Estimating Demand Systems with Bidding Data*

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

February 16, 2025

Abstract

We introduce a new framework for estimating demand across multiple assets, accommodating any degree of substitution or complementarity. Unlike existing methods, our approach does not require price instruments, which are often difficult to obtain. We outline the data requirements for implementing our approach and demonstrate its application using bidding data from Canadian Treasury bill auctions. Applying the traditional concept of substitutability—where demand for one good decreases as the price of another rises—we find that different Treasury bills are weak complements for primary dealers. Understanding these demand elasticities is a crucial factor in assessing potential changes to market rules.

Keywords: Demand estimation, multi-unit auctions, Treasuries

JEL classification: D44, C14, E58, G12

*The presented views are those of the authors, not necessarily of the Bank of Canada. All errors are our own. We thank the Treasury, auctions and settlement systems team and the debt-management team at the Bank of Canada for critical insights into debt management and auction operations. We also thank Nikhil Agarwal, Rui Albuquerque, Nina Boyachenko, Giulia Brancaccio, Markus Brunnermeier, Jens Christensen, Jeremy Fox, Matthew Gentry, Alessandro Gavazza, Pierre-Olivier Pineau, Eric Richert, Marzena Rostek, Paulo Somaini, Andreas Uthemann, Philippe van der Beck, Annette Vissing-Jorgensen, Ji Hee Yoon, and Yu Zhu, as well as seminar/conference participants at the Bank of Canada, Boston College, Chicago Booth, Columbia, EUI, Kellogg, Princeton, Penn, Stanford, Queen's University; EARIE 2018/2022, ECB workshop on money markets and central bank balance sheets 2018, EEA 2022, EWM 2021, Joint Bank of Canada-John Deutsch Institute Workshop on Financial Intermediation and Regulation 2018, MIT DSE conference 2022, CEPR Asset Pricing Meeting Gerzensee 2024. Correspondence to: [a]Jason Allen - Bank of Canada, Ottawa, ON K1A0G9, Canada, [b]Jakub Kastl (NBER and CEPR) - Economics, Princeton University, NJ 08540, USA, [c]Milena Wittwer (corresponding author) - Finance, Boston College, MA 02467, USA, Email: wittwer@bc.edu

1 Introduction

Each year, trillions of dollars' worth of related securities and commodities are sold concurrently. This includes Treasury securities of different maturities, various types of mortgage-backed securities, international carbon allowances, and diverse forms of renewable energy. Similarly, exchanges provide a wide array of stocks that investors can trade simultaneously. Technically, all these markets clear through auctions known as multi-unit auctions, wherein bidders demand or supply multiple units of the same good (e.g., Wilson (1979); Kyle (1985); Chen and Duffie (2021); Rostek and Yoon (2021b); Wittwer (2021); Budish et al. (2022)).

We exploit institutional features of parallel multi-unit auctions to develop a new framework for identifying demand schedules for multiple assets or commodities, and illustrate how to implement our method using data from Canadian Treasury auctions. Our demand estimation approach is well-suited for centralized markets where bidders submit packages of price-quantity orders for multiple assets, and the researcher observes these orders. Existing methods, such as the approach by Koijen and Yogo (2019), are better suited when one only observes holding positions or aggregate trade shares.

The key idea behind our identification of demand systems lies in combining two institutional features prevalent in many markets. First, parallel auctions take place under the same market rules, with the same set of participants, at the same time, and within the same economic environment. This allows us to control for unobserved heterogeneity across markets. Second, in multi-unit auctions, bidders submit full demand schedules, i.e., packages of price-quantity orders, which combines their willingness to pay with (in some cases) a strategic component. As a result, there is no need to aggregate data across different time periods and market participants. This approach ensures that, after eliminating the strategic component, variation in quantities is attributable to variation in prices and not something omitted that is potentially correlated with prices.

Our approach offers several advantages over existing demand estimation methods for financial assets, providing a practical solution to shortcomings of standard approaches, as recently formalized by Fuchs et al. (2024). First, there is no need for a price instrument. In contrast, the existing literature tackles the endogeneity issue by employing instrumental variables and making exogeneity and validity assumptions (following Berry et al. (1995) in

industrial organization, and Koijen and Yogo (2019) in finance). This is especially challenging in financial settings where prices and quantities are correlated across assets and markets. Second, our method can flexibly identify any degree of interdependency, ranging from perfect substitutes to perfect complements. Inspired by Chen and Duffie (2021), Rostek and Yoon (2021b), and Wittwer (2021), we allow demand for one asset to depend on the prices of all available assets. In this regard, our approach is closer to the Almost Ideal Demand System (e.g., Hausman et al. (1994)). Third, we do not need to impose any correlation of latent demand shocks for different goods, which is typically arbitrarily chosen and affects demand estimates (Gentzkow (2007)). Other methods for generalizing demand estimation for financial products, such as those by Koijen and Yogo (2020), Chaudhary et al. (2023), and Darmouni et al. (2023), which estimate nested logit models, require assumptions about the correlation structure. For example, assuming that bonds with similar credit ratings or maturities are close substitutes while others are not. Fourth, our method enables flexible time-variation in demand coefficients, providing a valuable tool for analyzing market stability or the transmission and effects of monetary policy at higher frequency. Fifth, our approach can discover latent demand heterogeneity. This is useful because it is rare for researchers to observe all demand factors, which include beliefs about market conditions or balance sheet constraints (Brunnermeier et al. (2021)).

In Section 2, we introduce our framework and explain how to apply it in different market environments. Although straightforward to allow bidders to buy and sell simultaneously, in our baseline model bidders only seek to buy two assets in parallel auctions. Each bidder observes a two-dimensional private signal that determines how much they are willing to pay for different amounts of each asset. They view the assets as substitutes (complements) if they demand less (more) of one asset when the price of the other asset increases. Based on this information, they submit bids which characterize the prices they would like to pay for different amounts of the assets in each auction. Each auction clears at the lowest price at which aggregate demand satisfies supply. Winning bidders either pay the market clearing price (uniform price auction) or the prices they offered to pay (discriminatory price auction).

To estimate demand systems using our framework the data set must include bidding information. This means that we can observe packages of price-quantity pairs linked to the same trade, a requirement that can be easily verified. Ideally, each bidder has an identifier,

which is often the case in one-sided auctions (see Clark et al. (2021)), and for some exchange markets (e.g., Foucault et al. (2007); Brogaard et al. (2019); Bogousslavsky and Muravyev (2024); Dix and Wittwer (2024)). If bidder identifiers are unavailable, as, for instance, with NYSE Trade and Quote (TAQ) data, one could construct demand curves by clustering orders using time-stamps.

We first explain how to recover demand systems for environments with price-taking bidders with bidding data. This is more suitable for markets with many small bidders, such as online platforms like Robinhood, where retail investors can trade on a limit order book. Price-taking bidders bid how much they are actually willing to pay. If assets are neither substitutes nor complements, we can identify each bidder's demand by leveraging variation across the price-quantity pairs which are associated with the bidder's order. We simply regress submitted bids on demanded quantities plus order fixed effects.

When bidders maximize portfolios, different assets are likely substitutable or complementary. Thus, their true demand for one asset not only depends on the price for that asset (own-asset demand coefficient), but also the price of the other asset (cross-asset demand coefficient). This means that bidders would ideally want to condition their bids for each of the assets on how much they win of the other asset. However, auction rules prevent this: each auction clears based on demand submitted for that asset, and not based on demand submitted for asset bundles. Therefore, bidders take expectations of how much they will win of the other asset. To identify demand systems, i.e., own-and cross-asset demand coefficients, we need to back out these expectations. We illustrate one way to do this by simulating market clearing through resampling bids.

Next, we explain how to recover demand systems for markets with few, large bidders. In these cases, bidders are not price-takers. We generalize our framework to these settings. The main difference to environments with price-taking bidders is that strategic bidders internalize that their own order will impact the market clearing price. They shade their bids to minimize the payments they need to make in equilibrium. To identify demand coefficients in this environment, we define an auction game that determines how bidders place their bids and what they win in all scenarios. This creates a mapping between each bidder's true demand and their submitted bids, enabling us to estimate the expected winnings for each bidder in each auction by repeatedly simulating market clearing based on the observed bids.

Doing so, we make a methodological contribution by extending techniques for identifying demand (or willingness to pay) from bidding data in multi-unit auctions by Guerre et al. (2000), Hortaçsu (2002) and Kastl (2011) to allow demand to depend not only on the allocation of the underlying asset, but also on prices of assets of other assets. This complements a handful of alternative estimation strategies for identifying product synergies in other auction formats, such as package auctions, sequential auctions, and ascending auctions (e.g., Cantillon and Pesendorfer (2006); Kim et al. (2014); Jofre-Bonnet and Pesendorfer (2003); Kong (2021); Fox and Bajari (2013)).

In Section 3, we demonstrate how to adapt our baseline model to a specific empirical context and apply our framework using data from Canadian Treasury auctions, which are comparable to Treasury auctions in many other developed economies. Our data set includes information on all Canadian T-bill auctions between 2002 until 2015.[1] We observe which security and maturity type is issued at what amount, all submitted bids, i.e., demand schedules, and where the market clears. We can also identify each bidder thanks to unique identifiers. We focus on primary dealers, the key participants in this market.

We find that the T-bill demand of primary dealers is price-elastic, and that bills of different maturities are weak complements. This means that a dealer's demand for one maturity rises when the price of another maturity falls. For example, a price increase of one basis point in 6-month bills reduces dealer demand in the 6-month auction by approximately 10 percentage points of 6-month supply, or C$ 212 million (own-price demand effect). Concurrently, this price increase raises dealer demand in the 3-month auction by about 1 percentage point of 3-month supply, or C$ 57 million, and by about 1 percentage point of 12-month supply, or C$ 21 million (cross-price demand effect). All cross-price effects are comparable in magnitude: approximately 1 for maturities three months apart (e.g., 3M and 6M, or 6M and 12M), and around 0.5 for 3M and 12M.

Importantly, T-bills exhibit weak complementarity despite the fact that the prices of all maturities—and the underlying bond valuations for dealers—tend to move in the same direction, in line with existing studies. Moreover, dealers are willing to pay less for one maturity when they acquire more of another maturity. This underscores that when assets

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

are interdependent it matters whether we analyze price changes relative to quantities or quantity changes relative to prices.

We conclude our analysis by highlighting the significance of understanding demand elasticities as being the fundamental first step in optimizing any objective function to improve market rules. In the context of Treasury auctions, the government's main aim is to issue debt at the lowest possible cost for taxpayers. We argue that the government should behave like a monopolist who price discriminates, and exploit differences in demand elasticities to obtain higher auction revenues.[2] In various contexts, such as analyzing limit order books or auctions on carbon allowances, the objectives differ, aiming to maximize total gains from trade or reduce total carbon emissions, respectively.

Our demand analysis complements two additional strands of literature beyond the recent work on demand estimation based on Koijen and Yogo (2019) that we have already discussed. The first is the broad and longstanding literature that quantifies how asset prices respond to changes in quantity across various asset classes, employing diverse methodological approaches. This literature considers the own-price effects, i.e., by how much the price of a security (or an asset class) changes when demand or supply of that security changes—similar to the own-demand effect we measure. This is true for early contributions, such as Shleifer (1986) or Hasbrouck (1991), as well as more recent contributions, including Chang et al. (2014), Frazzini et al. (2018), Schmickler (2020), and Pavlova and Sikorskaya (2023), where own-price impact is also referred to as a micro multiplier. Our main distinction from these and related papers is to highlight the importance of analyzing cross-asset effects, i.e., by how much does the price of a security change when demand or supply for other securities changes, and thereby uncover spill-over effects across assets. As such, we complement few studies that consider asset spill-overs, such as Chaudhary et al. (2023) for the U.S. corporate bond market.

[2]Our aim is not to attain the optimal debt allocation, but to propose a simple-to-implement policy modification. This approach avoids altering the predetermined debt issuance calendar at the fiscal year's outset and eliminates the need for adjustments to established auction formats, a typically challenging task. For example, Klemperer (2010) proposes a combinatorial auction format to sell multiple financial securities. Despite excellent theoretic properties the auction hasn't become popular in practice. Another example is that it took a long academic debate with early contributions by Bikhchandani and Huang (1989) and Back and Zender (1993), as well as field experiments to change the format of U.S. Treasury auctions from discriminatory price to uniform price. Other countries, such as China, are still debating which auction format to use (e.g., Barbosa et al. (2022)).

Second, our T-bill application relates to a literature that analyzes aggregate demand for government debt with market-level data of the secondary market. Most of this literature focuses on comparing long-term with short-term debt (e.g., Gagnon et al. (2011); D'Amico et al. (2012); Lou et al. (2013); Krishnamurthy and Vissing-Jørgensen (2011, 2012, 2015); Jansen et al. (2024)). We consider different types of short-term debt (i.e., bills), more similar to Nagel (2016), Greenwood et al. (2015), and Krishnamurthy and Li (2023). In contrast to these papers, we zoom in on the primary market, where we can identify demand of individual institutions which act as primary dealers. This is useful because primary dealers are key players in Treasury markets, and likely exhibit different preferences from the average investor in safe assets.

2 Methodology: Demand Estimation with Bidding Data

Our methodology enables the estimation of full demand systems that account for substitutabilities and complementarities without needing an instrument. It applies to environments where traders submit demand or supply curves, i.e., packages of price-quantity orders, for multiple assets simultaneously, and the market clears via a centralized mechanism.

Environments where this is the case include simultaneous one-sided auctions in which multiple units of securities or goods are exchanged and bidders are either buying or selling, not both. Example are the sale of Treasuries by governments and the purchase of Treasuries by central banks as part of quantitative easing. Similarly, mortgage-backed securities are bought and sold via auctions and reverse auctions. Other examples include the sale of commodities, international carbon allowances, and various types of renewal energy.[3]

Moreover, our approach is also applicable when market participants buy and sell on exchanges or online broker platforms. In these cases, investors can simultaneously trade multiple products via limit orders. The limit orders, or packages of them, represent net demand curves, which specify how much a trader seeks to buy net of how much they seek

[3]Treasury securities are sold in parallel in most large economies issue, including the U.S. (the auction schedule is available here). For details regarding auctions for mortgage-backed securities and quantitative easing, see, for example, Becker et al. (2023), and Song and Zhu (2018), respectively. For examples of simultaneous auctions of carbon allowances by the E.U. and U.K., click here or here, by Canada and Quebec, click here; of (spot) diamonds, see Cramton et al. (2013); of renewable energy, see Ryan (2022) or click here; of procurement products, see Somaini (2020). All websites to click on were accessed on 02/02/2025.

to sell at different prices. Market clearing of parallel auctions occurs when all traders who seek to trade at that moment have submitted their net-demand curves. In practice, trading occurs continuously over time. Thus, to operationalize our approach with exchange or broker-platform data, one would need to discretize time and assume that market-clearing occurs at specific intervals.

2.1 Baseline Framework

We introduce our baseline framework for traders who buy assets in parallel.[4] Additionally, this framework approximates trading environments where the group of investors whose demand we aim to estimate are buying from other investors, with the supply from these latter traders taken as given. To adjust the framework for settings where bidders buy and sell simultaneously, it is sufficient to replace demand with net demand. This approach is common in the related market microstructure literature following Kyle (1989), e.g., Rostek and Yoon (2021b); Wittwer (2021); Budish et al. (2022).

There are M assets. Total supply of these assets, $\boldsymbol{Q} = (Q_1, ..., Q_M)$, is unknown to all bidders. $\mathbf{Q}$ is highlighted in **bold**, because it is a vector—a convention we adopt throughout the paper. Supply is drawn from some continuous joint distribution. For simplicity, we normalize expected supply to 1, $\mathbb{E}[Q_m] = 1$, for all m.

There are N bidders who seek to buy multiple units of these assets. Each bidder observes a signal, $\boldsymbol{s_i} = (s_{i,1}, ..., s_{i,M})$. The signal determines how much they are willing to pay for the assets. It is comprised of a known common component, $\boldsymbol{\mu} = (\mu_1, ..., \mu_M) \in \mathbb{R}^M_+$, and a private component, $\boldsymbol{\nu_i} = (\nu_{i,1}, ..., \nu_{i,M})$. The common component captures the resale price of the asset or common balance sheet shocks, while the private component reflects idiosyncratic bidder-specific demand-shocks, such as client orders.[5] The private component is drawn from some commonly known atomless distribution with strictly positive density on bounded support. While we allow signals to be correlated within a bidder across assets, they are independently drawn across bidders. For simplicity, $\boldsymbol{s_i}$ is also independent from supply—an assumption

[4] Alternatively, we could formulate the baseline for auctions where all bidders sell, as in quantitative easing or procurement auctions. In these cases, demand curves become supply curves, but the concepts and results remain applicable.

[5] For simplicity, we do not explicitly model bidder inventory, but it is straightforward to introduce exogenous private inventory positions, $\boldsymbol{z_i}$. Signal $\boldsymbol{s_i}$ would then be a function of $\boldsymbol{\mu}, \boldsymbol{\nu_i}$, and $\boldsymbol{z_i}$.

that we can relax. In Section 2.6 we discuss approaches that take into account correlation in signals or private beliefs about the common component.

