# FTG Working Paper Series

Centralizing Over-The-Counter Markets?

by

Milena Wittwer
Jason Allen

Working Paper No. 00109-00

Finance Theory Group

www.financetheory.com

# Centralizing Over-The-Counter Markets?*

Jason Allen[▷] Milena Wittwer[◁]

March 16, 2023

**Abstract**

In traditional over-the-counter markets, investors trade bilaterally through intermediaries. We assess whether and how to shift trades on a centralized platform with trade-level data on the Canadian government bond market. We document that intermediaries charge a markup when trading with investors, and specify a model to quantify price and welfare effects from market centralization. We find that many investors would not use the platform, even if they could, because it is costly, competition for investors is low, and investors value relationships with intermediaries. Market centralization can even decrease welfare, unless competition is sufficiently strong.

**Keywords:** OTC markets, platforms, demand estimation, government bonds
**JEL:** D40, D47, G10, G20, L10

*The views presented are those of the authors and not necessarily the Bank of Canada. Any errors are our own. We thank experts at the Bank of Canada—in particular, Corey Garriott and the Triton project team for critical insights into debt management—and at CanDeal for providing detailed institutional information. We also thank Darrell Duffie, Liran Einav and Matthew Gentzkow in addition to Luis Armona, Markus Baldauf, Timothy Bresnahan, Eric Budish, Nina Buchmann, Jose Ignacio Cuesta, Piotr Dworczak, Jakub Kastl, David K. Levine, Negar Matoorian, Paul Milgrom, Monika Piazzesi, Tobias Salz, Chad Syverson, Or Shachar, Paulo Somaini, Alberto Teguia, Robert Wilson, Ali Yurukoglu, all anonymous referees, and participants in the IO, finance, and macro student workshops at Stanford. [▷]Jason Allen (Bank of Canada): jallen@bankofcanada.ca; [◁]Milena Wittwer (Boston College): wittwer@bc.edu.

# 1 Introduction

Each year, trillions of dollars worth of bonds, mortgage-backed securities, currencies, commodities, and derivatives are traded in over-the-counter (OTC) markets. In traditional OTC markets, buyers must contact sellers one by one in order to trade bilaterally. Most OTC markets, therefore, rely on large financial institutions—dealers—to intermediate between investors, such as firms, banks, public entities, or individuals.

A series of antitrust lawsuits that have accused dealers of abusing market power, combined with dramatic events during the COVID-19 crisis, raised the question of whether and how to centralize OTC markets. One approach is to centralize trade on an electronic platform on which investors can access multiple dealers simultaneously. Yet, even though these types of platforms already exist in many OTC markets, investors are reluctant to use them. As a result, a two-tiered market structure with bilateral and platform trading persists in most modern OTC markets (BIS (2016)).

We assess whether and how to centralize an OTC market with trade-level data on the Canadian government bond market. We document that most investors trade bilaterally rather than on a centralized platform, and estimate a structural model to understand why this is the case, and to assess welfare effects from shifting more investors on the platform.

Three features render the Canadian government bond market a particularly attractive setting for our research questions. First, government bonds are more homogeneous than other securities. We can therefore rule out confounding factors that might explain a decentralized market structure and identify dealer market power. Second, a centralized platform—which is like platforms in many other OTC markets, including the largest ones in the United States and Europe—exists, but not all investors have access to it. This is a useful institutional feature we exploit in our empirical strategy. Third, a reporting regulation allows us to observe trade-level data, so that we can zoom in on each individual trade, unlike most prior studies on government bond markets.[1]

The data cover essentially all trades that involve Canadian government bonds, as well as bidding data from all primary auctions in which the government issues bonds. The data set is unique in that it includes identifiers for market participants and securities, so that we can trace both through the market. We observe the time, price, and size of trades, and

[1] Relative to other OTC markets, we know little about government bond markets because trade-level data are not readily available. The U.S., for example, began collecting trade-level data in mid-2017 but does not make it accessible to academics. Some countries granted access to data on parts of their Treasury market (Dunne et al. (2015); Monias et al. (2017); de Roure et al. (2020); Kondor and Pintér (2022); Pintér and Semih (2022)).

know whether a trade was executed bilaterally or on the platform and whether an investor is institutional (with platform access) or retail (without platform access). In addition, we collect bid and ask quotes that dealers post to attract investors.

Our trade-level data allow us to document novel facts that motivate our structural model. Despite better prices on the platform, institutional investors mostly trade bilaterally—and do so with a single dealer (home dealer). Conditional on entering the platform, they also trade with other dealers. Retail investors—who cannot trade on the platform—face worse prices than institutional investors. This could just be because they are willing to pay more. The fact that institutional investors who lose platform access for regulatory reasons obtain worse prices, however, raises the possibility that making platform access universal could lead to better prices.

To assess price and welfare effects when centralizing the market, we introduce a model with dealers and investors who have different values for realizing trade. Dealers compete for investors with platform access by simultaneously choosing quotes at which they are willing to trade on the platform. Institutional investors can enter the platform at a cost to trade with any of the dealers at the posted quotes. Alternatively they trade bilaterally with their home dealer and earn a loyalty benefit. Retail investors can only trade bilaterally.

In estimating the model, we face the common challenge that prices (here quotes) are endogenous. Our solution is to construct a new cost-shifter instrument that changes the dealer's costs to sell but not the investor demand. For this, we use bidding data on primary auctions, in which dealers buy bonds from the government to sell them at a higher price to investors. When a dealer wins more than she expected to win when bidding, she can more cheaply satisfy investor demand, either because of how others bid in the auction or because the government issued more than the dealer expected. Thus, how much more the dealer wins relative to what she expected to win represents an exogenous cost-shifter, which we construct with estimation techniques from the multi-unit auctions literature.

Our findings are threefold. First, the model estimates suggest that many (institutional) investors don't use the platform because it is unattractive: platform costs are high and dealer competition is weak. Further, investors stay off the platform because they enjoy large loyalty benefits when trading with their home dealer. Only investors with high values for trading enter the platform in equilibrium.

Second, dealer market power plays a key role in explaining the observed market structure. To show this, we conduct various counterfactuals, in which we reduce the frictions that arise due to imperfect platform access and design. This shifts more investors on the platform. We find that total gains from trade—our measure for welfare—decrease the more investors trade on the platform. This is counterintuitive as it should become easier for high-valued buyers to match with low-valued sellers. The problem is that dealers with market power charge prices that distort the investor's trading decision away from the efficient outcome.

Third, dealer-investor relationships have large effects on welfare. To highlight this, we repeat all counterfactuals but shut off the loyalty benefit. This increases competition for investors as they more easily switch to other dealers. Quotes become more favorable for investors, making the platform more attractive. As a result, more investors enter the platform on which they trade with dealers that have high values to trade—for instance, because they seek to offload inventory. Welfare increases by up to 27%.

Taken together, our results have valuable policy implications for OTC markets. They help explain why the two-tiered market structure with bilateral and platform trading that we observe in most OTC markets is so persistent, and highlight three sources of inefficiency. First, dealer market power alone diminishes welfare by 12%. Second, dealer-investor relationships foster market power and are distortive when an investor trades with a dealer because of the relationship and not because the dealer has a high value for trading. Lastly, there are frictions that prevent dealers from buying and selling. These became apparent in the recent COVID-19 crisis, when dealers in several countries failed to absorb the excess supply of government bonds on their balance sheets. Market centralization can reduce all three types of frictions, but welfare gains are limited when investors are loyal to dealers. We expect this to be true for many other OTC markets for standardized financial products (such as simple interest rate swaps or credit derivative index products).

Finally, our findings highlight two general lessons beyond OTC markets. First, markets that seem efficient might not be. With trade-level data, we find that the Canadian government bond market—which is liquid and offers a homogeneous, cash-like good—does not achieve the first-best. The degree of efficiency is lower than aggregate statistics that approximate the degree of efficiency, such as bid-ask spreads, would suggest. Second, introducing platforms to reduce frictions in decentralized markets—which are pervasive throughout the economy—has limitations. Platforms by Uber, Airbnb, Amazon, and others might not be designed to foster competition, because they are owned by profit-maximizing firms.

**Related literature.** Our main contribution is to empirically assess whether and how to centralize OTC markets. This adds to a literature that analyzes different aspects of decentralized or centralized financial markets, typically via reduced-form analysis.[2] The most

[2]Examples include Barclay et al. (2006); Loon and Zhong (2016); Fleming et al. (2017); Brancaccio et al. (2017); Plante (2017); Abudy and Wohl (2018); Biais and Green (2019); Benos et al. (2020); Kozora et al. (2020); Riggs et al. (2020); Coen and Coen (2021); Fleming and Keane

closely related paper is Hendershott and Madhavan (2015) (HM). Similar to HM, we build a model in which investors choose between bilateral trading and trading on a multi-dealer platform; however, we highlight the trade-off dealers face when choosing quotes, which are exogenous in HM. Further, we structurally estimate our model, which allows us to conduct counterfactual analyses.

By creating a data set with trade-level information, we contribute to a steadily growing literature that analyzes trade-level data on financial markets—recently reviewed by Bessembinder et al. (2020). Our market differs from those previously studied because it is highly liquid and features relatively high price transparency with little uncertainty about the true value of the asset.

By showing that even in this market there is evidence of price discrimination, we add to the descriptive evidence of price dispersion in less liquid or more opaque OTC markets (e.g., Green et al. (2007); Friewald and Nagler (2019); Hau et al. (2021)). Our findings differ from de Roure et al. (2020), who document an OTC discount in the German government bond market; and complement Kondor and Pintér (2022) and Pintér and Semih (2022), who study pricing, liquidity, and welfare in the UK bond market.

By estimating the demand and demand elasticity of an individual investor for government bonds, we contribute to a large literature that studies government bond markets using aggregate data (e.g., Garbade and Silber (1976); Krishnamurthy and Vissing-Jørgensen (2012, 2015)) and a young literature that estimates demand for financial assets (e.g., Koijen and Yogo (2019, 2020)). For estimation, we exploit techniques used to study multi-unit auctions to construct a cost-shifter instrument for prices outside of the auction (e.g., Hortaçsu and McAdams (2010); Kastl (2011); Hortaçsu and Kastl (2012); Allen et al. (2020)). Further, we apply an approach by Bresnahan (1981) that is commonly used in the literature on demand estimation to infer the marginal costs of firms from observable behavior in a trade setting. Here, "marginal costs" are values for realizing trade.

Our theory lies in between the theoretical literature on OTC markets (following Duffie et al. (2005), recently reviewed by Weill (2020)) and a large theoretical literature that studies decentralized or fragmented financial markets (with recent work by Glode and Opp (2019); Yoon (2019); Chen and Duffie (2021); Rostek and Yoon (2021); Wittwer (2020, 2021)). Similar to a few other papers, our model focuses on the selection of investors into trading venues (e.g., Liu et al. (2018); Vogel (2019)).[3] Different from these papers, we highlight the

(2021); Hendershott et al. (2021); Holden et al. (2021); Kutai et al. (2021); Lehar and Parlour (2021); O'Hara and Zhou (2021); Brancaccio and Kang (2022); Pintér and Semih (2022).

[3]Following the literature, we assume exclusive participation per market segment. Only recently,

importance of benchmark prices, as in Duffie et al. (2017), but we endogenize them.

Unlike most papers in the OTC literature, we do not highlight price opaqueness, because the market we study is more liquid and price-transparent than other markets. This is similar to Babus and Parlatore (2022), who study market fragmentation in OTC markets when there is no centralized platform, and to Baldauf and Mollner (2020), Wang (2022), and Yueshen and Zou (2022), who show that it can be theoretically optimal for investors to limit the number of dealers they contact. This is in line with our empirical findings.

## 2 Institutional environment

Government bond markets are ideal for studying whether centralizing OTC markets can decrease dealer market power, leading to welfare gains. The reason is that government bonds offer greater safety and liquidity than other securities. They are closer to a perfectly homogeneous good with a public market price and quick settlement. Therefore, we can rule out confounding factors that might drive markups and explain a decentralized market structure in other settings (such as high illiquidity, asymmetric information, counterparty risk, and product differentiation).

**Market players.** Government bond markets are populated by a few (in Canada, 10) primary dealers, which we refer to as dealers, and many investors. There are also smaller dealers and brokers, but they play a minor role in our case. Dealers are large banks, such as RBC Dominion Securities. Investors come in two types: they are either institutional or retail. Whether an investor is classified as institutional or retail is set by the Industry Regulatory Organization of Canada (see IIROC Rule Book). The biggest classifying factor is how much capital an investor holds. Only if she holds enough does she qualify as an institutional investor.

To get a sense of who investors are, we manually categorize 1,459 investors we can identify by name. The largest investor groups are asset managers, followed by pension funds, banks, and firms that are members of IIROC. Then we have public entities (such as governments, central banks, and universities), insurance companies, firms that offer brokerage services, and non-financial companies (see Appendix Figure A1).

**Market structure.** A country's government bond market often makes up a large part of its total bond market (in Canada 70%). It splits into two parts. The first is the primary

Dugast et al. (2022) generalizes Atkeson et al. (2015) to allow for endogenous participation decisions in multiple market segments.

market, in which the government sells bonds via auctions, primarily to dealers. The second and larger part is the secondary OTC market. It is similar to other OTC markets, with one segment in which dealers trade with other dealers or brokers and one in which dealers trade with investors. We focus on the larger (for Canada) dealer-to-investor segment.

Trade realizes either via bilateral negotiation or on (an) electronic platform(s). These platforms are called alternative trading systems in the U.S. and Canada, multilateral trading facility in Europe, and dark pools for equities (see Bessembinder et al. (2020) for an overview). We focus on the most common type of platform in the dealer-to-investor segment, which matches investors to dealers but not to other investors.

Given the dealers' strong influence on OTC markets and the fact that it is not uncommon for dealers to own the platforms (as in Canada), there are reasons to believe that these platforms are not designed to maximize investor surplus. One indication of this is that, unlike dealer-investor trades, inter-dealer/broker trades can be executed on an anonymous limit order book. Further, some platforms are only accessible to institutional investors.

In Canada, until recently, there was only one multi-dealer platform: CanDeal. It operates similarly to most other platforms (described below), including the largest ones in the U.S. and Europe.[4] Yet, unlike some platforms, CanDeal does not offer central clearing, different times to settlement, or higher price transparency than the bilateral market. This is useful for us, as it rules out confounding factors that might drive differences to bilateral trading.

**How dealers and investors trade.** Dealers post indicative quotes to advertise bonds on multiple outlets, including Bloomberg and CanDeal. Most investors can see (at least) the Bloomberg quotes and thus have a sense of the bond's market value. CanDeal quotes advertise platform prices. No one observes trade prices. The posted quotes are valid for trades smaller than C$ 25 million. Investors who want to trade larger amounts have to contact dealers privately for a quote.

An investor who seeks to trade contacts a dealer, over the phone or via message, who makes an offer that the investor either accepts or declines. If the investor declines, the investor could contact other dealers to seek more bilateral offers, but this is rare. As an alternative, institutional investors can trade on the platform, CanDeal. To have platform access, they each have to pay a monthly fee which ranges roughly between C$ 2,500-3,500 for

[4] A non-exhaustive list of alternative trading systems includes MarketAxess (the leader in e-trading for global bonds), BGC Financial L.P. (which offers more than 200 financial products); BrokerTec Quote (which leads the European repo market); Tradeweb Institutional (global operator of electronic marketplaces for rates, credit, equities, money markets).

a typical institutional investor.[5] In return, on the platform investors have access to all dealers simultaneously which reduces search costs. Investors choose among different alternatives for how to trade, but the most common is to run a request for quote (RFQ) auction.[6]

In an RFQ auction, an investor typically sends a request for quote to multiple—in Canada maximally four—dealers (Hendershott and Madhavan (2015); Riggs et al. (2020)).[7] The request reveals to the dealers the name of the investor, whether it is a buy or sell, the security, the quantity, and the settlement date. Knowing how many—but not which—dealers are participating, dealers respond with a price (that may differ from the posted quote). The investor chooses the deal that she likes best and the trade is executed shortly after.

Running an RFQ auction differs from contacting multiple dealers bilaterally, because it is easier to make dealers compete. This is for three reasons. First, the auction reduces search costs that can hinder competition in the bilateral market. It is faster and requires less effort than contacting multiple dealers. Second, dealers have to respond simultaneously, which prevents costly renegotiation. Third, dealers see how many other dealers compete for the investor, which can increase competition.

# 3 Data

**Main data source.** The main source is the Debt Securities Transaction Reporting System, MTRS2.0, collected by IIROC since November 2015. Our sample contains trade-level information on all bond trades of registered brokers or dealers from 2016 to 2019. The sample spans all trading days and 278 securities. We observe security identifiers (ISINs), the time, the side (buy/sell), the price, and the quantity of the trade. We also know whether an investor trades bilaterally or on the platform, and can identify whether the investor is institutional or retail as part of the reporting.

A unique feature of the data is that all traders have a unique identifiers. Dealers, brokers and the largest investors carry a legal identifier (LEI). Other investors have an anonymous ID. Besides the ID, we observe no other characteristics of an institution, such as balance

[5]The fee depends on how many traders use CanDeal and whether the institution pays for the basic or full package. A typical institutional investor has 3-4 people working on the trading desk, and pays C$ 725 per trader, plus C$ 350 for admin-rights, and compliance for the basic package, or C$ 2,895, plus C$ 140 per trader for the full package, both per month.

[6]CanDeal estimates that about 95% of trades go via RFQ auctions. In rare occasions a dealer and investor pair negotiates the trade bilaterally but executes it on the platform to benefit from the straight-through electronic processing.

[7]According to CanDeal's annual reports available at `https://www.candeal.com/en/news`, on average 58% of auctions are with four, 30% with three and 12% with two dealers.

sheet information.

Similar to the TRACE data set, the MTRS2.0 data are self-reported and requires cleaning (see Appendix A). The cleaned sample includes almost all (cash) trades of Canadian government bonds but misses trades between investors which are unreported but rare according to market experts.[8] To get a sense of how many trades our sample misses because of this or due to misreporting, we compare the daily trading volume of Treasury Bills in MTRS2.0 with the full volume, which must be reported to the Canadian Depository for Securities. Our data cover approximately 90% of all trades involving Treasury Bills.

**Additional data sources.** We augment our trade-level data with three additional data sources. First, we obtain bidding data on all government bond auctions between 2016 and July 2019 from the Bank of Canada. We can see who bids (identified by LEI) and all winning and losing bids. Importantly, we can link how much a dealer won in the primary market to how she trades in the OTC market, which we use to construct an instrument in our demand estimation.

Second, we scrape ownership information from the public registrar of LEIs (`gleif.org`). This tells us whether a counterparty LEI is a subsidiary of a dealer so that we can exclude in-house trades, i.e., trades between a dealer and one of its subsidiaries.

Lastly, we collect indicative quotes that dealers post as advertisement. For one, we obtain bid and ask quotes that are posted on the platform (CanDeal) to advertise platform prices for each security up to the millisecond; we do not observe which dealer posts which quote, or which quotes dealers provide when participating in an RFQ auction. Further, we obtain the mid-quote (the average between the bid and the ask quote) that is posted on Bloomberg for each security per hour. We use this mid-quote as a proxy for a bond's true market value which is commonly known by everyone. We believe that this is a reasonable assumption, because the Bloomberg mid-quote is very close to the price at which dealers trade with one another, which is often taken as the true value of a security in the related theory literature. We do not use the inter-dealer price directly because the inter-dealer market isn't sufficiently liquid to generate a price for all securities in the dealer-investor market.

