Centralizing over-the-counter markets?*

Jason Allen[a] Milena Wittwer[b]

December 31, 2021

Abstract

In traditional over-the-counter markets, investors trade bilaterally through intermediaries (dealers). We assess whether and how to shift trades onto a centralized platform with trade-level data on the Canadian government bond market. We show that investors who have access to a platform pay better prices than investors who do not, and specify a model to quantify price and welfare effects from making platform access universal. We find that not all investors would use the platform—it is costly, dealer competition is low, and investors value close relationships with dealers. Removing platform fees or increasing dealer competition can increase welfare by 22%.

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, Or Shachar, Paulo Somaini, Alberto Teguia, Robert Wilson, Ali Yurukoglu, and all participants in the IO, finance, and macro student workshops at Stanford. Correspondence to: [a]Jason Allen (Bank of Canada): jallen@bankofcanada.ca, [b]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. Unlike centralized markets, such as stock exchanges, OTC markets are considered to be decentralized because buyers must search for sellers one by one in order to trade. 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 accused dealers of abusing market power when trading with investors, combined with dramatic events during the COVID-19 crisis, raised questions regarding whether and how to centralize OTC markets.[1] A popular proposal is to shift trading onto multi-dealer platforms, on which investors run auctions with dealers. Yet, even though this approach has already been adopted in some markets, it is unclear whether it has sizable effects on prices and welfare because of poor platform design.[2]

We assess dealer market power and evaluate price and welfare effects from centralizing OTC markets with trade-level data on the Canadian government bond market. This market is considered to be close to efficient because it is highly liquid and features low price uncertainty. Nevertheless, we document sizable price deviations from the market value of bonds, referred to as markups. We also show that large (institutional) investors, who have access to a platform, pay systematically lower prices than small (retail) investors, who do not. Our main analysis quantifies the role of platform access in driving the price gap and the changes in market outcomes and welfare that could result if platform access were universal.

Three features render the Canadian government bond market a particularly attractive setting for our research question. First, government bonds are more homogeneous than other securities. We can therefore rule out confounding factors that might drive markups and identify dealer market power. Second, a multi-dealer platform—which

[1] Logan (2020) discusses the events in the U.S. Treasury market during the COVID-19 crisis and possible market reforms. For an overview of antitrust litigation, see "The Manipulation Monitor" at www.schlamstone.com, accessed on 12/30/2021.

[2] Examples of legislations that promote multi-dealer platforms are the Dodd-Frank Act, which mandates that standarized derivatives must be traded on platforms called swap execution facilities (SEFs), and the European Union's Markets in Financial Instruments Directive II (MiFID II).

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 prior studies on government bond markets.[3]

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 know whether a trade was executed bilaterally or on the platform and whether an investor is retail or institutional. In addition, we collect bid and ask quotes that are posted on Bloomberg and therefore publicly available. They indicate at what prices investors can trade, and serve as market values of each bond.

Our trade-level data allow us to document novel facts. A typical investor trades bilaterally with a single dealer (the investor's home dealer) and with more than one dealer on the platform. Dealers charge markups over market value which vary across investors and are systematically smaller for institutional than for retail investors. To test whether this is because of platform access, we use an event study design to show that the prices of an investor who (exogenously) loses platform access drop by an amount that is eight times the bid-ask spread. This raises the possibility that making platform access universal could lead to better prices. It neglects the fact, however, that dealers might respond by adjusting platform quotes.

To assess all price effects and quantify welfare gains when centralizing the market, we introduce a model, in which dealers and investors have different values for realizing trade. They play a two-stage game. First, dealers simultaneously post indicative quotes at which they are willing to trade on the platform. Then, institutional investors can enter the platform (at a cost) and run an auction among dealers, which determines the platform trade price. Alternatively, institutional investors trade bilaterally. Retail investors can only trade bilaterally. In our benchmark model, we normalize the trade

[3]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 (2021)).

surplus an investor obtains in a bilateral trade to zero. In a model extension that helps capture dealer-investor relationships, we allow investors to obtain a loyalty benefit whenever trading with their home dealer.

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 main findings are threefold. First, to quantify how much of the price gap between retail and institutional investors comes from platform access, we conduct a counterfactual in which we allow all retail investors to trade on the platform. Using the benchmark model, we find that the price gap would reduce by 32%–47%. The magnitude of the price change is similar to the outcome of the event study, even though we do not use any information about the event study to estimate the model.

Second, in the status quo and the counterfactual many investors choose not to use the platform because dealer competition on the platform is weak and it is costly to trade on the platform. A back of the envelop calculation, which relies on the model extension, reveals that the platform cost splits into three main components. First, investors have to pay a monthly fee to use the platform (31%–43% of the total cost). Second, investors forgo the loyalty benefit when trading with a dealer that is not their home dealer on the platform (43% of the total cost). Lastly, sharing information about the trade with multiple dealers, as required by the platform, is costly because it harms future trade prices (13% of the total cost).[4]

Third, welfare improves substantially when reducing platform costs or increasing

[4]This is in line with the concern raised by industry experts that investors are deterred from using platforms because they are reluctant to reveal their names. Dealers are accused of "long [having] arranged trades bilaterally with investors away from platforms" (*Financial Times* (2015)). One worry is that "the loss of anonymity deters access to platforms in practice" (Managed Funds Association (2015), p. 2).

platform competition, both in the benchmark and extended model. As a starting point, when we allow all investors to trade on the platform at the estimated platform usage cost, the total gains from trade, our measure of welfare, increase by 5%–10%, or C$ 92–C$ 127 million per year. The reason is that more trades are intermediated by dealers who have high values to trade—for instance, because they seek to offload their inventory positions—as more investors trade on the centralized platform. This effect becomes stronger, the higher platform participation, which increases the lower the platform usage costs or the higher platform competition. For instance, eliminating the platform fee or the cost of information sharing by making the platform anonymous brings an additional welfare gain of 4%–12%, or C$ 81–C$ 156 million per year.

Taken together, our results have valuable policy implications for OTC markets. They help explain why the two-tier market structure with bilateral and platform trading that we observe in most OTC markets is so persistent.[5] Further, they highlight three sources of inefficiency. First, dealer market power alone would distort welfare by 34% if the market was otherwise frictionless. Second, dealer-investor relationships can be distortive. This happens when an investor buys from (sell to) a dealer because she keeps a close relationship with this dealer, even though this dealer receives a low benefit for realizing the trade, for instance, because the dealer is already short (long). Lastly, there are frictions that prevent dealers from buying and selling. These became visible 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. 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) in which welfare gains are likely larger.

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—is far from the first-best. The degree of efficiency is much 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

[5] A survey conducted by the Bank of International Settlements shows the split between bilateral versus platform trading in various OTC markets (see BIS (2016), Graph 3).

by Uber, Airbnb, Amazon, and others require sharing information, which consumers (here, investors) might dislike. Further, platforms might not be designed to foster competition, because they are owned by profit-maximizing firms (here, dealers).

Related literature. Our main contribution is to empirically assess price and welfare effects when OTC markets are centralized. Therefore, we add to a literature that analyzes different aspects of decentralized or centralized financial markets, typically via reduced-form analysis.[6] The most 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 market for German government bonds, and complement Kondor and Pintér (2021), who show that investors who trade with more dealers on a day obtain better prices in the UK government 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

[6] Examples include Barklay 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 (2021); Hendershott et al. (2021a,b); Kutai et al. (2021); Lehar and Parlour (2021); O'Hara and Zhou (2021).

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 become 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)).[7] Different from these papers, we highlight the 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 (2019), who study market fragmentation in OTC markets when there is no centralized platform, and to Baldauf and Mollner (2020), who show that it can be theoretically optimal for an investor to disclose information when running an auction. This is in line with our empirical findings.

Paper overview. The paper is structured as follows: Sections 2 and 3 describe the institutional environment and the data, respectively. Section 4 provides descriptive evidence that motivates the need for market reforms. Section 5 introduces the model, which is estimated in Section 6. Section 7 presents the estimation results that of the basis for the welfare analysis in Section 8, Section 9 summarizes the robustness analysis conducted in the Appendix, and Section 10 concludes.

[7]Following the literature, we assume exclusive participation per market segment. Only recently, Dugast et al. (2019) generalizes Atkeson et al. (2015) to allow for endogenous participation decisions in multiple market segments.

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[8]). 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. Many also work as asset managers. 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 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).

[8]Accessible at: www.iiroc.ca/industry/rulebook/Pages/default.aspx.

These platforms are called alternative trading systems in the U.S. and Canada, multilateral trading facility in Europe, and dark pools for equities.[9] 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 are 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.[10] 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 investors trade. If an investor is interested in trading a bond, she typically consults Bloomberg (or another information provider) to check dealer-advertised prices and prices that reflect the bond's current market value.

Next, the investor contacts her dealer, traditionally over the phone or by text. The dealer makes a take-it-or-leave-it 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. Most investors have a single dealer with whom they trade bilaterally, either because search costs are high as it takes time and effort to solicit offers from multiple dealers one-by-one or because keeping a close relationship with one dealer is valuable (see Figures 1a and 1b). We call the dealer with whom an investor trades most frequently, the investor's home dealer.

As an alternative, institutional investors can trade on the platform, CanDeal.

[9]For an overview of platforms see Bessembinder et al. (2020).

[10]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).

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

(a) # of dealers in bilateral trades

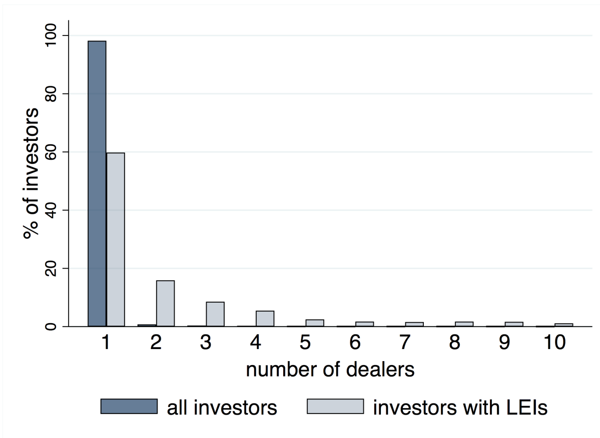

(b) Bilateral trades with the home dealer

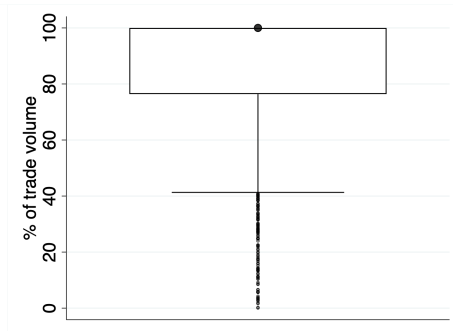

Figure 1a shows how many investor IDs trade bilaterally with 1,. . .,10 dealers, in blue as percentage of all IDs. In gray, we restrict the sample to the 1,459 largest investors that we can trace through the market because they are reported with a legal identifier (LEI). Figure 1b indicates that even these investors typically trade bilaterally with the same dealer (home dealer). It shows a box plot of the amount an investor with a LEI trades bilaterally with her home dealer relative to the total amount this investor trades bilaterally. The median is at 100% The graph looks similar with trade frequencies.