A bidder who observes signal $\boldsymbol{s_i}$ obtains gross utility, $V(\boldsymbol{q}, \boldsymbol{s_i})$, from owning bundle $\boldsymbol{q} = (q_1, ..., q_M)$. In our baseline model, we impose a simple utility function, which is common in the related theory literature (recently overviewed by Rostek and Yoon (2021a)):[6]

$$V(\boldsymbol{q}, \boldsymbol{s_i}) = \boldsymbol{s_i} \cdot \boldsymbol{q}' + \frac{1}{2}\boldsymbol{q} \cdot \boldsymbol{\Sigma} \cdot \boldsymbol{q}' \text{ with } \boldsymbol{\Sigma} = \begin{pmatrix} \lambda_1 & \delta_{2,1} & ... & \delta_{M,1} \\ \delta_{1,2} & \lambda_2 & ... & \delta_{M,2} \\ ... & & & \\ \delta_{1,M} & \delta_{2,M} & ... & \lambda_M \end{pmatrix}, \tag{1}$$

where $\lambda_m < 0, \delta_{m,-m} \in \mathbb{R}, |\delta_{m,-m}| < \lambda_m$, and $\delta_{m,-m} = \delta_{-m,m}$ for all $m \neq m$. This gross utility implies that in auction m, bidder i is willing to pay

$$\frac{\partial V(\boldsymbol{q}, \boldsymbol{s_i})}{\partial q_m} \equiv v_m(q_m, \boldsymbol{q_{-m}}, \boldsymbol{s_i}) = s_{m,i} + \lambda_m q_m + \boldsymbol{\delta_m} \cdot \boldsymbol{q_{-m}} \tag{2}$$

for amount q_m, conditional on getting $\boldsymbol{q_{-m}}$, which is an $(M-1) \times 1$ dimensional vector that summarizes the vector of quantities of all other assets $-m \neq m$. Equation (2) is bidder $\boldsymbol{s_i}$'s inverse demand for asset m. It depends on the bidder's signal and two types of demand coefficients: λ_m determines the own-price demand coefficient of asset m, and $\boldsymbol{\delta_m}$ is a $1 \times (M-1)$ dimensional vector of parameters that determines the cross-price demand coefficient. For example, with three assets, $\boldsymbol{\delta_1} = (\delta_{1,2}, \delta_{1,3})$.

The bidder's system of willingness to pay curves implies demand system $\boldsymbol{q}(\boldsymbol{p}, \boldsymbol{s_i}) = \boldsymbol{\Sigma}^{-1}(\boldsymbol{s_i} + \boldsymbol{p})$, where $\boldsymbol{q}(\boldsymbol{p}, \boldsymbol{s_i}) = (q_1(\boldsymbol{p}, \boldsymbol{s_i}), ..., q_M(\boldsymbol{p}, \boldsymbol{s_i}))'$, which we obtain by setting $v_m(q_m, \boldsymbol{q_{-m}}, \boldsymbol{s_i}) = p_m$ for all m, and solving for q_m.

Definition 1. *Asset m and $l \neq m$ are substitutes if the demand for asset m, $q_m(\boldsymbol{p}, \boldsymbol{s_i})$, strictly decreases in the price of asset l. They are complements if it is strictly increasing, and independent otherwise.*

Our definition for substitutes is closely related to standard definitions for substitutes in

[6]Utility function (1) arises, for example, in settings where bidders have utility with constant absolute risk aversion (CARA, as in Kyle (1989)), and take expectations over the future returns of the asset, which are jointly Normally distributed.

consumer theory using Marshallian demand (see, for example, Mas-Colell et al. (1995)). Under this definition, two assets m and $-m$ can be complements even though willingness to pay, and ultimately the market clearing price, for asset m decreases in the quantity of asset $-m$, i.e., $\frac{\partial v_m(q_m, \boldsymbol{q_{-m}}, \boldsymbol{s_i})}{\partial q_{-m}} < 0$. To see this, consider an example with $M = 2$ assets. Then

$$\begin{pmatrix} \frac{\partial q_1(\boldsymbol{p}, \boldsymbol{s_i})}{\partial p_1} & \frac{\partial q_1(\boldsymbol{p}, \boldsymbol{s_i})}{\partial p_2} \\ \frac{\partial q_2(\boldsymbol{p}, \boldsymbol{s_i})}{\partial p_1} & \frac{\partial q_2(\boldsymbol{p}, \boldsymbol{s_i})}{\partial p_2} \end{pmatrix} = \boldsymbol{\Sigma}^{-1} = \frac{1}{\lambda_1 \lambda_2 - \delta^2} \begin{pmatrix} \lambda_2 & -\delta \\ -\delta & \lambda_1 \end{pmatrix} \text{ with } \delta = \delta_{12} = \delta_{21} \text{ and } |\lambda_m| < \delta.$$

For $\delta < 0$, the willingness to pay for asset m decreases in the quantity of asset $-m$. However, the demand for asset m increases in the price of asset $-m$. Therefore, assets are complements.

Given their true demand system, each trader submits bidding curves which characterize the price they would like to pay for each amount of the asset for sale. In line with most real-world auctions, bidding functions must be step functions—Figure 1a shows an example. Formally, in auction m, each bidder submits up to $\overline{K}$ price-quantity tuples. Thus, the bidder's action set A_m is:

$$A_m = \begin{cases} (b_m, q_m, K_m) : dim\,(b_m) = dim(q_m) = K_m \in \{1, ..., \overline{K}\} \\ b_{m,k} \in [0, \infty) \text{ and } q_{m,k} \in [0, 1] \\ b_{m,k} > b_{m,k+1} \text{ and } q_{m,k} > q_{m,k+1} \forall k < K_m. \end{cases} \tag{3}$$

The auctioneer aggregates the bids. The market clears at the lowest price, P_m^c, at which aggregate demand satisfies aggregate supply. 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 quantity, $q_{m,i}^c$. In equilibrium, we refer to this quantity as $q_{m,i}^*$.

Depending on the auction rules, the bidder either pays the market clearing price for all units won (known as a uniform price auction), or the individual price offers (known as a discriminatory price auction). Uniform price auctions are commonly used to approximate trading on exchanges (following Kyle (1989)), while discriminatory price auctions are the most common way of selling government debt in primary markets (Brenner et al. (2009)).

When there are several prices at which total supply equals aggregate demand by all bidders, the auctioneer chooses the highest one. In the event of excess demand at the market clearing price, bidders are rationed pro-rata on-the-margin.

2.2 Data Requirements

To implement our methodology the data set must include all winning and losing bids of all bidders, that is, all submitted step functions. These data can often be obtained through collaboration with financial regulators or auctioneers (e.g., Hollifield et al. (2006); Foucault et al. (2007), Fleming et al. (2017), Budish et al. (2015), Brogaard et al. (2019), Korajczyk and Murphy (2019); Hortaçsu et al. (2018), Richert (2022), Becker et al. (2023)).

Ideally, each bidder carries an identifier so that the researcher can identify who submits which bids. This information is commonly available for one-sided auctions (Clark et al. (2021)). This information is also available for some limit-order book markets (e.g., Foucault et al. (2007); Brogaard et al. (2019); Brogaard and Garriott (2019); Korajczyk and Murphy (2019); Bogousslavsky and Muravyev (2024); Dix and Wittwer (2024); Garriott et al. (2024)).[7]

If bidder identifiers are unavailable, techniques from the empirical market microstructure literature could be used to construct demand curves by clustering orders, for instance, relying on time-stamps (e.g., Frazzini et al. (2018); Korajczyk and Murphy (2019)).[8] While the broader data availability makes this approach appealing, it requires additional assumptions, such as assuming that all active traders at a given moment aim to trade either all available securities or a known subset of them.

Our methodology is most suitable for applications where most bidders submit step functions with more than one step. This means observing packages of price-quantity tuples that can be associated with the same trade—a requirement that can be easily tested.[9] We expect this to be true, for instance, in settings where bidders are risk-averse, as this leads to

[7] One approach to learn whether trader identifiers might be available is to read the rules of the exchange. Some exchange markets, such as the Hong Kong Stock Exchange or the Australian Stock Exchange, disclose, for each limit order, the issuing broker's identifier. In other markets, including INET, Euronext, or the NYSE, brokers' identifiers are concealed (Foucault et al. (2007)).

[8] Using NYSE TAQ data, Conrad et al. (2015) and Conrad and Wahal (2020) highlight that large orders often interact with several smaller counterparty orders within a single limit order book. Additionally, large traders frequently submit orders across multiple trading venues in the U.S. These simultaneous orders account for 13% of volume in small stocks and 22% of volume in large stocks (Conrad et al. (2015)), fitting our framework for demand systems for multiple securities.

[9] When working with limit-order-book data one first verification is to count the number of limit orders versus market orders. On standard exchanges, the majority of trades are executed via limit orders by institutional investors (O'Hara (2015); Brogaard et al. (2019); van Kervel and Menkveld (2015)). Even retail investors rely heavily on limit orders (Bogousslavsky and Muravyev (2024)).

downward-sloping willingness to pay curves, as explained by Wittwer and Allen (2023). We also expect this to be true for sophisticated bidders uncertain about the market-clearing price, which incentivizes strategic demand reduction (Ausubel et al. (2014)), typically resulting in multi-step bidding functions.

2.3 Recovering Demand Systems for Price-Takers

We first explain how to recover demand systems for environments with price-taking bidders. This is more suitable for markets with many small bidders, such as online platforms where retail investors can trade on a limit order book, such as Robinhood.

Proposition 1 (Price-Takers)**.** *In auction m, a price-taking bidder i submits a step function that fulfills the following condition:*

$$b_{m,k} = s_{m,i} + \lambda_m q_{m,k} + \boldsymbol{\delta_m} \cdot \mathbb{E}[\boldsymbol{q^*_{-m,i}} | q_{m,k}] \text{ for all } k \text{ and all } m, -m \neq m. \tag{4}$$

When bidders are price-takers, they bid their actual willingness to pay for amount q_m in auction m. This depends on how much of the other securities they get allocated. Ideally, the bidder would want to condition their bid in auction m for amount $q_{m,k}$ on how much they will purchase of the other security in equilibrium, $\boldsymbol{q^*_{-m,i}}$. However, by the rules of nearly every auction, this is infeasible given that the step function in auction m can only depend on prices and quantities of security m, but not of other securities. Therefore, the bidder takes an expectation over how much they will win of the other security, conditional on winning $q_{m,k}$.

Since we observe all bids $\{b_{m,k}, q_{m,k}\}$ for all steps k, bidders, and securities m, Proposition 1 implies that we can identify all demand coefficients, $\lambda_m, \boldsymbol{\delta_m}$, when we either observe, or infer, how much each bidder expects to buy of the other asset, $\mathbb{E}[\boldsymbol{q^*_{-m,i}} | q_{m,k}]$. For concreteness, assume that we observe how much each bidder expects to win of the other asset for sale, $\hat{\mathbb{E}}[\boldsymbol{q^*_{-m,i}} | q_{m,k}]$, up to a random error, $\epsilon_{m,k}$. Then, we can identify demand coefficients by leveraging variation across the k steps within each bidding function, conditional on there being more than one step. To estimate the demand coefficients, we regress observed bids on demanded quantities and the expected quantities:

$$b_{m,k} = s_{m,i} + \lambda_m q_{m,k} + \delta_m \cdot \hat{\mathbb{E}}[\boldsymbol{q^*_{-m,i}}|q_{m,k}] + \epsilon_{m,k}, \tag{5}$$

such that $\delta_{m,-m} = \delta_{-m,m}$ for all m, and $-m \neq m$. When bidders observe different private signals, we need to include a bidding-function fixed-effect to absorb these private signals, $s_{m,i}$. When all bidders share the same information an OLS regression of bids on quantities and expected quantities suffices. The number of observed bidding functions (which are submitted by bidders with different signal realizations) together with the number of steps in these functions determines how many assets, M, one can include in the demand system.

Typically, we do not directly observe the conditional expectations of market-clearing quantities, $\hat{\mathbb{E}}[\boldsymbol{q^*_{-m,i}}|q_{m,k}]$. Therefore, we need to estimate them. Before discussing how to do this, it is useful to consider situations where such estimation is unnecessary. These are cases where we either observe the expectations directly or have reliable proxies. We explore two such scenarios.

Corollary 1 (Perfect Correlation)**.** *Let $M = 2$. Let private signals be, for all bidders, perfectly correlated across securities and have the same variance. Then, a price-taking bidder i submits a step function that fulfills the following condition:*

$$b_{m,k} = \zeta_{m,i} + (\lambda_m + \delta_m) q_{m,k}, \text{ with } \zeta_{m,i} = s_{m,i} + \mathbb{E}[q^*_{-m,i}] - \mathbb{E}[q^*_{m,i}] \text{ in auction } m. \tag{6}$$

Proof: When private signals are perfectly correlated across securities for all bidders, and have the same variance, the assets' prices are perfectly correlated across auctions and share the same variance—an implication that can easily be tested. Perfect correlation also implies that the winning quantity of one security is a linear function of the other security: $q^*_{-m,i} = \mathbb{E}[q^*_{-m,i}] - g\mathbb{E}[q^*_{m,i}] + g q^*_{m,i}$ with $g = \pm\sqrt{Var(q^*_{-m,i})/Var(q^*_{m,i})}$. In addition, identical variances imply $g = 1$. Therefore, $\mathbb{E}[q^*_{-m,i}|q_{m,k}] = \mathbb{E}[q^*_{-m,i}] - \mathbb{E}[q^*_{m,i}] + q_{m,k}$. With this, bidding function (4) becomes (6). $\square$

Thus, if we are interested in estimating the demand for two securities whose market-clearing prices are nearly perfectly correlated, we can use variations across steps within bidding functions to estimate the slope coefficient: $\lambda_m + \delta_m$ with a simple fixed-effect regression. If we incorrectly assume that securities are independent ($\delta_m = 0$), we would mistakenly conclude that the willingness to pay for security m declines more (less) sharply with the

quantity of security m when the δs are actually negative (positive). This would suggest that demand is less (more) responsive to price changes, and thus more inelastic (elastic), than it is in reality. Consistent with this finding, Fuchs et al. (2024) point out that cross-security spillovers lead to biased elasticity estimates in the Koijen and Yogo framework.

To separately identify λ_m from δ_m, we would need additional variation in the data. This can be achieved through cross-bidder or cross-time variation. For cross-bidder variation, if we are able to separate bidders into two groups, one with demand for both securities and another that only seeks to buy one security, then we identify the own-security demand coefficients, λ_m, from bids placed by the second group and the δ_m coefficients from the first group (assuming that the λ_m are common across groups).[10] For cross-time variation, we need to identify episodes of market clearing when one of the securities is unavailable for sale. From these instances, we can identify the λ_m coefficients. The δ_m coefficients can then be determined from instances when all securities are available for sale under the assumption that coefficients are common across time.

Corollary 2 (Zero Correlation)**.** *Let $M = 2$. When private signals are stochastically independent across securities and all bidders observe the same signals, $s_{m,i} = \mu_m$ for all i, m, a price-taking bidder i submits a step function that fulfills the following condition:*

$$b_{m,k} = \mu_m + \lambda_m q_{m,k} + \delta_m \frac{1}{N} \mathbb{E}[Q_m] \text{ in auction } m. \tag{7}$$

Proof: When private signals are stochastically independent across securities, and all bidders observe the same signals, $s_{m,i} = \mu_m$ for all i, m, the assets' prices are stochastically independent, so that $\mathbb{E}[q^*_{-m,i}|q_{m,k}] = \mathbb{E}[q^*_{-m,i}]$. Since all bidders share the same signal, they all win $1/N$ of the total supply in each auction, and thus $\mathbb{E}[q^*_{-m,i}] = \frac{1}{N}\mathbb{E}[Q_m]$. Bidding function (4) simplifies to expression (7). □

This simplification is useful because we observe at what point submitted step functions in auction m intercept with the price axis, that is, $\mu_m + \delta_m \frac{1}{N}\mathbb{E}[Q_m]$. If we also observe the bidder's common signal, μ_m, and the expected supply, $\mathbb{E}[Q_m]$, we can back out δ_m, since

[10] For institutional investors, this separation could be based, for instance, on investment mandates, following Koijen and Yogo (2019). For retail investors, we could infer their security investment set from their portfolio holdings, observable, for example, through Robintrack (`https://robintrack.net`) for Robinhood investors (e.g., Barber et al. (2022)).

we know how many bidders compete, N. Among these variables, we typically have a good sense of the supply (or average supply) that is available for sale (or purchase), $\mathbb{E}[Q_m]$. If using exchange-data, one proxy for the bidder's common signal could be the mid-price of the National Best Bid and Offer (NBBO) prior to the trade. In our empirical application—Treasury auctions—a natural candidate to approximate a common signal of the asset's value is the when-issued market price, which bidders observe prior to bidding. In other settings, price quotes provided by Bloomberg or Thomson Reuters data feeds could serve as proxies.

For the general case in which signals or supply are imperfectly correlated, leading to imperfect correlation of asset prices, we estimate expected quantities non-parametrically. To do this, we extend non-parametric resampling techniques from the empirical auction literature to take dependencies across auctions into account (following Guerre et al. (2000), Hortaçsu (2002) and Kastl (2011)).

We fix the M bidding functions simultaneously submitted by a bidder i. We then draw a random subsample of $N-1$ bidding functions in each of the M auctions with replacement from the sample of bids to construct M realizations of residual supply curves that the bidder faces in the auctions. Next, we determine at what points the residual supply curves intersect with the fixed bidding function. This gives one realization of the clearing prices, $\boldsymbol{P^c} = (P_1^c, ..., P_M^c)$, and the amount $\boldsymbol{q_i^*} = (q_{1,i}^*, ..., q_{M,i}^*)$ this bidder would have won for all $\boldsymbol{q_i^*}, \boldsymbol{P^c}$. Repeating this procedure many times provides a consistent estimate of the joint distribution of market clearing prices and the corresponding amount of each security i would win.[11] With these joint contributions, we can compute all expected quantities, up to some random error, using the fact that $\mathbb{E}[\boldsymbol{q_{-m,i}^*}|q_{m,k}] = \mathbb{E}[\boldsymbol{q_{-m,i}^*}|b_{m,k} \geq \boldsymbol{P_m^c} > b_{m,k+1}]$.

This resampling procedure achieves consistent estimates when bidder signals are drawn independently from one another, conditional on anything that bidders observe at the time of the auction. This can include information that is unavailable to the econometrician, for example, shocks to balance sheets or risk exposures that affect all bidders. Therefore, bidders can have common values, but all are equally good (or bad) in predicting the magnitudes of those common values.

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

2.4 Recovering Demand Systems for Strategic Bidders

Next, we address environments with sophisticated bidders that have price impact and therefore not price-takers. This is more likely to be true for markets that are characterized by a few large institutional investors, as in the case of our empirical application.