**Sample restrictions.** We exclude in-house trades because they are likely driven by factors that differ from those of a regular trade; for instance, tax motives or distributing assets

[8]In line with this, participation on all-to-all platforms—on which investors can directly trade with one another—remains low in markets in which these platforms already exist (Bessembinder et al. (2020)). One example is the U.S. government bond market, as was discussed at the 2020 U.S. Treasury Market Conference.

within an institution. In addition, we exclude trades that are outside of regular business hours (before 7:00 am and after 5:00 pm) because these trades are by foreign investors who might be treated differently.

For the estimation of our structural model, we impose some additional restrictions in order to construct an instrument for quotes using bidding data on the primary auctions (summarized in Appendix Table A1). We focus on primary dealers and drop trades after July 2019 because this is when our auction data ends. Due to data reporting, we exclude one dealer. Further, we exclude trades that are realized before the outcome of a primary auction was announced—10:30 am for bill auctions and 12:00 am for bond auctions.

**Unit of measurement.** The bonds in our sample differ with respect to when they mature and how much interest they pay until then (if any) in form of coupons. This implies that two different bonds may be traded at different prices, even though they bring comparable investment returns. To make different bonds more comparable, we convert each price into the yield-to-maturity (the annualized interest rate that equates the price with the present discount value of the bond) and report our findings in terms of yields rather than prices; a higher price implies a lower yield, and vice versa. All yields are expressed in bps; 1 bps is 0.01%.

**Key market features.** The typical trade between a dealer and one of the 572,760 investor IDs is small (see Appendix Figure A2). It involves a bond that is actively traded, because it was issued in a primary auction less than three months prior. Bid-ask spreads are narrow (0.5 bps at the median), and it takes only 0.11 (2.5) minutes between an investor who buys and an investor who sells (the same security from some dealer) on a day. Thus, the market is highly liquid. A yield difference of 1 bps is relatively large because of the low interest rate level throughout our sample—the median yield of a bond is about 150 bps.

# 4 Stylized facts

We document a series of stylized facts to motivate our structural model and show the variation in the data with which we can identify the main model parameters. The first fact is about dealer-investor relationships, which have been documented to also play a role in other market settings (e.g., Di Maggio et al. (2017); Hendershott et al. (2020a)).

**Fact 1** (Relationships)**.** *Most investors trade with a single dealer (see Figure 1). We call the dealer with whom the investor trades most frequently the home dealer.*

It could be that investors trade with the same dealer because this dealer offers better yields. To show that this is not the case, we regress

$$markup_{thsij} = \begin{cases} y_{thsij} - \theta_{ths} & \text{when the investor buys} \\ \theta_{ths} - y_{thsij} & \text{when the investor sells,} \end{cases} \tag{1}$$

on an indicator variable, $homedealer_{ij}$, that assumes value one if dealer $j$ is investor $i$'s home dealer, and investor fixed effects, $\zeta_i$. If the home dealer offered better yields, the $homedealer_{ij}$-coefficient would be positive by construction of $markup_{thsij}$. Instead, we cannot reject that the coefficient is zero (see Table 1).

It could also be costly for the investor to search for alternative dealers. However, the fact that investors still tend to trade with their home dealer on the platform, on which search frictions are minimal, suggests otherwise.

Therefore, we conjecture that it is valuable to trade with the home dealer because she offers "benefits" to her loyal clients. For example, the home dealer is frequently a custodian bank, which means that they hold their client's securities, manage financial accounts, and provide tax expertise. This could be beneficial for the client, but it could also just imply large switching costs, given that it requires resources to set up the infrastructure with another bank. The home dealer might also provide preferable discounts on other financial products, or serve only loyal clients in times of distress, and thereby grant insurance.

The second set of facts are about the degree and nature of dealer competition on the platform: Facts 2 (*i*) − 2 (*iii*) suggest that dealers need to compete for investors on the platform, and Fact 2 (*iv*) helps us understand when this competition takes place.

**Fact 2** (Platform competition)**.**
*(i) On the platform, an investor is more likely to trade with multiple dealers. (ii) On the platform, investors obtain better yields. (iii) Having platform access improves yields for investors. (iv) Platform yields are almost identical to the indicative quotes that dealers post on the platform.*

It is difficult to show Fact 2 (*i*) according to which an investor trades with more dealers on than off the platform, because platform participation is endogenous. Investors are less likely to choose the platform, so that there are fewer occasions to trade with different dealers—some investors only trade a single time. To control for the bias, we need an instrument.

Our idea is to leverage the benchmark status of a bond, which is achieved automatically once a certain amount of a bond has been issued (Berger-Soucy et al. (2018)). Benchmark bonds are more liquid than non-benchmark bonds, and existing studies show that liquid bonds tend to be more actively traded on platforms (e.g., Hendershott and Madhavan (2015)). At the same time, we wouldn't expect investors to trade benchmark bonds with more dealers than non-benchmark bonds once we control for how often the investor trades a bond. If anything, higher bond liquidity would decrease the number of trading partners in a sequential search model.[9]

To implement this idea, we regress the number of dealers an investor trades with in a month, $N\_dealers_{mi}$, on the fraction of all trades with security $s$ that the investor makes on the platform, $platform\_probability_{msi}$, in addition to year-month and investor fixed effects, $\zeta_m$, $\zeta_i$:

$$N\_dealers_{mi} = \alpha + \beta platform\_probability_{msi} + \zeta_m + \zeta_i + \epsilon_{msi}. \tag{2}$$

We instrument $platform\_probability_{msi}$ by whether the security has benchmark status in month $m$ or not, $benchmark_{ms}$. The estimates reported in Table 2 suggest that an investor trades with 1 additional dealer when her probability to trade on the platform increases by $1/0.348 \approx 2.8$ pp. The effect is bigger when we use the number of dealers an investor trades security $s$ with, $N\_dealers_{msi}$, as the dependent variable.

To derive Fact 2 (*ii*), we regress the markup defined in (1) on a platform indicator, $platform_{thsij}$. Even after including investor, day-hour, security and dealer fixed effects, and controlling for trade size, we find that the markup is higher on than off the platform (see Table 3). The difference is 0.12 bps, or roughly 20% of the bid-ask spread.

To show Fact 2 (*iii*), that an investor obtains better yields with platform access, we leverage the fact that during our sample 90 institutions lost their institutional status and therefore the right to trade on the platform. This happens, for instance, when investors no longer hold sufficient capital or are no longer willing to prove that they do. We regress $markup_{thsij}$ defined in (1) on an indicator variable equal to 1 $m$ months before/after $i$ loses access investor, $D_{mi}$, as well as hour-day, security and dealer fixed effects:

$$markup_{thsij} = \alpha + \sum_{m=Mi^-}^{Mi^+} \beta_m D_{mi} + \zeta_i + \zeta_{th} + \zeta_s + \zeta_j + \epsilon_{thsij}. \tag{3}$$

We let the fixed effects be pinned down by trade information on retail investors who never obtain access, since these investors are more similar to those who lost access than to (large) institutional investors. We find that investors who lose platform access realize worse yields (see Figure 2b). On average, the yield drops (per trade) by 1 bps in the first month and

[9] To see this, assume that each time an investor seeks to trade a security she first contacts her home dealer and then other dealers until some dealer can satisfy her demand. When it becomes more likely that each dealer can satisfy her demand, the investor trades with fewer dealers; in the extreme only with her home dealer.

decreases further by about 4 bps thereafter.[10]

The evidence for Fact 2 (*iv*) is shown in Figure 3. We see that platform yields are very close to the quotes dealers post on the platform. This suggests that dealers compete for the investors already when posting quotes to attract investors, which motivates a key feature of our model.

The third and fourth fact establish that most investors don't use the platform, even though it appears more competitive than bilateral trading.

**Fact 3** (Platform cost)**.** *Despite better yields on the platform, only about 30%–40% of institutional investors realize trades on the platform on a typical day (see Figure 4).*

This pattern has also been documented in other markets (e.g., McPartland (2016)). It suggests that it is costly for investors to use the platform. One cost factor is the platform subscription fee. Another factor could be that RFQ auctions are not anonymous, and investors might be reluctant to reveal information about their trade positions or strategies with multiple dealers (Managed Funds Association (2015)). This is especially true when trading large amounts.

**Fact 4** (Trade size)**.** *Investors are less likely to trade large amounts on the platform.*

To show this, we follow O'Hara and Zhou (2021) and split trades into categories, ranging from micro- ($<$ C\$ 100,000) to block trades ($\geqslant$ C\$ 25 million). We regress an indicator variable that tells us whether the trade is on the platform, $platform_{thsij}$, on trade size, $q_{thsij}$, for five trade-size categories, $\{(0, 0.1),[0.1, 1),[1, 5),[5, 25),[25, \infty)\}$, and investor fixed effects, $\zeta_i$. To show that the benchmark status no longer correlates with platform participation when controlling for trade size, we also include an indicator variable that is equal to 1 if the bond $s$ has benchmark status on day $t$, $benchmark_{ts}$:

$$platform_{thsij} = \alpha + \sum_{k=1}^{5} \beta_k \mathbb{I}(category = k) q_{thsij} + \beta_6 benchmark_{ts} + \zeta_i + \epsilon_{thsij}. \tag{4}$$

The estimates reported in Table 4 suggest that trade size affects the investor's decision to enter the platform, but only for trades larger than C\$ 25 million. This could be because

[10] This relationship would be causal if losing access to the platform were exogenous. This would be the case, for example, if an investor's regulatory capital fell marginally below the regulatory threshold due to shocks unrelated to trading demand. In this case, we would expect a change in investor-status but not a change in trading behavior. Alternatively, an investor might lose their status for potentially endogenous reasons, such as a change in business model or financial distress. Then we would expect to see changes in trading behavior. In Appendix Figure A5 we show that investors who lose access do not systematically change trading behavior after they lose access—providing us confidence that losing access is likely exogenous.

investors want to hide large trades, or because the indicative quotes that dealers post publicly are only valid for trades smaller than C$ 25 million.

**Take away.** Summarizing, Fact 1 highlights that dealer-investor relationships are important. Fact 2 suggests that dealers compete more aggressively for the investor on than off the platform and that at least some of this competition happens already when posting quotes. Fact 3 tells us that many investors stay off the platform, implying that it must be costly to trade on the platform. Fact 4 suggests that this cost is bigger for large block trades.

These facts motivate the main model parameters that we want to estimate: a loyalty benefit that an investor achieves whenever she trades with her home dealer, the degree of platform competition, and the cost of trading on the platform. In our benchmark model, we abstract from differences in trade sizes. In Appendix B.1 we present an extension in which traders either trade regular-sized trades or block trades to see whether our model correctly predicts that these traders face a higher cost of trading on the platform.

# 5 Benchmark model

Without loss of generality, we consider two separate games: one in which dealers sell to investors and one in which dealers buy from investors.[11] We explain the setting with buying investors; the other side is analogous. We use yields rather than prices in line with the rest of the paper and estimation. We denote a vector of quotes by $q_t = (q_{t1}...q_{tJ})$ and similarly for all other variables. All proofs are in Appendix C.

**Dealers and investors.** In period $t$, $J_t \geqslant 2$ dealers sell a bond to investors, bilaterally or on a platform. Each transaction is a single unit trade. The market (or fundamental) value of the bond is $\theta_t \in \mathbb{R}^+$. It is commonly known and exogenous, capturing macroeconomic factors that affect interest rates.

Each investor $i$ has a home dealer $d$, short for $d_i$. Each dealer, thus, has a home investor base. It consists of two investor groups, institutional and retail investors, indexed by $G \in \{I, R\}$. Each has a commonly known mass $\kappa^G$ of potential investors. W.l.o.g., we normalize $\kappa^I + \kappa^R = 1$. Of the potential investors in group $G$, $N_t^G$ investors actually seek to buy. This number is exogenous and unknown to the dealer until the end of the period.

[11]To see why considering separate games is without loss of generality, assume that investors may either buy or sell, but that the dealer does not know whether the investor is a buyer or seller. The dealer offers a bid and ask yield, such that the bid is optimal conditional on the investor's being a seller and the ask is optimal conditional on the investor's being a buyer. The ask (bid) yield is identical to the ask (bid) yield in our model with only buying (selling) investors.

Each dealer seeks to maximize profit from trading with investors. Ex post, dealer $j$ obtains profit $v_{tj}^D - y$ when selling one unit at yield $y$; $v_{tj}^D \in \mathbb{R}$ is the dealer's value for the bond. It may be driven by current market conditions, expectations about future demand, prices or inventory or market making costs. Counterfactually, if the market was frictionless and dealers neither derived value from holding bonds nor paid any costs for intermediating trades, $v_{tj}^D$ would equal the market value, $\theta_t$. In our estimation, we rely on cost-shifter instruments, which unexpectedly shift dealer values without directly affecting investor demand. Thus, strictly speaking, $v_{tj}^D$ is a function of dealer $j$'s cost-shifter.

An investor $i \in G$ obtains a surplus of $y - v_{tij}^G$ when buying from dealer $j$ at yield $y$, where

$$v_{tij}^G = \theta_t + \iota_{ti}^G - \Xi_{tij} \text{ with } \iota_{ti}^G \overset{iid}{\sim} \mathcal{F}_t^G \text{ and } \Xi_{tij} \sim \mathcal{G}_t$$

is the investor's value, also referred to as willingness to pay. It splits into four elements. The first is the commonly known market value of the bond, $\theta_t$.

The second is a liquidity shock, $\iota_{ti}^G$, which is drawn iid from a commonly known distribution with a continuous CDF $\mathcal{F}_t^G(\cdot)$ that has a strictly positive density on the support. It reflects individual hedging or trading strategies, balance-sheet concerns, or the cash needs of an institution.

The third element, $\Xi_{tij}$, is dealer-investor-specific, and depends on whether the dealer is the investor's home dealer or not:

$$\Xi_{tij} = \begin{cases} \xi_{tj} + r & \text{if the dealer is the investor } i\text{'s home dealer } (j = d) \\ \xi_{tj} & \text{if the dealer is not the investor } i\text{'s home dealer } (j \neq d). \end{cases}$$

We label $\xi_{tj}$ the dealer's base quality and $r \in \mathbb{R}^+$ the loyalty benefit. Both capture unobservable characteristics that makes trading with a specific dealer particularly attractive, independent of how the trade is realized—for instance, a high probability of delivery, immediacy, or ancillary services (such as offering investment advice on a broad range of securities). The key difference between the two is that the loyalty benefit captures services that are provided only to customers with whom the dealer has a close business relationship. For simplicity it is the same across dealers.

**Timing of events.** The game has two stages. First, dealers simultaneously post quotes at which they are willing to sell on the platform. Second, all investors draw their liquidity shocks. Institutional investors trade bilaterally or on the platform; retail investors can only trade bilaterally. In a bilateral trade, the dealer offers the yield that leaves the investor with the loyalty benefit. On the platform, the investor can buy at the posted quote from any of the dealers but has to pay a commonly known cost $c_t$ to enter.

To capture platform frictions, each investor $i$ draws an iid shock $\epsilon_{tij}$ from distribution $\mathcal{H}$ upon entry to the platform, which distorts the investor's behavior. It randomly increases her willingness to trade with dealer $j$ in period $t$ by shifting her utility by $\sigma\epsilon_{tij}$, where $\sigma \in \mathbb{R}^+$ measures the severeness of the friction. Inspired by Appendix B, where we model an RFQ auction, we refer to the platform friction as imperfect competition. Thus, we say that the platform becomes less competitive when $\sigma$ increases. In the extreme, when $\sigma \to \infty$, the investor chooses from which dealer to buy independently of the yield this dealer offers. When $\sigma = 0$, dealers compete à la Bertrand.[12]

In summary, the sequence of events is:

(1) Everyone observes the market value $\theta_t$, dealer qualities $\xi_t$, and the loyalty benefit $r$. Dealers observe their values $v_t^D$, and simultaneously post quotes $q_{tj} \in \mathbb{R}$.

(2) $N_t^G$ investors of both groups $G \in \{I, R\}$ each draw a liquidity shock $\iota_{ti}^G$.
Each investor contacts her home dealer $d$, who offers $y_{tid}^G = v_{tid}^G + r$.
A retail investor accepts the offer. An institutional investor can either accept or enter the platform. In the latter case, the investor pays cost $c_t$, observes the platform shock $\epsilon_{tij}$, and decides from which dealer to buy at $q_{tj}$.

A pure-strategy equilibrium can be derived by backward induction.

**Proposition 1.**
*(i) A retail investor with shock $\iota_{ti}^R$ buys bilaterally from home dealer $d$ at*

$$y_{tid}^R = \theta_t + \iota_{ti}^R - \xi_{td}. \tag{5}$$

*(ii) An institutional investor with shock $\iota_{ti}^I$ buys bilaterally from home dealer $d$ at*

$$y_{tid}^I = \theta_t + \iota_{ti}^I - \xi_{td} \quad \textit{if } \psi_{td}(q_t) \leqslant \iota_{ti}^I, \tag{6}$$

$$\textit{where } \psi_{td}(q_t) = \mathbb{E}[\max_{k \in \mathcal{J}_t} \tilde{u}_{tik}(\epsilon_{tik})] - \theta_t - c_t - r \tag{7}$$

$$\textit{with } \tilde{u}_{tij}(\epsilon_{tij}) = \xi_{tj} + \mathbb{I}(j = d)r + q_{tj} + \sigma\epsilon_{tij}. \tag{8}$$

*Otherwise, the investor enters the platform, where she observes $\epsilon_{tij}$ and buys from the dealer with the maximal $\tilde{u}_{tij}(\epsilon_{tij})$ at $q_{tj}$.*

[12] More broadly, in our estimation, $\sigma\epsilon_{tij}$ captures any idiosyncratic friction that prevents an investor from buying from the best dealer with the highest posted quote. This includes dealer inattention, meaning that the dealer does not actively monitor the platform when the investor requests a quote (similar to Liu et al. (2018)).

Proposition 1 characterizes where investors buy and at what yields. A retail investor always buys at a yield equal to her willingness to pay, excluding the loyalty benefit. An institutional investor trades on the platform if she expects that the surplus from buying on the platform minus the platform usage cost, $\mathbb{E}[\max_{k\in\mathcal{J}_t}\tilde{u}_{tik}(\epsilon_{tik})-(\theta_t+\iota^I_{ti})]-c_t$, will be higher than the loyalty benefit she receives in a bilateral trade. This is the case for urgent investors who are willing to pay a higher price, i.e., accept a low yield, due to a low liquidity shock. For them it is better to trade on the platform, because the platform quote is targeted to an investor with an average willingness to pay rather than to the investor's individual willingness to pay.

**Proposition 2.** *Dealer $j$ posts a quote $q_{tj}$ that satisfies*

$$q_{tj}\Big(1+\frac{1}{\eta^E_{tj}(q_t)}\Big(1-\frac{\partial\pi^D_{tj}(q_t)}{\partial q_{tj}}\Big/S_{tj}(q_t)\Big)\Big)=v^D_{tj}, \tag{9}$$

*where $\eta^E_{tj}(q_t)$ is the dealer's yield elasticity of demand on the platform and $\frac{\partial\pi^D_{tj}(q_t)}{\partial q_{tj}}$ is the marginal profit the dealer expects from bilateral trades with institutional investors. It is normalized by the size of the dealer's platform market share, $S_{tj}(q_t)$. When $\epsilon_{tij}$ are extreme value type 1 distributed, $\eta^E_{tj}(q_t)=q_{tj}\frac{\partial S_{tj}(q_t)}{\partial q_{tj}}/S_{tj}(q_t)$ where $S_{tj}(q_t)=S^j_{tj}(q_t)+\sum_{k\neq j}S^k_{tj}(q_t)$ with $S^l_{tj}=s^l_{tj}(q_t)\Pr(\iota^I_{ti}\leqslant\psi_{tl}(q_t))$ where $s^l_{tj}(q_t)=\frac{\exp(\frac{1}{\sigma}(q_{tj}+\xi_{tj}+\mathbb{I}(j=l)r))}{\sum_k\exp(\frac{1}{\sigma}(q_{tk}+\xi_{tk}+\mathbb{I}(k=l)r))}\forall l\in J_t$ and $\pi^D_{tj}(q_t)=\mathbb{E}[v^D_{tj}-(\iota^I_{ti}+\theta_t-\xi_{tj})|\psi_{tj}(q_t)\leqslant\iota^I_{ti}]$.*

Proposition 2 characterizes the quotes dealers post on the platform. Taking the quotes of the other dealers as given, each dealer chooses a quote that equals a fraction of her value, $v^D_{tj}$. To obtain an intuition regarding what determines the size of this fraction, it helps to abstract from the bilateral segment for a moment.