(c) # of dealers on the platform

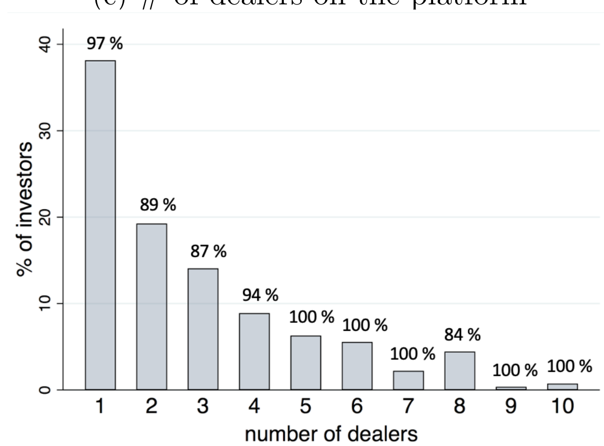

Figure 1c shows how many investors with a LEI who trade at least 10 times on the platform trade with 1,. . .,10 dealers on the platform. The numbers above each bar show the likelihood that the investor's home dealer is included in the set of dealer with whom the investor trades on the platform. We restrict the sample to highlight that investors trade with multiple dealers on the platform. Investors who trade x<10 times can only trade with x different dealers, even if these investors chose a different dealer each time.

Here, they 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.[11]

In an RFQ auction, an investor typically sends a request for quote to multiple (in Canada maximally four) dealers.[12] 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. The investor chooses the deal that she likes best and the trade is executed shortly after. As a result, over 60% of investors trades with more than one dealer on the platform (see Figure 1c).

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.

Platform usage costs. To use the platform, an eligible investor has to pay a monthly fee. However, industry experts have raised concerns that the actual costs are indirect, because—despite platforms appearing to be an attractive alternative to bilateral trading—a relatively small fraction of trades actually occurs on platforms in many markets (e.g., McPartland (2016)). In our case, only about 35% of the institutional investors realize trades on the platform on a typical day.

One cost factor might come from the fact that the RFQ auction is not anonymous. Industry experts have warned that the fact that investors have to reveal their name

[11] In rare occasions a dealer and investor pair negotiates the trade bilaterally but executes it on the platform so as to benefit from the straight through electronic processing that the platform offers. Our data does not allow us to identify these trades. We therefore asked CanDeal, who estimate that about 95% of trades go via RFQ auctions. This implies that less than 5% are negotiated bilaterally or executed on a limit order book.

[12] This is known from previous research (e.g., Hendershott and Madhavan (2015), Riggs et al. (2020)). We cannot verify this with our data since we only observe realized trades, similar to observing the winning but not the losing bids of an auction. Therefore, we looked into annual reports that CanDeal publishes: on average 58% of auctions are with four, 30% with three and 12% with two dealers in our sample period. The reports are available at: `https://www.candeal.com/en/news`.

and trade information when running the auction deters them from using the platform, for instance, because dealers can use such information against the investor and offer worse trade prices in the future (Managed Funds Association (2015)).

Another factor could be a relationship "cost": When trading with a new dealer, the investor might forgo services (including low margins or haircuts, account administration, transaction settlements, collection of dividends and interest payments, tax support, and foreign exchange management) that she would obtain if she traded with a dealer with whom she keeps a close relationship, such as a custodian bank. A strong relationship factor would be in line with existing research (e.g., Hendershott and Madhavan (2015); Hendershott et al. (2020a, 2021b))) and could explain why it is not uncommon for an investor to trade with the same dealer on and off the platform, even though search costs on the platform are minimal (see Figure 1c).

3 Data

Our main data source contains trade-level information on all government bond trades of registered brokers or dealers. We augment this data with additional data sources.

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 each dealer (and broker) carries a unique legal identifier (LEI) and thus can be traced throughout the market. Investors either have an LEI or an anonymous dealer-specific account ID. Trades with investors that have a LEI represent 25% of all trades with investors and capture 43% of the trade volume. Besides the LEI, we observe no other characteristics of an institution, such as balance 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.[13] 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 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.

Third, we collect information of the platform (CanDeal). We see the indicative quotes and execution prices on CanDeal for all Treasury Bills and both sides of the trade. Unfortunately, CanDeal does not report dealer identifiers so that we cannot merge their information with our base data set. We use their data only to show that a dealer's indicative ask (bid) quote is essentially identical to the average price at which an investor buys (sells) on the platform (see Appendix Figure A2). This motivates us to use the daily average trade price at which a dealer executes a buy/sell-side trade on the platform (reported in MTRS2.0) as a proxy for each dealer's daily indicative ask/bid quote when estimating our model.

Lastly, we collect the mid-quote (the average between the bid and the ask quote) posted on Bloomberg (BNG) 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 BNG mid-price 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-

[13]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.

dealer price directly because the inter-dealer market isn't liquid enough so that not all securities that are actively traded in the dealer-to-investor market also have an inter-dealer market price.

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

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. We focus on primary dealers and drop trades after July 2019 because we do not have auction data for the second half of 2019. 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. Appendix Table A1 summarizes all sample restrictions.

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%. This is a relatively large yield difference because of the low interest rate level throughout our sample. As comparison, the median yield of a bond is about 150 bps, and the median bid-ask spread is about 0.5 bps.

Normalization. The yield (and price) of a bond might be affected by many factors, and explaining all of them in a single model is intractable. Our approach, instead, is to control for factors that are not endogenous in our model. We do this by regressing the $yield_{thsij}$ 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, a flexible function of trade size, here $f(quantity_{thsij}) = \sum_p^3 (quantity_{thsij})^p$, an hour-day fixed effect, and a security-week fixed effect. We find that our results are not sensitive to how we control for trade size, and construct the residual from this regression. In addition, we normalize the Bloomberg yield by subtracting the estimated hour-day and security-week fixed effects.

We label the residualized trade yield y_{thsij} and the normalized Bloomberg yield θ_{ths}. For consistency, we use these normalized yields throughout the paper—that is, for the reduced-form evidence, as well as to estimate our model. However, our reduced-form evidence is robust to using the yields in the raw data rather than the normalized yields.

Key market features. The typical trade between a dealer and one of the 546,048 investor IDs is small (see Appendix Figure A3). It involves a bond that is actively traded, because it was issued in a primary auction less than three months ago and is on-the-run. Bid-ask spreads are narrow (0.5 bps at the median), and it takes only 0.13 (2.8) minutes between an investor who buys and an investor who sells (the same security) on a day. Taken together, the market is highly liquid.

4 Descriptive evidence: Why market reforms?

Our trade-level data allow us to document yield differences across investors, which suggests that the market is imperfect.

Markups and yield gap. To analyze whether dealers charge markups over the market value (θ_{ths}), we define the markup as:

$$(y_{thsij} - \theta_{ths})^+ = \begin{cases} y_{thsij} - \theta_{ths} & \text{when the investor buys} \\ \theta_{ths} - y_{thsij} & \text{when the investor sells.} \end{cases} \tag{1}$$

The higher $(y_{thsij} - \theta_{ths})^+$, the more favorable the yield for the investor, independent of whether she buys or sells.

Figure 2: Yield gap

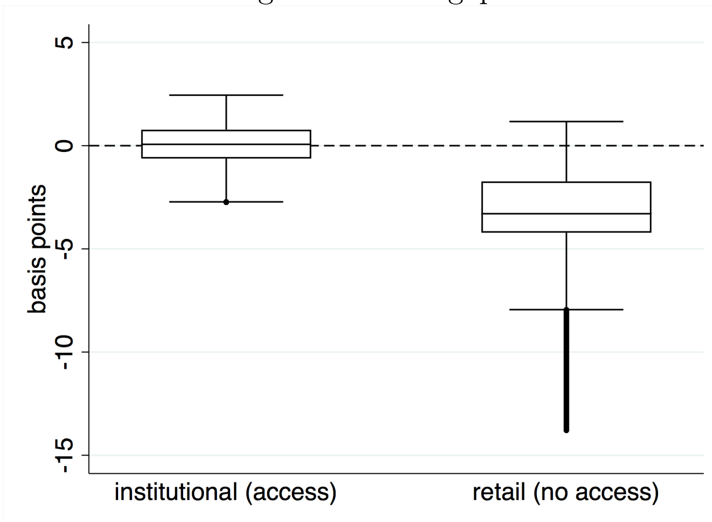

Figure 2 shows a box plot of markups for institutional and retail investors, excluding the upper and lower 5% of the distribution. To construct these markups, we regress $(y_{thsij} - \theta_{ths})^+$ as defined in (1) on indicator variables that distinguish retail from institutional investors ($retail_{thsij}$) and buy-side from sell-side trades (buy_{thsij}). In addition, we control for hour-day (ζ_{th}), security (ζ_s), and dealer (ζ_j) fixed effects. The residual measures how much worse the yield is relative to the market value.

Figure 2 shows that markups vary widely across investors, even when controlling for differences in trade size, security ID, time of trade, and dealer. Furthermore, these markups are systematically smaller for institutional investors—who have access to the platform—than for retail investors, who do not. At the median, a retail investor obtains a yield that is about 4 bps worse than an institutional investor. This amounts to a yearly monetary loss of roughly C$ 34,000 for the average retail investor who trades C$ 86 million per year.

Whether market reforms that shift trading onto the platform could have sizable effects on yields or welfare depends on what drives the yield gap. It could be that retail investors are willing to pay more and therefore realize worse yields. But it could also be that they obtain worse yields because they cannot trade on the platform. In that case, there is scope for market reforms.

Yields and platform access. The ideal but infeasible experiment to establish a causal link between platform access and yields would be to randomly assign platform access to some investors (treatment group) but not all investors (control group). Instead, we leverage the fact that 90 institutions lost the right to access the platform in our sample to conduct an event study.

Investors lose platform access when they lose their institutional status. This can happen for different reasons. First, the investor may no longer hold sufficient capital or may no longer be willing to prove that she does. For instance, the non-financial or financial assets of a firm could lose value so that a firm no longer has a net worth of C$ 75 million, which is the cutoff to classify as institutional investor. Second, the investor may terminate her membership with a regulated entity such as the Canadian Investor Protection Fund (CIPF), which protects investor assets in case of bankruptcy. Third, the investor may stop selling securities, offering investment advice, or managing a mutual fund. Our data do not allow us to disentangle these reasons because switchers are not reported with their LEI, and are therefore anonymous to us.