To estimate demand coefficients in this setting, we need to solve for an equilibrium that pins down market clearance. Since in typical empirical settings bidding functions are discrete, we follow Kastl (2011) and provide equilibrium conditions for pure-strategy Bayesian Nash Equilibria (BNE) in which bidders submit step functions. To formally define a BNE, note that a pure-strategy is a mapping from the bidder's signal to the action space (3) of all auctions, and denote the resulting bidding function in auction m by $b_{m,i}(\cdot, \boldsymbol{s_i})$.

Definition 2. *A BNE is a collection of functions* $\boldsymbol{b_i}(\cdot, \boldsymbol{s_i}) \equiv (b_{1,i}(\cdot, \boldsymbol{s_i}) \ ... \ b_{M,i}(\cdot, \boldsymbol{s_i}))$ *that for each bidder* i *and almost every signal* s_i *maximizes their expected gross utility (1) minus the total expected payment, where the expectation is taken over where the auction will clear. The BNE is symmetric if* $\boldsymbol{b_i}(\cdot, \boldsymbol{s_i}) = \boldsymbol{b}(\cdot, \boldsymbol{s_i})$ *for all* i.

Equilibrium conditions depend on the auction rules since the total payment differs across auction formats: In uniform price auctions, where every winner pays the market clearing price, the total expected payment is $\mathbb{E}\left[\sum_m P_m^c q_{m,i}^c\right]$, where P_m^c is the clearing price in auction m, and $q_{m,i}^c$ is the amount bidder i gets allocated of asset m when submitting $b_i(\cdot, s_i)$. In discriminatory price auctions, where winners pay their bids, the total expected payment is $\mathbb{E}\left[\sum_{m=1}^{M} \int_0^{q_{m,i}^c} b_{m,i}(x, \boldsymbol{s_i}) dx\right]$. We focus on symmetric BNE, which is natural given that bidders are ex ante identical.

Proposition 2 (Non-price-takers)**.** *Let*

$$\tilde{v}_{m,i}(q_{m,k}) = s_{m,i} + \lambda_m q_{m,k} + \boldsymbol{\delta_m} \cdot \mathbb{E}[\boldsymbol{q^*_{-m,i}} | q_{m,k}]. \tag{8}$$

In a uniform price auction, a strategic bidder i *submits a step function that fulfills the following condition at all steps in any symmetric BNE:*

$$\tilde{v}_{m,i}(q_{m,k}) = \mathbb{E}[P_m^c | b_{m,k} > P_m^c > b_{m,k+1}] + \frac{q_{m,k}}{\Pr(b_{m,k} > P_m^c > b_{m,k+1})} \frac{\partial \mathbb{E}[P_m^c \mathbb{I}(b_{m,k} \geq P_m^c \geq b_{m,k+1})]}{\partial q_{m,k}} \tag{9}$$

In a discriminatory price auction, the step function fulfills:

$$\tilde{v}_{m,i}(q_{m,k}) = b_{m,k} + \frac{\Pr\left(b_{m,k+1} \geq P_m^c\right)}{\Pr\left(b_{m,k} > P_m^c > b_{m,k+1}\right)}(b_{m,k} - b_{m,k+1}) \tag{10}$$

for all steps but the last step. At the last step, shading is zero, and the bidder demands the maximal amount that may be allocated in equilibrium.

Strategic bidders do not bid their true willingness to pay; instead, they shade their bids to minimize their payments. The extent of bid-shading depends on the auction rules and the bidders' beliefs about their competitors' bidding strategies, as established in the literature. In uniform price auctions shading is determined by the trade-off between the expected marginal utility, the LHS in expression (9), and the expected loss due to the price that needs to be paid and price impact, the RHS of expression (9); see Kastl (2011) or Hortaçsu et al. (2018) for more details. In discriminatory price auctions a bidder trades off the expected surplus on the marginal infinitesimal unit versus the probability of winning it, similar to first-pice auctions; see Kastl (2017), p. 237, and Clark et al. (2021) for more details.

To estimate demand coefficients, we need to estimate how much each bidder shades their bids. The shading factors depend on observed bids and the joint distribution of market clearing prices, which we can estimate by resampling from bids, as explained above. Using this distribution, we can compute the shading factors according to expressions (9) or (10), and point-identify the willingness to pay, $\tilde{v}_{m,i}(q_{m,k})$, at each step k (Kastl (2012)).[12]

Once we know the shading factors, we can estimate a linear fixed-effect regression (5) where we only need to replace submitted bids with the bidders' actual willingness to pay:

$$\tilde{v}_{m,i}(q_{m,k}) = s_{m,i} + \lambda_m q_{m,k} + \boldsymbol{\delta_m} \cdot \hat{\mathbb{E}}[\boldsymbol{q^*_{-m,i}}|q_{m,k}] + \epsilon_{m,k}, \tag{11}$$

such that $\delta_{m,-m} = \delta_{-m,m}$ for all $m, m \neq -m$. As with price-taking bidders, the error term in this regression represents noise coming from the resampling procedure.

[12]This is slightly more complicated for uniform-price auctions than for discriminatory price auctions, because we need to take the numerical derivative of $\mathbb{E}[P_m^c \mathbb{I}(b_{m,k} \geq P_m^c \geq b_{M,k+1})]$ with respect to quantity. To do this, take two quantity points that are close to one another, and the corresponding bids on the step function. Then compute $\mathbb{E}[P_m^c \mathbb{I}(b_{m,k} \geq P_m^c \geq b_{M,k+1})]$ at both points by relying on the price distribution we have already identified, and take the ratio between the difference in these expected values and the quantity points.

2.5 Comparison to Existing Approaches to Demand Estimation

Our approach to demand estimation differs on at least five dimensions compared to existing methods in industrial organization and finance, following Berry et al. (1995) and Koijen and Yogo (2019), respectively.

First, since we combine data from multi-unit auctions with institutional knowledge about the data generating process, our method does not require price instruments. In contrast, existing approaches to demand estimation require finding valid instruments and imposing exclusion restrictions to isolate exogenous variation in prices that shift demand. This is because researchers are typically linking together different price-quantity observations over time or the cross-section to recover the demand curve. This is challenging in financial settings where prices and demand are correlated across assets and markets, and often requires strong assumptions. For example, Koijen and Yogo (2019) assume that investors allocate their wealth across assets within an exogenously given investment set, and that these sets are perfectly observed.[13] The authors acknowledge that these are strong assumptions and call for future research (p. 1477), which has led to an active and ongoing search for better instruments (e.g., Gabaix and Koijen (2024)).

Second, our method can flexibly identify any type of interdependencies, substitutes and complements—which likely arise in many settings for various reasons, for instance, when bidders face budget, capacity, or liquidity constraints, or seek to hedge positions. In fact, we do not need to impose any relationship between own and cross-product demand elasticities given that we estimate separate slope parameters (λ's and δ's) in our demand systems—different from Koijen and Yogo (2019).[14]

In this regard, our approach is closer to the Almost Ideal Demand System (e.g., Hausman et al. (1994)), while Koijen and Yogo (2019) set up their model in the characteristic space of products. In order to derive their main regression specification, which correlates normalized portfolio shares of an asset with the price and other characteristics of that asset, Koijen and

[13] Under these assumptions they construct a Hausman-style instrument for each investor that exploits variation in the wealth and investment sets of other investors. The exclusion restriction requires the latent demand shock of an investor to be uncorrelated with the wealth and investment sets of other investors. This might not be the case for sophisticated investors or dealers who strategically choose which assets to invest in.

[14] Equations (14) and (15) in Koijen and Yogo (2019) specify that the relationship between own- and cross-asset demand elasticities is driven by portfolio weights and the coefficient on log market equity.

Yogo (2019) need to make the assumption that the mean and variance of expected returns of an asset only depend on characteristics of that asset, and not on any other characteristics (Assumption 2). However, this assumption is easily violated in equilibrium of portfolio choice models (Fuchs et al. (2024)). We do not need to make this assumption.

Third, our method does not require imposing a specific correlation structure for unobservable shocks that affect demand—the ϵ errors in standard BLP models or in Koijen and Yogo (2019). This is useful because we know from Gentzkow (2007) that these arbitrarily chosen correlation structures nontrivially affect the demand estimates. Other approaches to generalize demand estimation for financial products, e.g., Koijen and Yogo (2020), Chaudhary et al. (2023), and Darmouni et al. (2023), which estimate nested logit models, need to make assumptions about the correlation structure within and across assets.

Fourth, since our approach leverages variation within orders submitted at the same time (or within a short time frame), we can estimate demand coefficients that flexibly vary over time. In contrast, Koijen and Yogo (2019) relies on quarterly or annual holdings data, making it difficult to account for time variation. Therefore, our approach is especially useful for answering questions about market stability, or the transmission and effects of monetary policy or financial news, for which it is crucial to understand how demand reacts at high frequency.

Finally, our approach can discover latent demand heterogeneities of different individuals or institutions—as illustrated in Appendix C. This is useful, because it is rare for researchers to observe all potential factors that could drive differences in demand or preferences. Examples include balance sheets or liquidity requirements of institutions, and even diverging beliefs about market conditions (as discussed in Brunnermeier et al. (2021)).

2.6 Extensions of our Methodology

Our baseline framework builds on some assumptions which might not hold in all empirical settings that could be analyzed using our methodology. In Section 3, we show how to adjust the baseline framework to better fit the data generating process of one specific setting, namely Treasury auctions. To motivate future research applying the main ideas of our approach to a broad range of settings, we discuss potential methods for relaxing some of the main assumptions of our baseline framework.

Correlated Signals or Different Beliefs About the Common Value. One of our assumptions is that bidder signals are conditionally independent across bidders and that they all agree on the common (fundamental) value of the assets. In Section 3, we show how to extend the baseline model to allow for systematic disagreement between bidders by clustering bidders into groups, with the groups disagreeing about the common value, but each investor within a group sharing the same common value. The main challenge in relaxing the assumptions further is not solving for equilibrium conditions—analogues to Propositions 1 and 2—but rather, consistently estimating expected winning quantities and shading factors.

There are two approaches to accommodate correlated signals or different beliefs about the common value. The first is to extend the non-parametric resampling procedure to these settings, which, for example, means moving away from randomly resampling with replacement from observed bids. This is challenging and would represent a significant contribution to the empirical auction literature.

The second approach, which has been taken by Boyarchenko et al. (2021), among others, is to move away from non-parametric estimation. This involves specifying distributions for all random variables in the model. Additionally, for tractability, one would need to (incorrectly) assume that bidders submit continuous bidding functions and solve for an equilibrium with linear bidding functions. Under appropriate parametric assumptions, one could estimate demand coefficients via GMM, using the analogue of Proposition 1 for price-taking bidders and Proposition 2 for strategic bidders. The advantage of this method is greater flexibility in information or value structures. The drawbacks include imposing non-testable parametric assumptions and assuming an auction game that does not accurately reflect the data-generating process, as bids cannot be continuous in most empirical applications due to auction rules. For example, a low number of submitted steps, which is the case in our empirical setting, suggests that a continuous approximation of step functions may not be a good one.

Linear Demand. While it is common in industrial organization and finance to estimate linear demand specifications, it is valuable to relax this restriction in cases where the data is rich enough to do so—that is, includes sufficiently many bids and steps. In Appendix A, we explain how to non-parametrically estimate interdependencies between assets without imposing linearity.

Dynamic Effects. Our model is static, meaning it does not account for dynamic intertemporal bidding strategies. This is reasonable for bidders participating in only one auction or for those with a time horizon shorter than the auction frequency, due to high future uncertainty making precise strategic predictions impractical. The static model might also be the best feasible approximation when the auction is not the only market in which bidders are active, but the researcher only observes auction data.

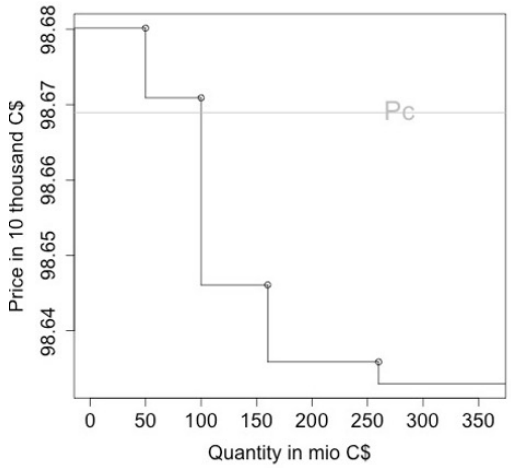
One example of this is our empirical application of Canadian Treasury auctions, where primary dealers are the most relevant bidders. These dealers are obligated to actively trade in the secondary market, requiring them to regularly fill their inventory positions at the auctions. A dynamic model of Treasury auctions that does not account for this obligation would incorrectly predict that dealers have strong incentives to refrain from bidding if they expect an upcoming auction to clear at a lower price. In reality, their market obligations prevent them from doing so. Moreover, the time between auctions is relatively long compared to the level of liquidity in the secondary market. Therefore, what happens in the secondary market prior to the auction is more important than how much a dealer won in the last auction.

In other contexts the static model might be inadequate. For instance, some large traders engage in continuous trading and strategically divide large trades into smaller ones, executed gradually over time, to minimize price impact. To extend our estimation approach to estimate demand for these traders, the first step would be to combine insights from the static auction literature on parallel uniform price auctions, such as Rostek and Yoon (2021b) or Wittwer (2021), with insights from dynamic single-asset uniform price auction models, such as Rostek and Weretka (2015). Empirically estimating such a dynamic model would be challenging but would represent a valuable contribution to the literature.

Summarizing, our approach to demand estimation can be applied to various centralized market settings where bidders submit (net) demand schedules for multiple assets, and we observe bidding data.

Figure 1: Bids in the Canadian T-Bill Market

(a) Competitive Bid for 12M



(b) Time to Deadline

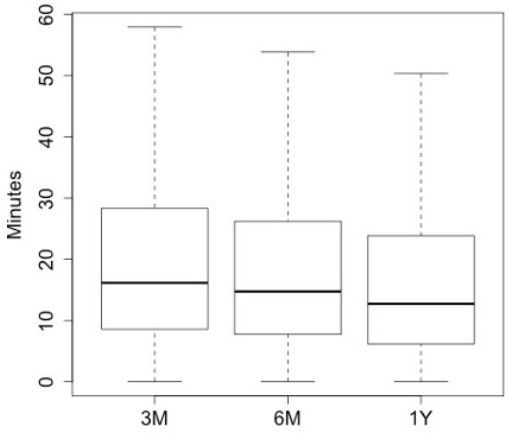


Figure 1a displays an example of a bidding step function. It is the one of the median dealer in a 12M auction, computed as follows: determine the median number of steps in all competitive bid functions submitted by dealers, and then take the median over all (price, quantity) tuples corresponding to each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

3 Empirical Application: Canadian Treasury Auctions

We use data on Canadian T-bill auctions to illustrate how to implement our approach to demand estimation when bidders are strategic.

3.1 Institutional Environment and Data

Canadian T-Bill Auctions. There are three types of Canadian T-bills (3M, 6M, 12M) which are sold every second Tuesday by the Bank of Canada in three separate, but parallel, discriminatory price auctions. T-bonds of different maturities are sold on different days. Two groups of bidders, dealers and customers, participate in the auctions. Customers can only submit bids through dealers, but like dealers, they tend to be large financial institutions.

Our focus lies on dealers, who buy more than 85% of the issued amount in an average auction. They play a key role in secondary markets, where they handle 96.7% of the daily trade volume in cash, and 86.5% of repo transactions.[15]

[15]The information on secondary market trades comes from MTRS2.0 data. It covers the period from 2016-2022. The data is self-reported and includes all trades with at least one firms that is registered with the Investment Industry Regulatory Organization of Canada (IIROC). Allen and Wittwer (2023) verify the fraction of missing data in MTRS2.0 is negligible for the cash market.

From the time the tender call opens until the auctions close, bidders may submit and update their bids, which come in two forms. The first form is a competitive bid. This is a step function with at most 7 steps. Quantities are denominated in multiples of C$1,000, while values are expressed in terms of yield to maturity. To simplify the analysis, we convert the yields into prices by employing a face value of C$1,000. This conversion ensures that the demand schedules are decreasing, rather than increasing (as shown in Figure 1a).

The second form of bidding is a non-competitive bid. This is a quantity order that the bidder will win for sure, at the average price of all accepted competitive bid prices. It is capped at $10 million for dealers and $5 million for customers, and hence trivial relative to the competitive order sizes—with one exception: the Bank of Canada itself. It utilizes non-competitive bids to reduce the previously announced supply and to purchase Treasuries (assets) to match its issuance of bank notes (liabilities).[16]

Bidding Data. Our data consists of all 366 Canadian T-bill auctions between 2002 and 2015, summarized in Table 1. We identify each bidder through a bidder ID and know whether the bidder is a dealer or a customer. The average auction has 11 to 12 dealers and 5 to 6 customers, and distributes roughly C$5.76 billion in the form of 3M and C$2.12 billion in the form of 6M and 12M bills (see Appendix Figure A1). Dealers tend to participate in all parallel auctions to keep their dealer status (see Bank of Canada (2016), p. 12, for details).

We observe all bids submitted from the opening of the tender call until the auction closes. The updating period lasts one week, although most bids arrive within 10 to 20 minutes prior to closing (see Figure 1b). Typically, a dealer updates their bid once. The higher average number of updates (2.26) is driven by outliers. In contrast, customers typically don't update their bids, with an average number of 0.1 and a median of no updates.

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

[16]For more details regarding the Canadian Treasury auction process, and related regulations, see Hortaçsu and Kastl (2012) and Allen et al. (2023).

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

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

Table 1 displays summary statistics of our auction sample, which goes from January 2002 until December 2015. There are 366 auctions per maturity. The face value is C$ 1,000. The total number of competitive bids (including updates) in the 3M, 6M, 12M auctions is 66,382, 48,927, and 56,721, respectively. These individual steps make up 18,272, 15,514, and 17,077 different step functions. The total number of non-competitive bids is 2,477, 2,378, and 1,932. From the raw data we drop competitive bids with missing bid price (133) and competitive or non-competitive bids with missing quantities (69). The variable "Global part." is the probability of attending the remaining auctions, conditional on bidding for one maturity. Comp and Non-comp dem. are the competitive and non-competitive demand of a bidder, respectively, excluding the Bank of Canada.