If the market only consisted of the platform, the dealer's quote would satisfy the classic markup rule of profit-maximizing firms: $q_{tj}(1+1/\eta^E_{tj}(q_t))=v^D_{tj}$. The price would equal the marginal cost multiplied by a markup, which depends on the price elasticity of demand, $\eta^E_{tj}(q_t)$. In our setting, the marginal cost is $v^D_{tj}$, and since the dealer chooses quotes in yields rather than prices, the markup is actually a discount.

When the market splits into the platform and a bilateral segment, there is an additional term. It captures the fact that a quote also affects how much profit the dealer expects to earn from bilateral trades, given that investors select where to buy based on these quotes. If the dealer decreases the quote, more investors buy bilaterally because they earn a higher yield there; how many depends on the cross-market (segment) elasticity between bilateral and platform trading. If this elasticity is high, investors easily switch onto the platform. To prevent this from happening, the dealer decreases the quote to make the platform less attractive.

In summary, when choosing the quote the dealer trades off the profit from selling bilaterally, where she extracts a higher trade surplus, with the profit she earns on the platform when stealing investors from other dealers.

**Discussion.** Our model builds on several simplifying assumptions. First, we assume that the number of dealers and investors who trade in a period is exogenous and that no trade between them fails. This is motivated by empirical evidence that suggests that trades of safe assets rarely fail (e.g., Riggs et al. (2020); Hendershott et al. (2020b)); and the fact that (primary) dealers have an obligation to actively trade—the least active dealer trades on 98% of dates.

Second, our game does not connect multiple periods. In particular, we assume that dealers' and investors' values for the bond are independent of prior trades. This implies that we set aside dynamic trading strategies. Dealers and investors can still trade every period and their values can capture continuation values, which may vary in time. However, when changing the market rules, we cannot account for changes in their continuation values.

Third, we abstract from an inter-dealer market. This is because dealers primarily trade with investors rather than dealers (see Appendix Figure A3) and because dealers do not sizably re-balance their inventory positions by trading with one another or brokers (see Appendix Table A2).[13] To validate that this is a good approximation, we let dealer valuations converge towards the hypothetical value $\theta_t$ that would arise with a perfectly frictionless inter-dealer market as part of our robustness analyzes in Appendix E.

Fourth, we assume that the home dealer chooses the bilateral yield that lets the investor earn the loyalty benefit, motivated by Fact 1 and Table 1. This implies that investors have no incentive to search for other dealers when trading bilaterally.[14] An alternative would be to allow the dealer to be strategic and set the bilateral yield so that an institutional investor is indifferent between trading bilaterally or entering the platform. Then, in contrast

[13]One reason for this is that investor demand or supply is likely correlated across dealers. This implies that there are some periods in which all dealers seek to sell more than they seek to buy in the inter-dealer market, and vice versa for other periods. Another reason could be that there are frictions, such as balance sheet constraints, that hinder dealers from absorbing inventory from each other.

[14]The assumption implies that dealers do not adjust the yields they charge in bilateral trades depending on how costly it is for them to realize the trade. We test whether this implication holds in our data and find supporting evidence (see Appendix Table A3). From an econometric perspective, you may think of this as a normalization similar to the way we normalize the utility of one choice in discrete choice models.

to the data, all investors who trade with the same dealer in a period would obtain the same bilateral yield unless they faced different platform costs, earned different loyalty benefits, or had different expectations over platform conditions. In our model, yield dispersion arises more naturally because investor's have different liquidity demands, i.e., shocks.

Fifth, motivated by Fact 2, we assume that on the platform investors trade at the posted quotes. Thus, dealers do not revise their quotes when learning information about the investor, perhaps to ensure that investors believe quotes to be informative when deciding where to trade. For the estimation, this implies that we can use the average platform trade price—which we observe—for the quotes that dealers post—which we do not observe. This is a common approach in the industrial organization literature whenever we do not observe information about the way prices are formed, such as the bargaining process, or in our case the RFQ auction. In settings in which the indicative quotes differ from the platform trade prices, we recommend estimating the model variation presented in Appendix B.2.

Sixth, we assume that there is a single bond because all actively traded bonds in our sample are relatively similar in terms of liquidity and risk.[15] Further, we have seen that the liquidity (i.e., benchmark) status of a bond does not determine the investor's decision to trade on the platform when controlling for trade size (recall Table 4).

Lastly, we abstract from order splitting by assuming that investors trade either bilaterally or on the platform but not both since we observe that investors typically maximally trade once per day or even week and do not split orders (see Appendix Figure A4).

# 6 Estimation

Including both sides of the market, we have four investor groups, indexed by $G$: retail and institutional investors who buy ($R$ and $I$) and sell ($R^*$ and $I^*$). We want to estimate the distribution of the liquidity shocks of each group ($\mathcal{F}_t^G$), the dealers' base qualities ($\xi_{tj}$), the loyalty benefit ($r$), the degree of competition on the platform ($\sigma$), the cost of using the platform, and the dealer's value, which may depend on whether the investor buys ($c_t$, $v_{tj}^D$) or sells ($c_t^*$, $v_{tj}^{*D}$).

Most of the parameters are period-specific to non-parametrically account for variation and correlation across periods that are driven by unobservable trends. This is important because trading and pricing in the Canadian government bond market is largely affected by

[15]To apply our framework to other OTC markets, in which assets are more heterogeneous, we recommend extending the model to include more than one asset, so as to take asset-specific factors that determine trade choices into account. Kozora et al. (2020) determine such factors for the U.S. corporate bond market.

global macroeconomic trends. In our benchmark model with loyalty benefits a period is a week. As robustness, we estimate the model with and without loyalty benefits with a period being a day in Appendix E. Note that in neither case do we have enough data to allow for heterogeneous dealer quality or platform costs, for instance, by including investor fixed effects.

## 6.1 Identifying assumptions

Our estimation builds on three identifying assumptions, a normalization, and two parametric assumptions that are not crucial for identification.

**Assumption 1.** *Within a period $t$, the liquidity shocks $\iota_{ti}^G$ are iid across investors $i$ in the same group $G$.*

This assumption would be violated if an investor trades more than once in a period and jointly decides whether and at what price to trade for all such trades. However, given that we observe very few investors who trade several times within the same period, this is unlikely (recall Appendix Figure A4).

Dealer qualities $\xi_{tj}$ may be correlated across dealers $j$. W.l.o.g. we decompose $\xi_{tj}$ into a part that is persistent over time and a part that might vary: $\xi_{tj} = \xi_j + \chi_{tj}$. Further, dealers may change quotes in response to demand shocks that are unobservable to the econometrician. To eliminate the implied endogeneity bias, we need an instrument for the quotes. Our solution is to extract unexpected supply shocks, $won_{\tilde{t}j}$, from bidding data in the primary auctions in which the government sells bonds to dealers:

$$\begin{aligned} won_{\tilde{t}j} = &\text{ amount dealer } j \text{ won in the last auction } \tilde{t} \\ &- \text{amount she expected to win when placing her bids.} \end{aligned} \tag{10}$$

These shocks work as cost-shifter instruments, shifting dealer values, because it is cheaper for dealers to satisfy investor demand when unexpectedly winning a lot at auction, given that auction prices are systematically lower than prices in the OTC market.[16]

Importantly, we use the expected rather than the actual winning amounts, since dealers anticipate or even know investor demand when bidding in the auction—for example, because investors place orders before and during the auction (as in Hortaçsu and Kastl (2012)). This information affects how dealers bid and, consequently, how much they win, which creates

[16]In reality, it is unlikely that dealers observe the cost-shifters of other dealers. Therefore, our model, which assumes that dealers know each others' values, is only an approximation.

a correlation between the unobservable demand shocks and the actual, but not expected, winning amounts.

To compute the expected amount and control for anything the dealer knows at the moment she places her bids, we model the bidding process in the auction and use techniques from the empirical literature on multi-unit auctions. In a nutshell, we fix a dealer in an auction, randomly draw bids (with replacement) from the other bidders, and let the market clear. This generates one realization of how much the dealer wins. Repeating this many times generates the empirical distribution of winning amounts, from which we compute the expectation (see Appendix D.2 for details).

**Assumption 2.** *Conditional on unobservables that drive aggregate demand and supply in period $t$, $\zeta_t$, and the time-invariant quality of the dealer, $\xi_j$, the demand shocks, $\chi_{tj}$, are independent of the unexpected supply shocks, $won_{\tilde{t}j}$:* $\mathbb{E}[\chi_{tj}|won_{\tilde{t}j}, \zeta_t, \xi_j] = 0$.

To better understand whether this assumption is plausible, it helps to think through where the surprise—and, with that, the identifying variation—comes from. For one, the dealer is surprised when the Bank of Canada issues a different amount to bidders than the dealer expected. However, the period fixed effect absorbs most of this effect. What is left is the surprise the dealer faces when other bidders bid differently than the dealer expected.

With this in mind, the biggest threat to identification is the following scenario: One dealer is hit by a negative shock and bids less, so that the other dealers win more than expected. If investors substitute from the unlucky dealer toward those who won more, the exclusion restriction would be violated. However, in our data, we see relatively little substitution of investors across dealers. Therefore, we are less worried that this is a first-order concern.

The exclusion restriction would also be violated if the dealer changed her quality based on how much she won at auction. We believe that this is unlikely to happen (often) for at least two reasons. First, dealers have incentives to smooth out irregular shocks to maintain their reputation and their business relationships with investors in the longer run. Second, dealers would risk revealing information about their current inventory positions if they changed the service they provide based on how much they win at auction.

**Assumption 3.** *Platform shocks $\epsilon_{tij}$ are iid across $t, j, i$.*

This assumption is less stringent than it might sound because we allow investor values (and hence liquidity demand) to be correlated across dealers by allowing dealer base qualities to be correlated. If dealer qualities were independent from one another, Assumption 3 would restrict how investors substitute across dealers. In particular, if the base quality of dealer

$j$ dropped, but the base qualities of the other dealers remained unchanged, investors would substitute from dealer $j$ to any of the other dealers with equal likelihood. However, with correlated dealer qualities $\xi_{tj}$, this scenario would not arise.

We normalize the quality of one dealer to 0, because—as is common in demand estimation—we cannot identify the size of the dealers' qualities but we can estimate the quality differences between dealers.

**Normalization 1.** *The benchmark dealer ($j = 0$) provides zero quality: $\xi_{t0} = 0\ \forall t$.*

Finally, we impose a standard functional form on the distribution of $\epsilon_{tij}$ to obtain market shares in closed-form solutions, and let liquidity shocks be normally distributed. The latter assumption is inspired by the shape of the histogram of shocks, $\hat{\iota}^G_{ti} = y^G_{tij} - \theta_t + \hat{\xi}_{tj}$, for investors who choose to trade bilaterally. It resembles a normal distribution, similar to Figure 8.

**Parametric Assumptions.**
*(i) Platform frictions $\epsilon_{tij}$ are extreme value type 1 (EV1) distributed.*
*(ii) Liquidity shocks $\iota^G_{ti}$ are drawn from a normal distribution $N(\mu^G_t, \sigma^G_t)$ for all $g, t$.*

Without imposing a distributional assumption on the liquidity shocks, we could nonparametrically estimate the (truncated) distribution of the liquidity shocks of investors who trade bilaterally—as explained in more detail below—as well as bounds on the cost parameters.

## 6.2 Identifying variation

We estimate the model separately for each investor group. Here, we focus on buying institutional investors and leave the other groups for Appendix D, which explains the estimation procedure in detail.

**Key variables.** The main variables used in the estimation are dealer market shares and yields. We use each dealer $j$'s period-specific market share on the platform of investors who enter the platform, $s_{tj}$. This share decomposes into what is bought by home investors, $s^j_{tj}$, and what is bought by home investors of other dealers $k \neq j$, $s^k_{tj}$. In addition, we use each dealer's period-specific bilateral market share relative to her platform market share, $\rho_{tj}$.

To be conservative, we control for factors that might affect the yield in reality but which are not endogenous in our model. We do this by constructing the residuals from regressing the yield of a trade on day $t$ in hour $h$ of security $s$ between dealer $j$ and investor $i$ on an indicator variable that separates trades in which the investor buys from trades in which she sells, an hour-day fixed effect, and a security-week fixed effect. In addition, we normalize the Bloomberg yield by subtracting the estimated hour-day and security-week fixed effects. We use this normalized Bloomberg yield, averaged across securities and hours in period $t$, to approximate the bond's market value observable to everyone, $\theta_t$.[17]

We measure the platform quote at which dealer $j$ sells in period $t$, $q_{tj}$, using the average yield at which she sells on the platform in that period—motivated by Fact 2. Thus, we construct a dealer-specific platform quote from dealer-specific platform trade yields to work around the data limitation that we don't observe dealer-specific quotes.

**Identification.** The main identifying variation for the competition parameter, dealers' qualities and the loyalty benefit comes from how dealers split the platform market in a period.

The competition parameter $\sigma$ is mainly identified from the within-week correlation between dealers' weekly platform market shares and their (cost-shifter) supply shocks, as shown in Figure 6. To derive an intuition for this, assume for a moment that dealers do not differ in quality ($\Xi_{tj} = 0 \; \forall j$). If the platform is perfectly competitive ($\sigma = 0$), a single dealer—namely the one with the most favorable supply shock and with it the best quote $q_{tj}$—captures the entire platform market share in that week. As $\sigma$ increases, this dealer loses more and more of her market share to the other dealers. How much of the market share each dealer gains depends (besides $\sigma$) on the dealers' supply shocks. Hence, the correlation between these market shares and the supply shocks pins down $\sigma$.

The loyalty benefit $r$ is identified from how many home investors buy from their home dealer relative to how many buy from another dealer in a week, so for a fixed set of quotes.

The dealers' base qualities $\xi_{tj}$ are determined by how the dealers split the platform market when posting similar quotes, conditional on how many home investors of each dealer enter the platform and the size of the loyalty benefit. Dealers with higher qualities capture a higher market share. Even without conditioning on platform prices and the loyalty benefit, we see how the size of the market share and dealer qualities are related (see Figure 7). For instance, the median quality and market share of d6 is above d5 and d7, while dealer qualities and market shares d4 and d8 are towards the bottom.

[17]An alternative would be to approximate the market value by the average trade yield, which is identical to the Bloomberg mid-quote for institutional investors (recall Figure 2a). The main difference would be that the market value would then be endogenous, and thus change when we change the market structure in our counterfactual analyses. We think that this is unlikely to happen in our setting since the market value of a Canadian government bond is strongly determined by macroeconomic factors, including global trends. The Bloomberg mid-quote captures this exogenous variation.

The distribution of the liquidity shocks and the platform usage costs are, for any given week, mainly identified from how bilateral yields vary across investors and how many investors choose to trade bilaterally rather than on the platform. This is illustrated in Figure 8. It shows the distribution of yields that institutional buyers realize and a black line. Investors who draw liquidity shocks that would imply a bilateral yield that lies below the line buy on the platform, according to Proposition 1. Therefore, the position of the line—and, with it, the size of $c_t$—is determined by the fraction of investors who buy bilaterally rather than on the platform. Further, the shape of the yields' distribution above the black line pins down the distribution of the liquidity shocks.

Finally, we back out the dealer's values $v_{tj}^D$ from the markup equation (9) of Proposition 2. We pick the $v_{tj}^D$ for which the equation holds, given all the estimated parameters. This is similar to a classic approach adopted in industrial organization to infer the marginal costs of firms from firm behavior.

# 7 Estimation results

We estimate the benchmark model and a model extension that distinguishes between small and large trades. An overview of all estimates is in Table 5. In Appendix E, we report information about the instrument and verify that the estimates are biased in the expected direction when we do not instrument the quotes.

## 7.1 Benchmark model

**Investor values.** The amount an investor is willing to pay can be decomposed into her value for the bond (the liquidity shock), her value for the dealer's base quality (shown in Figure 7b), and the loyalty benefit she receives whenever trading with her home dealer. The split between these components relies on the normalization that the benchmark dealer provides zero quality. Therefore, we focus our discussion on differences across investor groups and their total values.

A buying retail investor is willing to pay about 2 bps more than institutional investors. When selling, the difference is smaller—about 1 bps—perhaps because retail investors who sell are more active than retail investors who only buy. The loyalty benefit is large—2.37 bps—compared to the investors' total values. For instance, an institutional investor is on average wiling to accept a yield that is 3.89 bps below the market value when buying, and 4.03 bps above it when selling.

**Dealer values.** A dealer is typically willing to sell at a lower price than market value and buy at a higher price than market value, indicating that liquidity provision is costly. Relative to the median value of an institutional investor, the dealer's value is roughly 1 bps closer to the market value on both sides of the trade.

Since different dealers attach different values to realizing trade in any given period (see Figure 5), there might be welfare gains in matching investors to dealers with higher values. We quantify these gains below.

**Platform competition.** Competition is relatively low for at least two reasons ($\hat{\sigma} = 0.56$). First, the differences in quality of different dealers are sizable compared to the average liquidity shock (see Figure 7b). In consequence, a dealer who offers the best quote only captures a fraction of the platform market share. Second, investors are still likely to buy from their home dealer on the platform. In line with this, $\hat{\sigma}$ is larger (1.46), meaning that competition is even lower, when we ignore the loyalty benefit.

Low platform competition implies that demand on the platform is rather inelastic. The average yield elasticity of demand on the platform is about 110–114. This means that a dealer's demand (supply) increases on average by 1.10% (1.14%) if this dealer increases (decreases) her quote by 1 bps. Thus, even if the dealer were willing to sell at a price at which she usually buys (which is about 0.5 bps higher), she would sell less than 0.6% more.[18]

**Platform usage cost.** We estimate a median platform usage cost of 1.72-1.81 bps. This is only slightly above the actual fee an investor has to pay to trade on the platform, which is approximately 1.1-1.5 bps.[19] In contrast, if we don't account for loyalty benefits, the median platform cost is over 4 bps, which highlights that loyalty benefits drive switching costs that prevent investors from using the platform.

## 7.2 Extended model

Inspired by Fact 4, we estimate an extended model, described in Appendix B.1, that distinguishes between regular-sized small ($<$ C\$ 25 million) and large ($\geqslant$ C\$ 25 million) trades for institutional investors. Retail investors essentially never trade large amounts.

[18]In comparison, Krishnamurthy and Vissing-Jørgensen (2012) estimate that the spread between corporate and government bond yields would increase by 1.5–4.25 bps if the debt/GDP ratio would rise by 2.5%. Our estimate implies a 1.65–4.67% increase in demand when the yield increases by 1.5–4.25 bps.

[19]The platform fee is between C\$ 2,500–3,500 per month, and since a typical institutional investor trades 22 million per month on the platform, the monthly fee implies a per-unit cost of 1.1-1.5 bps.

Our model more closely mimics the data generating process of small rather than large trades given that platform quotes are not valid for large trades. We nevertheless would like to see if our estimates move in the expected direction when separating by trade size.