We define the event of investor i as the first time we observe this investor as a retail investor. We bucket time into months and pool the buy and sell side of the trade to obtain sufficient statistical power. We then test whether the investor obtains a worse yield (y_{thsij}) relative to the market value (θ_{ths}) when losing platform access by regressing

$$(y_{thsij} - \theta_{ths})^{+} = \zeta_i + \sum_{m=Mi^-}^{Mi^+} \beta_m D_{mi} + \zeta_{th} + \zeta_s + \zeta_j + \epsilon_{thsij}, \tag{2}$$

where D_{mi} is an indicator variable equal to 1 m months before/after i loses access and $\zeta_i, \zeta_{th}, \zeta_s, \zeta_j$ are investor, hour-day, security, and dealer fixed effects, respectively.

Our parameters of interest (the β's) are identified from how the trade yields of an investor who loses access change over time when controlling for time, security, and dealer-specific unobservables. The sizes of the hour, security, and dealer fixed effects are pinned down by trade information on retail investors who never obtain access (in our sample), since these investors are likely more similar to those who lost access than to institutional investors. This is because investors with access throughout tend to be large players and clearly meet the regulatory requirements.

We find that investors who lose platform access realize worse yields (see Figure 3). On average the yield drops (per trade) by 1 bps in the first month and decreases

Figure 3: Event study: Yield drop when losing platform access

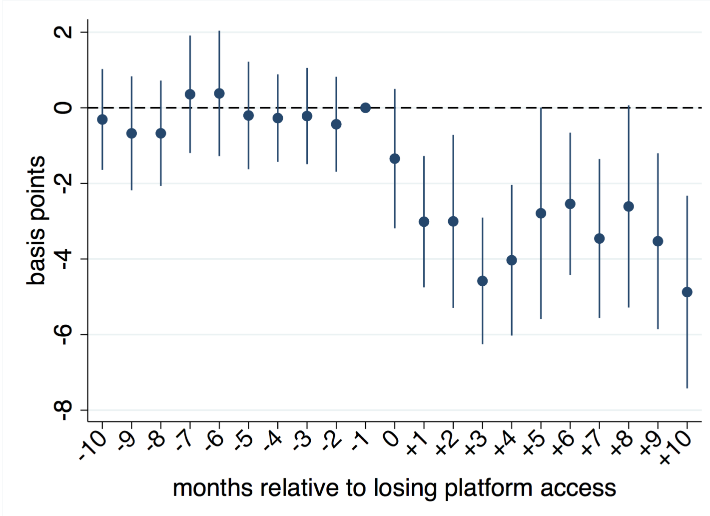

Figure 3 shows the β_m estimates and the 95% confidence intervals of event study regression (2) for 10 months before and after an investor i switched from having to not having access. Each β_m measures by how much the markup, $(y_{thsij} - \theta_{ths})^+$, for investor i differs m months before/after this event relative to 1 month prior to it. Standard errors are clustered at the investor level.

further by about 4 bps thereafter. This translates into sizable monetary losses for investors. The average institutional investor loses about C$ 1.7-6.8 million per year, suggesting that having platform access matters for yields and has a sizable monetary impact on returns.

The relationship would be causal if losing access to the platform were exogenous. This would be the case if an investor's regulatory capital fell 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 the status for potentially endogenous reasons, such as a change in business model. 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 exogenous.

Summary. We have gathered novel evidence that platform access matters for yields. This implies that there is scope to increase investor yields by centralizing the market. We now introduce and estimate a structural model to gain additional insights.

5 Model

Estimating a structural model has three advantages. First, we can quantify total gains from trade, our measure of welfare. For this we need to know the values for realizing trade. In the data, we only observe transaction prices, which lie somewhere between the value of the buyer and the value of the seller. Second, we can account for institutional investors selecting between trading bilaterally and on the platform. This selection problem can bias estimates of OLS regressions. Lastly, we can take into account how dealers and investors respond when centralizing the market.

Without loss of generality, we consider two separate games: one in which dealers sell to investors and one in which dealers buy from investors.[14] 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. To make the price-yield conversion, it helps to keep in mind that the yield is like a negative price. We denote a vector of quotes by $q_t = (q_{t1}...q_{tJ})$ and similarly for all other variables. Random variables are highlighted in **bold**. All proofs are in Appendix C.1.

5.1 Bechmark model

Dealers sell a bond to institutional and retail investors in a two-stage game, which is inspired by how investors trade in this market (see Section 2).

First, dealers simultaneously set quotes to maximize the expected profit from trading with investors. The quotes are posted publicly, and inform investors about the yields they may expect to realize when buying on the platform. This is motivated by the empirical fact whereby individual trade yields on the platform are on average identical to the quotes dealers post on the platform. Then, given these quotes, institutional investors decide whether to buy on the platform or bilaterally with a dealer, while retail investors can only buy bilaterally.[15]

[14]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.

[15]We let all institutional investors consider the platform because we observe no systematic difference between investor types on versus off the platform (see Appendix Figure A1).

In choosing a quote, a dealer faces a trade-off: On the one hand, when the dealer decreases the quote, the platform becomes less attractive for investors. As a result, more investors stay off the platform and buy bilaterally. On the other hand, when she increases the quote, more investors entering the platform buy from her, increasing her platform market share.

An institutional investor also has a trade-off: When she buys bilaterally she has to leave all surplus to the dealer because the dealer discovers her willingness to pay. When entering the platform the investor can extract positive surplus thanks to (more direct) competition between dealers but has to pay a cost to use the platform. As a result, in equilibrium only investors with a high willingness to pay enter the platform.

Formal details. On a fixed day t, $J_t \geqslant 2$ dealers sell a bond to infinitely many 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 κ^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 on any particular day. This number is exogenous and unknown to the dealer until the end of the day.

Each dealer seeks to maximize profit from trading with investors. Ex post, dealer j obtains a profit of $\pi_{tj}(y) = v_{tj}^D - y$ when selling one unit at yield y. Here, $v_{tj}^D \in \mathbb{R}$ is the dealer's value for the bond. It may be driven by current market conditions, expectations about future demand, prices or inventory costs. 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, θ_t. However, since this is unlikely in reality—for instance, because it is costly to hold inventory—we refrain from imposing $v_{tj}^D = \theta_t$.

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

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

is the investor's value, also referred to as willingness to pay. It splits into three elements. The first is the commonly known market value of the bond, θ_t. The second is a privately known liquidity shock, ν^G_{ti}, which is drawn iid from a commonly known distribution with a continuous CDF $\mathcal{F}^G_t(\cdot)$ that has a strictly positive density $f^G(\cdot)$ on the support. It reflects individual hedging or trading strategies, balance-sheet concerns, or the cash needs of an institution. The third element is dealer-specific, ξ_{tj}. It is drawn from an arbitrary distribution with CDF $\mathcal{G}_t(\cdot)$ and absorbs unobservable dealer characteristics (similar to a fixed effect in a linear regression). We label this term "dealer quality", as it captures anything that makes trading with a specific dealer particularly attractive, independent of how the trade is realized. It could, for example, reflect the dealer's probability of delivery, the speed of processing the trade, her ability to hold, or release large quantities or to provide ancillary services.

The game has two stages. In the first stage, dealers simultaneously post indicative quotes at which they are willing to sell on the platform, which is accessible to institutional investors. When choosing the quote, each dealer maximizes the expected profit from selling to institutional investors. Each supposes that they can sell on the platform at the posted quote and forms expectations over how much institutional investors are willing to pay.

In the second stage of the game, all trades realize. In a bilateral trade, the dealer discovers the investor's willingness to pay and offers a yield equal to that.[16] On the platform, each investor runs an auction with all dealers (or, alternatively, with a random subset of dealers) which determines by how much the platform trade yield of an investor differs from the posted quotes.

We formalize such an auction-game in Appendix C.2. Here we only give the main idea, because we cannot estimate this game without bidding data from the platform. Before bidding in an auction for investor i, each dealer j updates her value for the bond, drawing a signal, ϵ_{tij}, from a commonly known distribution with CDF $\mathcal{H}_t(\cdot)$. Each dealer then decides by how much to move away from the posted quote. In equilibrium, her bid is equal to her posted quote plus a stochastic term that depends on the signal, ϵ_{tij}, and a parameter σ:

[16]This implies that the dealer may occasionally accept to trade at a loss. This happens when the investor draws an extreme liquidity shock that lies above the dealer's value, and captures the idea that a dealer is willing to occasionally help an investor in need to sustain their bilateral relationship.

$$q_{tj} + \sigma\epsilon_{tij} \text{ where } \epsilon_{tij} \sim \mathcal{H}_t \text{ and } \sigma \in \mathbb{R}. \tag{3}$$

This auction-game is only one possible micro-foundation for the yield an investor achieves on the platform (3). More broadly, in our estimation, ϵ_{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)) .

Parameter σ measures the degree of competition on the platform and depends, for instance, on the number of dealers who bid in the auction.[17] When $\sigma = 0$, all investors buy from the dealer who offers the best quote and quality, i.e., the one with $\max_j\{\xi_{tj} + q_{tj}\}$. In that case, the platform is perfectly competitive and dealers compete à la Bertrand. If $\sigma \to \infty$, investors buy from the dealer for which the realization of ϵ_{tij} is the highest, regardless of the dealer's quote or quality. In this case, each dealer acts as a monopolist on the platform.

To use the platform and choose from any of the dealers, the investor has to pay a commonly known cost c_t. This represents any obstacle to access the platform, including privacy concerns or relationship costs, and is motivated by the empirical fact that even though platform yields are better than bilateral yields (see Appendix Table A2), on a typical day only 35% of institutional investors use the platform.[18]

In summary, the sequence of events is:

(1) Dealers observe the market value θ_t, their qualities ξ_t and values v^D_{tj}.
They simultaneously post quotes $q_{tj} \in \mathbb{R}^+$.

(2) N^G_t investors of both groups $G \in \{I, R\}$.
Each investor observes $\theta_t, q_{tj} \forall j$ and draws its liquidity shock ν^G_{ti}.
She contacts her home dealer, who observes ν^G_{ti} and offers $y^G_{tid} = v^G_{tid}$.
A retail investor accepts the offer. An institutional investor can accept or enter the platform. In the latter case, the investor pays cost c_t, observes the platform shock ϵ_{tij}, and decides from which dealer to buy at $q_{tj} + \sigma\epsilon_{tij}$.

[17] σ could be shaped by a variety of factors, including limited information- or risk-sharing (as in Boyarchenko et al. (2021)).

[18] The cost can also absorbs differences in the service a dealer provides on versus off the platform. Although uncommon, a bilateral trade could be part of a package or come with additional investment advice.