3.2 Model Extension for Canadian Treasury Auctions

Canadian Treasury auctions are more complex than our baseline setting in Section 2.1. We enhance the model to better fit the data generating process. Since many features of Canadian Treasury auctions, such as multiple bidder groups and bid-updating, are common to Treasury auctions in general, our model can offer guidance on estimating demand for Treasuries in other countries.

We introduce two groups of bidders: dealers (d) and customers (c). The total number of bidders in each group is commonly known.[17] In addition, all or some bidders could carry a latent type, as illustrated in a model extension in Appendix C. Given the presence of non-competitive bids, supply to these (competitive) bidders is random, and thus unknown to them when placing their bids.

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

Each bidder i of group g draws a private, potentially multi-dimensional signal $\boldsymbol{s_i^g} = (s_{1,i}^g, \ldots, s_{M,i}^g)$. We assume that dealers' and customers' private signals $\boldsymbol{s_i^d}$ and $\boldsymbol{s_i^c}$ are for all bidders i independently drawn from common atomless distribution functions F^d and F^c with support $[0,1]^M$ and strictly positive densities f^d and f^c. The bidder's signal affects their true willingness to pay, which is given by equation (2), but now with bidder-group specific signals and demand-coefficients.

A bidder's signal is private and independent from the other bidders' signals conditional on anything that every bidder (but not necessarily the econometrician) knows at the time of the auction. This includes a reference price-range provided by the auctioneer (see Appendix Figure A2), in addition to all public information that is available in the active when-issued market. The presence of this market implies that most, if not all, information relevant for price-discovery is aggregated prior to the auction and that any private information about future resale value can be arbitraged away.

Supporting this idea, we demonstrate in Table 2 that the average secondary market yield at delivery (one day after the auction) can be almost perfectly predicted using the when-issued yield from the day before the auction, with an R^2 of over 99%. This suggests that bidders are equally informed about their common value for the bond, and that their signals are private.[18] The small heterogeneity of information at the time of the auction is likely driven by idiosyncratic factors, such as the structure of individual dealer balance sheets, client order flows, investment opportunities, or repo needs.[19] Given that the market is highly segmented and many clients trade with a single dealer (as shown in Allen and Wittwer (2023), Figure 1), it is a reasonable first-order approximation to abstract from correlations in these factors across bidders.

To provide additional support for this assumption, we conduct a battery of statistical

[18]Consistent with this notion, Hortaçsu and Kastl (2012) and Allen et al. (2023) conduct formal statistical tests for private signals in Canadian Treasury auctions using the same type of dataset as ours and covering overlapping time periods. The idea is to investigates whether dealers are simply learning about competition in the upcoming auction without updating their own information (which is the case if signals are private), or if they are also learning about fundamentals and thus updating their own signals after observing customers' bids. Both studies fail to reject the null hypothesis that dealers only learn about the degree of competition in the auction, providing support for the assumption of private signals.

[19]The presence of multi-step bidding functions, or downward-sloping curves, indicates that bidders experience residual uncertainty. If bidders were certain about the auction's clearing price they would have incentive to submit single-step bids at that price to minimize their payments.

Table 2: Post-Auction Prices are Predictable

	3M Bills		6M Bills		12M Bills	
Lag WI ask	$+0.998^{***}$	.	$+0.996^{***}$	.	$+0.996^{***}$	.
	(0.001)	.	(0.001)	.	(0.001)	.
Lag WI bid	.	$+0.996^{***}$	.	$+0.995^{***}$	.	$+0.995^{***}$
	.	(0.001)	.	(0.001)	.	(0.001)
Constant	+0.001	-0.02^{***}	+0.010	−0.009	$+0.008^{**}$	-0.014^{***}
	(0.003)	(0.004)	(0.004)	(0.009)	(0.004)	(0.004)
R^2	0.999	0.999	0.999	0.996	0.999	0.999
N	398		398		359	

Table 2 tests whether post-auction yields are predictable. Our dependent variable is the average secondary market Treasury bill yield on the day immediately following the auction. We regress this yield on the when-issue ask and bid quotes of the auctioned security on the day before the auction (Lag WI ask, and Lag WI bid). We collect these quotes from CanPX, which is the main data feed for inter-dealer brokers. In unreported results, we estimate specifications with t-2 WI yields and the results are similar. Standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

tests in Appendix B. The idea for verifying conditional independence is to randomly split bids in an auction into two parts and check for correlations across these sets. We compute four different test statistics and all of them lead us to fail rejecting the null hypothesis of conditional independence at a 5% significance level, apart from occasional exceptions.

Finally, we account for the fact that dealers may update their own bids when observing a customer's bid. We let each dealer place (competitive) bids at a discrete number of time slots $\tau = 1, ..., \Gamma$ over the course of the auction. To rationalize the fact that dealers do not wait until the very last moment to submit all of their bids, we follow Hortaçsu and Kastl (2012).[20] Each time a dealer places a bid, they think that this might be their last bid, and therefore the one that counts for market clearing. This assumption reflects the fact that it takes some time to compute bids, enter them manually into the auction interface and transmit them to the auctioneer. Dealers are therefore uncertain about whether a last-minute bid will make it in time, and have an incentive to place "early" bids before auction closure (see Figure 1b).

Since the typical customer does not update their bid, we assume that customers only bid once. Each customer is matched to a dealer who observes their bid. This provides the dealer

[20]Formally, at $\tau = 1$, a dealer draws an iid random variable $\Psi_i \in [0, 1]$ which is one dimension of the dealer'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 dealer's later bids will make it in time to be accepted by the auctioneer. Thus, for $\tau > 1$, the dealer'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.

with additional private information about the degree of competition, which we denote by $Z_{m,i,\tau}$. Thus, if the dealer observed a customer's bid in all M auctions, their information set is $\boldsymbol{\theta_i^g} = (\boldsymbol{s_i^g}, Z_{1,i,\tau}, ..., Z_{M,i,\tau})$. If the dealer only has a customer in one auction, say for maturity 1, $\boldsymbol{\theta_{i,\tau}^g} = (\boldsymbol{s_i^g}, Z_{1,i,\tau})$, and so on.

A bidding strategy now maps from a bidder's set of information sets to the action space (3) for all auctions and for each time the bidder has a chance to bid. We focus on group-symmetric BNE, in which bidders within the same group choose the same bidding function when observing the same information set. For dealers, which are in the focus of our analysis, the equilibrium conditions are analogous to Proposition 2 for each τ with the only difference that the signal $\boldsymbol{s_i}$ is replaced by the dealer's information set at that time $\boldsymbol{\theta_{i,\tau}^d}$.

3.3 Estimation of Demand for Canadian T-Bills

We estimate demand coefficients under various assumptions about the data-generating process. To take customer bidding and dealer updating into account when estimating the joint-distribution of winning quantities, and thus the expected amounts, we implement an augmented resampling procedure, which is described in Appendix D.

For all regression specifications, we use a subsample of competitive dealer bids with at least two steps, which encompasses nearly all dealer bids (see Appendix Figure A3). Our findings remain robust when restricting our analysis to bids with at least three (or even more) steps, as illustrated in Appendix Table A1. To compare bids in auctions with different supply levels, we normalize quantities to a range of 0 to 100, representing the share of total supply in percentage points. Prices are expressed in cents, i.e., C\$1/100, and all bills have a face value of C\$1,000.

Slopes of Bidding Functions. Before allowing for interdependencies across maturities and bid-shading, we analyze the observed slope in submitted bidding functions. If the prices of different maturities were independent from one another ($\delta_m = 0$), and bidders were price-takers, we would identify the own-demand coefficients, λ, from regressing the observed bids on demanded quantities and bid-function fixed effects.

The coefficients in Table 3a say that a dealer decreases their bid for 3M bills by 0.499 cents (or 0.0499 bps of the face value) when demanding one percentage point more of the

Table 3: Demand Coefficients

(a) Slopes of Bidding Functions

	3M Bill Auction		6M Bill Auction		12M Bill Auction	
Slope	-0.499^{***}	(0.0025)	-0.792^{***}	(0.0046)	-1.565^{***}	(0.0081)
N	58,542		42,282		50,410	

(b) With Bids as Independent Variables

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	-0.391^{***}	(0.0030)	λ_{6M}	-0.714^{***}	(0.0046)	λ_{12M}	-1.542^{***}	(0.0099)
$\delta_{3M,6M}$	$+0.023^{***}$	(0.0046)	$\delta_{6M,3M}$	$+0.023^{***}$	(0.0046)	$\delta_{12M,3M}$	$+0.024^{***}$	(0.0054)
$\delta_{3M,12M}$	$+0.024^{***}$	(0.0054)	$\delta_{6M,12M}$	$+0.048^{***}$	(0.0075)	$\delta_{12M,6M}$	$+0.048^{***}$	(0.0075)
N	58,542			42,282			50,410	

(c) With Values as Independent Variables

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	-0.518^{***}	(0.0043)	λ_{6M}	-1.009^{***}	(0.0072)	λ_{1Y}	-2.271^{***}	(0.0161)
$\delta_{3M,6M}$	-0.045^{***}	(0.0070)	$\delta_{6M,3M}$	-0.045^{***}	(0.0070)	$\delta_{12M,3M}$	-0.070^{***}	(0.0082)
$\delta_{3M,12M}$	-0.070^{***}	(0.0082)	$\delta_{6M,12M}$	-0.216^{***}	(0.0122)	$\delta_{12M,6M}$	-0.216^{***}	(0.0122)
N	58,542			42,282			50,410	

Table 3a regresses the competitive bids on the demanded quantities and bid-function fixed effects for each maturity. Table 3b reports the coefficients for equation (5), that is, with the observed competitive bids as independent variables rather than the estimated values. Table 3c reports the coefficients when using values, as in equation (11). In Table 3b and 3c we impose symmetry in the cross-demand coefficients, i.e., $\delta_{m,-m} = \delta_{-m,m}$. In all cases, the sample is restricted to functions by dealers with at least two steps; bids and values are in cents, i.e., C$1/100, and quantities in percentage of auction supply—when expressing bids in yields and quantities C$, the findings are qualitatively the same. The bills' face value is C$1,000. The first three columns show the estimates for the 3M bill auction, the next three for the 6M bill auction and the last three for the 12M bill auction. The point estimates are in the second, fifth and eight column. Standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

supply in a 3M auction; in the 6M auction the bid drops by 0.792 cents (or 0.0792 bps), and in the 12M auction by 1.565 cents (or 0.1565 bps). These price drops are not negligible when compared to the difference between the maximum and minimum bid in the average bidding function, which is 8.264 cents for 3M bills, 13.167 cents for 6M bills, and 26.135 cents for 12M bills.[21] However, in absolute terms, the bidding curve is rather flat, implying that dealer demand is highly price-elastic, and becomes steeper in maturity length. This is in line with Vayanos and Vila (2021) models, according to which yields for longer maturities are more affected by demand changes because of duration risk.

[21] A different way of interpreting these number is to recognize that bills are sold at a discount from par and the interest earned is the size of the discount. A decrease in the price by 0.499 cents for 3M bills is the same as an increase in the discount of 0.499 cents. Across maturities, this increase is roughly 6% of the average range in the bidding function.

Demand Systems with Price-Takers. As a next step, we relax the assumption that different maturities are independent, but still impose that dealers are price-takers. We estimate equation (5), that is, regress the observed bids in one auction on the corresponding demanded quantities in that bidding function and the expected winning quantities for the two other maturities plus bid-function fixed effects.

We report the estimates in Table 3b. The own-demand coefficients, λ, are similar to those of Table 3a. The cross-demand coefficients, δ, are positive and statistically significant, which indicates that dealers are willing to pay higher prices when they win more of the other maturities.[22] This implies, for example, that the auction for 3M bills is expected to clear at a higher price when the supply of 6M bills increases, contradicting the conventional notion in the literature.

Demand Systems with Strategic Bidders. Next, we allow for interdependencies and strategic bidding and re-estimate the regressions using the estimated values, which we derive based on Proposition 2. Results from regressions (11) are shown in Table 3c. Given that bid shading decreases with quantity, as depicted in Appendix Figure A5, the willingness-to-pay curve is steeper than the bidding curve. Accordingly, the own-demand coefficients, λ, are larger when accounting for bid shading.

All cross-demand coefficients, δ, turn negative, highlighting the importance of eliminating bid shading and using true values to accurately identify interdependencies, even when shading factors appear small. The results now indicate that the auction for 3M bills is expected to clear at a lower price when the supply of 6M bills increases, and similarly for all other maturities, aligning with the literature, unless bid shading changes dramatically in response to supply changes.

The price drops due to cross-maturity independencies are small in absolute terms but not negligible compared to the own-demand effects. For example, a dealer's willingness to pay for 6M bills decreases by $\lambda_{6M} = 1.009$ cents (or 0.1009 bps) when getting 1 percentage point more of the 6M supply. In addition, since bills are not independent from another, dealers' willingness to pay for 3M bills drops by $\delta_{3M,6M} = 0.045$ cents (or 0.0045 bps) and for 12M bills by $\delta_{12M,6M} = 0.216$ cents (or 0.0216 bps). The estimated δ coefficients imply

[22]This is also the case when regressing the observed bids on the amounts bidders actually win—which we observe in the raw data—rather than the amounts they expect to win when bidding.

Table 4: Complements or Substitutes?

	m=3M	m=6M	m=12M
$\frac{\partial q_{3M}}{\partial p_m}$	−19.44	+0.75	+0.53
$\frac{\partial q_{6M}}{\partial p_m}$	+0.75	−10.15	+0.94
$\frac{\partial q_{12M}}{\partial p_m}$	+0.53	+0.94	−4.51

Table 4 shows the own-and cross price coefficients of the inverted demand system that is implied by the dealer's willingness to pay of Table 3c. It tells us by how many percentage points of supply dealer demand for maturity $l \in \{3M, 6M, 12M\}$ changes in response to a price change of 1 bps of maturity $m \in \{3M, 6M, 12M\}$. For example, the second column ($m = 6M$) tells us that a price of 1 bps increase in 6M bills decreases dealer demand in the 6M auction by 10.15 percentage points of 6M-supply (own-price demand effect). Simultaneously, this price increase boosts dealer demand in the 3M and 12M auctions by about 0.75 percentage point of 3M-supply and 0.94 percentage point 12M-supply, respectively (cross-price demand effect).

that the dealer's value for the 6M bill decreases by approximately 1.82 cents when obtaining the average amount of the 3M (7.3% of supply) and 12M bills (6.9% of supply), compared to receiving nothing. In the 3M auction the corresponding price decrease is about 0.78 cents and in the 12M auctions about 1.96 cents.

Substitutes or Complements. To determine whether bills are substitutes or complements according to Definition 1, we compile all demand coefficients into the matrix $\mathbf{\Sigma}$ and then invert it. We express price changes in basis points (1/100th of a percent) rather than in whole percentage points because the latter represents a change that is too large for the context of our sample period.

The results, shown in Table 4, indicate that, on average, bills are weak complements. For example, a price increase of 10 cents—i.e., 1 bps of the face value—in 6M bills decreases dealer demand in the 6M auction by approximately 10 percentage points of 6M-supply, or C$ 212 million (own-price demand effect). Simultaneously, this price increase boosts dealer demand in the 3M by about 1 percentage point of of 3M-supply, or C$ 57 million, and by about 1 percentage point of 12M-supply, or C$ 21 million (cross-price demand effect). All cross-price effects are comparable in magnitude: approximately 1 for maturities three months apart (e.g., 3M and 6M, or 6M and 12M), and around 0.5 for 3M and 12M.

Implied Market Demand Elasticities. To convert the demand coefficients in Table 3 into market demand elasticities—quantifying the percentage change in market demand per 1% (or in our case 1/100%) price increase—we need to take a stance on the size of the underlying demand. The standard approach in the literature is to approximate demand

using the total amount outstanding of the asset (or asset class).[23] To follow this method, we collect data on the daily amounts outstanding of bills and bonds from 2003 to 2015 (excluding 2002 due to data unavailability). We group securities by remaining maturity into clusters: 3M ± 2 weeks, 6M ± 2 weeks, and 12M ± 2 weeks. Assuming that securities within each cluster are perfect substitutes, and that dealers act as intermediaries considering the entire market's demand, we can calculate the percentage change in demand for each cluster in response to a 1 bps increase of the auction-clearing price (details are in Appendix E). If one is willing to assume that the demand slope estimated from price changes on the order of fractions of a basis point can be extrapolated to price changes corresponding to a 1% change, our demand elasticities could be scaled by a factor of 100 to represent the percentage change in demand for a 1% change in price.

To understand the magnitudes, we first compute demand elasticities under the assumption that dealers act as price-takers and bills of different maturities are independent (all δs are zero). In this case, demand coefficients can be inferred directly from submitted bids (Table 3a). We find that elasticities are very high in absolute values, as shown in Table 5a, ranging from 12 for 12M bills to 49 for 3M bills. Two factors contribute to these high values. First, our clustering approach is conservative. Expanding clusters, for example, including securities with the same maturity ± 1 month instead of ± 2 weeks, reduces elasticities, simply because the total demand within each cluster—the denominator—increases. Second, a 1 bps price increase is relatively large given the low-interest-rate environment during our sample period. For instance, a 1 bps price increase for 3M bills represents a 2% reduction in the yield to maturity.

Table 5b incorporates strategic dealer behavior and interdependencies among bills (using coefficients from Table 3c). Since all δ coefficients are positive, the own-price elasticities become smaller in absolute terms. In comparison, cross-demand elasticities are moderate, ranging from 2% (for the 12M demand response to a 3M price change) to 23% (for the 6M demand response to a 12M price change) of the corresponding own-price elasticities.

[23] In our setting, an alternative approach is to use the average quantity a dealer demands at the market-clearing price in the auction as the underlying demand. Computing the implied demand elasticities is straightforward in light of Table 4, and with the information that the average dealer demand at market-clearing is 7.3%, 6.7%, and 6.9% of supply, and the clearing prices are C$994.898, C$990.570, and C$979.8145 in the 3M, 6M, and 12M auction, respectively. However, since our objective is to align our results with the demand elasticities reported in the existing literature, we do not adopt this approach.