Since investors are less likely to trade large amounts on the platform—perhaps, because they are averse to sharing trade information with multiple dealers—we expect to estimate larger platform costs for these trades despite the actual platform fee being lower per unit of trade. Further, given that trade size correlates with how active and connected an investor is to multiple dealers, the loyalty benefit should be lower. In addition, if large trades are executed by more sophisticated investors, they should be more similar to dealers than retail investors in terms of values. They might even switch more easily to different dealers on the platform.

Our estimates, reported in Table 5c, confirm all conjectures. We find that platform costs are about 1 bps higher for large trades. The demand elasticity on the platform is about 1.6 times higher than for small-trade investors, meaning that large-trade investors more easily switch across dealers on the platform. This is because they earn a lower loyalty benefit than small-trade investors. In addition, they draw less extreme liquidity shocks, making them more similar to dealers than retail investors. Platform competition appears slightly lower when accounting for differences in trade-size, yet the difference is not statistically significant.

### 7.3 Model fit

Before assessing welfare effects from centralizing the market, we validate whether our parsimonious model can replicate patterns in the data that will drive the welfare effects, most of which we haven't targeted explicitly in our estimation. To be conservative, we display findings using the benchmark model with all trades. The extended model with only small trades fits the data even better.

Inspired by Facts 1-3, we compare model predicted to observed investor trading behavior and yields in Figure 9. In addition, Appendix Figure A6 shows observed and predicted market concentration.

Our model correctly predicts that it is likely for an investor to trade with her home dealer. An institutional investor is predicted to trade with 90% probability with her home dealer. Conditional on entering the platform, the likelihood drops to 80%. In the data, institutional investors are slightly less likely to trade with their home dealer. This is because in the data some of the large investors with LEIs trade with multiple dealers bilaterally (recall Figure 1a), while our model assumes that all investors trade with one dealer bilaterally. On the platform, institutional investors are slightly more likely to trade with their home dealer in the data. They are also slightly more likely to enter the platform in the data. However, all differences are small.

In the data and our model, institutional investors obtain significantly better yields than retail investors. In addition, platform yields are slightly better than bilateral yields. Finally, when losing the right to trade on the platform, an investor obtains worse yields (recall Figure 2b). Our model predicts a similar drop in yields when we eliminate platform access in a counterfactual.

# 8 Counterfactual exercises

We now conduct counterfactuals to analyze the welfare effects from market centralization. For this, we use the extended model with small trades which fits the data generating process more closely than the model with large trades. As such, our focus is qualitative. Quantitively, all welfare findings are small relative to the total size of the market, not only because we focus on small trades, but also because we keep the number of trades and the trade volume fixed. This implies that we do not account for a potential growth in investor demand in response to more favorable market conditions. We present all findings for investors who buy, but our findings generalize to selling investors, since the buy and sell sides of the market are close to symmetrical.

Crucially, we take into account how dealers and investors respond to the changes in the market rules: as investors enter the platform, dealers adjust their quotes, which in turn affects the trading decisions of investors. A new equilibrium arises, in which all investors select onto the platform, as in (*ii*) of Proposition 1, and dealers set quotes that are valid for all investors according to Proposition 2. Doing so, dealers behave as if there were a representative investor who draws liquidity shocks from a normal distribution with mean $\mu_t = \kappa^R \mu_t^R + \kappa^I \mu_t^I$ and standard deviation $\sigma_t = \kappa^R \sigma_t^R + \kappa^I \sigma_t^I$ where $\kappa^R = 0.1$ is the fraction of trades by retail investors and $\kappa^I = 0.9$ those by institutional investors on an average day.[20]

**Types of frictions.** Before making welfare statements, it helps to think through what types of frictions there might be and which of them move around in our counterfactual analyses. First, frictions stem from poor platform design: platform access is limited, it is costly to use, and platform competition is imperfect. Second, frictions arise when not all investors and dealers consider trading with one another so that some welfare-improving trade

[20]We can show that these quotes are numerically equivalent to those that arise when dealers take into account that retail investors may more strongly select onto the platform than institutional investors.

matches cannot form. Lastly, there are frictions that arise from the fact that dealers must use their own inventory for the vast majority of trades (known as principal trading). Limited balance sheet capacity or an illiquid inter-dealer market lead dealers to have different values for trading.

In our counterfactuals, we fix the dealer's values (and with that the frictions driving heterogeneous dealer values) and change the frictions that arise due to imperfect platform access and design. Since this shifts more trades onto the centralized platform on which it is easy to access all dealers simultaneously and search costs are close to zero, this reduces search frictions.

To highlight how welfare changes as market centralization increases, we express all welfare gains or losses relative to a world in which all trades are executed bilaterally in a fully decentralized market. Parting from this, we first allow all institutional investors and then all investors to trade on the platform at the estimated platform usage costs. Then, we remove platform fees. Finally, we shift all trades on the platform. The first-best is achieved when all investors trade on a frictionless centralized platform—a useful theoretic benchmark that is likely hard to achieve in reality. An overview of welfare gains in all counterfactuals is presented in Table 6. The respective trade probabilities are in Appendix Table A4.

**When does welfare change?** We measure welfare by the total expected gains from trade. Therefore, welfare effects come entirely from re-matching who trades with whom. If all market participants have the same value for the bond, such re-matching would not affect welfare.

**Definition 1.** *The expected welfare is* $W_t = \sum_G \kappa^G W_t^G$*, where*

$$W_t^G = \sum_j \mathbb{E}[v_{tj}^D - v_{tj}^G(\iota_{ti}^G) | \textit{investor } i \in G \textit{ buys from dealer } j] \tag{11}$$

*is the expected welfare from trading with investors of group* $G \in \{I, R\}$ *with dealer value,* $v_{tj}^D$*, and investor value,* $v_{tj}^G(\iota_{ti}^G) = \theta_t + \iota_{ti}^G - \xi_{tj} - \mathbb{I}(j = d)r$*. Proposition 1 specifies which dealer the investor buys from.*

Welfare increases when investors are more likely to buy from more efficient dealers, i.e., dealers with higher $v_{tj}^D + \xi_{tj} + \mathbb{I}(j = d)r$. This happens when more investors buy from dealers with high values or higher base qualities, or from home dealers.

Due to market power it is not clear whether welfare increases when more investors trade on the platform. An investor chooses the dealer with the highest $q_{tj} + \sigma\epsilon_{tij} + \xi_{tj} + \mathbb{I}(j = d)r$ on the platform, but this dealer may not be the efficient one. This is because dealers compete imperfectly when posting quotes and charge a markup over their true values ($v_{tj}^D \neq q_{tj}$), and because the platform itself is imperfect ($\sigma \neq 0$).

**How important is market power?** To assess the importance of market power, we make the platform perfectly competitive ($v_{tj}^D = q_{tj}$ and $\sigma = 0$) at different degrees of market centralization.

We find that welfare decreases as more investors trade on the platform—unless the platform is perfectly competitive. The overall welfare effect is small, ranging between -3% and 2%, in all counterfactuals because two opposing effects almost entirely cancel each other out. On the one hand, welfare decreases because fewer investors trade with their home dealer, and therefore lose the loyalty benefit. On the other hand, welfare increases because on the platform more investors trade with higher-valued dealers or with dealers that offer better base quality. Only when the platform is perfectly competitive, and dealers no longer charge prices that distort the investor's decision, does the positive effect dominate.

This suggests that increasing dealer competition is more effective in increasing welfare than eliminating platform entry barriers. One way to foster competition would be to allow investor to request quotes from all, instead of only four, dealers.

**How important are relationships?** Due to the large loyalty benefit, many investors buy from their home dealer, independent of the market structure. This has two effects. First, loyalty reduces competition for investors. To see this, consider the counterfactual in which all investors trade on the platform but competition is imperfect. When we remove the loyalty benefit quotes increase by 1.8 bps, which is more than double the observed bid-ask spread, thanks to higher competition for non-loyal investors. Second, the loyalty benefit affects welfare because it increases the investor's value, and because loyal investors are less likely to trade with high-valued or high-quality dealers.

To highlight the welfare importance of relationships, we repeat all counterfactuals, but shut off the loyalty benefit in two ways. First, we eliminate the direct effect that the loyalty benefit $r$ has on welfare, by excluding $r$ from the gains from trade. Formally, $v_{tj}^G(\iota_{ti}^G) = \theta_t + \iota_{ti}^G - \xi_{tj}$ in Definition 1, so that $r$ can be interpreted as switching cost that distorts investor behavior but has no direct effect on utility. Second, we eliminate this distortion by shutting off the loyalty benefit entirely.[21]

When the loyalty benefit distorts behavior without affecting the investor's value directly,

[21] In both cases, given that the loyalty benefit $r$ enters the investor's value in Definition 1, setting $r$ to zero decreases welfare. This is a mechanical effect—similar to obtaining profit gains when eliminating costs—and should not be taken as policy recommendation.

it is welfare improving to shift more trades on the platform, even if the platform is imperfectly competitive. This is because the negative welfare effect coming from fewer trades with the home dealer is eliminated. However, since many investors still trade with their home dealer on the platform (as they do bilaterally), welfare gains are moderate, ranging between 3% and 12% depending on the degree of centralization and platform competition.

When shutting off the loyalty benefit, more investors enter the platform and on the platform they no longer buy from their home dealer—unless that dealer offers the best quote/base quality. Therefore, welfare gains are larger in all counterfactuals, ranging between 10% and 27%. This highlights that relationships may have large welfare effects.

**What drives the welfare gain?** In all specifications, almost the entire welfare gain from centralization comes from matches to dealers with higher values, $v_{tj}^D$, rather than to dealers with better qualities, $\xi_{tj}$. Since dealer values would likely be more similar if dealers didn't have to carry inventory, this highlights the fact that dealers cannot freely sell and buy as much they would like. For instance, a dealer who unexpectedly took a long inventory position might be more pressed to sell than a dealer who is short, but she might not be able to sell as much as she would like until the end of the day. Similar frictions triggered dramatic events in bond markets in March 2020. When dealers failed to absorb enough bonds onto their balance sheets to meet the extraordinary supply of investors, the Federal Reserve System purchased trillions of U.S. government bonds and temporarily relaxed balance sheet constraints to rescue the market (Duffie (2020); Schrimpf et al. (2020); He et al. (2022)). Our findings suggest that market centralization can reduce these frictions.

**Will the market fix itself?** Dealers who own the platform in Canada could realize that they can more easily offload their inventory positions when more investors trade on the platform and make the platform more attractive, for instance, by eliminating the platform fee. This is unlikely to happen because dealers earn lower trade profits as more investors trade on the platform and obtain more investor friendly prices.[22] For example, on average a dealer would lose roughly C$ 5.4 million in trade profits from small-trade investors per year if platform access was universal and free, without accounting for the profit loss due to eliminated platform fees. A dealer would suffer an additional C$ 22 million per year if the loyalty benefit was eliminated, which highlights how much dealers benefit from loyal investors. These losses are sizable compared to the average annual profit a bank generates from its market-making activities in all financial markets combined, which is roughly C$ 413

[22]This raises the question why dealers founded the platform in the first place. They followed the global trend away from OTC markets.

million during our sample period.[23]

This highlights that dealers—who are the key market players in most government bond markets—have strong incentives to prevent changes in the market structure, even if they are welfare-improving. This also suggests that dealers might exit the market for government bonds if they can no longer profit from loyal investors, and raises the question whether alternative liquidity providers would step in to make markets. Allen et al. (2022) and Hendershott et al. (2021) have taken the first steps in addressing these questions.

# 9 Conclusion

We use trade-level data on the Canadian government bond market to study whether to centralize OTC markets by shifting bilateral trades onto a multi-dealer platform on which dealers compete for investors. We show that even in a seemingly frictionless market, platform access can lead to better prices for investors, but only when dealer market power is limited. Further, we estimate welfare gains because more trades are intermediated by dealers who urgently seek to trade. We expect this to be true for many other OTC markets.

Shifting more bilateral trading onto multi-dealer platforms is a first feasible step towards full market centralization, which would be achieved if all trades cleared on a central limit order book. Our findings highlight two reasons that render it challenging to break up existing market structures. First, dealers have market power and incentives to prevent market reforms. Second, investors keep long-lasting relationships with dealers, and might therefore be reluctant to adapt their trade behavior. Future research could assess the full welfare effects from changing existing relationship structures.

[23]We obtain this information for the eight largest primary dealers from the regulatory financial reports (Form 1 of the IIROC Rule Book). The profit is the revenue minus costs (including taxes) associated with activities that the bank conducts in their role as a dealer, i.e., trading for their own books and on behalf of clients.

# References


Abudy, M. M. and A. Wohl (2018). Corporate bond trading on a limit order book exchange. *Review of Finance 22*(4), 1413–1440.

Allen, J., A. Hortaçsu, E. Richert, and M. Wittwer (2022). Entry and exit in Treasury auctions. Manuscript.

Allen, J., J. Kastl, and M. Wittwer (2020). Estimating demand systems for Treasuries. Manuscript.

Atkeson, A. G., A. L. Eisfeldt, and P.-O. Weill (2015). Entry and exit in over-the-counter markets. *Econometrica 83*(6), 2231–2292.

Babus, A. and C. Parlatore (2022). Strategic fragmented markets. *Journal of Financial Economics 145*(3), 876–908.

Baldauf, M. and J. Mollner (2020). Competition and information leakage. Working paper.

Barclay, M. J., T. Hendershott, and K. Kotz (2006). Automation versus intermediation: Evidence from Treasuries going off the run. *Journal of Finance 66*(5), 2395–2414.

Benos, E., R. Payne, and M. Vasios (2020). Centralized trading, transparency, and interest rate swap market liquidity: Evidence from the implementation of the Dodd-Frank Act. *Journal of Financial and Quantitative Analysis 55*(1), 159–192.

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

Bessembinder, H., C. Spatt, and K. Venkataraman (2020). A survey of the microstructure of fixed-income markets. *Journal of Financial and Quantitative Analysis 55*(1), 1–45.

Biais, B. and R. Green (2019). The microstructure of the bond market in the 20th century. *Review of Economic Dynamics 33*, 250–271.

BIS (2016). Electronic trading in fixed income markets. Markets Committee Papers no. 7.

Brancaccio, G. and K. Kang (2022). Search frictions and product design in the municipal bond market. Manuscript.

Brancaccio, G., D. Li, and N. Schürhoff (2017). Learning by trading: the case of U.S. market for municipal bonds. Manuscript.

Bresnahan, T. F. (1981). Departures from marginal-cost pricing in the American automobile industry: Estimates for 1977–1978. *Journal of Econometrics 17*(2), 201–227.

Chen, D. and D. Duffie (2021). Market fragmentation. *American Economic Review 111*(4), 2247–2274.

Coen, J. and P. Coen (2021). A structural model of liquidity in over-the-counter markets. Manuscript.

de Roure, C., E. Moench, L. Pelisson, and M. Schneider (2020). OTC discount. Discussion paper Deutsche Bundesbank no. 42/2019.

Di Maggio, M., A. Kermani, and Z. Song (2017). The value of trading relations in turbulent times. *Journal of Financial Economics 124*(2), 266 – 284.

Duffie, D. (2020). Still the world's safe haven? Redesigning the U.S. Treasury market after the COVID-19 crisis. Report.

Duffie, D., P. Dworczak, and H. Zhu (2017). Benchmarks in search markets. *Journal of Finance 72*(5), 1983–2044.

Duffie, D., N. Gârleanu, and L. H. Pedersen (2005). Over-the-counter markets. *Econometrica 73*(6), 1815–1847.

Dugast, J., S. Üslü, and P.-O. Weill (2022). A theory of participation in OTC and centralized markets. *The Review of Economic Studies 89*(6), 3223–3266.

Dunne, P. G., H. Hau, and M. J. Moor (2015). Dealer intermediation between markets. *Journal of the European Economic Association 13*(5), 770–804.

Fleming, M. and F. Keane (2021). The netting efficiencies of marketwide central clearing. Federal Reserve Bank of New York Staff Reports no. 964.

Fleming, M. J., B. Mizrach, and G. Nguyen (2017). The microstructure of a U.S. Treasury ECN: The BrokerTec platform. Federal Reserve Bank of New York Staff Reports no. 381.

Friewald, N. and F. Nagler (2019). Over-the-counter market frictions and yield spread changes. *Journal of Finance 74*(6), 3217–3257.

Garbade, K. D. and W. L. Silber (1976). Price dispersion in the government securities market. *Journal of Political Economy 84*(4), 721–740.

Glode, V. and C. C. Opp (2019). Over-the-counter versus limit-order markets: The role of traders' expertise. *Review of Financial Studies 33*(2), 866–915.

Green, R. C., B. Hollifield, and N. Schürhoff (2007). Dealer intermediation and price behavior in the aftermarket for new bond issues. *Journal of Financial Economics 86*, 643–682.

Hau, H., P. Hoffmann, S. Langfield, and Y. Timmer (2021). Discriminatory pricing of over-the-counter derivatives. *Management Science 67*(11), 6660–6677.

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

Hendershott, T., D. Li, D. Livdan, and N. Schürhoff (2020a). Relationship trading in OTC markets. *Journal of Finance 75*(2), 683–734.

Hendershott, T., D. Li, D. Livdan, and N. Schürhoff (2020b). True cost of immediacy. Working paper.

Hendershott, T., D. Livdan, and N. Schürhoff (2021). All-to-all liquidity in corporate bonds. Working paper.

Hendershott, T. and A. Madhavan (2015). Click or call? Auction versus search in the over-the-counter market. *Journal of Finance 80*(1), 419 – 447.

Holden, C. W., D. Lu, V. Lugovsky, and D. Puzello (2021). What is the impact of introducing a parallel OTC market? Theory and evidence form the Chinese interbank FX market. *Journal of Financial Economics 140*(1), 270–291.

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

Hortaçsu, A. and D. McAdams (2010). Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the Turkish Treasury auction market. *Journal of Political Economy 118*(5), 833–865.

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

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

Koijen, R. S. and M. Yogo (2020). Exchange rates and asset prices in a global demand system. Manuscript.

Kondor, P. and G. Pintér (2022). Clients' connections: Measuring the role of private information in decentralized markets. *Journal of Finance 77*(1), 505–544.

Kozora, M., B. Mizrach, M. Peppe, O. Shachar, and J. Sokobin (2020). Alternative trading systems in the corporate bond market. Federal Reserve Bank of New York Staff Reports no. 938.

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

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

Kutai, A., D. Nathan, and M. Wittwer (2021). Exchanges for government bonds? Evidence during COVID-19. Manuscript.

Lehar, A. and C. A. Parlour (2021). Decentralized exchanges. Manuscript.

Liu, Y., S. Vogel, and Y. Zhang (2018). Electronic trading in OTC markets vs. centralized exchange. Swiss Finance Institute Reserach Paper no. 18-19.

Loon, Y. C. and Z. K. Zhong (2016). Does Dodd-Frank affect OTC transaction costs and liquidity? Evidence from real-time CDS trade reports. *Journal of Financial Economics 119*(3), 645–672.

Managed Funds Association (2015). Why eliminating post-trade name disclosure will improve the swaps market. MFA position paper.

McPartland, K. (2016). Corporate bond electronic trading continues growth trend. Industry report.

Monias, S., M. Nguyen, and G. Valentex (2017). Funding constraints and market illiquidity in the European Treasury bond market. TSE Working Paper no. 17-814.

O'Hara, M. and X. Zhou (2021). The electronic evolution of corporate bond dealers. *Journal of Financial Economics 140*(2), 368–390.

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

Plante, S. (2017). Should corporate bond trading be centralized? Manuscript.

Riggs, L., E. Onur, D. Reiffen, and H. Zhu (2020). Swap trading after Dodd-Frank: Evidence from index CDS. *Journal of Financial Economics 137*, 857–886.

Rostek, M. J. and J. H. Yoon (2021). Exchange design and efficiency. *Econometrica 89*(6), 2887–2928.