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

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

$$y_{tid}^R = \theta_t + \nu_{ti}^R - \xi_{td}. \tag{4}$$

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

$$y_{tid}^I = \theta_t + \nu_{ti}^I - \xi_{td} \quad \text{if } \psi_t(q_t) \leqslant \nu_{ti}^I, \tag{5}$$

$$\text{where } \psi_t(q_t) = \mathbb{E}[\max_{k \in \mathcal{J}_t} \tilde{u}_{tik}(\boldsymbol{\epsilon_{tik}})] - \theta_t - c_t \text{ with } \tilde{u}_{tij}(\boldsymbol{\epsilon_{tij}}) = \xi_{tj} + q_{tj} + \sigma\boldsymbol{\epsilon_{tij}}. \tag{6}$$

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

This proposition characterizes where investors buy and at what yields. A retail investor always buys at a yield that equals her willingness to pay (statement *(i)*). An institutional investor trades on the platform if she expects the surplus from buying on the platform minus the platform usage cost, $\mathbb{E}[\max_{k\in\mathcal{J}_t} \tilde{u}_{tik}(\boldsymbol{\epsilon_{tik}}) - (\theta_t + \nu_{ti}^I)] - c_t$, will be higher than the zero surplus she receives in a bilateral trade (statement *(ii)*). 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.

The proposition highlights the fact that yields for institutional investors are higher than for retail investors because they have access to the platform: Those who obtain better yields on the platform buy on the platform; others buy bilaterally.

Proposition 2 (Dealers)**.** *Dealer j posts a quote q_{tj} that satisfies*

$$q_{tj}\Big(1 + \frac{1}{\eta_{tj}^E(q_t)}\Big(1 - \frac{\partial \pi_{tj}^D(q_t)}{\partial q_{tj}}\Big/ S_{tj}(q_t)\Big)\Big) = v_{tj}^D, \tag{7}$$

where $\eta_{tj}^E(q_t)$ is the dealer's yield elasticity of demand on the platform and $\frac{\partial \pi_{tj}^D(q_t)}{\partial q_{tj}}$ 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, $S_{tj}(q_t)$. Formally, $\eta_{tj}^E(q_t) = q_{tj}\frac{\partial S_{tj}(q_t)}{\partial q_{tj}}/S_{tj}(q_t)$ with $S_{tj}(q_t) = \sum_{j\in J_t} \Pr(\boldsymbol{\nu_{ti}^I} \leqslant \psi_t(q_t)) \Pr(\tilde{u}_{tki}(\boldsymbol{\epsilon_{tki}}) < \tilde{u}_{tij}(\boldsymbol{\epsilon_{tij}})\ \forall k \neq j)$, where $\tilde{u}_{tij}(\boldsymbol{\epsilon_{tij}})$, $\psi_t(q_t)$ as in (6), and $\pi_{tj}^D(q_t) = \mathbb{E}[v_{tj}^D - (\boldsymbol{\nu_{ti}^I} + \theta_t - \xi_{tj})|\psi_t(q_t) \leqslant \boldsymbol{\nu_{ti}^I}]$.

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 consisted of the platform only, the dealer's quote would satisfy $q_{tj}(1 + 1/\eta^E_{tj}(q_t)) = v^D_{tj}$. This is equivalent to the classic markup rule of firms that set prices to maximize profit. Each chooses a price that equals its 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.

5.2 Model extension: Loyalty benefit

Our benchmark model abstracts from the possibility that investors might prefer to trade with dealers with whom they have a relationship. To take this feature into account, we extend our model in Appendix C.3 and allow investors to earn a loyalty benefit whenever trading with their home dealer. Similar to the quality of a dealer, the loyalty benefit may capture dealer services related to the trade. The key difference is that the loyalty benefit captures services that are provided only to customers with whom the dealer has a close business relationship.

The extended model helps rationalize why we observe that investors oftentimes trade with only one dealer when trading on the platform (recall Figure 1c).[19] However,

[19]This model predicts that no investor wants to trade bilaterally with a dealer that is

with our data, we can only estimate this model on a small subset of investors who are reported with a LEI. Further, to obtain sufficient power we have to pool trades within a week (rather than a day) which implies that we can no longer account for unobservable movements of prices and demand/supply within the same week. Therefore, we use the model without loyalty benefit—in which the platform cost absorbs any reason for which an investor prefers to trade bilaterally, including a loyalty benefit—as our benchmark. Unless specified otherwise, we refer to the benchmark model when we talk about our model.

5.3 Discussion of simplifying assumptions

Our model builds on several simplifying assumptions. First, because we do not observe failed trades, we assume that the number of dealers and investors who trade on a day is exogenous and that no trade between them fails. We believe that this is not problematic for two reasons. First, empirical evidence suggests that trades of safe assets rarely fail (e.g., Riggs et al. (2020); Hendershott et al. (2020b)). Second, (primary) dealers have an obligation to actively trade: The least active dealer trades on 98% of dates. We can, thus, abstract from market entry and exit of dealers.

Second, our game does not connect multiple days. 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 day 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 Figure 4) and because dealers do not sizably balance their inventory positions by trading with one another or brokers (see Appendix Table A3).[20] To validate that this is a good approximation, we let dealer

not the investor's home dealer if the loyalty benefit is positive. We think that this is a good first-order approximation of the data generating process given we that observe investors trading bilaterally with non-home dealers in rare occasions. To rationalize these occasions, we could assume that dealers are hit by random shocks that prevent them from realizing trade. When a home dealer cannot trade, the investor has to contact a different dealer.

[20]One reason for this is that dealers face similar demand (or supply) from their investors. This implies that there are some days on which all dealers seek to sell more than they seek to

Figure 4: Trade volume per market segment

Figure 4 shows the distribution of daily trade volume (sum of all trades quantities) per security in the dealer-to-investor market (in black) and the inter-dealer market (in gray), excluding the upper 5th percentile.

valuations converge towards the hypothetical value θ_t that would arise with a perfect inter-dealer market as part of our robustness analyzes.

Fourth, in our benchmark model we assume that the dealer offers a bilateral yield that equals the investor's full willingness to pay, leaving the investor with zero trade surplus. You may think of this as a normalization similar to the way we normalize the utility of one choice in discrete choice models since it is typically not possible to identify all utilities attached to each choice but only relative utility differences. 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 A4). In addition, we show that our findings are robust to allowing investors to earn a loyalty benefit whenever trading with her home dealer, which implies that the investor extracts a positive trade surplus in a bilateral trade.[21]

An alternative pricing assumption would be to assume that the dealer sets the bilateral yield in order to render an institutional investor indifferent between trading bilaterally or entering the platform, i.e., $y_t^I = \mathbb{E}[\max_{k \in \mathcal{J}_t} \tilde{u}_{tik}(\boldsymbol{\epsilon_{tik}})] - c_t$. Then, in con-

buy in the inter-dealer market, and vice versa for other days. Another reason could be that there are frictions, such as balance sheet constraints, that hinder dealers from absorbing inventory from each other.

[21] In previous versions of the paper, we also showed that the findings are robust to allowing investors to extract an exogenous fraction $\phi \in [0, 1]$ of surplus in the bilateral trade.

trast to the data, all investors would obtain the same bilateral yield unless they faced different platform entry costs or had different expectations over platform conditions. In our model, yield dispersion arises more naturally because investor's have different liquidity needs, even if costs and expectations are not heterogeneous.

Fifth, motivated by the empirical feature that investors tend to trade with a single dealer bilaterally, we assume in our benchmark model that an investor does not search for other dealers when trading bilaterally. This is equivalent to assuming that (latent) search costs in the bilateral market are too high to render search profitable. In the extended model, it is optimal for the investor to avoid searching for other dealers as long as the loyalty benefit is sizable.

Sixth, given that most trades in our sample are small and of similar size, we assume that all trades have the same size, normalized to one (recall Appendix Figure A3).[22] This implies that our findings are expressed in terms of unit of the trade on an average bond and that we cannot analyze whether and how trade sizes and volume change as we change the market rules.

Finally, we assume there is a single bond because all actively traded bonds in our sample are similar in terms of liquidity and risk.[23] Further, 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 and do not split orders (see Appendix Figure A4).

6 Estimation of the benchmark model

Including both the buy and sell 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^*), respectively. For all of them, we want to estimate the daily distribution of the liquidity shocks (F_t^G), in addition to the daily dealer qualities (ξ_{tj}) and the degree of competition on the platform (σ). We allow the cost of using the platform and the

[22] Appendix Tables A5 and A6 provide additional supporting evidence. In the related theoretical literature, following Duffie et al. (2005), this assumption is common.

[23] 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.

dealer's value to depend on whether the investor buys (c_t, v_{tj}^D) or sells (c_t^*, v_{tj}^{*D}).

Notice that most of the parameters are day-specific. This allows us to nonparametrically account for variation and correlation across days that are driven by unobservable market trends. This is important because the yields and demands for Canadian government bonds are largely affected by global macroeconomic trends. To obtain sufficient statistical power, we also cannot allow the platform usage cost or the taste for dealer quality to be heterogeneous across investors. Instead, we estimate both for the average investor. Details of the estimation of the benchmark and the extended model are in Appendix D.

6.1 Identifying assumptions

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

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

This assumption would be violated if an investor trades more than once in a day 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 day, this is unlikely (recall Appendix Figure A4).

Assumption 2. *W.l.o.g. we decompose dealer quality, ξ_{tj}, into a part that is persistent over time and a part that might vary: $\xi_{tj} = \xi_j + \chi_{tj}$. The time-varying parts, χ_{tj}, capture day- and dealer-specific demand shocks that are unobservable to the econometrician. They are drawn iid across dealers j within a day t.*

Crucially, Assumption 2 does not rule out that dealers may change quotes in response to demand shocks that are unobservable to the econometrician. Formally, ξ_{tj} may be correlated with q_{tj}. 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:

$$\begin{aligned} won_{\tilde{t}j} = &\text{ amount dealer } j \text{ won at the last auction day } \tilde{t} \\ &- \text{amount she expected to win when placing her bids.} \end{aligned} \tag{8}$$

These shocks work as cost-shifter instruments 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.

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 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.1 for details).

Assumption 3. *Conditional on unobservables that drive aggregate demand and supply on day t, ζ_t, and the time-invariant quality of the dealer, ξ_j, the demand shocks, $\boldsymbol{\chi_{tj}}$, are independent of the unexpected supply shocks, $won_{\tilde{t}j}$:* $\mathbb{E}[\boldsymbol{\chi_{tj}}|won_{\tilde{t}j}, \zeta_t, \xi_j] = 0$.