Table 5: Implied Market Demand Elasticities

(a) Assuming Price-Takers and Bill Independence (based on Table 3a)

1 bps price increase of	3M bills	6M bills	12M bills
Percentage change in 3M demand	−49.35	0	0
Percentage change in 6M demand	0	−26.53	0
Percentage change in 12M demand	0	0	−11.77

(b) With Strategic Dealers and Bill Interdependence (based on Table 3c)

1 bps price increase of	3M bills	6M bills	12M bills
Percentage change in 3M demand	−47.87	+1.84	+1.28
Percentage change in 6M demand	+1.58	−22.23	+1.95
Percentage change in 12M demand	+0.99	+1.75	−8.66

Table 5a converts the demand coefficients of Table 3a into market demand elasticities, which gives the percentage change in market demand of maturity $l \in \{3M, 6M, 12M\}$ (the row) in response to 1/100% price increase of the average clearing prices of maturity $m \in \{3M, 6M, 12M\}$ (the column). The average clearing prices are C$994.898 for 3M bills, C$990.570 for 6M bills, and C$979.8145 for 12M bills. We use the daily amount outstanding of bills and bonds between 2003 and 2015 to approximate total market demand: C$25.6 bn for 3M ± 2 weeks securities, C$11.0 bn for 6M ± 2 weeks securities, and C$12.4 bn for 12M ± 2 weeks securities. Table 5b uses the demand coefficients of Table 3c, which correctly account for the fact that dealers are strategic and that different maturities are not independent from each other.

Comparison to the Existing Literature. Our analysis underscores the importance of considering cross-asset interdependencies and clarifies the definition of substitutes for assets.

Early studies by Park and Reinganum (1986), Simon (1991, 1994), and Duffee (1996), as well as more recent work by Nagel (2016) and Krishnamurthy and Li (2023), examined the relationships between different T-bills or government debt more broadly. Traditionally, two assets are considered "substitutes" when the price (or spread) of one asset decreases as the supply or traded quantity of the other asset increases. Although this is true in our data, we conclude that bills are weak complements. This is because we use a different notion of substitutability—Definition 1—which is inspired by the standard definition of "gross substitutes" with Marshallian demand (Mas-Colell et al. (1995)).

The key take-away is that when assets are interdependent it matters whether we analyze how prices change relative to quantities or how quantities change relative to prices. The two-asset example following Definition 1 illustrates that we draw different conclusions about the assets' substitutability depending on the mapping we consider. Given the growing interest in estimating demand as a function of prices, following Koijen and Yogo (2019), we hope our notion of substitutability proves useful.

The focus of the growing demand-based asset pricing literature has been on the own-demand effect. For example, Koijen and Yogo (2020) estimate the (annual aggregate) own-demand elasticity of investors for short-term debt following the approach developed in Koijen and Yogo (2019). This approach assumes that the prices of one asset are independent of the characteristics of other assets (Assumption 2 in Koijen and Yogo (2019)) and that market participants are price-takers. Under these assumptions, our most comparable estimate is based on the slope of the (inverted) bidding curves, which we can directly infer from the data. If we calculate the change in demand in a typical auction (averaged across maturities) as a percentage of the total outstanding government debt with a maturity of less than one year, in response to a 1% price increase—disregarding concerns about extrapolation—our bidding data suggest a demand elasticity of approximately 275 in absolute terms (see Appendix E for details). In comparison, Koijen and Yogo (2020) report an elasticity of roughly 25.

Our demand elasticity estimates differ from those in Koijen and Yogo (2020) for several reasons. First, we focus on demand for Canadian T-bills, while Koijen and Yogo (2020) estimate the elasticity for all short-term debt, which includes more than just government debt, of various countries. These securities are denominated in various currencies and subject to differing inflation and default risks. Therefore, the finding that aggregate elasticity across such a diverse pool is lower than for a single class of securities aligns with recent work by Chaudhary et al. (2023).[24] Second, Koijen and Yogo (2020) focus on the demand of investors who hold these assets, whereas our demand coefficients reflect the preferences of intermediaries who actively trade them. Our findings are consistent with intermediaries being more sensitive to price changes than households. Third, Koijen and Yogo (2020) derive demand elasticities from annual holdings data, while our demand coefficients are based on intermediaries' willingness to pay at a specific point in time, namely when dealers place bids. Although there is no established framework for aggregating demand elasticities over time, we anticipate that the temporal scope—short-term versus long-term—affects elasticity measures. Finally, during our sample period, interest rates were lower than those analyzed in Koijen and Yogo (2020), which may also influence the demand elasticity estimates.

[24] Chaudhary et al. (2023) demonstrate that security-level multipliers—which are inverse demand elasticities, like the slope estimates of Table 3a—are much smaller than multipliers of security portfolios, implying that demand elasticities are smaller when aggregating across securities than at the security-level.

Extensions and Future Research. We offer two extensions to our baseline estimation to motivate future research. First, we allow demand coefficients to vary flexibly over time. Since we treat auction days separately, we can estimate demand coefficients over time. Figure 2a shows the time series of own-demand λ-coefficients, while Figure 2b displays the cross-maturity δ-coefficients, focusing only on those that are statistically significant at the 10% level. During the financial crisis of 2007-2009, we observe that the own-demand coefficients are larger in absolute value. This finding aligns with Albuquerque et al. (2024), who show that bidding functions in Portuguese Treasury auctions become steeper when market volatility is high. While the cross-maturity coefficients do not exhibit systematic changes during the crisis, they show substantial variability across auction days. Time-variation is so significant that bills may be complements or substitutes according to Definition 1, depending on the day.

The second extension adds latent dealer heterogeneity. To do so, we extend our model in Appendix C. Each dealer now has a business type that is unobservable to the econometrician. We uncover that, in the average auction, bills are complementary for the average dealer because they are complementary for the larger dealers. For smaller dealers, bills are essentially independent, perhaps because these dealers cater to specific investors with tastes for specific maturities. This finding sheds light on the nuanced perspectives of different market participants, and highlights dealer heterogeneity, in line with a growing literature on heterogeneous intermediary asset pricing (e.g., Kargar (2021); Brunnermeier et al. (2021)).

To better understand the drivers behind substitution patterns across dealers or over time, we propose a micro-foundation for dealer demand in the primary market, detailed in Appendix F. This framework suggests that cross-maturity elasticities are influenced by several factors: the ease with which investors substitute between maturities in the secondary market, the cost of turning down a client in the secondary market, and the likelihood of investors arriving with preferences for different maturities.

We hope that our estimation findings, together with the micro-foundation, motivates future research using data from both primary and secondary markets to test these hypotheses and explore the factors driving demand heterogeneity over time or in the cross-section.[25]

[25]In Appendix G, we take a first step by analyzing a unique dataset on secondary market cash and repo trades. A key limitation is that the transaction-level data available (post-2016) does not overlap with the auction sample. As a result, we cannot definitively attribute primary market behavior to what we observe

Figure 2: Time-varying demand estimates

(a) Own-demand λ coefficients (b) Cross-maturity δ coefficients

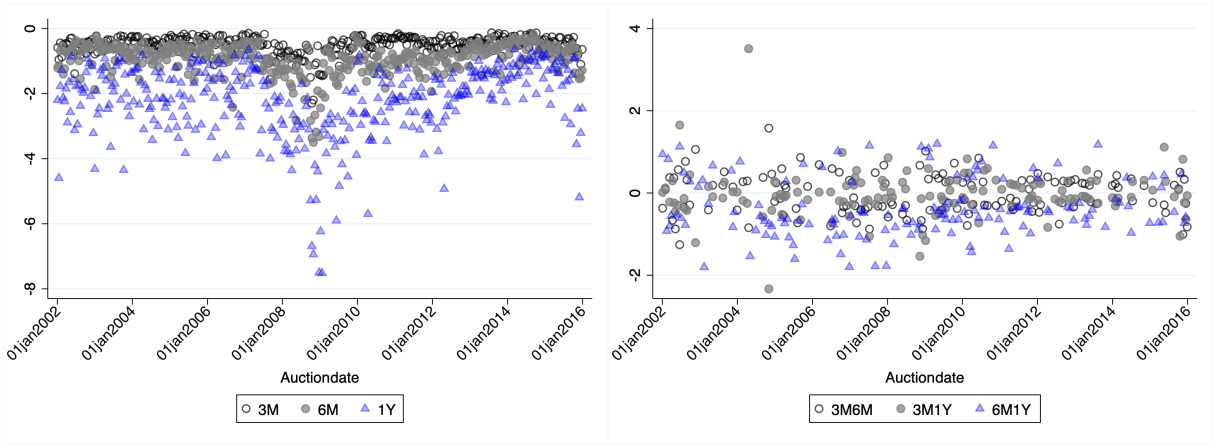


Figure 2a shows the own-demand coefficients when estimated for each auction date separately, while Figure 2b displays the cross-maturity coefficients that are statistically significant at 10%. The data sample is restricted to functions by dealers with at least two steps; bids and values are in cents, i.e., C$1/100, and quantities in percentage of auction supply—when expressing bids in yields and quantities C$, the findings are qualitatively the same. The bills' face value is C$1000.

3.4 Policy Take Away for Treasury Auctions

Knowing demand elasticities is a crucial input to proposing better market rules. In our application, debt managers aim to minimize the cost of debt funding, which is equivalent to maximizing auction revenues. In other settings, the objective might be different. For example, with data on limit order books, where packages of limit orders constitute demand step functions, we would want to maximize total gains from trade; with data on auctions of carbon allowances it might be more natural to reduce total carbon emissions.

A starting point to achieve the government's objective is to split the amount of debt to be issued over a short window (in our case the same day) more strategically across different auctions to achieve higher revenue. To achieve this, the government could behave like a monopolist who price discriminates, and exploit differences in demand elasticities. When the auction format is uniform price, it is generally revenue-increasing to issue more of the security whose demand responds less strongly to price changes (here the shorter bills) and less of the other securities. This is because everyone pays the market prices, and the market

in the secondary market, nor link both markets in our structural model. However, the data allows us to provide evidence that different types of T-bills serve distinct purposes and attract investors with specific maturity preferences (preferred habitat). This diversity in usage and investor preference leads to imperfect substitution across maturities.

price of the longer bills increases faster than the market price of the shorter bill decreases. When the auction format is discriminatory price, shifting supply towards the short bill may decrease total revenue. This is because the revenue of one auction is determined by the area underneath the entire aggregate demand curve, and not only the market clearing price. Given that bidders endogenously respond to changes in supply, and demand schedules are step functions, it is an empirical question as to whether total revenue increases or decreases when reshuffling supply.

Due to the low price sensitivity in the demand for Canadian bills during our sample period (2002–2015), cost savings from altering the current maturity split are small. This is evident from the nearly flat bidding functions and small demand coefficients (see Table 3). With flat curves, large changes in supply would be required to move prices in a way to result in meaningful cost savings for taxpayers. In markets where demand is more responsive to price changes the savings could be sizable. In those cases, our framework can be used as a foundation to quantify the effects, and help split the total supply of T-bills more strategically across maturities for any given auction day. Future research could further expand on this by incorporating inter-temporal trade-offs, such as debt rollover, into the allocation strategy.

4 Conclusion

We exploit the institutional features of parallel multi-unit auctions to develop a new methodology to identify demand systems that can accommodate any degree of substitution or complementarity between assets. We compare our method to existing approaches to demand estimation and demonstrate its implementation using data from Canadian Treasury auctions. Using the traditional notion of substitutability—where two goods are substitutes if the demand for one decreases when the price of the other increases—we find that different T-bills are weak complements for primary dealers. This is true even though the price of any maturity decreases when the supply of other maturities increases. Lastly, we emphasize the importance of understanding demand elasticities to effectively change market rules.

References

Albuquerque, R., Cardoso-Costa, J., and Faias, J. A. (2024). Price elasticity of demand and risk-bearing capacity in sovereign bond auctions. *Review of Financial Studies*, 37(10):3149–3187.

Allen, J., Hortaçsu, A., Richert, E., and Wittwer, M. (2023). Entry and exit in Treasury auctions. Working Paper.

Allen, J. and Wittwer, M. (2023). Centralizing over-the-counter markets? *Journal of Political Economy*, 131(12):3267–3540.

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

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

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

Barber, B. M., Huang, X., Odean, T., and Schwarz, C. (2022). Attention-induced trading and returns: Evidence from Robinhood users. *Journal of Finance*, 77(6):1341–3190.

Barbosa, K., Silva, D. G. D., Yang, L., and Yoshimoto, H. (2022). Auction mechanisms and Treasury revenue: Evidence from the Chinese experiment. *American Economic Journal: Microeconomics*, 14:394–419.

Becker, J., Hendricks, K., Houde, J.-F., and Raisingh, D. (2023). Asymmetric information and the supply-chain of mortgages: The case of Ginnie Mae loans. Working paper.

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

Bikhchandani, S. and Huang, C.-F. (1989). Auctions with resale markets: An exploratory model of Treasury bill markets. *The Review of Financial Studies*, 2(3):311–339.

Blum, J. R., Kiefer, J., and Rosenblatt, M. (1961). Distribution free tests of independence based on the sample distribution function. *The Annals of Mathematical Statistics*, 32(2):485–498.

Bogousslavsky, V. and Muravyev, D. (2024). An anatomy of retail option trading. Working paper.

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

Brenner, M., Galai, D., and Sade, O. (2009). Sovereign debt auctions: Uniform or discriminatory? *Journal of Monetary Economics*, 56(2):267–274.

Brogaard, J. and Garriott, C. (2019). High-frequency trading competition. *Journal of Financial and Quantitative Analysis*, 54(4):1469–1497.

Brogaard, J., Hendershott, T., and Riordad, R. (2019). Price discovery without trading: Evidence from limit orders. *The Journal of Finance*, 74(4):1621–1658.

Brunnermeier, M., Farhi, E., Koijen, R. S. J., Krishnamurthy, A., and Ludvigson, S. C. (2021). Review article: Perspectives on the future of asset pricing. *The Review of Financial Studies*, 34:2126–2160.

Budish, E., Crampton, P., and Shin, J. (2015). The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics*, 130(4):1547–1621.

Budish, E., Cramton, P., Kyle, A. S., Lee, J., and Malec, D. (2022). Flow trading. NBER working paper 31098.

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

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

Chang, Y.-C., Hong, H., and Liskovich, I. (2014). Regression discontinuity and the price effect of stock market indexing. *Review of Economic Studies*, 28(1):212–246.

Chaudhary, M., Fu, Z., and Li, J. (2023). Corporate bond multipliers: Substitutes matter. Working paper.

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

Clark, R., Houde, J.-F., and Kastl, J. (2021). The industrial organization of financial markets. In: The Handbook of Industrial Organization, K. Ho, A. Hortaçsu, and A. Lizzeri (eds.), North-Holland Press, Volume 5, Chapter 15.

Conrad, J. and Wahal, S. (2020). The term structure of liquidity provision. *Journal of FInancial Economics*, 136:239–259.

Conrad, J., Wahal, S., and Xiang, J. (2015). High-frequency quoting, trading, and the efficiency of prices. *Journal of Financial Economics*, 116:271–291.

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

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

Darmouni, O., Siani, K., and Xiao, K. (2023). Nonbank fragility in credit markets: Evidence from a two-layer asset demand system. Working paper.

Dix, R. and Wittwer, M. (2024). Estimating asset demand with trade data. mimeo.

Duffee, G. R. (1996). Idiosyncratic variation of Treasury bill yields. *The Journal of Finance*, 51(2):527–551.

Fleming, M. J., Mizrach, B., and Nguyen, G. (2017). The microstructure of a U.S. Treasury ECN: The BrokerTec Platform. *Federal Reserve Bank of New York, Staff Report No. 381*, pages 1–52.

Foucault, T., Moinas, S., and Theissen, E. (2007). Does anonymity matter in electronic limit order markets? *The Review of Financial Studies*, 20(5):1707–1747.

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

Frazzini, A., Israel, R., and Moskowitz, T. J. (2018). Trading costs. Working paper.

Fuchs, W., Fukuda, S., and Neuhann, D. (2024). Demand-system asset pricing: Theoretical foundations. Working paper.

Gabaix, X. and Koijen, R. S. J. (2024). Granular instrumental variables. *Journal of Political Economy*, 132:2274–2303.

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

Garriott, C., van Kervel, V., and Zoican, M. (2024). The queuing friction in limit order book markets. Working paper.

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

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

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

Hasbrouck, J. (1991). Measuring the information content of stock trades. *The Journal of Finance*, 46(1):179–207.

Hausman, J., Leonard, G., and Zona, J. D. (1994). Competitive analysis with differentiated products. *Annales d'Economie et de Statistique*, 34:159–18.

Hollifield, B., Miller, R. A., Sandås, P., and Slive, J. (2006). Estimating the gains from trade in limit-order markets. *Journal of Finance*, 61(6):2753–2804.

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

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

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

Jansen, K. A. E., Li, W., and Schmid, L. (2024). Granular Treasury demand with arbitrageuers. Working paper.

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

Kargar, M. (2021). Heterogeneous intermediary asset pricing. *Journal of Financial Economics*, 141(2):505–532.

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

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

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

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

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

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

Koijen, R. S. J. and Yogo, M. (2020). Exchange rates and asset prices in a global demand system. NBER working paper 27342.

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

Korajczyk, R. A. and Murphy, D. (2019). High-frequency market making to large institutional trades. *The Review of Financial Studies*, 32(3):1034–1067.

Krishnamurthy, A. and Li, W. (2023). The demand for money, near-money, and Treasury bonds. *Review of Financial Studies*, 36:2091–2130.

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

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

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

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

Kyle, A. S. (1989). Informed speculation with imperfect competition. *Review of Economic Studies*, 56(3):317–355.

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

Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press.

Mudholkar, G. S. and Wilding, G. E. (2003). On the conventional wisdom regarding two consistent tests of bivariate independence. *Journal of the Royal Statistical Society, Ser. D (The Statistician)*, 52(1):41–57.

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

O'Hara, M. (2015). High frequency market microstructure. *Journal of Financial Economics*, 116(2):257–270.

Park, S. Y. and Reinganum, M. R. (1986). The puzzling price behavior of Treasury bills that mature at the turn of calendar months. *Journal of financial Economics*, 16(2):267–283.