Schrimpf, A., H. S. Shin, and V. Sushko (2020). Leverage and margin spirals in fixed income markets during the COVID-19 crisis. BIS Bulletin no. 2.

Vogel, S. (2019). When to introduce electronic trading platforms in over-the-counter markets? Manuscript.

Wang, C. (2022). The limits of multi-dealer platforms. Manuscript.

Weill, P.-O. (2020). The search theory of OTC markets. *Annual Review of Economics 12*, 747–773.

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

Wittwer, M. (2021). Connecting disconnected financial markets? *American Economic Journal: Microeconomics 13*(1), 252–282.

Yoon, J. H. (2019). Endogenous market structure: Over-the-counter versus exchange trading. Manuscript.

Yueshen, B. Z. and J. Zou (2022). Less is more. Manuscript.

Table 1: Yields are not better with the home dealer

| | Bilateral | | Platform | |
|---|---|---|---|---|
| homedealer | -0.0257 | (0.0516) | -0.0257 | (0.0234) |

Table 1 shows the estimates when regressing $markup_{thsij}$, as defined in (1), on an indicator variable that assumes value 1 if the trade realizes with the home dealer, $homedealer_{ij}$, and investor fixed effects, $\zeta_i$, separately for bilateral trades, using 191,570 observations, and for platform trades, using 110,977 observations. The adjusted $R^2$ is 0.085 and 0.021, respectively. Standard errors are in parentheses and clustered at the investor level. Estimates shrink further towards zero when including hour-day, $\zeta_{th}$, security, $\zeta_s$, or dealer, $\zeta_j$ fixed effects.

Table 2: On the platform, an investor trades with more dealers

| | First stage | | Second stage | |
|---|---|---|---|---|
| benchmark | 0.537 | (0.224) | | |
| platform_probability | | | 0.348 | (0.054) |

Table 2 shows the estimation results of IV regression (2), implemented via two stage least squares. In the first stage, we regress investor $i$'s share of trades with bond $s$ that execute on the platform in a month $m$, $platform\_probability_{mis}$, on an indicator variable that is 1 if security $s$ has benchmark status in that month, $benchmark_{ms}$. In the second stage, we regress the number of dealers an investor trades with in a month, $N\_dealers_{mi}$, on the predicted variable of the first stage. In both stages, we use 140,611 observations and include month-year, $\zeta_m$, and investor, $\zeta_i$, fixed effects. The adjusted $R^2$ is 0.447 and 0.941, respectively. Standard errors are in parentheses and clustered at the investor level.

Table 3: Platform yields are better than bilateral yields

| | Coefficient | SE |
|---|---|---|
| platform | 0.122 | (0.019) |

Table 3 shows the estimates when regressing $markup_{ths}$ as defined in (1) on an indicator variable that is 1 if the trade realizes on the platform, $platform_{thsij}$, trade size, $q_{thsij}$, and investor, $\zeta_i$, hour-day, $\zeta_{th}$, security, $\zeta_s$, and dealer, $\zeta_j$, fixed effects, using 1,191,477 observations. The adjusted $R^2$ is 0.403. Standard errors are in parentheses and clustered at the investor level.

Table 4: Large block trades are traded off the platform

| | Coefficient | SE |
|---|---|---|
| $\mathbb{I}(category = 1)q$ | $-0.056300$ | (0.2800000) |
| $\mathbb{I}(category = 2)q$ | $+0.019100$ | (0.0128000) |
| $\mathbb{I}(category = 3)q$ | $-0.002620$ | (0.0039900) |
| $\mathbb{I}(category = 4)q$ | $-0.000734$ | (0.0006370) |
| $\mathbb{I}(category = 5)q$ | $-0.000379$ | (0.0000875) |
| benchmark | $+0.001140$ | (0.0300000) |

Table 4 shows the estimates of regression (4), which regresses an indicator variable that is 1 if the trade is on the platform, $platform_{thsij}$, on trade size, $q_{thsij}$, for five size categories, $\{(0, 0.1)$,$[0.1, 1)$,$[1, 5)$,$[5, 25)$,$[25, \infty)\}$ million, an indicator for whether the traded security has benchmark status, $benchmark_{ts}$, and investor fixed effects, $\zeta_i$, using 784,809 observations. The adjusted $R^2$ is 0.462. The coefficients mechanically increase as trade size increases, but only the coefficient of the largest category, $\beta_5$, is statistically significant. To avoid the mechanical effect, we could normalize trade size by the average in each category. Then $\beta_5$ becomes the largest coefficient. Standard errors are in parentheses and clustered at the investor level.

Table 5: Overview of the estimates for different versions of our model

(a) Benchmark model with loyalty benefits

| $\hat{\mu}^I$ | $\hat{\sigma}^I$ | $\hat{\mu}^R$ | $\hat{\sigma}^R$ | $\hat{c}$ | $1/\hat{\sigma}$ | $\hat{r}/\hat{\sigma}$ |
|---|---|---|---|---|---|---|
| −1.14 | +2.92 | −3.45 | +5.72 | +1.72 | +1.76 | +4.18 |
| (0.07) | (0.05) | (0.35) | (0.62) | (0.20) | (0.59) | (0.06) |
| $\hat{\mu}^{I^*}$ | $\hat{\sigma}^{I^*}$ | $\hat{\mu}^{R^*}$ | $\hat{\sigma}^{R^*}$ | $\hat{c}^*_t$ | $1/\hat{\sigma}$ | $\hat{r}/\hat{\sigma}$ |
| +1.34 | +3.13 | +2.38 | +4.77 | +1.81 | +1.76 | +4.18 |
| (0.08) | (0.31) | (0.31) | (0.05) | (0.24) | (0.59) | (0.06) |

(b) Benchmark model with no loyalty benefits

| $\hat{\mu}^I$ | $\hat{\sigma}^I$ | $\hat{\mu}^R$ | $\hat{\sigma}^R$ | $\hat{c}$ | $1/\hat{\sigma}$ | $\hat{r}/\hat{\sigma}$ |
|---|---|---|---|---|---|---|
| −1.14 | +2.97 | −3.45 | +5.72 | +4.77 | +0.68 | 0 |
| (0.06) | (0.05) | (0.35) | (0.62) | (0.10) | (0.24) | |
| $\hat{\mu}^{I^*}$ | $\hat{\sigma}^{I^*}$ | $\hat{\mu}^{R^*}$ | $\hat{\sigma}^{R^*}$ | $\hat{c}^*_t$ | $1/\hat{\sigma}$ | $\hat{r}/\hat{\sigma}$ |
| +1.34 | +3.16 | +2.38 | +4.77 | +4.91 | +0.68 | 0 |
| (0.06) | (0.05) | (0.31) | (0.59) | (0.11) | (0.24) | |

(c) Extended model with small and large trades

| $\hat{\mu}^{IS}$ | $\hat{\sigma}^{IS}$ | $\hat{\mu}^{IL}$ | $\hat{\sigma}^{IL}$ | $\hat{\mu}^R$ | $\hat{\sigma}^R$ | $\hat{c}^S$ | $\hat{c}^L$ | $1/\hat{\sigma}$ | $\hat{r}^S/\hat{\sigma}$ | $\hat{r}^L/\hat{\sigma}$ |
|---|---|---|---|---|---|---|---|---|---|---|
| −1.25 | +3.01 | −1.03 | +2.68 | −3.49 | +5.75 | +1.63 | +2.87 | +1.44 | +4.07 | +2.46 |
| (0.10) | (0.05) | (0.15) | (0.12) | (0.35) | (0.62) | (0.27) | (0.39) | (0.48) | (0.06) | (0.10) |
| $\hat{\mu}^{IS^*}$ | $\hat{\sigma}^{IS^*}$ | $\hat{\mu}^{IL^*}$ | $\hat{\sigma}^{IL^*}$ | $\hat{\mu}^{R^*}$ | $\hat{\sigma}^{R^*}$ | $\hat{c}^{S*}_t$ | $\hat{c}^{L*}_t$ | $1/\hat{\sigma}$ | $\hat{r}^S/\hat{\sigma}$ | $\hat{r}^L/\hat{\sigma}$ |
| +1.49 | +3.23 | +1.09 | +2.60 | +2.58 | +4.78 | +1.81 | +2.87 | +1.44 | +4.07 | +2.46 |
| (0.12) | (0.05) | (0.18) | (0.13) | (0.31) | (0.58) | (0.32) | (0.45) | (0.48) | (0.06) | (0.10) |

Table 5a shows the median over all periods of the point estimates of the benchmark model per investor group $G$. In (b) we re-estimate the model but shut off the loyalty benefits by setting $r$ to 0. In (c) we show the estimates of the extended model that separates between small and large institutional trades. The corresponding medians of the standard errors are in parentheses. All estimates are in bps. The implied elasticity of demand/supply on the platform, averaged across dealers, are +114/−110 in (a), +96/−96 in (b), and +98/−96 and +156/−156 for small and large trades, respectively.

Table 6: Average welfare gain compare to a fully decentralized market

| **With loyalty** | **Acc I** | **Acc IR** | **No fee** | **Central** |
|---|---|---|---|---|
| Imperfect competition | −1.29 | −1.50 | −2.07 | −3.02 |
| Perfect competition | +1.30 | +1.45 | +1.64 | +2.12 |
| **Switching cost** | **Acc I** | **Acc IR** | **No fee** | **Central** |
| Imperfect competition | +3.37 | +3.94 | +5.58 | +9.77 |
| Perfect competition | +7.07 | +7.92 | +9.10 | +11.89 |
| **No loyalty** | **Acc I** | **Acc IR** | **No fee** | **Central** |
| Imperfect competition | +9.72 | +16.72 | +17.89 | +19.99 |
| Perfect competition | +21.40 | +23.68 | +24.79 | +26.71 |

Table 6 displays the average welfare gain (in %) in all counterfactuals relative to a fully bilateral market. In With loyalty the loyalty benefit enters the investor's value; in Switching cost the loyalty benefit does not directly affect the investor's value but distorts her trading decisions; in No loyalty the loyalty benefit is entirely eliminated. Counterfactual Acc I allows institutional investors to access the platform at estimated platform usage costs; Acc IR allows all investors to access the platform at estimated costs; No fee eliminates the platform fee; Central shifts all trades are on the platform. Dealer (platform) competition is either imperfect or perfect, i.e., $v_{tj}^D = q_{tj}$ and $\sigma = 0$. The baseline welfare of the fully bilateral market is about C$ 737 million per year when the loyalty benefit enters welfare directly and about C$ 441 million per year when it does not.

Figure 1: With how many dealers does an investor trade?

(a) Number of dealers (b) Home dealer

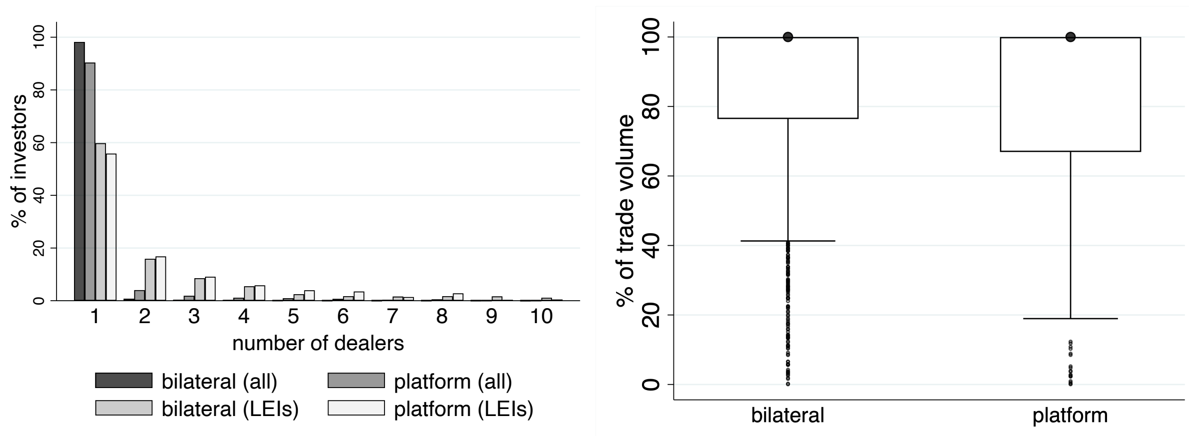


Figure 1a shows the fraction of investors that trade with 1,. . .,10 dealers—bilaterally in black and on the platform in gray. In lighter gray, we restrict the sample to the 1,459 largest investors with a LEI. Figure 1b shows a box plot of the amount an investor with a LEI trades bilaterally or on the platform with her home dealer relative to the total amount this investor trades in each market segment.

Figure 2: Platform access and yields

(a) Yield gap: retail vs. institutional investors

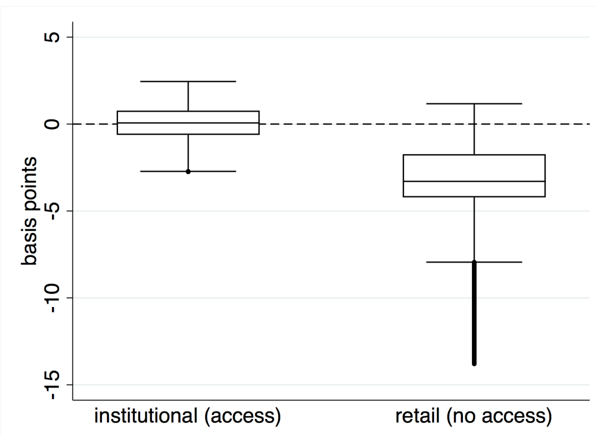


(b) Yield drop when losing platform access

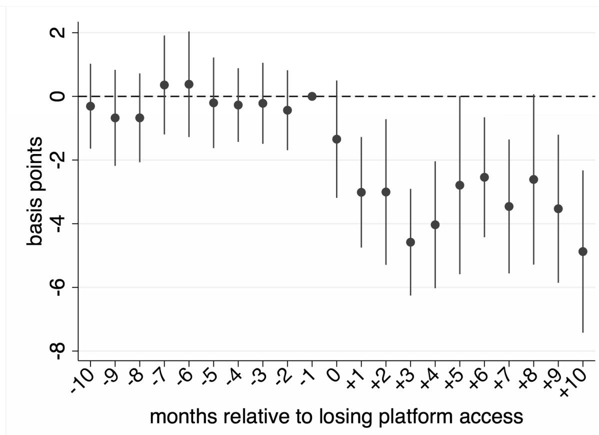


Figure 2a shows a box plot of markups for institutional and retail investors, excluding the upper and lower 5% of the distribution. For this, we regress $markup_{thsij}$ as defined in (1) on indicator variables that distinguish retail from institutional investors, $retail_{thsij}$, and buy-side from sell-side trades, $buy_{thsij}$, as well as hour-day, $\zeta_{th}$, security, $\zeta_s$, and dealer, $\zeta_j$, fixed effects. Figure 2b shows the $\beta_m$ estimates and the 95% confidence intervals of event study regression (3) for 10 months before and after an investor $i$ switched from having to not having platform access. Each $\beta_m$ measures by how much the markup, $markup_{thsij}$, for investor $i$ differs $m$ months before/after this event relative to 1 month prior to it. Standard errors are clustered at the investor level. Both graphs look similar when also controlling for trade size.

Figure 3: Platform quotes versus platform trade yields

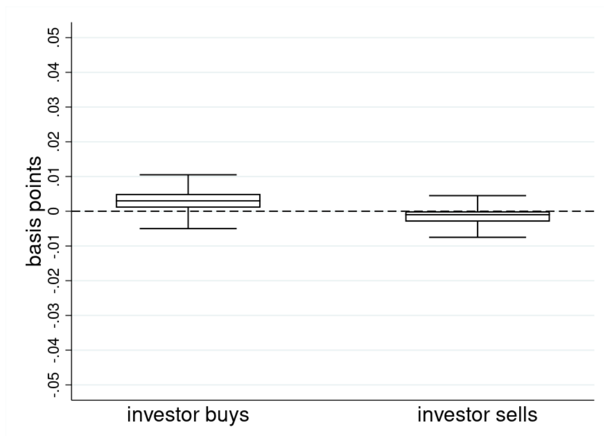


The left-handed box plot of Figure 3 depicts the distribution of the difference between the yield at which a security was bought on the platform and the security's ask quote that is shown on the platform at the time of the trade or shortly before (within 10 seconds). The right-handed box plot shows the analogous for when the investor sells with the bid quote replacing the ask quote. The differences is in basis points. The x-axis is scaled up to 0.05 bps which is 1/10th of the median bid-ask spread to visualize the small magnitude of these differences.

Figure 4: Platform market share

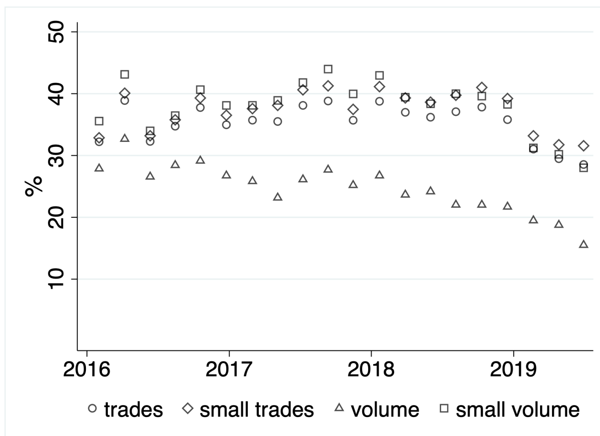


Figure 4 visualizes the platform's market share over time. The circles capture the percentage of trades by institutional investors that are executed on the platform on a day. The triangles show the daily trade volume on the platform relative to the total daily trade volume of institutional investors. The diamonds and squares compute the analogue shares but only using trades that are smaller than C$ 25 million.

Figure 5: Dealer values

(a) Identifying variation

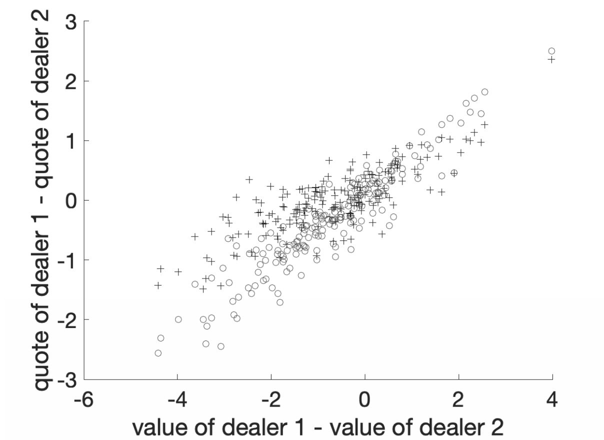


(b) Estimated dealer values for 10 weeks

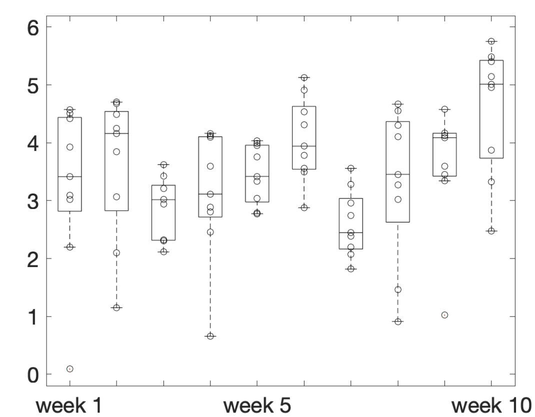


Figure 5a illustrates what drives two dealers to have different values in the same week. The crosses show the correlation between the difference of the observed quotes of two dealers ($q_{t1} - q_{t2}$) and the difference of the estimated values of those dealers ($\hat{v}^D_{t1} - \hat{v}^D_{t2}$). The circles replace the quotes by the markup equation that would arise if there was no bilateral market segment: $q_{tj}(1 + 1/\eta^E_{tj}(q_t))$. Figure 5b displays box plots of the estimated dealer values, $\hat{v}^D_{tj}$, net of the bond's market value, $\theta_t$, in bps for the last ten weeks in our sample.