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 date 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 4. *Platform shocks $\boldsymbol{\epsilon_{tij}}$ are iid across t, j, i.*

This assumption is frequent in demand estimation. In our case, it implies that a drop in dealer j's daily quality, ξ_{tj}, leads investors to substitute to any of the other dealers $k \neq j$ with equal likelihood if their daily qualities, ξ_{tk}, remain unchanged but not otherwise. In our model extension, this assumption is further relaxed since the probability that an investor buys from a specific dealer then depends on whether the dealer is the investor's home dealer or not.

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 rely on two parametric assumptions. The first imposes a functional form on the distribution of $\boldsymbol{\epsilon_{tij}}$ that is standard in demand estimation. It implies that the dealer's market shares (on the platform) have a closed-form solution. The second assumption is inspired by the shape of the histogram of shocks, $\hat{\nu}^G_{ti} = y^G_{tij} - \theta_t + \hat{\xi}_{tj}$, for investors who choose to trade bilaterally. It resembles a normal distribution, similar to Figure 6.[24]

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

[24]Without imposing a distributional assumption on the liquidity shocks, we could non-parametrically 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.2.

Key variables. The first set of variables used in the estimation are each dealer's daily market share on the platform (among investors who enter the platform), s_{tj}, and each dealer's daily bilateral market share relative to her platform market share, ρ_{tj}. The second set of variables are the normalized yields of Section 3. We approximate the bond's daily market value observable to everyone, θ_t, by the normalized Bloomberg yield, averaging across securities and hours of the day and the quote at which dealer j sells on a day, q_{tj}, by the average yield at which she sells on the platform on that day, as explained above.

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 2). The main difference would be that the market value would become 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.

Identification. The main identifying variation for the competition parameter and the dealers' qualities comes from how dealers split the platform market on a day.

The competition parameter (σ) is mainly identified from the within-day correlation between dealers' daily platform market shares and their (cost-shifter) supply shocks (see Figure 5). 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 on that day. As σ 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 σ) on the dealers' supply shocks. Hence, the correlation between these market shares and the supply shocks pins down σ.

Figure 5: Dealers' daily platform market shares and their cost shifters

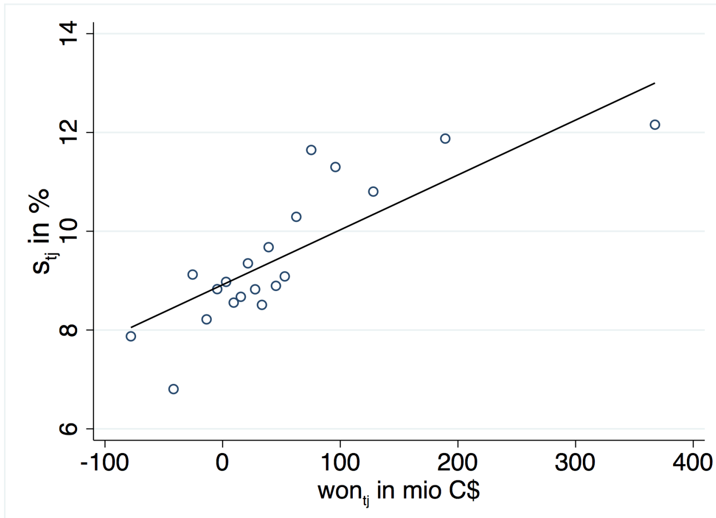

Figure 5 shows a binned scatter plot to visualize the correlation between dealers' daily market shares on the platform (s_{tj}) and their unexpected supply shocks ($won_{\tilde{t}j}$) when partialling out day fixed effects.

Figure 6: Identifying variation for c_t, μ_t^I, σ_t^I

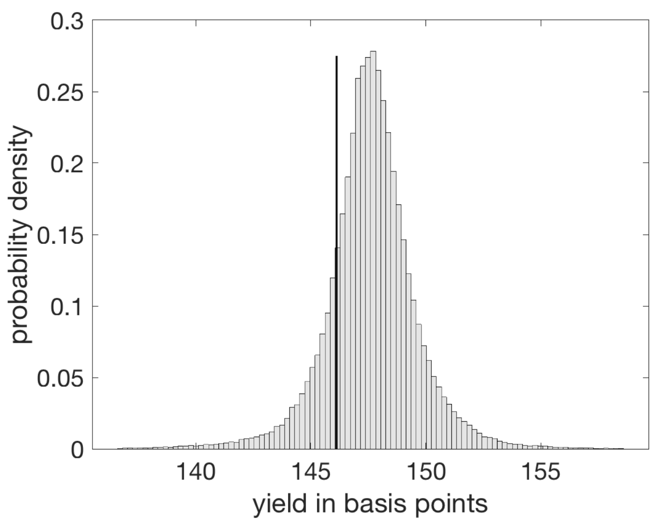

Figure 6 shows a probability density histogram of the yields (in bps) that institutional buyers realize, excluding the upper and lower 0.1 percentile of the distribution, and a black line. This line is the average cutoff (across days and dealers) that determines whether an institutional investor buys bilaterally or on the platform, according to Proposition 1.

The dealers' qualities (ξ_{tj}) are determined by how the dealers split the platform market when posting the same or very similar quotes: Dealers with higher qualities capture a higher market share.[25]

The distribution of the liquidity shocks and the platform usage costs are, for any given day, 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 6. 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. This is because the investor realizes a yield $y^I_{tid} = \theta_t + \nu^I_{ti} - \hat{\xi}_{td}$ when buying bilaterally. Since we observe the trade yield (y^I_{tid}) and market value (θ_t) and we have already estimated dealer qualities ($\hat{\xi}_{td}$), we can solve for the liquidity shock (ν^I_{ti}) pointwise.

Finally, we back out the dealer's value (v^D_{tj}) from the markup equation (7) of Proposition 2. We pick the v^D_{tj} 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 now report the estimates of our benchmark model, and analyze what factors deter investors from using the platform. In the end, we validate that our parsimonious model can replicate the event study—which we have not used to estimate the model—before moving on to the counterfactual exercises.

[25] We expect dealer quality to vary across dealers, because we observe systematic differences in how much of the market each dealer captures when posting the best quote (see Figure 7a).

Table 1: Estimates (median across days)

Buys	$\hat{\mu}^I$	$\hat{\mu}^R$	$\hat{\sigma}^I$	$\hat{\sigma}^R$	$\hat{c}$	$1/\hat{\sigma}$	$\hat{\eta}^E$
	-0.82	-2.92	2.81	5.12	3.46	1.29	$+174.68$
	(0.13)	(0.75)	(0.10)	(0.94)	(0.16)	(0.25)	
Sells	$\hat{\mu}^{I^*}$	$\hat{\mu}^{S^*}$	$\hat{\sigma}^{I^*}$	$\hat{\sigma}^{S^*}$	$\hat{c}_t^*$	$1/\hat{\sigma}$	$\hat{\eta}^{*E}$
	$+0.93$	$+1.95$	2.88	4.52	3.54	1.29	-179.28
	(0.14)	(0.66)	(0.10)	(0.96)	(0.17)	(0.25)	

Table 1 shows the median over all days of the point estimates per investor group G, in addition to the implied elasticity of demand ($\hat{\eta}^E$) and of supply ($\hat{\eta}^{*E}$) on the platform, averaged across dealers. The corresponding medians of the standard errors are in parentheses. All estimates are in bps.

7.1 Estimates of the benchmark model

We report the estimates for a median day in Table 1—for example, $\hat{\mu}^I = median_t(\hat{\mu}_t^I)$. In all box plots, the upper and lower 1st percentile of the distribution are excluded.

Investor values. The amount the investor is willing to pay splits into her value for the bond (the liquidity shock) and her value for the dealer's quality (shown in Figure 7b). The split between these two components relies on the normalization that the benchmark dealer provides zero quality. Therefore, we focus on the differences across investor groups and dealers rather than the absolute values.

A buying retail investors are 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. This suggests that the yield gap of 4 bps between retail and institutional investors is not driven entirely by platform access, but that differences in the willingness to pay account for some of it. Below, we quantify how much.

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. Since different dealers attach different values to realizing trade on any given day (see Figure 8), there might be welfare gains in matching investors to dealers with higher values on that day. We quantify these gains below.

Platform competition. Competition is relatively low ($\hat{\sigma} = 0.78$) for at least two reasons. 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 (see Figure 7a). Second, investors are likely to buy from their home dealer also on the platforms. This is explored in the model extension.

Low platform competition implies that demand on the platform is rather inelastic.[26] The average yield elasticity of demand on the platform is about 174–179. This means that a dealer's demand (supply) increases on average by 1.74% (1.79%) if this dealer increases (decreases) her quote by 1 bps. 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 1% more.

The elasticity of demand of an individual investor is similar, although not directly comparable, to the aggregate elasticity of demand in the U.S. government bond market: 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 2.6% increase in demand when the yield increases by 1.5 bps.

Platform usage cost. In line with concerns that have been raised by industry experts, we find that a high cost of about 3.5 bps prevents investors from using the platform. In the next section, we analyze what drives this cost.

7.2 Platform cost decomposition

We conjecture that three factors manly drive the cost. The first is the actual fee an investor has to pay to enter the platform, the second is a relationship component, and the last is a cost of sharing information about trade sizes and prices. To estimate how much of each of these two factors contribute to the total cost, one would

[26] To see this, re-write $\eta^E_{tj}(q_t)$ of Proposition 2 assuming that the platform shocks are iid EV1 distributed: $\eta^E_{tj}(q_t) = \frac{1}{\sigma}(1 - s_{tj}(q_t))q_{tj} + rest(q_t)$, where $s_{tj}(q_t) = \exp(\frac{1}{\sigma}(\xi_{tj} + q_{tj}))/\sum_{k\in\mathcal{J}_t}\exp(\frac{1}{\sigma}(\xi_{tk} + q_{tk})$ is dealer j's daily platform market of investors who enter the platform, and $rest(q_t) = f^I(\psi_t(q_t))/\mathcal{F}^I(\psi_t(q_t))s_{tj}(q_t)q_{tj}$ comes from the fact that only a mass of $\mathcal{F}^I(\psi_t(q_t))$ of investors (per dealer) choose to enter the platform. This term would be zero if all investors traded on the platform, and is small given our estimates.