Pavlova, A. and Sikorskaya, T. (2023). Benchmarking intensity. *The Review of Financial Studies*, 36(3):859–903.

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

Romano, J., Shaikh, A., and Wolf, M. (2010). Multiple testing. *The New Palgrave Dictionary of Economics*.

Rostek, M. and Weretka, M. (2015). Dynamic thin markets. *Review of Financial Studies*, 28(10):536–557.

Rostek, M. and Yoon, J. (2021a). Equilibrium theory of financial markets: Recent developments. prepared for The Journal of Economic Literature.

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

Ryan, N. (2022). Holding up green energy. NBER working paper 29154.

Schmickler, S. (2020). Identifying the price impact of fire sales using high-frequency surprise mutual fund flows. Working paper.

Shleifer, A. (1986). Do demand curves for stocks slope down? *The Journal of Finance*, 41(3):579–590.

Simon, D. P. (1991). Segmentation in the Treasury bill market: Evidence from cash management bills. *Journal of Financial and Quantitative Analysis*, 26(1):97–108.

Simon, D. P. (1994). Further evidence on segmentation in the Treasury bill market. *Journal of Banking & Finance*, 18(1):139–151.

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

Song, Z. and Zhu, H. (2018). Quantitative easing auctions of Treasury bonds. *Journal of Financial Economics*, 128:103–124.

van Kervel, V. and Menkveld, A. J. (2015). High-frequency trading around large institutional orders. *Journal of Finance*, 74(3):1091–1137.

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

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

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

Wittwer, M. and Allen, J. (2023). Market power and capital constraints. Working paper.

ONLINE APPENDIX

Estimating Demand Systems with Bidding Data

by Jason Allen, Jakub Kastl, and Milena Wittwer

Section **A** explains how to estimate non-linear demand.

Section **B** provides empirical evidence in favor of conditional independent private values.

Section **C** presents a model extension that allows bidders to carry latent types.

Section **D** describes the resampling procedure we implement for our empirical application.

Section **E** explains how we compute market demand elasticities.

Section **F** introduces a micro-foundation for dealer demand in Treasury auctions.

Section **G** provides empirical evidence to understand what drives substitution patterns.

Section **H** presents proofs that we have not discussed in the body of the paper.

A Identifying Non-Linear Demand

In line with the traditional assumption in the demand estimation literature, we assume in the main text that demand is linear in prices. However, with rich enough auction data, we do not need to impose that assumption. In fact, it is possible to non-parametrically identify the shape of all willingness to pay curves. This can be done different ways, depending on details of the auction setting.

We illustrate one approach for the case when there are three assets, $M = 3$, and bidders are strategic. We assume that the distribution of a bidders' winning quantities has finite support. However, the insights from this example generalize to settings with more maturities and continuous winning quantities.

The main data requirement is that we need to observe data from L auction days that differ in their supply distribution (or some other observable characteristic that shifts bidder expectations), but are otherwise ex ante identical. This means that true demand for a given signal, $v_m(q_{m,k}, \boldsymbol{q_{-m}}, \boldsymbol{s_i})$, is the same across auctions, but the expected marginal utility from

winning at auction—the analogue of expression (8)—

$$\tilde{v}_m(q_{m,k}, \boldsymbol{s_i}|\theta_{i,l}) = \mathbb{E}[v_m(q_{m,k}, \boldsymbol{q^*_{-m,i}}, \boldsymbol{s_i})|\theta_{i,l}] \tag{12}$$

differ since information, $\theta_{i,l}$, differs across auctions l. The idea is to leverage the fact that the expected marginal utility varies across auctions only because the expectation changes but not because the true underlying willingness to pay changes.

To see how this works, we first need to estimate the joint distribution of market clearing prices and the distribution of winning quantities, as explained in Section 2.3, or in Appendix D for the case of Canadian Treasury auctions. We further back out how much bidders expect to marginally benefit from winning each submitted step in all auctions, by relying on the equilibrium conditions, as specified in Proposition 2. This gives us $\tilde{v}_m(q_{m,k}, \boldsymbol{s_i}|\theta_{i,l})$ for all maturities, bidders, steps and auction days.

Given that $\tilde{v}_m(q_{m,k}, \boldsymbol{s_i}|\theta_{i,l})$ is strictly increasing in $\boldsymbol{s_i}$ for any l by Lemma 1 in Appendix H, any quantile ψ of the distribution of expected values at some step k corresponds to the quantile of the signal distribution: $Quantile_\psi(\tilde{v}_m(q_{m,k}, \boldsymbol{s_i}|\theta_{i,l})) = Quantile_\psi(\boldsymbol{s_i})$.

Thus, using variation across auctions l, by holding the quantile of the expected values at the same quantity point $q_{m,k}$ fixed, we can identify $v_m(\cdot, \cdot, \boldsymbol{s_i})$ at all quantities. For this note that:

$$\tilde{v}_1(q_{1,k}, \boldsymbol{s_i}|\theta_{i,l}) = \sum_{h=1}^{H}\sum_{r=1}^{R} p_{1,i,l}^{k,h,r} v_1\left(q_{1,k}, (q_{2,h}, q_{3,r}), \boldsymbol{s_i}\right) \text{ for } l = 1, ..., L, \tag{13}$$

$$\text{where } p_{1,i,l}^{k,h,r} = \Pr(q_{2,i}^* = q_{2,h} \cap q_{3,i}^* = q_{3,r}|b_{1,i,k} \geq P_1^c > b_{1,i,k+1}, \theta_{i,l}) \tag{14}$$

is the probability that $q_{2,h}$ and $q_{3,r}$ will be the quantities awarded in the auction for asset 2 and 3, respectively, conditional on winning the k'th step in the auction for asset 1, and H and R are the numbers of points in the support of winning quantities for asset 2 and 3. This is a system of $L \times K$ linear equations in $H \times R \times K$ unknowns. If we have $H \times R$ bids at the same $q_{1,k}$ for each step k, and the matrix of coefficients of this system of linear equations is invertible (i.e., if it is full rank), then there is a unique solution.

Once we know the shape of $v_m(\cdot, \cdot, s_i)$, it is straightforward to back out the degree of interdependency. For instance, the dependency between asset 1 and 2 when winning (q_1, q_2, q_3)

is $\delta_{1,2}(\boldsymbol{q}, \boldsymbol{s_i}) = \frac{\partial v_1(q_1, \boldsymbol{q_{-1}}, \boldsymbol{s_i})}{\partial q_2}$. With finite data, we can point-identify finite points on $v_m(\cdot, \cdot, \boldsymbol{s_i})$ and approximate $\delta_{1,2}(\boldsymbol{q}, \boldsymbol{s_i})$ by

$$\delta_{1,2}^{k,h,r}(\boldsymbol{s_i}) = \frac{v_1(q_{1,k}, (q_{2,h+1}\ q_{3,r}), \boldsymbol{s_i}) - v_1(q_{1,k}, (q_{2,h}\ q_{3,r}), \boldsymbol{s_i})}{q_{2,h+1} - q_{2,h}} \ \forall k, h, r. \tag{15}$$

In our empirical application we would like to avoid pooling data across auctions because the distributions of signals may easily shift over time, especially after important macroeconomic events, or monetary policy announcements. Therefore, we identify linear demand curves. However, in other settings, it might be beneficial to choose a non-parametric approach.

B Testing Independence of Private Information

Similar to Hortaçsu and Kastl (2012), we conduct a battery of tests for (conditional) independence of private information. We test whether dealers' bids are conditionally independent by first randomly splitting the (quantity-weighted) bids within each auction into two halves (dropping one observation if the number of bids in an auction is odd). We then construct 4 test statistics based on the two vectors of bids: the Pearson coefficient, Spearman rank correlation, Kendall's tau and, finally, the Blum et al. (1961) nonparametric test for independence.[26] This battery of tests is conducted separately for each of 1,098 auctions in our sample, hence we test for *conditional* independence. We considered 100 random splits of bids.

Using all bids in our sample, we fail to reject the null hypothesis of independence at 5% significance level in all 4 tests in over than 95% of auctions. In between 49 and 56 auctions (out of 1,098) on average (across the 100 splits) we reject the null hypothesis.

To account for potential correlation in bids that is due to dealers updating their bids within an auction, we also conduct these tests based on using only first submitted bid by a dealer and also using only the last submitted bid. For first-bids only, the test rejects independence in between 34 auctions (based on Kendall-tau test) and 64 (based on Pearson) out of 1,098 auctions across 100 permutations. When only last bids are considered, the Kendall-tau test rejects for 33 auctions, Spearman for 46, Pearson for 58 and BKR for 53.

[26] Mudholkar and Wilding (2003) critical values are used for the Blum-Kiefer-Rosenblatt test statistic.

Finally, since customer bids might also induce correlation in bids—dealers can and do react after routing such a bid—we also ran these tests after removing customer bids. Spearman and Kendall-tau test reject only for 18 auctions (out of 1,098) on average, while the Pearson test rejects for 56.

It might be expected that we do not always fail to reject the null hypothesis, since occasionally (and certainly quite infrequently) there could be information that arrives during the time that bids are being submitted, in which case bids before and after such information might appear correlated. In addition, multiple testing leads to over-rejection of the null hypothesis (Romano et al. (2010)). Overall, we conclude that conditional independence seems like a reasonable assumption to impose in our context.

C Heterogeneous Dealers

We consider an extension of the auction game introduced in Section 3.2 in which dealers have a latent business type, χ. Each dealer is either a market-maker ($\chi = mm$) or a niche-customer ($\chi = nc$).[27] The econometrician does not know which dealer is of which type, but the bidders do. In this case, private signals are drawn independently from three distributions $F^{d,mm}$, $F^{d,nc}$ and F^{c}, and the marginal willingness to pay may be bidder specific: $v_{m,i}(q_m, \boldsymbol{q_{-m}}, \boldsymbol{s^g_{m,i}}) = s^g_{m,i} + \lambda_{m,i} q_m + \boldsymbol{\delta_{m,i}} \cdot \boldsymbol{q_{-m}}$. All other assumptions remain unchanged.

To estimate demand coefficients, we extend the resampling procedure as in Cassola et al. (2013) to account for asymmetric latent types. We proceed in three steps: (i) partition dealers into the two groups; (ii) estimate a model where resampling is conditional on that assignment; (iii) use the estimated dealer-specific demand coefficients to classify dealers into types—we assign a dealer to the market-maker group if at most two of the δ coefficients in the demand system are positive. We repeat until (iii) yields the same assignment as we started with in (i) for all but one dealer.[28] Finally, we estimate regression (11) for each dealer group separately, identifying the group-specific average δ^{χ} and λ^{χ} parameters.

We find that there are two dealer groups with different preferences (see Appendix Table

[27] In theory one could allow for more than two types. In the estimation, this is feasible only if there are sufficiently many bidders that participate in an auction, which is not the case in our setting.

[28] We allow one dealer to switch groups, as it is theoretically possible that dealers do not fall definitively into one of the two groups; there is no formal argument that our procedure converges for all dealers.

A3). For the 11 dealers in group one, bills are (in most cases) slightly more interdependent than for the average dealer in our benchmark model. For the 4 dealers in the second group, bills are mostly independent. Dealers in the first group win on average larger amounts in the auctions than dealers in the second group. This is interesting since we do not use this information to classify dealers into groups.

D Resampling Procedure for Canadian T-Bill Auctions

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

Three complications arise in this setting. First, bids may not be submitted at the exact same time given electronic or human delays. We define bids to be "simultaneous" if they are the closest bids placed by a bidder within a 200-second window or if they are the final bids made before the auction deadline. The choice of a 200-second upper bound seems reasonable based on the observed time differences between bids across different maturities, where bidders did not update their bids during the auctions (see Appendix Figure A4). Second, a customer might bid via different dealers for different maturities, and third, two bids for the same maturity but by different customers might go through the same dealer. Neither of these cases happens often. Therefore, we assume that the information set of dealers who observe the same customer is independent across maturities, conditional on his own signal. In addition, we restrict the number of possible observed customer bids to two given that most customers only submit one bid and that there are many more dealers than customers in a typical auction.

With this, our procedure goes as follows: draw N_c customer bids from the empirical distribution of customer bids at date t. If a customer did not participate in one auction, replace his bid by 0. For each customer, find the dealer(s) who observed this customer's bid(s). If the customer submitted only one bid, we take the dealer who observed it. If the customer submitted more than one bid, draw uniformly over dealer-bids having observed this customer. Finally, if the total number of dealers drawn is at this point lower than the total number of potential dealers, draw the remaining bids from the pool of uninformed dealers,

i.e., those who do not observe a customer bid in any of the three auctions. Note that—while theory allows for many updates—we restrict the number of possible observed customer bids to two in order to simplify our resampling algorithm. This includes most cases as most bidders only update once or twice.

To correct for outliers that occasionally occur due to small values of the denominator in the estimated (marginal) hazard rate of the market clearing price, we trim our estimated values. Specifically, we restrict each to be lower than the bidder's maximal bid plus a markup of about 5 bps (C$0.5 for 12M, C$0.25 for 6M, C$0.125 for 3M). This approach is conservative in light of the distribution of how bidders shade the untrimmed estimated values per step (see Appendix Figures A5). The less we trim, the larger, in absolute value, are all demand coefficients (see Appendix Table A2).

E Computing Market Demand Elasticities

To convert our demand coefficients into market demand elasticities, which tell us by how many percent T-bill holdings change in response to a bps price change, we collect data on the daily amount outstanding of bills from 2003 to 2015 (excluding 2002 due to unavailability). We group securities based on their remaining maturity into clusters: 3 months +/1 2 weeks (3M), 6M ± 2 weeks (6M), and 12M ±2 weeks (12M), and refer to these clusters as 3M, 6M, and 12M bills, respectively. Given these clusters, the average amount outstanding for 3M is C$25.6 bn, for 6M it is C$11.0 bn; for 12M it is C$12.4 bn.[29]

This conversion, while helpful for comparison with the existing literature, relies on several assumptions that influence the magnitude of the elasticities. First, we assume that bills within a cluster—selected arbitrarily—are perfect substitutes. Second, we assume that dealers act as intermediaries representing the aggregate market demand, disregarding any balance sheet constraints they may face.

[29] Alternatively, if we pool more securities into one cluster, the total amount outstanding increases. For example, we could pool all Canadian government bills and bonds with a maturity of less than 3 months (C$88.4bn amount outstanding), securities with a maturity between 3M and 6M (C$48.6bn amount outstanding), and securities with 6M to 12M to maturity (C$67.9 bn amount outstanding). This would decrease all demand elasticities, even though we would be using the same underlying demand coefficients of Table 3.

Calculations for Table 5a. Assuming that dealers are price-takers and bills are independent, the demand coefficients of Table 3a tell us by how much demand increases when the price increases by 1 cent. Since the average 3M auction clears at C$994.898, a C$10 increase raises the price by 1.005%. Therefore, a 1 bps price increase leads to a reduction in demand of (10/0.499)/1.005 percentage points of supply per dealer, i.e., C$1.149 bn given average supply is C$5.76 bn. Since there are 11 dealers in a typical auction, total demand decreases by 11*C$1.149 bn, which is 49.35% of total 3M bill demand (C$25.6 bn). In the 6M auction, which clears on average at a price of C$990.57, the analogue demand decrease is 11*C$0.265 bn, which implies a reduction in total demand (which is C$11.0 bn) of 26.50%. In the 12M auction, which clears on average at a price of C$979.8145, the demand reduction is 11*C$0.138 bn, which is a decline of 12.24% in demand (which is C$12.4 bn).

Calculations for Table 5b. Table 4 (which uses the coefficients of Table 3c) tells us by how much demand in each auction changes in response to a price change of 1 bps. Further note that a price increase of 1 percentage point is an increase of 1.005% for 3M bills, 1.009% for 6M bills, and 1.021% for 12M bills given average auction prices. A 1 bps price increase for 3M bills leads to a demand change of (19.44)/(100*1.005)*C$5.76 bn = C$1.114 bn per dealer. Given that in a typical auction 11 dealers participate, total demand decreases by 11*C$1.114 bn, which is 47.87% of the daily amount outstanding of 3M bills (C$25.6 bn). Total demand for 6M bills increases by 11*(0.75)/(100*1.005)*C$2.12 bn = C$0.174 bn, which is a 1.58% change in 6M bill holdings (C$11.0 bn); for 12M bills total demand increases by 11*(0.53)/(100*1.005)*C$2.12 bn, which is a 0.99% change in 12M bill holdings (C$12.4 bn).

Calculation for Comparison with Koijen and Yogo (2019). The demand coefficients of of Table 3a—where bids are expressed in cents—imply that demand per dealer decreases by C$1000/100*(1/0.499*5.76 + 1/0.792*2.12+ 1/1.565*2.12)/3 bn ≈ C$51.91 bn in an average auction (averaging across maturities) when the price increases by 1 pp of the bills' face value, or, equivalently, when the average auction clearing price increases by 1.012%. With 11 dealers, who typically participate in the auction, total demand decreases by 11*C$51.91/1.012 bn per 1% price increase. This represents 275% of the total daily amount outstanding of bills and bonds with a maturity less than one year (C$205 bn).

F Micro-Foundation of Dealer Demand

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

Motivated by the evidence presented in Appendix Figure A8, our model features market segmentation in the spirit of Vayanos and Vila (2021). Dealers buy securities with different maturities in the primary market to sell them to their clients who are investors that have preferences for specific maturities (preferred habitat) in the secondary market. We restrict the number of maturities to $M = 2$, and drop the superscript g and the subscripts i, τ for simplicity. Generalizing the model to more than two maturities is straightforward but mathematically cumbersome and brings little additional insight.

Each dealer has a type $\boldsymbol{s}$, which decomposes into $\boldsymbol{\nu}$ (known by all dealers) and $\boldsymbol{t}$ (private information): $\boldsymbol{s} = (\boldsymbol{t}, \boldsymbol{\nu})$ with $\boldsymbol{t} = (t_1, t_2)$ and $\boldsymbol{\nu} = (a, b, e, \gamma, \kappa_1, \kappa_2, \rho)$.