Figure 6: Dealers' weekly platform market shares and their cost shifters

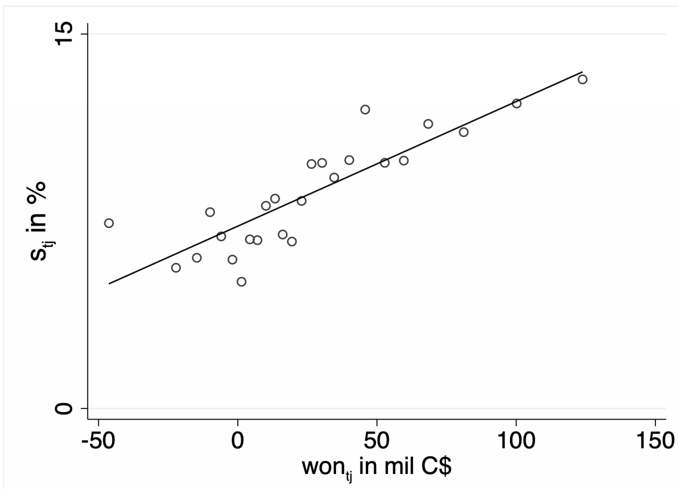


Figure 6 shows a binned scatter plot to visualize the correlation between dealers' weekly total market shares on the platform and their unexpected supply shocks when partialling out week fixed effects. We exclude outliers of the cost shifters for better visualization.

Figure 7: Identifying variation for dealer base qualities

(a) Observed platform market shares (b) Estimated base qualities

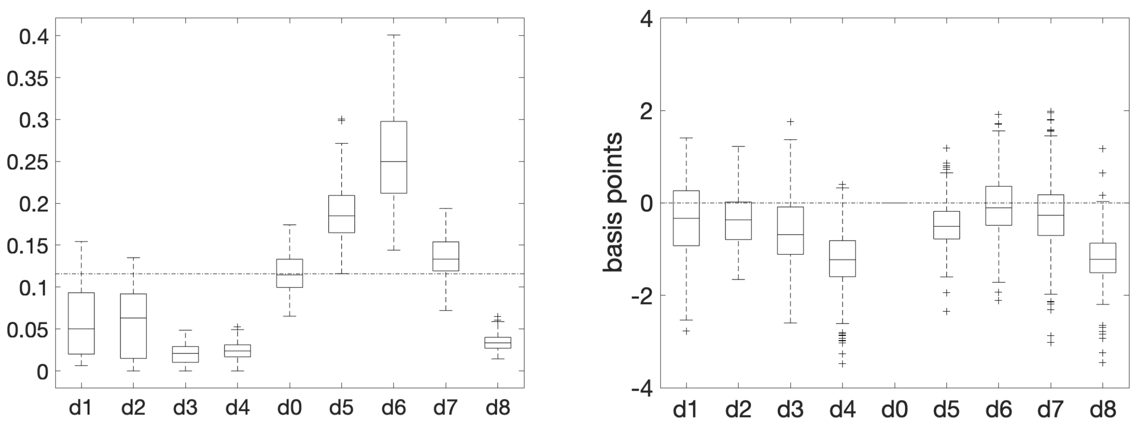


Figure 7a displays a box plot for each dealer, labeled d0 (benchmark) to d8. Each shows the distribution of how much of the total platform market this dealer captures (in %) in a week. Figure 7b shows a box plot of the estimated base qualities, $\hat{\xi}_{tj}$, in bps for each dealer.

Figure 8: Identifying variation for $c_t, \mu_t^I, \sigma_t^I$

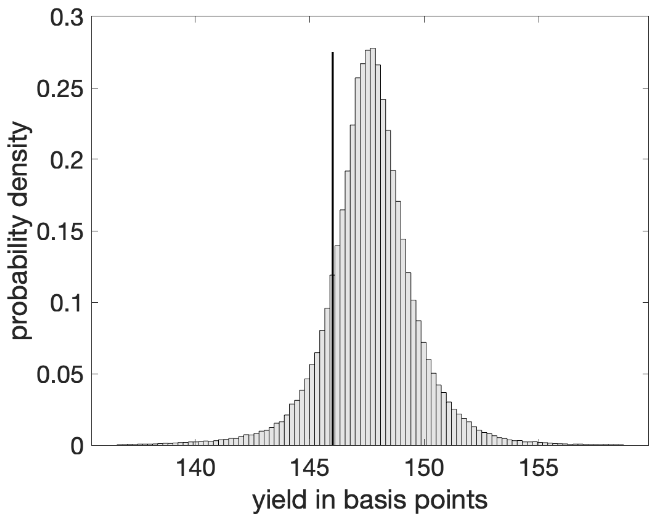


Figure 8 shows a probability density histogram of the yields (in bps) that institutional buyers realize, excluding the upper and lower 0.1 percentile of the distribution, and a black line. This line is the average cutoff (across weeks and dealers) that determines whether an institutional investor buys bilaterally or on the platform, according to Proposition 1.

Figure 9: Model fit

(a) Trading

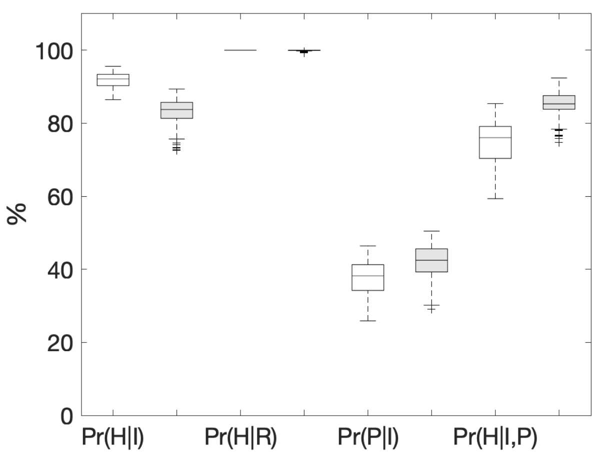


(b) Yields

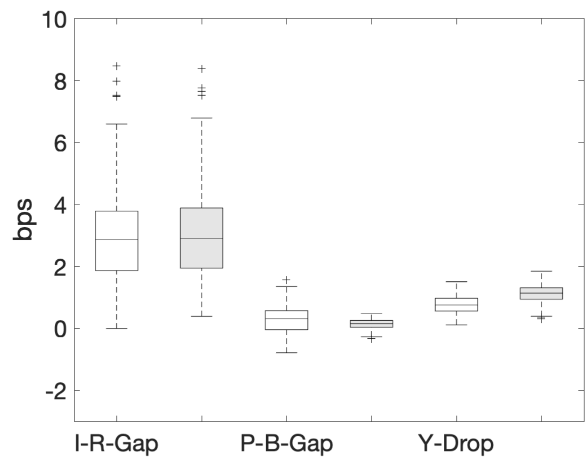


Figure 9 shows box plots of trade probabilities (in %) and yields (in bps) in the data (in gray) and in the model (in white). The distribution is taken over periods. Figure 9a shows the probabilities to trade with a home dealer for institutional and retail investors, Pr(H|I) and Pr(H|R), the probability to enter the platform, Pr(P|I), and to trade with a home dealer on the platform for institutional investors, Pr(H|I,P). Figure 9b displays the yield gap between retail and institutional investors, I-R-Gap, the yield difference between bilateral and platform yields for institutional investors, P-B-Gap, and the amount by which the yield drops when losing platform access, Y-Drop. Mathematical details are in Appendix D.3.

# Online Appendix

**Centralizing Over-The-Counter Markets?**

By Jason Allen and Milena Wittwer

Appendix A explains how we clean the raw data.
Appendix B describes model extensions and our micro-foundation for platform quotes.
Appendix C provides all proofs.
Appendix D explains our estimation.
Appendix E summarizes our robustness analyses.

# A Data cleaning

For 0.15% of trades out of 3,755,901 observations in the raw data, we correct the execution time, the date, or the settlement date. We correct 100 cases in which subsidiaries of reporting dealers or brokers are labeled retail investors, and drop 20 observations that were reported without retail/institutional indicator. Of the investors who switch from retail to institutional or vice versa, we drop investors who do not permanently switch. For example, CIBC Investor Services Inc., a subsidiary of the Canadian Imperial Bank of Commerce (CIBC), classifies as a retail investor according to the rules, even though CIBC is one of the biggest banks in Canada. A reporting dealer who trades with CIBC Investor Services Inc. might falsely believe that this investor is institutional and report it as such.

We exclude trades that exhibit yields that are extreme relative to the public Bloomberg mid-yield since it is difficult to rationalize why anyone would be willing to accept these trades. They could be reporting errors or part of a larger investment package that we do not observe. To detect these outliers, we analyze the distribution of the markup, defined in (1). It has extremely long but very thin tails. We drop the upper and lower 1% of this distribution for each investor group.

We focus on CanDeal or bilateral trades only, which means that we ignore 0.41% of the observations with incorrect trading venues. In these rare cases, the dealer makes a mistake and typically reports the ID of her counterparty as the trading venue.

In rare cases in which a Bloomberg quote for a security is missing in an hour of the day, we use the daily average Bloomberg quote of this security.

# B Model extensions and micro-foundation

## B.1 Accounting for trade size

Motivated by Fact 4, we extend the model to differentiate between small and large trades: $g \in \{S, L\}$. Given that retail investors don't trade large amounts, there are three groups of investors $G$: small institutional investors, large institutional investors and retail investors. Like in the benchmark model, dealers observe the investor group. In contrast to before, she now offers a size-specific loyalty benefit, $r^g$, base quality, $\xi_t^g = \xi_j + \xi_t + \chi_{tj}^g$, and quote $q_{tj}^g$. The sequence of events and equilibrium for each group remain unchanged.

## B.2 Modeling platform competition

In the benchmark model, the platform shock $\epsilon_{tij}$ distorts investor behavior because it affects investor utilities directly. Here, we present a model variation in which the investor's utility always equals her value $v_{tij}^G$, and $\epsilon_{tij}$ affects the utility indirectly via the price the investors pay on the platform.

The main difference is that the posted quotes are only indicative. Investors still observe these quotes to decide where to trade, but then run an auction on the platform—with all dealers or with a random subset of dealers—that determines by how much dealers adjust the indicative quotes, i.e., the actual trade price. In Appendix B.3, we formalize a stylized platform auction-game. We show that in equilibrium, a dealer's bid equals her posted quote plus a stochastic term that depends on the signal, $\epsilon_{tij}$, and a parameter $\sigma$:

$$q_{tj} + \sigma\epsilon_{tij} \text{ where } \epsilon_{tij} \sim \mathcal{H} \text{ and } \sigma \in \mathbb{R}. \tag{12}$$

Investors in this model behave like in the benchmark model. Thus, Proposition 1 goes through with the only difference being that investors now pay $q_{tj} + \sigma\epsilon_{tij}$ on the platform. The reason is that the investor cares about her surplus which is affected in the same way by $\sigma\epsilon_{tij}$ in both models. It doesn't matter whether $\sigma\epsilon_{tij}$ is part of the investor's utility or part of her payment. In contrast, for dealers who collect the payments, this difference matter.

**Proposition 3.** *Dealer $j$ posts a quote $q_{tj}$ that satisfies*

$$\begin{aligned} 0 = &(v_{tj}^D - (\theta_t + \psi_{tj}(q_t) - \xi_{tj}))\frac{\partial \rho_{tj}(q_t)}{\partial q_{tj}} \\ &+ \sum_{k=1}^{J_t} \frac{\partial S_{tj}^k(q)}{\partial q_j}(v_{tj}^D - q_{tj} - \sigma E_{tj}^k(q_t)) - \sum_{k=1}^{J_t} S_{tj}^k(q)\left[1 + \sigma\frac{\partial E_{tj}^k(q_t)}{\partial q_{tj}}\right] \end{aligned} \tag{13}$$

*with* $\rho_{tj}(q) = \Pr(\iota_{ti}^I \geqslant \psi_{tj}(q_t))$, $S_{tj}(q_t) = S_{tj}^j(q_t) + \sum_{k \neq j} S_{tj}^k(q_t)$ *with* $S_{tj}^l = s_{tj}^l(q_t)\Pr(\iota_{ti}^I \leqslant$

$\psi_{tl}(q_t))$ *and* $s^l_{tj}(q_t) = \Pr(\xi_{tj} + q_{tj} + \mathbb{I}(j = l)r + \sigma\epsilon_{tij} > \xi_{tk} + q_{tk} + \mathbb{I}(k = l)r + \sigma\epsilon_{tik}\ \forall k \neq j) \forall l \in J_t$ *and* $E^j_{tj}(q_t) = \mathbb{E}[\epsilon_{tij}|\xi_{tj} + r + q_{tj} + \sigma\epsilon_{tij} > \xi_{tk} + q_{tk} + \sigma\epsilon_{tik}\ \forall k \neq j]$, $E^l_{tj}(q_t) = \mathbb{E}[\epsilon_{tij}|\xi_{tj} + q_{tj} + \sigma\epsilon_{tij} > \xi_{tk} + \mathbb{I}(l = k)r + q_{tk} + \sigma\epsilon_{tik}\ \forall k \neq j]$. *When* $r = 0$,

$$q_{tj}\left(1 + \frac{1}{\eta^E_{tj}(q_t)}\left(\left(1 + \sigma\frac{\partial E(q_t)}{\partial q_{tj}}\right) - \frac{\pi^D_{tj}(q_t)}{\partial q_{tj}}\right)\right) = v^D_{tj} - \sigma E_{tj}(q_t) \tag{14}$$

*with* $E_{tj}(q_t) = \mathbb{E}[\epsilon_{tji}|\tilde{u}_{tj}(\epsilon_{tji}) > \tilde{u}_{tk}(\epsilon_{tki})\ \forall k \neq j]$ *where* $\tilde{u}_{tij}(\epsilon_{tij}) = \xi_{tj} + q_{tj} + \sigma\epsilon_{tij}$.

Dealer $j$ anticipates that the investor pays $q_{tj} + \sigma\epsilon_{tij}$ on the platform, but doesn't know the platform shocks. Therefore, she takes an expectation of these shocks. When investors don't obtain loyalty benefits ($r = 0$), the condition simplifies to a markup equation that is similar to (9). The only difference to the benchmark model is that the dealer expects to earn platform price $q_{tj} + \sigma E_{tj}(q_t)$ rather than $q_{tj}$.

## B.3 Micro-foundation for platform prices

Here, we solve for an equilibrium in an RFQ auction to justify why it can be reasonable to assume that investor $i$ obtains a yield of $q_{tj} + \sigma\epsilon_{tij}$ when trading with dealer $j$ on the platform in period $t$. To make the auction tractable, we let dealers be ex ante identical—i.e., abstract from dealer qualities and loyalty benefits. This implies that all dealers post the same quote $q_t$ in the first stage of the game. We drop the day subscript for convenience and focus on formalizing an auction that occurs on the platform between an investor and dealers.

$K \geqslant 2$ dealers seek to sell one unit of the bond in a first-price auction. For simplicity, we let all dealers participate in the auction ($J = K$). Alternatively, we could let the investor select a subset of dealers at random ($K < J$).

Each dealer aims at maximizing profit. Ex post, the dealer obtains a profit of $v - y$, when selling one unit at yield $y$ and the reservation value is $v$. The dealer's reservation value on the platform splits into two parts. The first part is the publicly posted platform quote. This is commonly known to all dealers. The second part $v_2$ is unknown, drawn iid from a commonly known normal distribution:

$$v = q + v_2 \text{ with } q \in \mathbb{R}^+ \text{ and } v_2 \overset{iid}{\sim} N(\mu_v, \sigma_v^2).$$

Before running the auction with investor $i$, each dealer $j$ draws a private signal $x_{ij}$ about the common value component of her value:

$$x_{ij} = v_2 + s\omega'_{ij}, \text{ where } \omega'_{ij} \overset{iid}{\sim} N(0, 1) \text{ and } s \in \mathbb{R}^+.$$

The signals may be correlated across dealers, which is more general than in our structural estimation where platform shocks are assumed to be iid. To achieve independence in the theoretical model, we could assume that the dealers draw independent signals, conditional on $q$.

Given these signals, each dealer submits her bid, and the dealer with the highest bid wins the auction and trades with the investor at a yield equal to the winning bid.

**Proposition 4.** *In a symmetric equilibrium the dealer with the highest signal, $x_{ij}$, wins the auction, and the investor obtains the following yield on the platform:*

$$y_{ij}^E = q + \sigma\epsilon_{ij} \text{ where } \epsilon_{ij} = (x_{ij}/\sigma + s) \text{ and} \tag{15}$$

$$\sigma \text{ solves } 0 = \int_{-\infty}^{\infty} \left[-\Phi(-z)^{J-1} + (z-\sigma)(J-1)\Phi(-z)^{J-2}\phi(-z)\right]\phi(z)\, dz. \tag{16}$$

*$z \sim N(0,1)$ and $\Phi(\cdot)$, $\phi(\cdot)$ are the CDF, PDF of the standard normal distribution.*

# C Proofs

The equilibrium is derived by backward induction. For notational convenience, we drop the subscript $t$ and the superscript $I$.

**Proposition 1.** Statement $(i)$ holds by assumption. To derive statement $(ii)$, note that conditional on entering the platform and observing $\epsilon_{ij}$, investor $i$ buys from dealer $j$ if $u_{ij} > u_{ki}\ \forall k \neq j$, where $\tilde{u}_{ij}(\epsilon_{ij}) = \xi_j + q_j + \mathbb{I}(j = d)r + \sigma\epsilon_{ij}$. When choosing whether or not to enter the platform, the investor compares how much she obtains when trading bilaterally, which is the loyalty benefit $r$, with how much she expects to earn when entering the platform, which is $-(\theta + \iota_i) + \mathbb{E}[\max_{k\in\mathcal{J}} \tilde{u}_{ki}(\epsilon_{ki})] - c$. She buys bilaterally if $\psi(q) \leqslant \iota_i$ with $\psi(q) = \mathbb{E}[\max_{k\in\mathcal{J}} \tilde{u}_{ki}(\epsilon_{ki})] - \theta - c - r$. □

**Proposition 2.** In choosing quote $q_j$, a dealer $j$ takes the quotes of the other dealers as given and anticipates how investors will react according to Proposition 1. She does not know which liquidity shocks investors will draw, and therefore maximizes the expected profit from platform and bilateral trading. Formally, $\max_{q_j} \pi_j(q) = \max_{q_j}\{\pi_j^D(q) + \pi_j^E(q)\}$. Here, $\pi_j^D(q) = \mathbb{E}[v_j^D - (\iota_i + \theta - \xi_j)|\psi_j(q) \leqslant \iota_i]$ is the expected profit from bilateral trades; and $\pi_j^E(q) = S_j(q)(v^D - q_j)$, where $S_j(q) = S_j^j(q) + \sum_{k\neq j} S_j^k(q)$ with $S_j^l = s_j^l(q)\Pr(\iota_i \leqslant \psi_l(q))$ and $s_j^l(q) = \Pr(\xi_j + q_j + \mathbb{I}(j = l)r + \sigma\epsilon_{ij} > \xi_k + q_k + \mathbb{I}(k = l)r + \sigma\epsilon_{ik}\ \forall k \neq j)\forall l \in J$ is the expected profit from platform trades. When $\epsilon_{ij}$ are EV1 distributed, platform market shares simplify to $s_j^l(q) = \frac{\exp(\frac{1}{\sigma}(q_j+\xi_j+\mathbb{I}(j=l)r))}{\sum_k \exp(\frac{1}{\sigma}(q_k+\xi_k+\mathbb{I}(k=l)r))}$. Setting the partial derivative w.r.t. $q_j$ to 0, and rearranging gives the markup equation. □

**Proposition 3.** The proof is analogous to the proof of Proposition 2. The difference is that the dealer expects a different trading profit on the platform, namely: $\pi_j^E(q) = S_j^j(q)(v_j^D - q_j - \sigma E_j^j(q)) + \sum_{l \neq j} S_j^l(q)(v_j^D - q_j - \sigma E_j^l(q))$ with $S_j^k(q)$ and $E_j^k(q)$ for all $k \in J$ as defined in Proposition 3. When $r = 0$, the first-order condition simplifies to markup equation (14). □

# D Details about the estimation

In Appendix D.1 we describe the estimation procedures for different versions of our model, assuming that we have cost-shifter instruments. In Appendix D.2 we explain how to construct these instruments from bidding data of primary auctions.