Figure 7: Dealer qualities

(a) Platform market shares with best quote

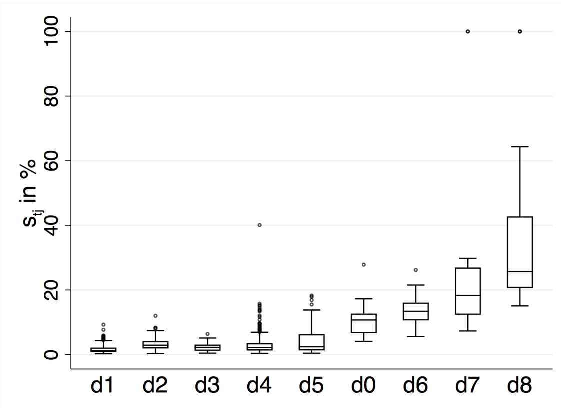

(b) Estimated dealer qualities

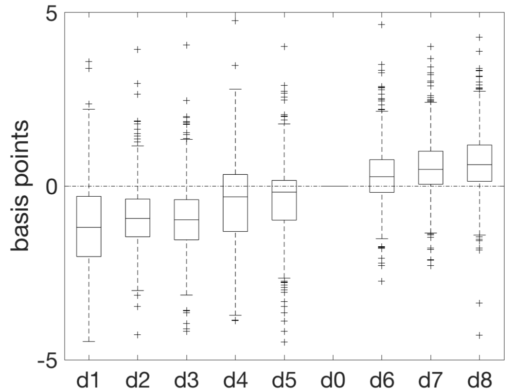

Figure 7a displays a box plot for each dealer, labeled d0 (benchmark) to d8. Each shows the distribution of how much of the total platform market this dealer captures (in %) on days on which she posts the best quote relative to other dealers. Figure 7b shows a box plot of the estimated qualities, $\hat{\xi}_{tj}$, in bps for each dealer.

Figure 8: Dealer values

(a) Identifying variation across dealers

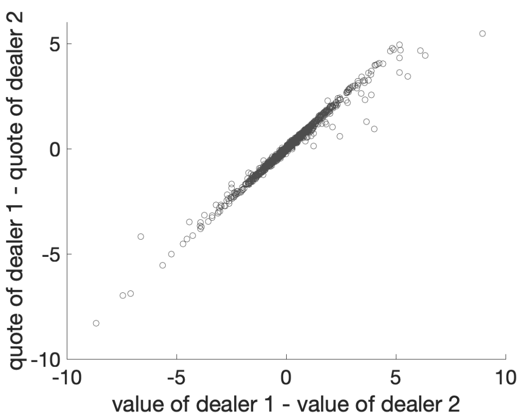

(b) Estimated dealer values for 10 days

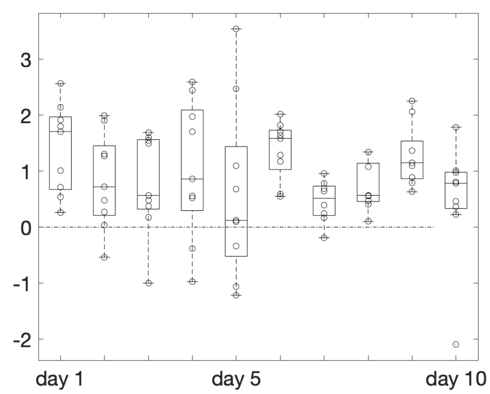

Figure 8a illustrates what drives two dealers to have different values on the same day. It shows the correlation between the difference of the observed quotes of two dealers $(q_{t1} - q_{t2})$ and the difference of the estimated values of those dealers $(\hat{v}^D_{t1} - \hat{v}^D_{t2})$. Figure 8b displays box plots of the estimated dealer values, $\hat{v}^D_{tj}$, net of the bond's market value, θ_t, in bps for the last ten days in our sample. The graphs look similar when picking different dealer pairs and days, respectively.

ideally observe how investors and dealers behave on the platform. We don't have this information but still would like to provide a rough back of the envelop calculation.

Platform fee. We know that on average the actual fee to trade on the platform ranges roughly between C$ 2,500-3,500 per month for a typical institutional investor.[27] This translates into a per-unit cost of about 1.1-1.5 bps per trade for the typical institutional investor who trades 22 million (units) per month on the platform.

Relationship component. To get a sense of how important the relationship factor is, we focus on the subsample of investors who are reported with a LEI and re-estimate the benchmark model. As comparison, we estimate the extended model in which the investor obtains a loyalty benefit whenever she buys from her home dealer. The loyalty benefit is identified from how many home investors buy from their home dealer relative to how many buy from another dealer on a day, so for a fixed set of quotes (see Appendix D.2.2 for details).

We find that the loyalty benefit is significant and accounts for a large part of the total platform usage cost (see Appendix Table A7). The platform usage cost drops to about 2 bps when accounting for the relationship component, i.e., when going from the benchmark model to the extended model. This suggests that about 1.5 bps of the platform usage cost of the benchmark model are benefits that an investor gives up when trading with other dealers than her home dealer on the platform. We conjecture that this is true also for investors without LEIs because the estimates of the benchmark model using the full and restricted sample are similar.

Cost of information sharing. Estimating the cost of information sharing is a challenging task. An ideal setting to gather evidence that information sharing might be costly, would be one in which investors share information with more dealers but—unlike in RFQ auctions—competition is kept constant so that it cannot not affect the terms of trade. This would allow us to separate the cost of information sharing from

[27]The cost depends on how many traders use CanDeal and whether the institution pays for the basic or full package. Consider a small institutional investor who typically has 3-4 people working on the trading desk. For the basic package, this institution pays C$ 725 per month for three traders, plus a fee for admin-rights (C$ 150) and a compliance fee (C$ 200). This amounts to C$ 2,525–C$ 3,250 and is a little less than the full package costs C$ 3,315–C$ 3,455 per month (C$ 2,895 per institution + C$ 140 per trader).

the benefit of higher competition in an auction.

An institutional feature renders behavior in the Canadian primary auctions close to this ideal setting: investors can participate in these auctions but must place their bids with one or several dealers. This is called order splitting. The chosen dealer(s) observe(s) how much the investor would like to buy of a security and at what prices, similar to what dealers observe in an RFQ auction. However, unlike in the RFQ auctions, the dealers only pass on this information to the auctioneer and do not compete to trade with the investor.[28]

In Appendix B, we show that conditional on the size of the bid, investors who share information with four dealers (as in a typical RFQ auction) relative to one dealer (as in a typical bilateral trade) obtain a yield that is about 0.5 bps worse when trading post-auction.[29] This suggests that it is costly to share information because it negatively affects future terms of trade.[30]

Summary. Taken together, this evidence suggests that roughly 31%–43% of the estimated platform usage costs represent the actual platform fee, 43% a relationship cost and 14% a cost of sharing information.

7.3 Model fit

Before assessing the price and welfare effects from centralizing the market, we validate whether our parsimonious (benchmark) model can replicate the event study in Section 4. Recall that 90 institutional investors lost platform access in our sample. Crucially, we did not use any information on how yields change when this happens to estimate the model. Instead, we use this information to test whether our model predicts a similar impact on yields.

We find that the model's prediction is very similar to the reduced-form estimate

[28]The fact that the dealer observes the investors' bids might affect the degree of competition in the auction. However, compared to other factors, such as the total number of bidders in the auction, this effect is likely second-order.

[29]A natural question that arises is why investors split orders with several dealers if this has negative consequences. One explanation is that dealers have limits on how much they can demand. Consistent with this, we find that an investor splits her order with more dealers when she demands large amounts at auction (see Appendix Table A10).

[30]Our static model incorporates this effect in reduced from, as a cost. We encourage future research on how information sharing shapes trade dynamics and prices.

Table 2: Model fit

	Event study		Model prediction	
Change in yield	-1.15	(0.340)	[-0.95, -0.98] for I	[-1.15, -1.24] for R

Table 2 compares by how much yields drop for investors who lose platform access with the standard error in parentheses (first column) with what the model predicts (second column). The former is the estimate of the event study regression but collapses the time before and after the event: $(y_{thsij} - \theta_{ths})^+ = \zeta_i + \beta access_{thi} + \zeta_{th} + \zeta_s + \zeta_j + \epsilon_{thsij}$, where $(y_{thsij} - \theta_{ths})^+$ is defined as in (1) and $access_{thi}$ assumes value 1 if the investor has platform access and 0 otherwise. To compute the drop in the expected yield (before observing the liquidity shock) according to our model, we rely on Proposition 1. We keep the quotes constant at the observed levels, as in the event study.

(see Table 2). On average, an institutional investor who loses platform access obtains a yield that is 1.15 bps worse in the data. Our model predicts that the yield drops by 0.95–1.24 bps. This similarity reassures us that the model makes adequate predictions about what happens when we make platform access universal in our counterfactual exercise.

These differences in yields translate into sizable monetary losses for investors. For an average institutional investor, who trades about 17,391 million units of bonds per year, a yield reduction of 1 bps implies an annual loss of C$ 1.7 million. For an average retail investor, who trades about 86 million units, the annual loss is about C$ 8,622. These numbers imply that the value for having platform access for all investors combined amounts to about C$ 640 million per year.

8 Counterfactual exercises

We now conduct counterfactuals using the benchmark model—unless stated other—to assess how much of the gap between retail and institutional investors' yields is due to platform access and to quantify welfare gains when centralizing the market.

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.[31] 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.[32] See Appendix D.3 for formal details.

8.1 What drives the yield gap?

When a retail investor obtains (costly) platform access, she expects a yield increase of about 1 bps. This implies that the gap between retail and institutional investors decreases on average by roughly 32% when the investor is buying. When the investor is selling, the percentage change is larger, 47%, because the yield gap in the status quo is smaller.

The yield gap does not close completely, because many retail investors stay off the platform: Only 52%—60% of the retail investors would trade on the platform. The remaining would trade bilaterally. These investors obtain worse yields than institutional investors in bilateral trades because they are, on average, willing to pay more.

Platform participation is weak because it is costly to use the platform and because the platform is not perfectly competitive. To separate these two factors we eliminate the platform usage cost. As a result more retail investors (83%) would use the platform, and the yield gap would close by 52% for buying and 82% for selling investors. Some investors would still stay off the platform because of their low willingness to pay. For them, the platform quotes are not attractive.

8.2 Welfare analysis

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. An overview of all counterfactuals is in Table 3.

[31] The findings are similar if we allow dealers to discriminate between investors on the platform, and post a quote that is investor-group specific rather than a single quote for both groups.

[32] 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.

Types of frictions. First, frictions stem from poor platform design: platform access is limited, it is costly to use the platform 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 matches cannot form—either, because search frictions in the bilateral market are high, or because keeping close relationships is valuable. 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 capacities 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 reduce 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. 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.

When does welfare increase? It is not obvious that there are sizable welfare gains when improving platform access or design, because any welfare gain in our setting comes entirely from re-matching who trades with whom, as we keep the number of trades and the trade volume fixed. In the extreme, if all market participants have the same value for the bond, such re-matching would not affect welfare. To see this, let us define welfare formally.