Rather than assuming that dealers are risk-averse, we assume that they face a cost of not meeting investor demand.[30] A dealer who draws type $\boldsymbol{s}$ obtains the following gross benefit from "consuming" amounts $(1 - \kappa_1)q_1$ and $(1 - \kappa_2)q_2$:

$$U(q_1, q_2, s) = t_1(1 - \kappa_1)q_1 + t_2(1 - \kappa_2)q_2. \tag{16}$$

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

In the secondary market a mass of investors with random demand $\{x_1, x_2\}$ arrives to the dealer. Equivalently, you may imagine that there are two types of investors, each with a random demand for one of the two maturities. We assume that each of $\{x_1, x_2\}$ is on-the-margin uniformly distributed on $[0, 1]$ but allow both amounts to be correlated. Specifically, $\{x_1, x_2\}$ assumes the following (Farie-Gumbel-Morgenstern cupola) density $f(x_1, x_2) = 1 +$

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

$3\rho(1-2F_1(x_1))(1-2F_2(x_2))$ with marginal distributions $F_m(x_m) = x_m$ and correlation parameter $\rho \in \left[-\frac{1}{3}, +\frac{1}{3}\right]$.

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

$$p_{i,1}(x_1, x_2|q_1, q_2) = \begin{cases} a - bx_1 - ex_2 & \text{for } x_1 \leq \kappa_1 q_1 \text{ and } x_2 \leq \kappa_2 q_2 \\ a - bx_1 & \text{for } x_1 \leq \kappa_1 q_1 \text{ and } x_2 > \kappa_2 q_2 \\ 0 & \text{for } x_1 > \kappa_1 q_1 \text{ and } x_2 > \kappa_2 q_2. \end{cases} \tag{17}$$

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

Overall, resale prices are characterized by three parameters $\{a, b, e\}$. A higher intercept $a > 0$ increases the dealer's bargaining power, and with it the price they can charge for each unit sold. Parameter $b > 0$ governs the price-sensitivity of investors. Large investors (who demand more) have more negotiating power and can drive down the price. Parameter $e \in \mathbb{R}$ drives the cross-price sensitivity.

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

$$revenue(x_1, x_2|q_1, q_2) = p_1(x_1, x_2|q_1, q_2)x_1 + p_2(x_1, x_2|q_1, q_2)x_2. \tag{18}$$

Turning down investors is costly for the dealer. An unhappy investor is, for instance, less likely to contact the dealer again in the future. In reality, a dealer might even want to source the security an investor demands in the secondary market so as to avoid losing their investor in the longer run. This is costly for the dealer because it is expensive to borrow or buy additional T-bills on the secondary market when demand is high. In our model, dealers face

the following cost function:

$$cost(x_1, x_2|q_1, q_2) = \begin{cases} 0 & \text{if } x_1 \leq \kappa_1 q_1 \text{ and } x_2 \leq \kappa_2 q_2 \\ \gamma x_1 & \text{if } x_1 > \kappa_1 q_1 \text{ and } x_2 \leq \kappa_2 q_2 \\ \gamma x_2 & \text{if } x_1 \leq \kappa_1 q_1 \text{ and } x_2 > \kappa_2 q_2 \\ \gamma x_1 x_2 & \text{if } x_1 > \kappa_1 q_1 \text{ and } x_2 > \kappa_2 q_2. \end{cases} \tag{19}$$

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

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

$$V(q_1, q_2, s) = U(q_1, q_2, s) + \mathbb{E}\left[revenue(x_1, x_2|q_1, q_2) - cost(x_1, x_2|q_1, q_2)\right], \tag{20}$$

where the expectation is taken over $\{x_1, x_2\}$. This determines how much the are willing to pay on-the-margin: $v_m(q_m, q_{-m}, s) = \frac{\partial V(q_m, q_{-m}, s)}{\partial q_m}$.

Proposition 3. *The marginal willingness to pay for amount q_m of a dealer with signal $s = (s_m, s_{-m})$, conditional on winning q_{-m} in the other auction $-m$ can be approximated by*

$$v_m(q_m, q_{-m}, s) = f_m(s_m) + \lambda_m q_m + \delta_m q_{-m} \tag{21}$$

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

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

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

When buying only for its own account a dealer is willing to pay the marginal value that the bill brings to their own institution, t_1. When they anticipate that they will sell at least some of maturity 1, their demand in auction 1 decreases in q_1 as long as their investors are sufficiently price-elastic, i.e., b is sufficiently high. If they sell all of both maturities, the demand is independent of their private type t_1, but depends on the amount they win of the other maturity.

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

Corollary 3. *The cross-maturity demand coefficient in the dealer's demand, δ_m, increases when (i) the inverse demand in the secondary market (17) for asset m decreases more strongly in response to larger quantities of asset $-m$ ($e \downarrow$), (ii) it is more costly to turn down investors ($\gamma \uparrow$), or (iii) it is more likely that investors with demand for different maturities arrive ($\rho \uparrow$).*

The corollary yields noteworthy implications that could be tested with data that combines information from the primary and secondary market.

G Potential Sources of Interdependencies of T-Bills

To understand what might drive imperfect substitution of bills in the primary market, we exploit a unique secondary market cash and repo data set (MTRS 2.0). For each traded security, we observe the time, price and quantity, and trade-type (buy/sell/repo/reverse-repo). We also observe who trades, e.g., bank, pension fund, or hedge fund. The data provide us with a unique look into the use of Treasury bills by different groups of investors. The main caveat is that the time period during which transaction-level data is available (post-2016) does not overlap with the auction sample. Therefore, we cannot definitively attribute primary market behavior to what we observe in the secondary market, or link both markets in our structural model.

We focus on two hypotheses that address substitutability: (i) different types of bills are used in different ways, and (ii) different types of investors prefer different maturities (preferred habitat). We define bills based on the number of days left to maturity: 3M bills are bills with less than 91 days left to maturity, 6M bills have 92-184 days left to maturity and 12M bills are all others. Data are from January 2016 to January 2022. We summarize the results graphically.

First, Appendix Figure A6 shows that 12M bills are more heavily used in repo agreements than 3M bills.[31] One explanation for this pattern is that the money-like premium on bills decreases in the length to maturity (see, for example, Greenwood et al. (2015)), and therefore 3M bills are the most expensive. As a result, 3M bills are generally avoided as collateral in repurchase agreements—bills are not perfect substitutes. Furthermore, Appendix Figure A7 highlights heterogeneity across investor groups in the extent they use different bills as repo-collateral. Governments and central banks, as well as pension funds and insurance companies tend to use 3M bills more often than other investor groups, while hedge funds and mutual funds mostly use 12M bills.

Finally, Appendix Figure A8 highlights that the trading frequency of different maturities varies by investor type. To better visualize trade-patterns, we separate the buy-from the sell-side. We find that governments and central banks are more likely to both buy and sell longer maturities compared to other groups. This could be because they seek to reduce search costs of replacing bills that mature. On the contrary, asset managers and mutual funds tend to trade 3M bills, perhaps to fulfill liquidity requirements.

H Proofs

Proof of Proposition 1. When bidders are price-takers, they bid their actual willingness to pay for amount q_m in auction m. This depends on how much of the other securities they get allocated. Ideally, the bidder would want to condition their bid in auction m for amount $q_{m,k}$ on how much they will purchase of the other security in equilibrium, $\boldsymbol{q^*_{-m,i}}$. However, by the rules of the auction, this is infeasible given that the step function in auction m can only depend on prices and quantities of security m, but not of other securities. Therefore, the

[31] Approximately 15% of repo transactions involve bills. The remainder is bonds.

bidder takes an expectation over how much they will win of the other security, conditional on winning $q_{m,k}$. □

Proof of Proposition 2. Fix an auction format, and omit the bidder i-subscript for notational convenience. Take the perspective of dealer with signal $\boldsymbol{s}$, and let all other agents $j \neq i$ play a symmetric equilibrium. In this equilibrium it must be optimal for the dealer to choose the same set of functions $\{b_1(\cdot, \boldsymbol{s}), ..., b_M(\cdot, \boldsymbol{s})\}$ as all other dealers with signal $\boldsymbol{s}$. These M functions must jointly maximize the dealer's expected total surplus. It must therefore be the case that each of the functions $b_m(\cdot, \boldsymbol{s})$ maximizes the expected total surplus separately when fixing all the other bidding functions $-m$ at the optimum. To determine the necessary conditions of an equilibrium we can consequently fix the dealer's strategy in all but one auction at the equilibrium.

The remainder of the proof extends Kastl (2012)'s proof for a K-step equilibrium of a discriminatory price auction that takes place in isolation; and similarly for a uniform price auction. Hortaçsu et al. (2018) use the simpler version of Kastl (2012)'s equilibrium condition for uniform price auctions since ties occur with zero probability with supply uncertainty and when the restriction on the number of steps is not binding. To avoid redundancies, we do not copy and paste the original proof, but only highlight the necessary changes.

Kastl (2012)'s proofs rely on a series of lemmas. To show that they extend to our setting it suffices to replace a bidder's marginal willingness to pay for amount $q_{m,k}$ in auction m with how much the bidder expects to benefit from winning amount $q_{m,k}$ in auction m and taking the expectation over how much they will win in the other parallel auctions, $\tilde{v}_m(q_{m,k})$ (the LHS in the expressions of Proposition 2). The only thing we need to guarantee is that $\tilde{v}_m(q_{m,k})$ fulfills all necessary properties that the marginal utility is assumed to fulfill in Kastl (2012) and Hortaçsu et al. (2018). This is established in Lemma 1, and it generalizes to the case with dealer updating in our extended model.

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

Proof of Lemma 1. Most properties of $v_m(q_m, q_{-m}, s)$ carry over to $\tilde{v}(\cdot)$ in a straightforward way. Statement (i) holds directly by assumption.

For (ii) and (iv): Continuity of $\tilde{v}_m(q_m)$ in q_m is implied if $\mathbb{E}\left[\boldsymbol{q^*_{-m,i}}|q_m\right]$ is continuous in q_m. Since the bidding strategy in auctions of other maturities does not depend on q_m, the expected allocation also cannot change with quantity demanded in auction for maturity m. Hence, $\tilde{v}\left(\cdot\right)$ is both weakly decreasing and continuous in q_m.

(iii) By assumption, $v_m\left(q_m, q_{-m}, s\right)$ is strictly increasing in s_m for any q_{-m} and q_m, and by the Envelope Theorem, q^*_m must be increasing in s_m for all m. Further, given that s_m and s_{-m} are affiliated, a bidder expects to win more of the other maturity $-m$ conditional on winning more of maturity m when drawing a high s_m. Therefore $\tilde{v}_m(\cdot)$ must be strictly increasing in s_m at all q_m. $\square$

Proof of Proposition 3. Given the aggregate inverse demand of investors, (17), the dealer expects the following payoff (20) from owning q_1, q_2:

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

Insert the functional forms (16), (17), and $f(x_1,x_2) = 1 + 3\rho(1-2F_1(x_1))(1-2F_2(x_2))$, integrate and take the partial derivative w.r.t. q_1 for illustration. Then do a Taylor expansion around the mean of (x_1,x_2), which is $\left(\frac{1}{2},\frac{1}{2}\right)$, to obtain:

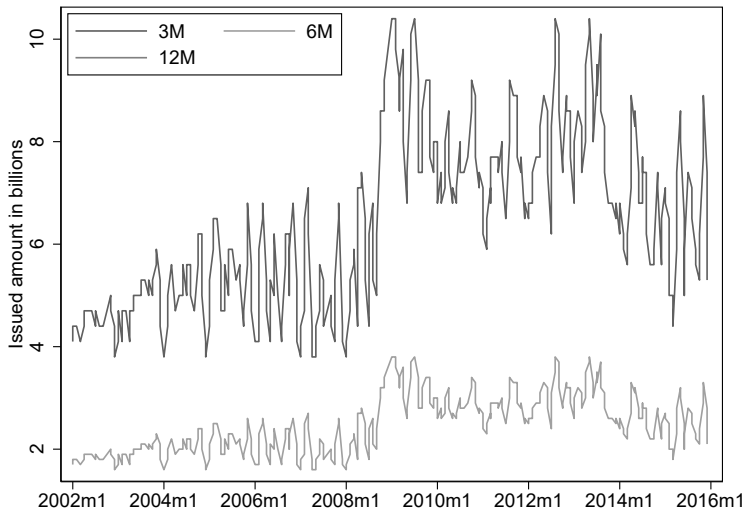
$$v_1(q_1,q_2,s_1) = (1-\kappa_1)t_1 + h_0(\kappa_1,\kappa_2,\gamma,\rho) + h_1(\kappa_1,\kappa_2,\gamma,a,b,e,\rho)q_1 + h_2(\kappa_1,\kappa_2,e,\rho)q_2$$

with

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

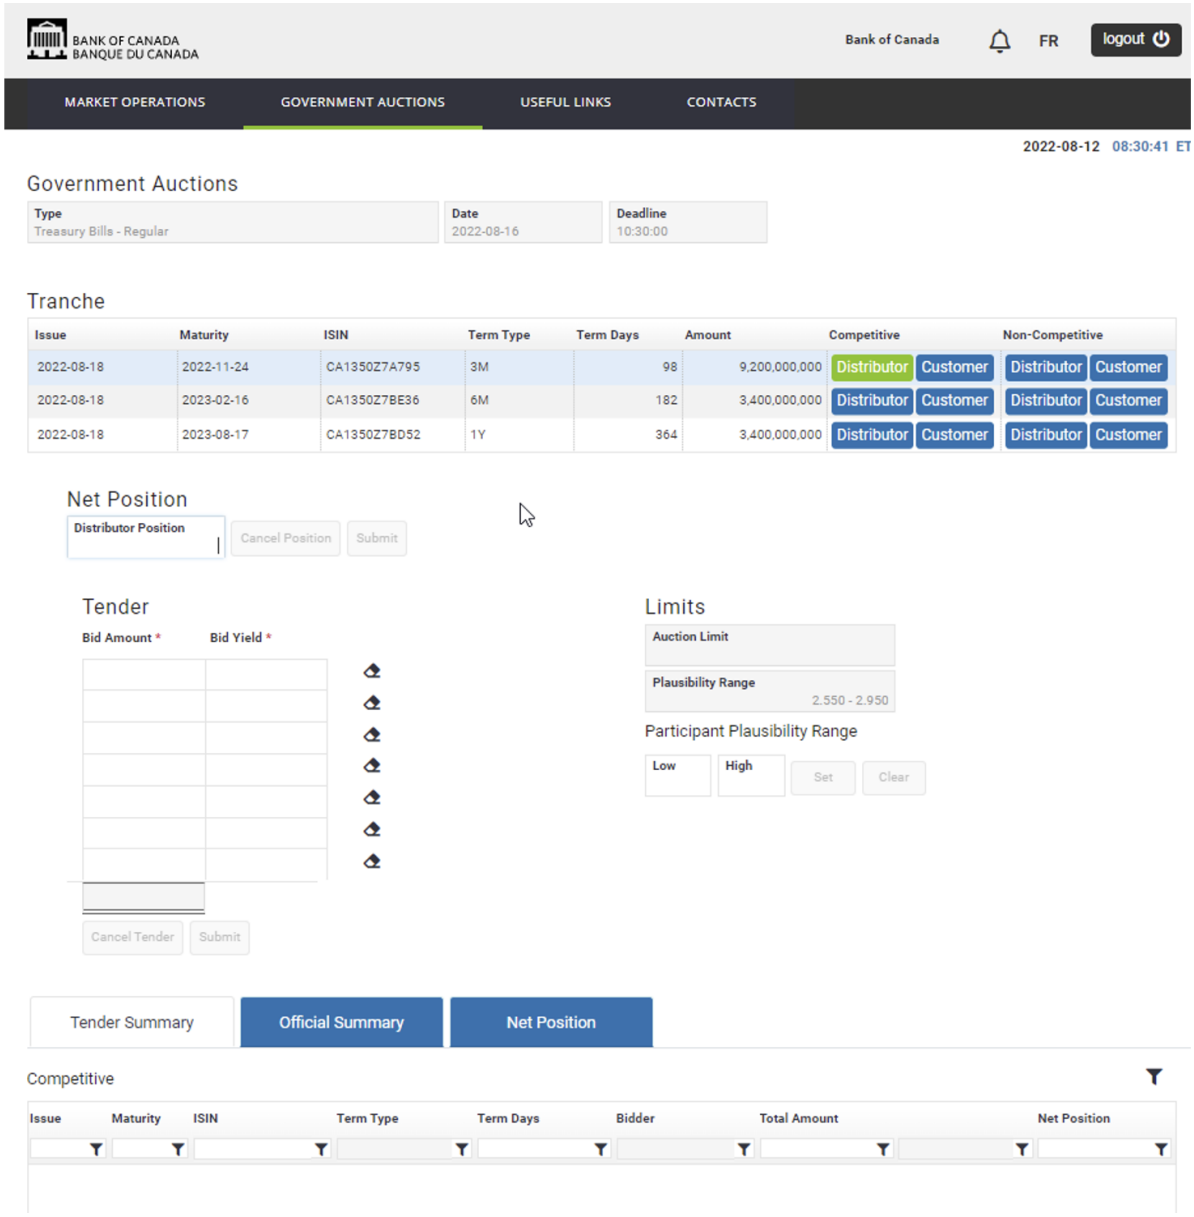
Proof of Corollary 3. Given our parameter restrictions, $\frac{\partial h_2(\kappa_1,\kappa_2,\gamma,e,\rho)}{\partial e} \leq 0$, $\frac{\partial h_2(\kappa_1,\kappa_2,\gamma,e,\rho)}{\partial \gamma} \geq 0$, $\frac{\partial h_2(\kappa_1,\kappa_2,\gamma,e,\rho)}{\partial \rho} \geq 0$ for any $\kappa_m \in [0,1]$ and any $\rho \in [-1/3, 1/3]$. $\square$

Appendix Figure A1: Issuance of Canadian 3M, 6M, 12M T-Bills



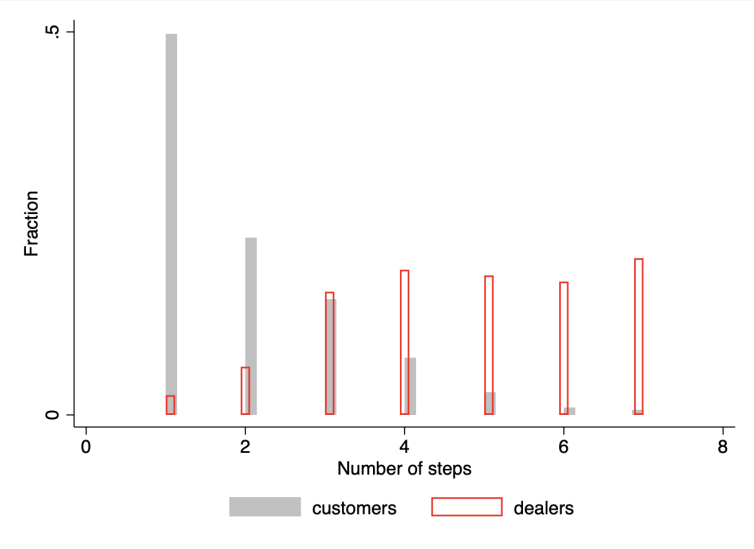
Appendix Figure A1 displays a time series of the issued supply of the 3M, 6M, and 12M bills, where the 6M issuance do not appear in the graph because they are identical to 12M issuance. Over time, the amounts issued of the different maturities are perfectly correlated.