## D.1 Estimation procedure

**Benchmark model.** We explain the estimation for buying investors. The estimation is analogous for selling.

Denote the share of investors who have home dealer $l$ and buy from dealer $j$ on the platform by $s_{tj}^l(q_t, \xi_t, r, \sigma)$. According to Proposition 1,

$$s_{tj}^l(q_t, \xi_t, r, \sigma) = \frac{\exp(\frac{1}{\sigma}(q_{tj} + \xi_{tj} + \mathbb{I}(j = l)r))}{\sum_k \exp(\frac{1}{\sigma}(q_{tk} + \xi_{tk} + \mathbb{I}(k = l)r))} \text{ for all } l \in \mathcal{J}_t \tag{17}$$

since $\epsilon_{tij}$ are EV1 distributed.

We abbreviate $s_{tj}^l(q_t, \xi_t, r, \sigma)$ by $s_{tj}^l$ for all $j$, divide this expression for all $j \neq 0$ by the equivalent expression for the benchmark dealer ($j = 0$), take logs, and use Normalization 1 to obtain

$$\log(s_{tj}^l / s_{t0}^l) = \zeta_j + \mathbb{I}(l = j)\frac{r}{\sigma} + \zeta_t + \frac{1}{\sigma}\tilde{q}_{tj} + rest_{tj}, \tag{18}$$

where $\tilde{q}_{tj} = q_{tj} - q_{t0}$, $\zeta_j = \frac{1}{\sigma}\xi_j, \zeta_t = mean_j\left(\frac{1}{\sigma}\chi_{tj}\right), rest_{tj} = \frac{1}{\sigma}\chi_{tj} - \zeta_t$.

Under Assumption 2, which implies $\mathbb{E}[rest_{tj} | \tilde{w}on_{\tilde{t}j}, \zeta_t, \zeta_j] = 0$, we can identify $r$ and $\sigma$ from expression (18), when instrumenting $\tilde{q}_{tj}$ by $\tilde{w}on_{\tilde{t}j} = won_{\tilde{t}j} - won_{\tilde{0}j}$. This can be done in various ways. We estimate an IV regression of log market share ratios on relative quotes that includes home dealer, dealer and date fixed effects.

With this, we compute $\hat{\xi}_{tj}$ for all $j \neq 0$ and how much utility the investor expects to receive on the platform

$$\mathbb{E}[\max_{k \in \mathcal{J}_t} \tilde{u}_{tik}(\epsilon_{tik})] = \hat{\sigma} \ln \left( \sum_{k \in \mathcal{J}_t} \exp\left(\frac{1}{\sigma}(\hat{\xi}_{tk} + \mathbb{I}(k = d)\hat{r} + q_{tk})\right)\right)$$

given $\tilde{u}_{tik}(\epsilon_{tik})$ in (8) and $\epsilon_{tij} \overset{\text{iid}}{\sim} EV1$; in addition to cutoff (7) that determines whether the investor enters the platform.

Given these estimates, we estimate the remaining parameters for each period $t$ and each investor group $G \in \{I, R\}$ via GMM. Here, we match the observed to the predicted expectation and variance of the bilateral yield of a buying investor of group $G$, and the probability that an institutional investor buys bilaterally. To compute the predicted moments, we rely on $\iota_{ti}^G \sim N(\mu_t^G, \sigma_t^G)$ and Proposition 1.

**Extended model.** We estimate the extended model with small and large trades analogously to the benchmark model. The only difference is that there are now two groups of institutional investors. Therefore, each of the dealer platform market shares (17) splits into two, the share of small trade investors and the share of large trade investors: $s_{tj}^{lg}(q_t, \xi_t, \sigma)$ for $g \in \{S, L\}$. Equation (18) adjusts accordingly, so that the adjusted IV regression identifies $r^S$ and $r^L$ in addition to $\sigma$ and dealer base qualities. The remaining parameters are estimated as before, for each investor group $G$ and period $t$ separately.

**Alternative model of Appendix B.2.** The alternative model differs from our benchmark model in two ways which affect the estimation.

The first difference comes from the fact that the platform trade yields no longer equal the indicative platform quotes. We therefore either have to observe both, or, if we only observe platform yields, back out which indicative quotes generate these platform yields.

Backing out indicative quotes from trade yields is challenging because it involves finding a complex fixed point. To illustrate this, consider the model without loyalty benefits ($r = 0$). Assume that we observe $I_{tj}$ investors $i$ buying from dealer $j$ on the platform in period $t$ at yields $y_{tij}^p$, and denote the share of investors buying from dealer $j$ on the platform among all investors who enter the platform in period $t$ by $s_{tj} = \frac{I_{tj}}{\sum_j I_{tj}}$. To back out indicative quotes, guess some initial values for $wedge_{tj}^\tau$ $\forall t, j$ and compute

$$q_{tj}^\tau = \frac{1}{I_{tj}} \sum_i y_{tij}^p - wedge_{tj}^\tau \text{ for all } t, j,$$

in addition to $q_{tj}^\tau - q_{t0}^\tau$, where $j = 0$ is the benchmark dealer. Then, estimate dealer qualities $\xi_{tj}^{\tau+1}$ and $\sigma^{\tau+1}$ by regressing the observed log-ratio of dealer platform market shares, $\log(s_{tj}/s_{t0})$, on the instrumented relative quotes, $q_{tj}^\tau - q_{t0}^\tau$, dealer and period fixed effects. This step is like in the benchmark model without loyalty benefits. With the estimates, compute

$$wedge_{tj}^{\tau+1} = \sigma^{\tau+1}\hat{\mathbb{E}}[\epsilon_{tij}|\xi_{tj}^{\tau+1} + q_{tj}^{\tau} + \sigma^{\tau+1}\epsilon_{tij} > \xi_{tk}^{\tau+1} + q_{tk}^{r} + \sigma^{\tau+1}\epsilon_{tik}\ \forall k \neq j],$$

where $\epsilon_{tij}$ are drawn iid from the EV1 distribution with variance 1. Check whether the guess was correct, i.e., $wedge_{tj}^{\tau} = wedge_{tj}^{\tau+1}$, which implies $q_{tj}^{\tau} = q_{tj}^{\tau+1}$, for all $t, j$. Repeat if this isn't the case, and hope for convergence.

The second difference is to back out the dealer's value from a markup equation (9) that includes expectation terms, $E_{tj}^{k}(q_t)$ and $\frac{\partial E_{tj}^{k}(q_t)}{\partial q_{tj}}$. These terms have no closed-form, but we can compute them numerically by drawing many EV1 draws for all dealers. The rest is as before.

## D.2 Supply shock instruments

To construct the supply shock instruments, we model the primary auction and back out how much each dealer expects to win when placing its bid. For this, we rely on Hortaçsu and Kastl (2012), who introduce a model that captures the main institutional features of Canadian Treasury auctions with $N_d$ (potential) dealers and $N_c$ (potential) customers, who bid via dealers. A bid is a step function with maximally 7 steps that specifies how much the bidder seeks to buy at different prices. See Hortaçsu and Kastl (2012) for details of the auction game.

**Constructing supply shocks for each auction.** We first estimate the distribution of the residual supply curve a dealer faces in equilibrium. We draw $N_c$ customer bids from the empirical distribution of customer bids in the auction, replacing bids by customers who did not bid in the auction with 0. We then find the dealer(s) who observed each of the customer bids and draw their bids. When the customer submits more than one bid, we draw bids uniformly from all dealers who observed this customer. If at that point the total number of dealers we have drawn is still lower than the number of potential dealers minus one, we draw the remaining dealer bids from the pool of dealers who do not observe a customer bid.

We then let the market clear for each realization of the residual supply curve. This gives us the distribution of how much the dealer won in the auction, $q_j^c$, at market clearing price $P^c$. It also specifies, for each step of the dealer's bidding function, how likely it is that the market clears at that step, i.e., $b_k \geqslant P^c > b_{k+1}$.

With that, we compute how much the dealer expects to win when bidding:

$$\hat{\mathbb{E}}[\text{amount dealer } j \text{ wins}|\text{bids}] = \sum_{k}^{K_j} \hat{\Pr}(b_k \geqslant P^c > b_{k+1}|\theta_j^d) * \hat{\mathbb{E}}[q_j^c|b_k \geqslant P^c > b_{k+1}, \theta_j^d],$$

where $K_j$ are the steps in dealer $j$'s bidding function, and $\theta_j^d$ is the dealer's information set at the time of the bid. To obtain our instrument $won_{j\tilde{t}}$ of auction $\tilde{t}$, we subtract $\hat{\mathbb{E}}$[amount dealer $j$ wins |bids] from the (observed) amount that bidder $j$ actually won.

**Constructing the instrument.** When a period is a day, we use the supply shocks of the last primary auction, which takes place 2-3 days before a typical trade. When the period is a week, there is a complication that arises from the fact that primary auctions don't occur regularly. Ideally we would want to have an auction each Monday. Then we could construct market shares using all trades that happen after auction announcement on Monday until Friday evening, and instrument each dealer's weekly quote by the supply shock received before trading in that week. In reality, most auctions-times are irregular.

In our main specification, we average dealer quotes within the week and construct the instrument by averaging over the supply shocks a dealer received in that week. As robustness, we also use a day and the time between two subsequent auctions (see Appendix E).

## D.3 Details regarding the model fit

Figure 9 and Appendix Figure A6 show that the model fits several moments in the data. Here we explain how these moments are computed for the buy-side for which all functions are defined in Propositions 1 and 2. The formulas for the sell-side are analogous.

**Pr(H|I).** The model-predicted probability for an institutional investor to buy from her home dealer $j$ is $1 - \hat{\mathcal{F}}_t^I(\hat{\psi}_{tj}(q_t)) + \hat{S}_{tj}^j(q_t)$. The data analogue is the ratio between the number of times an institutional investor with home dealer $j$ buys from $j$ in period $t$ and the total number of times an institutional investor with home dealer $j$ buys in period $t$.

**Pr(R|I).** The model-predicted probability for a retail investor to buy from the home dealer 100% given that retail investors cannot trade on the platform and all investors trade with their home dealers bilaterally. The data analogue is computed like for institutional investors.

**Pr(P|I).** The model-predicted probability for institutional investors to buy on the platform is $\hat{\mathcal{F}}_t^I(\hat{\psi}_{tj}(q_t))$. The data analogue is the ratio between the number of times an institutional investor with home dealer $j$ buys on the platform in period $t$ and the total number of times an institutional investor with home dealer $j$ buys in period $t$. For retail investors, the probability of trading on the platform in the model and the data is 0.

**Pr(H|I,P).** The model-predicted probability for institutional investors to trade with their home dealer on the platform, conditional on platform entry, is $\hat{s}_{tj}^j(q_t)$. The data analogous is the ratio between the number of times an institutional investor with home dealer $j$ buys from $j$ on the platform in period $t$ and the total number of times an institutional investor with home dealer $j$ buys on the platform in period $t$.

**I-R-Gap.** The model-predicted expected yield gap between institutional and retail investors is $\hat{\mathcal{F}}_t^I(\hat{\psi}_{tj}(q_t))q_{tj} + (1 - \hat{\mathcal{F}}_t^I(\hat{\psi}_{tj}(q_t)))\hat{\mathbb{E}}_t[\theta_t + \iota_{tj}^I - \hat{\xi}_{tj} | \iota_{tj}^I \geqslant \hat{\psi}_{tj}(q_t)] - \hat{\mathbb{E}}_t[\theta_t + \iota_{tj}^R - \hat{\xi}_{tj}]$. The data analogue is $\text{avg\_yield}_{tj}^I - \text{avg\_yield}_{tj}^R$, where $\text{avg\_yield}_{tj}^I$ is the average (residualized) yield at which institutional investors buy from dealer $j$ in period $t$ and $\text{avg\_yield}_{tj}^R$ is the same for retail investors.

**P-B-Gap.** The model-predicted difference between the platform yield and the expected bilateral yield is $q_{tj} - \hat{\mathbb{E}}_t[\theta_t + \iota_{tj}^I - \hat{\xi}_{tj} | \iota_{tj}^I \geqslant \hat{\psi}_{tj}(q_t)]$. The data analogue is $\text{avg\_platform\_yield}_{tj} - \text{avg\_bilateral\_yield}_{tj}$ where $\text{avg\_platform\_yield}_{tj}$ is the average (residualized) yield at which investors buy from dealer $j$ on the platform in period $t$ and $\text{avg\_bilateral\_yield}_{tj}$ is the same for bilateral trades.

**Y-Drop.** The expected yield drop for retail-like investors when losing platform access according to the model is the difference between the yield a retail investor expects in the counterfactual in which she has access to the platform and the yield she expects in the status quo, with no access. The yield drop according to the data is the point estimate of the static version of the event study (recall Figure 2b):

$$markup_{thsij} = \zeta_i + \beta access_{thi} + \zeta_{th} + \zeta_s + \zeta_j + \epsilon_{thsij}.$$

To create a box plot, we draw $T$ times from a normal distribution with mean equal to the point estimate $\hat{\beta} = 1.16$ and standard deviation equal to the standard error of the estimate, which is 0.340.

**HHI.** The model-predicted average Herfindahl–Hirschman Index is $\sum_j market\_share_{tj}^2$ with $market\_share_{tj}^2 = \frac{\kappa_I[1 - \hat{\mathcal{F}}_t^I(\hat{\psi}_{tj}(q_t)) + \hat{S}_{tj}(q_t)] + \kappa_R}{J_t}$. The observed market shares are the ratios between the number of times an investor buys from dealer $j$ in period $t$ and the total number of times an investor buys in period $t$.

**P-HHI.** The model-predicted average Herfindahl–Hirschman Index on the platform is $\sum_j market\_share\_p_{tj}^2$ with $market\_share\_p_{tj}^2 = \frac{\hat{s}_{tj}(q_t)}{\sum_j \hat{s}_{tj(q_t)}}$. The observed market shares are the ratios between the number of times an investor buys from dealer $j$ in period $t$ on the platform and the total number of times an investor buys in period $t$ on the platform.

# E Robustness analysis

**Instrument.** Given IV regression (18), the instrument mostly affects the estimates of the competition parameter $\hat{\sigma}$.

In Appendix Table A6, we show the first stage of the IV regression to validate that our instrument $\tilde{won}_{\tilde{t}j}$ explains sufficient variation in the endogenous quotes; and check whether $\hat{\sigma}$ is biased in the expected direction when we do not instrument the quotes by $\tilde{won}_{\tilde{t}j}$ and replace Assumption 2 with $\mathbb{E}[\chi_{tj}|\zeta_t, \xi_j] = 0$. The OLS estimate implies an elasticity that is close to zero. The endogeneity bias goes in the expected direction. It comes from a misspecified estimate of $\sigma > 0$, which is biased downward if dealers decrease the yield quote (i.e., increase the price) in response to higher demand for reasons that are unobservable to the econometrician.

In addition, we use a different instrument for the quote; namely, the amount a dealer won in the most recent auctions rather than the amount she won unexpectedly. The advantage of this instrument is that it is model-free. The disadvantage is that it does not address the concern that dealers might anticipate investor demand and bid accordingly in the auction. The estimates are similar, but, as expected, slightly downward biased.

Further, we verify that the estimates of IV regression (18) are similar when using a day instead of a week as a period. The advantage of using a day is that our instrument is cleaner, as discussed in Section D.2. The disadvantage is that we quickly lose statistical power because we don't have sufficiently many trades in a day to construct all market shares defined in (17) with sufficient precision. This is similar to the zero market share problem faced by many applications of demand estimation. To gain some power, we omit dealer fixed effects. We find that the point estimate of the $\frac{1}{\sigma}$ coefficient is smaller (implying that competition is lower), but not statistically different from the estimate of the analogous weekly regression. The implied loyalty benefit is slightly higher in the daily regression (with roughly 3) than in the weekly regression (roughly 2.3). This suggests that, if anything, we might be underestimating the importance of market power and dealer-investor relationships in our baseline specification. The same conclusion can be drawn when using the time between two subsequent auctions as a period.

**Inter-dealer market.** To study the role of the inter-dealer market, assume for a moment that this market is frictionless and dealers can therefore immediately offload bonds bought (sold) from investors by selling (buying) them to (from) other dealers. This is often the case in related theory work (following Duffie et al. (2005)). Then, all dealers would value the bond at its market value $\theta_t$. In our data, dealers only offload 16% of what they accumulate from trades with investors by trading with one another or brokers (see Appendix Table A2). This fraction is rather low but might still affect the dealers' values $\hat{v}_{tj}^D$—which are estimated under the simplifying assumption that dealers do not trade with one another—and with that the welfare findings.

To approximate by how much our welfare findings change when we let dealers balance out a small fraction of their positions by trading with one another, we shrink each dealer's value $\hat{v}_{tj}^D$ towards the bond's market value $\theta_t$. Specifically, we compute welfare in the status quo and all counterfactuals under the assumption that the dealer's actual value is $\theta_t + (1-x)(\hat{v}_{tj}^D - \theta_t)$ rather than $\hat{v}_{tj}^D = \theta_t + (\hat{v}_{tj}^D - \theta_t)$ for different values of $x$. Appendix Table A5 reports the numbers for $x = 16\%$. We find the percentage increase in welfare to be extremely robust. This is true across counterfactuals and specifications, even when increasing $x$ for $x < 100$. Mechanically, the absolute value of the welfare gain decreases. For $x = 16\%$, the decrease is small.

Appendix Figure A1: Investor types and where they trade

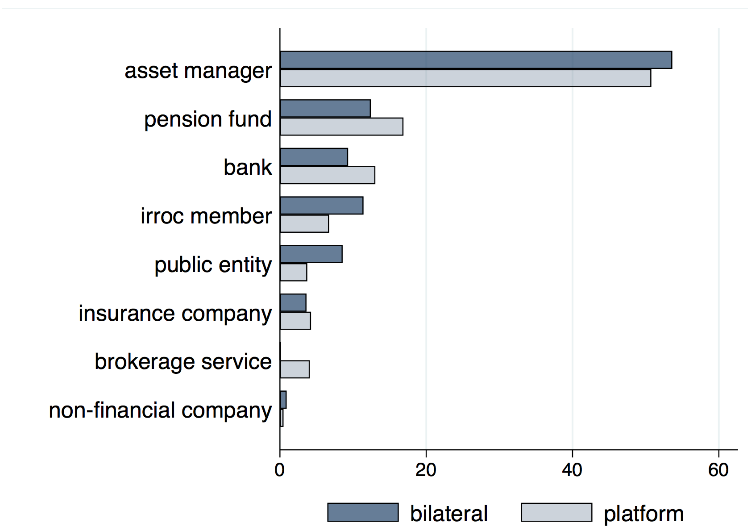


Appendix Figure A1 shows how much each investor type trades on the platform versus bilaterally as a percentage of the total amount investors trade.