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

$$W_t^G = \sum_j \mathbb{E}[v_{tj}^D - v_{tj}^G(\boldsymbol{\nu_{ti}^G}) | \textit{investor } i \in g \textit{ buys from dealer } j] \tag{9}$$

is the expected welfare from trading with investors of group $G \in \{I, R\}$ *with dealer value,* v_{tj}^D*, and investor value,* $v_{tj}^G(\nu_{ti}^G) = \theta_t + \nu_{ti}^G - \xi_{tj}$*. Proposition 1 specifies which dealer the investor buys from in the benchmark model, and Proposition 4 in the extended model with a loyalty benefit.*[33]

[33]Note that, in the extended model, the loyalty benefit influences the investor's decision with whom to trade but does not directly improve the gain associated with a particular trade, because this benefit captures dealer services that are outside of government bond

Given Definition 1, welfare increases when investors become more likely to buy from more efficient dealers, i.e., dealers with higher $(v_{tj}^D + \xi_{tj})$.[34] To see this, compute the welfare change when going from the status quo to the counterfactual world:

$$\Delta W_t = \sum_G \kappa^G \sum_j \Delta\gamma_{ij}^G * (v_{tj}^D + \xi_{tj}), \tag{10}$$

where $\Delta\gamma_{tj}^G$ is the change in the probability that $i \in G$ buys from dealer j on day t.

Removing platform entry barriers. In the first set of counterfactuals, we reduce entry barriers to the platform. We keep dealer market power fixed, in that we allow dealers to charge a markup when setting quotes and fix the degree of competition on the platform, σ at the estimated value.

In the benchmark model, we find that welfare increases by about 10% on an average day when all investors can trade on the platform (Figure 9). The reason is that more investors match with more efficient dealers when platform participation is higher. The cheaper it becomes for investors to trade on the platform, the higher platform participation and the larger is the welfare gain. For example, when we remove the platform fee and make the platform anonymous, eliminating the cost of information sharing, welfare increases by 22%.

The welfare gain translates into a sizable monetary gain of C$ 127–426 million per year, or roughly 0.8–2.7 bps of GDP.[35] The magnitude is smaller, yet within the same ballpark of the aggregate value of investors to have platform access (C$ 640 million per year). This is reassuring because the aggregate value for platform access represents the rent distribution from dealers to investors, which should be larger than the welfare gain since it does not account for dealer losses. Further, the aggregate value for platform access was obtained via a back-of-the-envelop calculation using the event study estimates, independently from the structural model.

market, such as account administration or tax support.

[34]In theory, the welfare effect is ambiguous. On the platform, an investor chooses the dealer with the highest $(q_{tj} + \sigma\epsilon_{tij}) + \xi_{tj}$, but this dealer must not be more efficient than the dealer chosen in the status quo because the platform is not perfectly competitive $(v_{tj}^D \neq q_{tj})$.

[35]In absolute terms, the welfare increases by 1.8-6 bps per day. Since trade-sizes are normalized to one, this holds under the assumption that the total trade volume is 9 units (1 unit per dealer). To compute the monetary gain, we re-weight the number by the actual trade volume between dealers and investors, which is on average 25 billion (units) per day and 6,4 trillion (units) per year.

Figure 9: Welfare gain

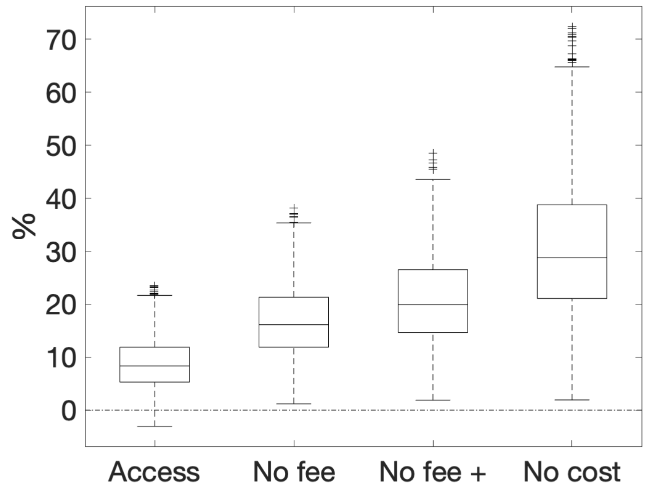

Figure 9 illustrates by how much welfare increases when allowing all investors to trade on the platform at the full estimated cost (Access), when removing the platform fee (No fee), when removing the information cost in addition (No fee +), and when eliminating all costs. In all cases, it shows the distribution of the percentage change in welfare, $100 * \frac{\Delta W_t}{W_t}$, over days.

Table 3: Average welfare gain in all counterfactuals

Benchmark model	**Access**	**No fee**	**No fee +**	**No cost**
Imperfect competition	9.95% (C$127)	18.00% (C$230)	22.17% (C$283)	33.39% (C$426)
Perfect competition	38.36% (C$489)	50.00% (C$638)	55.37% (C$706)	67.65% (C$863)
Extended model	**Access**	**No fee**	**No fee +**	**No cost**
Imperfect competition	4.60% (C$92)	8.62% (C$173)	10.47% (C$210)	11.94% (C$239)
Perfect competition	17.78% (C$356)	21.70% (C$435)	23.38% (C$469)	24.68% (C$495)

Table 3 displays the average welfare gain, $mean_t(\Delta W_t)$ for each counterfactual—universal platform access at estimated cost (Access), at no platform fee (No fee), at no platform fee and no information cost (No fee+), at no cost (No cost) when platform competition is imperfect and perfect—and model specification. The welfare gain is shown in % of welfare achieved in the status quo and as annual monetary gain in million C$. To obtain the latter, we re-weight the welfare gain that our model in which all trades are normalized to one by the actual trade volume that dealers and investors exchange in an average year.

Removing dealer market power. To quantify the distortion that arises from dealer market power, we shut down dealer market power on the platform and repeat the counterfactuals that remove platform entry barriers. We force dealers to post quotes that are equal to their true values, which eliminates the markup and, in addition, make the platform perfectly competitive (or frictionless) so that investors trade at the posted quotes by sending $\sigma \to \infty$. The first-best is achieved when all investors trade on this frictionless platform—a useful theoretic benchmark that is likely hard to achieve in reality.

Our findings highlight that distortions due to market power are substantial (see Table 3). For example, welfare increases by 38.36%-9.95%≈ 28% (or C$362 million per year) more when all investors can trade on a perfectly competitive platform than when they can trade on the imperfect platform, both at the estimated platform usage cost.

How important are relationships? To analyze what effect relationships may have on welfare, we repeat the welfare analysis using our extended model with a loyalty benefit (on the restricted sample). Relative to the benchmark model, investors are more likely to trade with their home dealer, independent of whether they trade bilaterally or on the platform. Therefore, fewer new trade matches form as more investors enter the platform. As a result, all welfare gains are smaller than in the benchmark model but still sizable with C$ 92–C$ 495 million per year (see Table 3).[36]

The difference between the welfare gain in the benchmark model and the extended model gives us a sense of how important relationships are. Take for instance the counterfactual in which we remove all platform fees and make the platform perfectly competitive. This counterfactual achieves the first-best in the benchmark but not in the extended model. A comparison of the two models, thus, tells us how strongly the gains from trade are distorted only because investors prefer trading with specific dealers in an otherwise frictionless market (C$ 368 million per year).

Future research could study why relationships are sticky and how they affect total welfare generated from all trading and investment activities between dealers

[36] If we assume that the loyalty benefit enters welfare, welfare can decrease when an investor trades with a dealer who offers a better price on the platform than the home dealer. Since this becomes more likely as more and more investors enter the platform, the welfare gain becomes smaller and smaller as platform entry becomes less and less costly.

and investors. Answering this question is beyond the scope of this paper as it would require observing how dealers and investors interact in all markets, not just the cash market for government bonds.

What drives the welfare gain? Interestingly, in all specifications, almost the entire welfare gain, ΔW_t in (10), comes from matches to dealers with higher values, v^D_{tj}, rather than to dealers with better quality, ξ_{tj}. This implies that a dealer's quality or services (such as fast and reliable delivery of the trade) is less important relative to the dealer's value for realizing trade (which are are shaped by inventory positions or expectations about future prices and demand are crucial).

This highlights the fact that in the status quo, 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); He et al. (2021); Schrimpf et al. (2020)). Our findings suggest that market centralization would reduce these frictions.

Will the market fix itself? Given the evidence in favor of market centralization, it is natural to ask whether the market might fix itself. For instance, the dealers who own the platform in Canada could realizes that they can more easily offload their inventory positions when more investors trade on the platform and eliminate 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.[37] For example, on average a dealer would lose roughly C$ 26 million in trade profits per year if platform access was universal and free. In addition, the dealers who own the platform would no longer earn profits from collecting the platform fee.

This highlights, that dealers—who are the key market players in most government

[37]This raises the question why dealers founded the platform in the first place. They followed the global trend away from OTC markets twenty years ago to preserve market power.

bond markets—have strong incentives to prevent changes in the market structure, even if they are welfare-improving.

Summary. Taken together, our findings suggest that even in a government bond market—which is commonly viewed as one of the most well-functioning financial markets—there is large potential to increase welfare by centralizing the market.

9 Robustness

We conduct several tests to verify the robustness of our findings in Appendix E. First, we test the robustness of our parameter estimates. For example, we check whether the estimates are biased in the expected direction when we do not instrument the quotes or use the amount a dealer won as an instrument. In addition, we restrict the sample to exclude occasionally large trades to verify that our estimates are not driven by outliers. Second, we verify that our estimates and welfare findings are robust to the assumption that dealers don't trade with one another. For this, we let dealer valuations converge towards a hypothetical dealer value that would arise with a perfect inter-dealer market.[38] Taken together, our robustness tests confirm our expectations and suggest that our main findings are qualitatively robust.

10 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. Further, we estimate large 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 feasible first step in the right direction, but there are many other ways how to further centralize the

[38]In previous versions of the draft, we verified that that measurement errors in the dealers' qualities (ξ_{tj}) do not significantly bias the distribution of the liquidity shocks, and allowed for dealer-specific platform usage costs (c_{tj}). In addition, we checked that the results are similar when we allow the investor to capture a fraction ϕ of trade surplus in the bilateral trade.

market. 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.

References

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

Allen, J., Kastl, J., and Wittwer, M. (2020). Primary dealers and the demand for government debt. Working paper.

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

Babus, A. and Parlatore, C. (2019). Strategic fragmented markets. Working paper.

Baldauf, M. and Mollner, J. (2020). Principal trading procurement: Competition and information leakage. Working paper.

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

Benos, E., Payne, R., and Vasios, M. (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.

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

Biais, B. and Green, R. (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. Report submitted by a Study Group established by the Markets Committee of the Bank of International Settlements.

Boyarchenko, N., Lucca, D., and Veldkamp, L. (2021). Taking orders and taking notes:

Dealer information sharing in treasury auctions. *Journal of Political Economy*, forthcoming.