Appendix Figure A2: Auction Interface



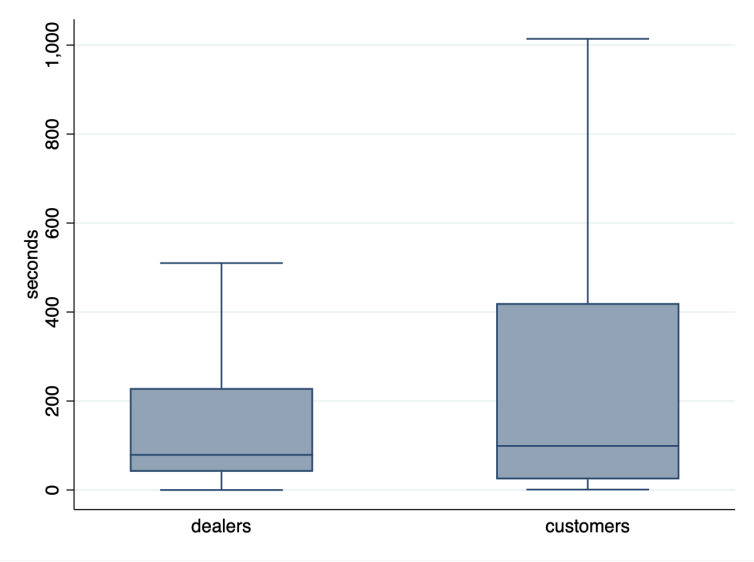
Appendix Figure A2 shows a screen shot of what a bidder sees when bidding. He sees three "Tranches", listing the three different securities for sale. If he is a dealer, placing a competitive bid for his own account, he clicks on "Distributor" (in green). He fills in what "Net Position" he currently holds of the security, and his bid "Tender" of maximally 7 steps. The bidder sees a non-binding "Plausibility Range" suggesting yields at which the auction might clear. When the dealer is content with his bid, he clicks "Submit". Then he can choose a different "Tranche" to bid in a different auction. At the bottom of the page, he sees an overview of all his submitted bids for all auctions.

Appendix Figure A3: Steps by Bidder Groups



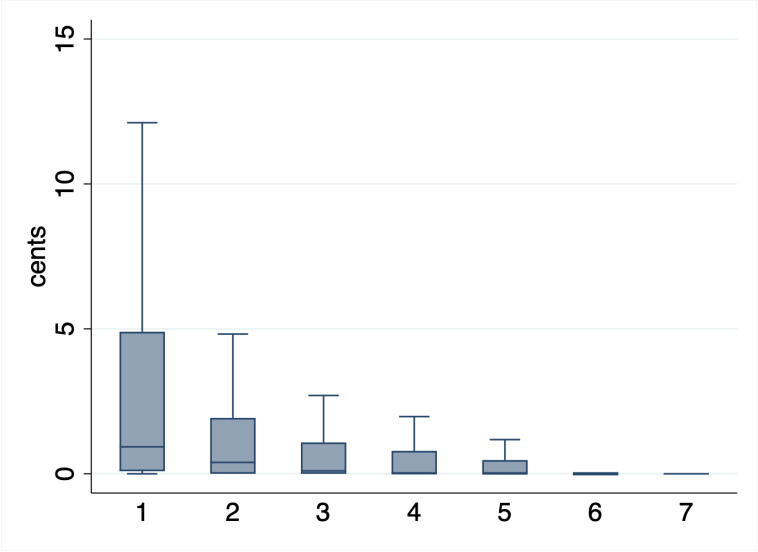
Appendix Figure A3 shows a histogram of the number of steps customers (in gray) and dealers (in red) submit. The fraction is measured in percentage points.

Appendix Figure A4: Time Between Bids of Those who do not Update



Appendix Figure A4 shows the distribution of the time difference (measured in seconds) between the bids that a dealer and a customer who do not update the bids placed in different auctions. Outliers are excluded.

Appendix Figure A5: Distribution of the Untrimmed Shading Factor



Appendix Figure A5 shows box plots of the untrimmed shading factor per step in a bidding function. For each step, the distribution is taken over dealers, time, and maturities. The shading factor is in cents; the face value is C$ 1,000.

Appendix Figure A6: Volume Traded in the Repo Market

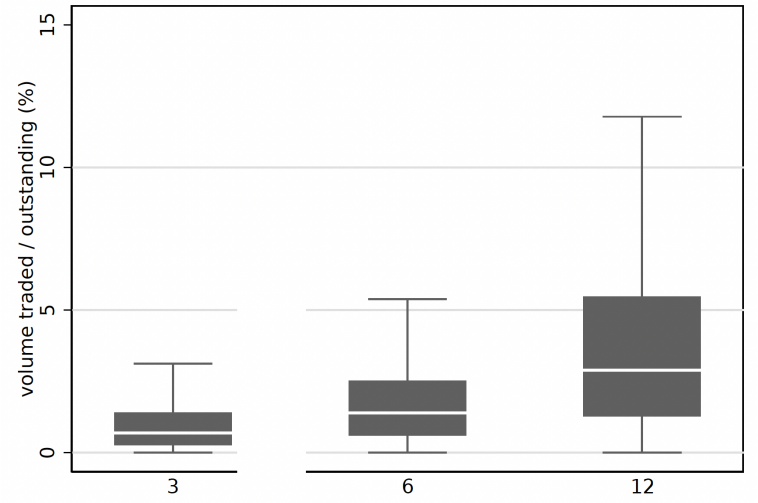
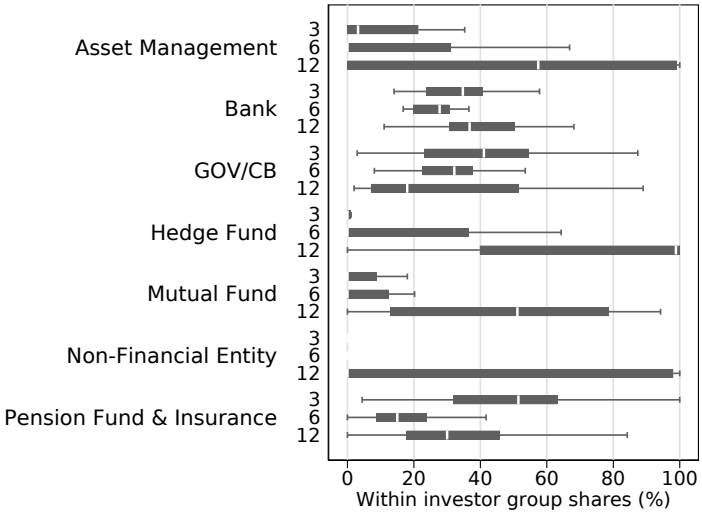


Figure A6 shows the daily amount exchanged in repo transactions that involve 3M, 6M, and 12M bills, as a percentage of total amount outstanding of bills in the respective categories. This information is constructed from MTRS2.0 data, using transactions between 01/2016 until 01/2022.

Appendix Figure A7: Repo Trade Shares per Investor Group

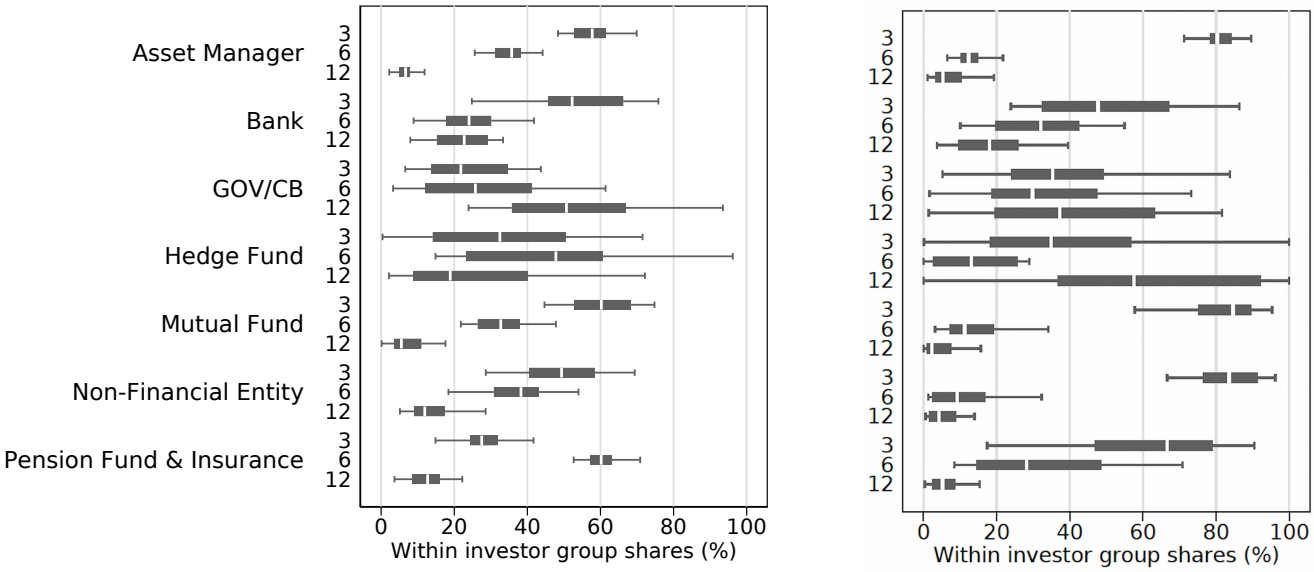


Appendix Figure A7 shows the distribution of the within investor group share of quarterly gross repo volume by maturity class, measured in percentage, including both sides of the contract. For example, one observation in the 3M box plot, labeled GOV/CB, shows what percentage of the total quarterly repo trade volume by government entities and central banks uses 3M bills. Insurance stands for insurance company. 3M bills are bills with less than 91 days left to maturity, 6M bills have 92-184 days left to maturity and 12M bills are all others. This information is constructed from MTRS2.0 data, using transactions between 01/2016 until 01/2022.

Appendix Figure A8: Cash Market Trade Shares per Investor Group

(a) Buy-Side

(b) Sell-Side



Appendix Figure A8a shows the distributions of how much each investor group buys in form of 3M, 6M, and 12M bills in percentage of the total amount bought by this investor group in the T-bill cash market, on average within a quarter. GOV/CB includes government entities and central banks, and Insurance stands for insurance companies. Figure A8b shows the analogue, but for the sell-side. 3M bills are bills with less than 91 days left to maturity, 6M bills have 92-184 days left to maturity and 12M bills are all others. This information is constructed from MTRS2.0 data, using transactions between 01/2016 until 01/2022.

Appendix Table A1: Demand Coefficients—Robustness w.r.t. Number of Steps

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	-0.534^{***}	(0.0047)	λ_{6M}	-1.054^{***}	(0.0078)	λ_{12M}	-2.366^{***}	(0.0175)
$\delta_{3M,6M}$	-0.048^{***}	(0.0077)	$\delta_{6M,3M}$	-0.048^{***}	(0.0077)	$\delta_{12M,3M}$	-0.066^{***}	(0.0089)
$\delta_{3M,12M}$	-0.066^{***}	(0.0089)	$\delta_{6M,12M}$	-0.222^{***}	(0.0132)	$\delta_{12M,6M}$	-0.222^{***}	(0.0132)
N	55,822			38,856			46,778	

Appendix Table A1 is analogous to Tables 3c but estimated on a subsample of step functions with at least three steps. Values are in cents, i.e., C$1/100, and quantities in percentage of auction supply. The bills' face value is C$1,000. The first three columns show the estimates for the 3M bill auction, the next three for the 6M bill auction and the last three for the 12M bill auction. The point estimates are in the second, fifth and eight column. Standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table A2: Demand Coefficients—Robustness w.r.t. Value Trimming

(a) 3M Bill Auction

markup	4 bps		10 bps		20 bps		40 bps	
λ_{3M}	-0.499^{***}	(0.0040)	-0.602^{***}	(0.0064)	-0.746^{***}	(0.0107)	-0.996^{***}	(0.0196)
$\delta_{3M,6M}$	-0.035^{***}	(0.0064)	-0.095^{***}	(0.0107)	-0.182^{***}	(0.0185)	-0.327^{***}	(0.0341)
$\delta_{3M,12M}$	-0.053^{***}	(0.0074)	-0.146^{***}	(0.0124)	-0.275^{***}	(0.0213)	-0.511^{***}	(0.0392)
N	58,542		58,542		58,542		58,542	

(b) 6M Bill Auction

markup	4 bps		10 bps		20 bps		40 bps	
λ_{6M}	-0.968^{***}	(0.0065)	-1.194^{***}	(0.0108)	-1.512^{***}	(0.0188)	-2.062^{***}	(0.0347)
$\delta_{6M,3M}$	-0.034^{***}	(0.0064)	-0.095^{***}	(0.0107)	-0.182^{***}	(0.0185)	-0.327^{***}	(0.0341)
$\delta_{6M,12M}$	-0.171^{***}	(0.0109)	-0.427^{***}	(0.0188)	-0.807^{***}	(0.0332)	-1.420^{***}	(0.0618)
N	42,282		42,282		42,282		42,282	

(c) 12M Bill Auction

markup	4 bps		10 bps		20 bps		40 bps	
λ_{12M}	-2.160^{***}	(0.0145)	-2.780^{***}	(0.0247)	-3.669^{***}	(0.0431)	-5.264^{***}	(0.0805)
$\delta_{12M,3M}$	-0.053^{***}	(0.0074)	-0.146^{***}	(0.0124)	-0.275^{***}	(0.0213)	-0.511^{***}	(0.0392)
$\delta_{12M,6M}$	-0.171^{***}	(0.0109)	-0.427^{***}	(0.0188)	-0.807^{***}	(0.0332)	-1.420^{***}	(0.0618)
N	50,410		50,410		50,410		50,410	

Appendix Table A2 reports the coefficients for equation (11), estimated using competitive dealer bids with at least two steps, and using different markups to trim value estimates (4 bps, 10 bps, 20 bps, 40 bps). The estimates for a markup of 5 bps, our favorite specification in light of the distribution of shading factors, shown in Appendix Figure A5, are in the main text. Values are in cents, i.e., C$1/100, and quantities in percentage of auction supply. The bills' face value is C$1,000. Standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table A3: Demand Coefficients for the Extended Model of Appendix C

(a) Average Dealer

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	-0.530^{***}	(0.0045)	λ_{6M}	-1.029^{***}	(0.0073)	λ_{12M}	-2.290^{***}	(0.0162)
$\delta_{3M,6M}$	-0.047^{***}	(0.0073)	$\delta_{6M,3M}$	-0.047^{***}	(0.0073)	$\delta_{12M,3M}$	-0.081^{***}	(0.0084)
$\delta_{3M,12M}$	-0.081^{***}	(0.0084)	$\delta_{6M,12M}$	-0.250^{***}	(0.0124)	$\delta_{12M,6M}$	-0.250^{***}	(0.0124)
N	58,542			42,282			50,410	

(b) Dealer Group 1

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	-0.496^{***}	(0.0046)	λ_{6M}	-0.968^{***}	(0.0074)	λ_{12M}	-2.168^{***}	(0.0164)
$\delta_{3M,6M}$	-0.051^{***}	(0.0074)	$\delta_{6M,3M}$	-0.051^{***}	(0.0074)	$\delta_{12M,3M}$	-0.092^{***}	(0.0085)
$\delta_{3M,12M}$	-0.092^{***}	(0.0085)	$\delta_{6M,12M}$	-0.290^{***}	(0.0124)	$\delta_{12M,6M}$	-0.290^{***}	(0.0124)
N	45,405			33,464			40,958	

(c) Dealer Group 2

	3M Bill Auction			6M Bill Auction			12M Bill Auction	
λ_{3M}	-0.760^{***}	(0.0138)	λ_{6M}	-1.443^{***}	(0.0246)	λ_{12M}	-3.045^{***}	(0.0524)
$\delta_{3M,6M}$	$+0.011$	(0.0241)	$\delta_{6M,3M}$	$+0.011$	(0.0241)	$\delta_{12M,3M}$	-0.009	(0.0272)
$\delta_{3M,12M}$	-0.009	(0.0272)	$\delta_{6M,12M}$	$+0.081$	(0.0452)	$\delta_{12M,6M}$	$+0.081$	(0.0452)
N	13,137			8,818			9,452	

Appendix Tables A3 is analogous to Table 3c. They report the coefficients for equation (11), but for the extended model with two dealer groups, presented in Appendix C. In Table A3a we show demand estimates for the average dealer, in Table A3b and A3c we separate between dealers in group 1 and group 2. Values are in cents, i.e., C$1/100, and quantities in percentage of auction supply. The bills' face value is C$1,000. The first three columns show the estimates for the 3M bill auction, the next three for the 6M bill auction and the last three for the 12M bill auction. The point estimates are in the second, fifth and eight column. Standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.