Appendix Figure A2: Distribution of trade sizes

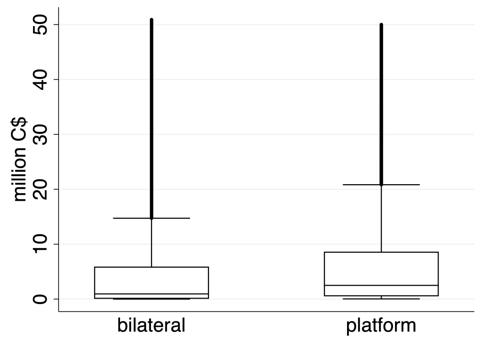


Appendix Figure A2 shows the distribution of trade sizes (in million C$) for bilateral and platform trades, excluding the lower and upper 5th percentile.

Appendix Figure A3: Trade volume per market segment

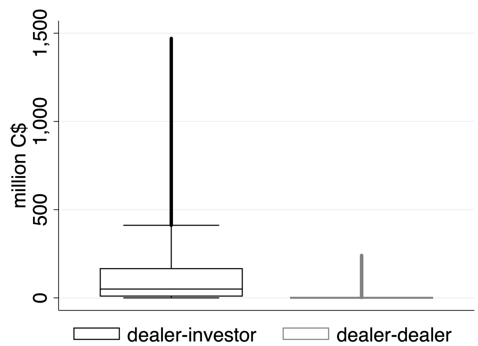


Appendix Figure A3 shows the distribution of daily trade volume per security between dealers and investors (in black) and between dealers (in gray), excluding the upper 5th percentile.

Appendix Figure A4: Number of times an investor trades in a week

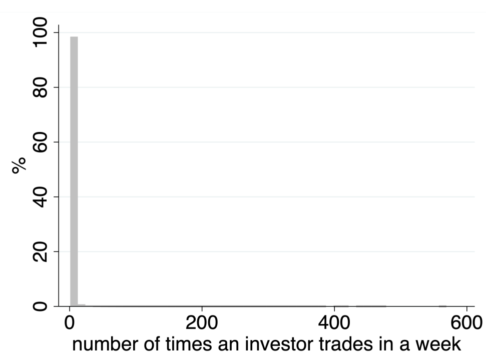


Appendix Figure A4 shows a probability density histogram of the number of times an investor trades—i.e., either sells or buys—in a week. The graph is similar when counting how many trading venues (bilateral vs. platform) an investor uses in a week.

Appendix Figure A5: Event study—observable trade behavior

(a) Trade size

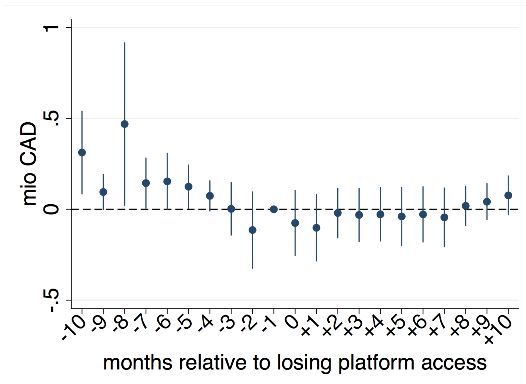


(b) Weeks-to-maturity of traded bonds

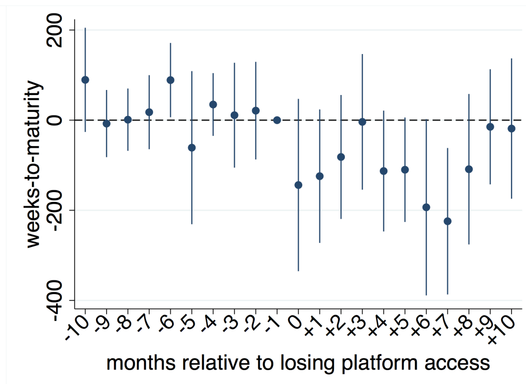


(c) Duration of traded bonds

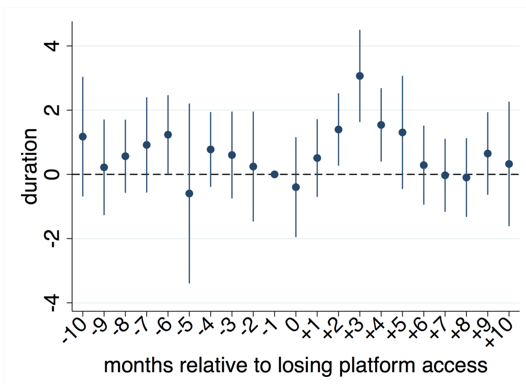


(d) Convexity of traded bonds

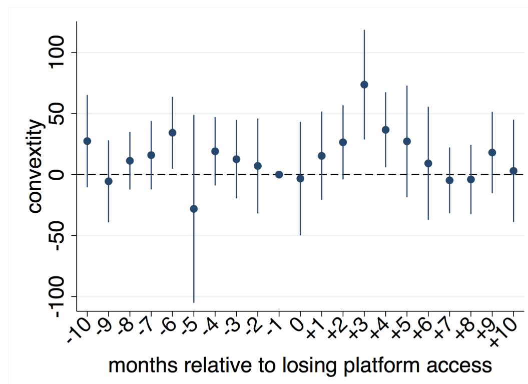


Appendix Figures A5a–A5d visualize changes in observable trade behavior when the investor loses platform access. They show the $\beta_m$ estimates and the 95% confidence intervals of regressions that are similar to the event study regression (3), but with outcome variables that capture trade behavior. For Figure A5a we regress $q_{thsij} = \zeta_i + \sum_{m=Mi^-}^{Mi^+} \beta_m D_{mi} + \zeta_{th} + \zeta_s + \zeta_j + \epsilon_{thsij}$ to see whether the amount traded changes. We see that the average trade size four months before the investor loses platform access ($m = -3$) is statistically indistinguishable from the trade sizes afterwards. Figures A5b to A5d illustrate whether the investor trades bonds with different characteristics; namely, the length to maturity, the duration (which approximates the bond's price sensitivity to changes in interest rates), and the convexity (which measures by how much the duration of the bond changes as interest rates change). In these regressions, we exclude the security fixed effect that absorbs any security specific unobservable in the main event study regression because it would absorb any characteristic of the bond. All standard errors are clustered at the investor level. The graphs look similar when we look at the number of dealers with whom investors trade and how often or how much investors trade in a month.

Appendix Figure A6: Model fit–Market concentration

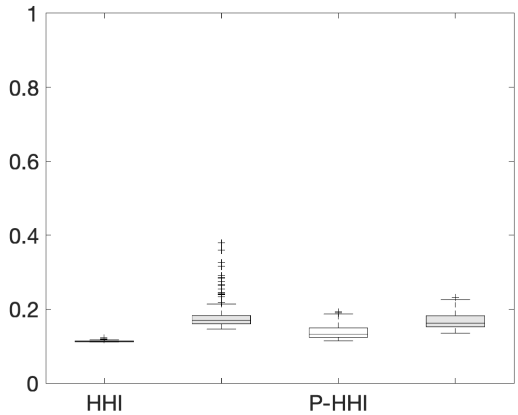


Appendix Figure A6 shows box plots of the Herfindahl Index of the entire market (HHI) and on the platform (P-HHI) in the data (in gray) and in the model (in white), both in percentage points. The distribution is taken over periods. Given the market structure, in which each dealer has its own investor base, the over-all HHI is relatively low both on and off the platform. This does not imply that dealer market power is low. In the extreme in which all investors traded bilaterally with only their home dealer, the HHI would suggest that the market is perfectly competitive since the market is split equally across dealers, even though each dealer has monopoly power over its investors. Our model slightly under-predicts the HHI because we assume that each dealer's home investor base has the same size, while in the data some dealers have a larger investor base than others.

Appendix Table A1: Sample restrictions

| Restrictions | Sample size | Size ↓ in % |
|---|---|---|
| All dealer-to-investor trades | 1,948,764 | |
| w/o extreme yields | 1,914,031 | 1.78% |
| w/o in-house trading | 1,668,520 | 12.82% |
| w/o errors in trading venue | 1,620,148 | 2.89% |
| w/o out of business hours | 1,523,037 | 5.99% |
| w/o false investor-type indicator | 1,517,714 | 0.34% |
| w/o trades after July 2019 (model only) | 1,346,462 | 11.28% |
| w/o non primary dealers (model only) | 1,252,718 | 6.96% |
| w/o one of the primary dealers (model only) | 1,139,412 | 9.04% |
| w/o trades prior announcement (model only) | 1,003,542 | 11.92% |

Appendix Table A1 summarizes how we restrict the raw data. We exclude extreme yields and trades by institutions that are likely reported as institutional investors but are retail or vice versa. Further, we exclude in-house trades, trades that are not realized on CanDeal or bilaterally, and trades that occur out of business hours. For the structural estimation, we focus on trades with primary dealers only. We exclude trades after July 2019 because our auction data do not cover the second half of 2019. Lastly, we exclude trades on auction dates prior to the auction announcement and drop one dealer due to data reporting.

Appendix Table A2: Balancing positions in the inter-dealer market

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| $\text{net\_supply}^C_{tsj}$ | $-0.161$ | $-0.161$ | $-0.162$ | $-0.162$ |
| | (0.002) | (0.045) | (0.045) | (0.045) |
| dealer fixed effect ($\zeta_j$) | – | ✓ | ✓ | ✓ |
| date fixed effect ($\zeta_t$) | – | – | ✓ | ✓ |
| security fixed effect ($\zeta_s$) | – | – | – | ✓ |
| Constant | 0.259 | 0.259 | 0.260 | 0.261 |
| | (0.255) | (0.105) | (0.105) | (0.112) |
| Observations | 177,165 | 177,165 | 177,164 | 177,157 |
| Adjusted $R^2$ | 0.026 | 0.027 | 0.022 | 0.021 |

Appendix Table A2 shows how much dealers balance out their inventory positions by trading with other dealers or brokers. We regress a dealer's net supply in the inter-dealer market, $\text{net\_supply}^D_{tsj}$, on the dealer's net supply in the dealer-investor market, $\text{net\_supply}^C_{tsj}$. The net supply (in million C\$) is the total amount the dealer $j$ sold of a security $s$ on a day $t$ minus the total amount this dealer bought in each market segment, respectively. Column (1) shows an OLS regression: $\text{net\_supply}^D_{tsj} = \text{net\_supply}^C_{tis} + \epsilon_{tis}$. The coefficient tells us that a dealers buys on average C\$ 0.16 million more in the inter-dealer market when she supplies C\$1 million more in the dealer-investor market. The result is robust when adding fixed effects in columns (2)-(4). Standard errors are clustered at the dealer level in columns (2)-(3) and at the dealer-security level in column (4).

Appendix Table A3: Yields and supply shocks

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| won | $-0.034$ | $+0.007$ | $+0.198$ | $+0.375$ |
| | (0.062) | (0.052) | (0.028) | (0.030) |
| Observations | 152,871 | 158,310 | 79,943 | 85,183 |
| Adjusted $R^2$ | 0.730 | 0.685 | 0.315 | 0.355 |

Appendix Table A3 shows the $\gamma$ coefficient of regression $y_{thsij} = \alpha + \beta\theta_{ths} + \gamma won_{\tilde{t}j} + \zeta_j + \zeta_i + \zeta_t + \epsilon_{thsij}$—which regresses the bilateral yield of a trade (in bps) between investor $i$ and dealer $j$ in hour $h$ of date $t$ with security $s$ on the bond's market value, the dealers' unexpected supply shocks (in billion C\$), dealer, investor, and date fixed effects—for buying investors in column (1) and for selling investors in column (2). The point estimate is close to 0 and statistically insignificant, suggesting that dealers don't adjust bilateral yields in response to supply shocks. Instead, if we replace bilateral yields with platform yields, the coefficient increases and becomes statistically significant, as shown in columns (3) and (4) for buying and selling investors, respectively. Both findings are in line with our model assumptions.

Appendix Table A4: Average trade probabilities

| **With loyalty/Switching cost** | | **Acc I** | **Acc IR** | **No fee** | **Central** |
|---|---|---|---|---|---|
| Imperfect competition | Pr(P\|I) | 40 | 41 | 63 | 100 |
| | Pr(H\|I) | 91 | 91 | 87 | 78 |
| | Pr(P\|R) | 0 | 60 | 72 | 100 |
| | Pr(H\|R) | 100 | 86 | 84 | 78 |
| Perfect competition | Pr(P\|I) | 79 | 79 | 88 | 100 |
| | Pr(H\|I) | 91 | 91 | 90 | 87 |
| | Pr(P\|R) | 0 | 80 | 85 | 100 |
| | Pr(H\|R) | 100 | 91 | 90 | 87 |
| **No loyalty** | | **Acc I** | **Acc IR** | **No fee** | **Central** |
| Imperfect competition | Pr(P\|I) | 62 | 87 | 93 | 100 |
| | Pr(H\|I) | 45 | 23 | 18 | 11 |
| | Pr(P\|R) | 0 | 84 | 88 | 100 |
| | Pr(H\|R) | 100 | 25 | 21 | 11 |
| Perfect competition | Pr(P\|I) | 92 | 92 | 96 | 100 |
| | Pr(H\|I) | 19 | 19 | 16 | 11 |
| | Pr(P\|R) | 0 | 87 | 91 | 100 |
| | Pr(H\|R) | 100 | 22 | 19 | 11 |

Appendix Table A4 displays the average trade probabilities in percentage of an investor in all counterfactuals: Pr(P|I) is the probability for an institutional investor I to trade on the platform. Pr(H|I) is the probability of trading with the home dealer, similarly for retail investors R. In With loyalty the loyalty benefit enters the investor's value; in Switching cost the loyalty benefit does not directly affect the investor's value but distorts her trading decisions; in No loyalty the loyalty benefit is eliminated. Counterfactual Acc I allows institutional investors to access the platform at estimated platform usage costs; Acc IR allows all investors to access the platform at estimated costs; No fee eliminates the platform fee; Central shifts all trades are on the platform. For all counterfactuals, we distinguish between imperfect and perfect (platform) competition.

Appendix Table A5: Average welfare gain compare to a fully decentralized market with shrunk dealer values

| **With loyalty** | **Acc I** | **Acc IR** | **No fee** | **Central** |
|---|---|---|---|---|
| Imperfect competition | −1.50 | −1.76 | −2.54 | −3.81 |
| Perfect competition | +0.97 | +1.08 | +1.22 | +1.57 |
| **Switching cost** | **Acc I** | **Acc IR** | **No fee** | **Central** |
| Imperfect competition | +3.64 | +4.24 | +5.83 | +10.14 |
| Perfect competition | +7.30 | +8.18 | +9.40 | +12.28 |
| **No loyalty** | **Acc I** | **Acc IR** | **No fee** | **Central** |
| Imperfect competition | +10.91 | +17.50 | +18.72 | +20.91 |
| Perfect competition | +22.17 | +24.55 | +25.70 | +27.68 |

Appendix Table A5 is analogous to Table 6 but when we shrink dealer values towards the market value, assuming that the dealer's value is $\theta_t+(1-0.16)(\hat{v}_{tj}^D-\theta_t)$. The table displays the average welfare gain (in %) in all counterfactuals relative to a fully bilateral market. In With loyalty the loyalty benefit enters the investor's value; in Switching cost the loyalty benefit does not directly affect the investor's value but distorts her trading decisions; in No loyalty the loyalty benefit is entirely eliminated. Counterfactual Acc I allows institutional investors to access the platform at estimated platform usage costs; Acc IR allows all investors to access the platform at estimated costs; No fee eliminates the platform fee; Central shifts all trades are on the platform. Dealer (platform) competition is either imperfect or perfect, i.e., $v_{tj}^D = q_{tj}$ and $\sigma = 0$. The baseline welfare of the fully bilateral market is about C$ 691 million per year when the loyalty benefit enters welfare directly and about C$ 381 million per year when it does not.

Appendix Table A6: Estimates of IV regression (18) for different specifications

(a) Benchmark model with loyalty benefits

| | First stage-1 | | First stage-2 | | IV-1 | | IV-2 | | OLS | |
|---|---|---|---|---|---|---|---|---|---|---|
| won | +0.528 | (0.132) | +0.159 | (0.029) | | | | | | |
| $r/\sigma$ | −0.036 | (0.021) | −0.034 | (0.021) | +4.184 | (0.061) | +4.162 | (0.059) | +4.127 | (0.057) |
| $1/\sigma$ | | | | | +1.765 | (0.599) | +1.172 | (0.443) | +0.111 | (0.043) |
| Observations | 3,375 | | 3,375 | | 3,375 | | 3,375 | | 3,375 | |
| Adjusted $R^2$ | 0.284 | | 0.287 | | 0.788 | | 0.787 | | 0.787 | |

(b) Benchmark model with no loyalty benefits

| | First stage-1 | | First stage-2 | | IV-1 | | IV-2 | | OLS | |
|---|---|---|---|---|---|---|---|---|---|---|
| won | +0.647 | (0.188) | +0.191 | (0.043) | | | | | | |
| $1/\sigma$ | | | | | +0.686 | (0.245) | +0.406 | (0.192) | −0.158 | (0.019) |
| Observations | 1,897 | | 1,897 | | 1,897 | | 1,897 | | 1,897 | |
| Adjusted $R^2$ | 0.220 | | 0.223 | | 0.775 | | 0.775 | | 0.779 | |

(c) Extended model with small and large trades

| | First stage-1 | | First stage-2 | | IV-1 | | IV-2 | | OLS | |
|---|---|---|---|---|---|---|---|---|---|---|
| won | +0.586 | (0.143) | +0.180 | (0.032) | | | | | | |
| $r^S/\sigma$ | −0.079 | (0.028) | +0.116 | (0.029) | +2.468 | (0.105) | +2.536 | (0.0974) | +3.968 | (0.055) |
| $r^L/\sigma$ | +0.115 | (0.029) | −0.077 | (0.028) | +4.073 | (0.066) | +4.028 | (0.0611) | +2.637 | (0.088) |
| $1/\sigma$ | | | | | +1.443 | (0.484) | +0.853 | (0.352) | +0.015 | (0.031) |
| Observations | 5,541 | | 5,541 | | 5,541 | | 5,541 | | 5,541 | |
| Adjusted $R^2$ | 0.171 | | 0.173 | | 0.738 | | 0.737 | | 0.735 | |

(d) Benchmark model with loyalty benefits—daily versus weekly

| | First stage-w | | First stage-d | | IV-w | | IV-d | |
|---|---|---|---|---|---|---|---|---|
| won | +0.530 | (0.131) | +0.403 | (0.106) | | | | |
| $r/\sigma$ | +0.020 | (0.022) | +0.069 | (0.021) | +4.131 | (0.063) | +2.879 | (0.041) |
| $1/\sigma$ | | | | | +1.780 | (0.625) | +0.953 | (0.375) |
| Observations | 3,375 | | 12,990 | | 3,375 | | 12,990 | |
| Adjusted $R^2$ | 0.213 | | 0.217 | | 0.730 | | 0.696 | |

Appendix Tables A6a–A6c show the estimates of IV regression (18), which regresses platform market shares on quotes, for our benchmark model with and without loyalty benefits, and for the extended model with small and large trades. In column (OLS), we do not instrument the quotes. In columns (IV-1) and (IV-2), we use our constructed supply shocks and the amount a dealer actually won at primary auction as instruments, respectively. The first stage estimation results, from regressing the quotes on the instrument, are shown in columns (First stage-1) and (First stage-2), respectively. In Appendix Tables A6d we omit dealer fixed effects and run the IV regression for the benchmark model with loyalty benefits with daily market shares and quotes in columns (First stage-d) and (IV-d), and weekly market shares and quotes in columns (First stage-w) and (IV-w). The instruments are in billion C\$, quotes are in bps. Standard errors are in parentheses.