444***	(0.120)
Constant	995725.1***	(0.539)	995715.5***	(0.520)
Observations	27656		27656	

Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 13: Dealer group 2 - 3M Bill Auction (using bidding functions with > 2 steps)

	v_k		b_k	
λ_{3M}	-4.795***	(0.0295)	-4.512***	(0.0279)
$\delta_{3M,6M}$	-0.346***	(0.0711)	-0.184**	(0.0672)
$\delta_{3M,1Y}$	0.275***	(0.0741)	0.429***	(0.0701)
Constant	995671.3***	(0.471)	995661.2***	(0.445)
Observations	28147		28147	

Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$