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

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 Duffie, D. (2021). Market fragmentation. *American Economic Review*, forthcoming.

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

De Roure, C., Moench, E., Pelisson, L., and Schneider, M. (2020). OTC discount. Working paper.

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

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

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

Dugast, J., Üslü, S., and Weill, P.-O. (2019). A theory of participation in OTC and centralized markets. Working paper.

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

Financial Times (2015). Wall street banks to settle CDS lawsuit for $1.9bn. Ben McLannahan and Joe Rennison.

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

Fleming, M. J., Mizrach, B., and Nguyen, G. (2017). The microstructure of a U.S. Treasury ECN: The BrokerTec platform. Federal Reserve Bank of New York Staff Reports.

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

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

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

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

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

He, Z., Nagel, S., and Song, Z. (2021). Treasury inconvenience yields during the COVID-19 crisis. *Journal of Financial Economcis*, page forthcoming. Working paper.

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

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

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

Hendershott, T., Livdan, D., and Schürhoff, N. (2021b). Do we need dealers in OTC markets? Working paper.

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

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

Hortaçsu, A. and McAdams, D. (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 Yogo, M. (2019). A demand system approach to asset pricing. *Journal of Political Economy*, 127(4):1475–1515.

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

Kondor, P. and Pintér, G. (2021). Clients' connections: Measuring the role of private information in decentralized markets. *Journal of Finance*, forthcoming.

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

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

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

Kutai, A., Nathan, D., and Wittwer, M. (2021). Central exchanges for government bonds? Evidence during COVID-19. Working paper.

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

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

Logan, L. K. (2020). Speech: Treasury market liquidity and early lessons from the pandemic shock. Accessed on 05/05/2021 at: `https://www.newyorkfed.org/newsevents/speeches/2020/log201023`.

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

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., Nguyen, M., and Valentex, G. (2017). Funding constraints and market illiquidity in the European Treasury bond market. Working paper.

O'Hara, M. and Zhou, X. (2021). The electronic evolution of corporate bond dealers. *Journal of Financial Economics*, forthcoming.

Plante, S. (2017). Should courporate bond trading be centralized? Working paper.

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

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

Schrimpf, A., Shin, H. S., and Sushko, V. (2020). Leverage and margin spirals in fixed income markets during the COVID-19 crisis. BIS Bulletin.

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

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. Working paper.

Online Appendix

A Data cleaning

For a few trades of the 3,755,901 observations in the raw data, we change the execution time, the date, or the settlement date. First, 296 trades were reported on a weekend. We count them as Monday 7:00 am trades if reported on a Sunday and Friday 5:00 pm trades if reported on a Saturday. In all other cases, we keep the time of day and only change the date. 162 observations are reported to settle after maturity. We replace their settlement date with the maturity date. 5,355 trades settled before they were executed. We replace the reported settlement date with the date on which the trade would settle according to trade conventions.

We correct 100 cases in which subsidiaries of reporting dealers or brokers are labeled retail investors, and we 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. This excludes trades with investors who are in a gray area. 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 $(y_{thsij} - \theta_{ths})^+$, 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.

Finally, 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 Empirical evidence: Cost of sharing information

The goal is to whether sharing information in primary auctions leads to worse trade yields in the future. For this, we compare yields of investors who trade similar amounts of the same security within the same hour and with the same dealer in the secondary market, but who have shared information with a different number of dealers (from 0 = no dealer to 7 = seven dealers, the upper bound according to the auction rules) in the most recent primary auction (which on average took place 3 days before the trade). We do this separately for investors who sell and investors who buy because the effect from sharing information about buying a bond might affect sell and buy-side trades differently.

As a starting point, we regress the trade yields investors realize after a primary auction ($yield_{thsji}$) on the number of dealers with whom an investor shared information in this auction ($N\_dealers_{ai}$), the size of her bid in that auction (bid_size_{ai}), hour-security (ζ_{ths}), dealer (ζ_j) and similar trade size (ζ_q) fixed effects. To generate ζ_q, we split trade sizes into buckets with equal numbers of observations ranging form from micro-trades of less than C\$ 0.05 million, up to large trades of more than C\$ 50 million.

$$yield_{thsji} = \alpha + \beta N\_dealers_{ai} + \gamma bid\_size_{ai} + \zeta_{ths} + \zeta_j + \zeta_q + \epsilon_{tsji}$$

Controlling for the size of the bid is similar to including a fixed effect to absorb unobservable differences between investors who bid in the primary auction ($bid\_size >$ 0) and those who did not ($bid\_size = 0$) with the difference that the effect is weighted by the bid size for investors who bid in the auction. The idea is that the size of the bid might be correlated with the degree of an investor's eagerness to sell/buy, which might affect the yield independently of whether the investor has shared information or not.

We run two specifications. First, we include all trades. Second, we restrict the sample to bilateral trades of investors with LEIs who trade with a single dealer in our sample. These investors are less likely to negotiate trades with multiple dealers or to share trade information in a different way than by splitting orders in primary auctions. Therefore, the effect we estimate can more directly be linked to sharing information in the primary auction.

In both specifications, our parameter of interest is identified from variation in the number of dealers with whom an investor shared information in the primary auction. The parameter tells us by how much the yield of an investor who shared information with more dealers in the primary auction differed from the yield of investors who shared information with less dealers.

We find that investors who have shared information with more dealers in the primary auction sell at a yield that is about 0.5 bps higher, i.e., lower price (see Appendix Table A8). The effect is weaker for investors who buy. This could be because it is harder to use information about how much and at what prices an investor wanted to buy bonds in the past when the investor wants to buy more, rather than when she now seeks to sell.

Next, we analyze how costly it is to share information with a single dealer (as in a bilateral trade) relative to sharing information with four dealers (as in a typical RFQ auction). For this, we estimate the following regression:

$$yield_{thsji} = \alpha + \sum_{k=0}^{7} \beta_k * \mathbb{I}(N\_dealers_{ai} = k) + \gamma * bid\_size_{ai} + \zeta_{ths} + \zeta_j + \zeta_q + \epsilon_{tsji},$$

and compare $\hat{\beta}_4$ with $\hat{\beta}_1$. The difference approximates the cost of information sharing in a typical RFQ auction relative to trading bilaterally with a single dealer. Since investors who trade with a single dealer never share information with more than two dealers, we can compute this difference only for our first specification that uses all trades. We include all estimates in Appendix Table A9 for completeness.

We find that selling investors who have shared information with four dealers obtain a yield that is 1.502–0.960 bps $\approx$ 0.5 bps worse relative to investors who have shared information with a single dealer. For buying investors, the estimates are too noisy to find any effect. On both side of the trade, the relationship between the number of dealers and the change in the yield is not increasing monotonically. This is likely due to the fact that there are not that many occasions in which an investor shares information with two, five or six dealers.

Taken together, the statistically significant estimates of both regressions point towards a cost of roughly 0.5 bps.

C Mathematical appendix

C.1 Proof of Propositions 1 and 2

The equilibrium can be derived by backward induction. For notational convenience, we drop the subscript t and the superscript I throughout the proof.

Proposition 1. Statement (i) holds by assumption. To derive statement (ii), begin in the last stage. Conditional on entering the platform and observing ϵ_{ij}, investor i buys from dealer j if $u_{ij} > u_{ki} \forall k \neq j$, where $\tilde{u}_{ij}(\epsilon_{ij}) = \xi_t + q_j + \sigma\epsilon_{ij}$. Ex ante, dealer j's market share on the platform (of investors on the platform) is

$$s_j(q) = \frac{\exp(\delta_j)}{\sum_k \exp(\delta_k)} \text{ with } \delta_j = \frac{1}{\sigma}(\xi_j + q_j) \text{ given } \epsilon_{ij} \sim EV1 \; . \tag{11}$$

By assumption, a home dealer d offers $y_{id} = \theta + \nu_i - \xi_d$ in a bilateral trade and is always willing to trade. The investor obtains no surplus when buying bilaterally and expects to earn $-(\theta + \nu_i) + \mathbb{E}[\max_{k\in\mathcal{J}} \tilde{u}_{ki}(\boldsymbol{\epsilon_{ki}})] - c$ when entering the platform. She decides to buy bilaterally if $\psi(q) \leqslant \nu_i$ with $\psi(q) = \mathbb{E}[\max_{k\in\mathcal{J}} \tilde{u}_{ki}(\boldsymbol{\epsilon_{ki}})] - \theta - c$. □

Proposition 2. Consider home dealer d. In choosing the quote, the dealer anticipates how investors will react, but does not know which liquidity shocks investors will draw. She supposes that she can sell on the platform at the posted quote, q_d, and chooses this quote to $\max_{q_d} \pi_d(q) = \max_{q_d}\{\pi_d^D(q) + \pi_d^E(q)\}$, where $\pi_d^D(q) = \int_{\psi(q)}^{\infty}(v^D - (\theta + \nu - \xi_d))f(\nu)d\nu$ given that $y_d = \theta + \nu - \xi_d$ is the expected profit from bilateral trades, and $\pi_d^E(q) = S_d(q)(v^D - q_d)$, where $S_d(q) = \sum_j F(\psi(q))s_d(q)$ and $s_d(q)$, as in (11), is the expected profit from platform trades.[39] Taking the partial derivative w.r.t. q_d, and rearranging gives the markup equation. □

[39]If the dealer didn't think that she sold at the posted quote, the markup equation would be: $E(q_d)(1 + 1/\eta_d^E(q)(1 - \frac{\partial \pi_d^D(q)}{\partial q_d}/S_d(q))) = v_d^D$. The yield the dealer expects to receive conditional on winning the RFQ auction, $E(q_d) = q_d + \sigma \int ... \int [\mathbb{I}(\tilde{u}_{id}(\epsilon_{id}) > \tilde{u}_{ij}(\epsilon_{ij}) \; \forall j \neq d)\epsilon_{id}] dH(\epsilon_{i1}) .... dH(\epsilon_{iJ})$, would replace the posted quote q_d. This markup equation would render the estimation of the model computationally intense.

C.2 Micro-foundation

Here, we solve for an equilibrium in an RFQ auction to justify why it is reasonable to assume that investor i obtains a yield of $q_{tj} + \sigma\epsilon_{tij}$ when trading with dealer j on the platform on day t. To make the auction tractable, we let dealers be ex ante identical—i.e., abstract from dealer qualities. 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.

Auction game. $K \geqslant 2$ 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 3. *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:*

$$y_{ij}^E = q + \sigma\epsilon_{ij} \text{ where } \epsilon_{ij} = (x_{ij}/\sigma + s) \text{ and} \tag{12}$$

$$\sigma \text{ solves } 0 = \int_{-\infty}^{\infty} [-\Phi(-z)^{J-1} + (z-\sigma)(J-1)\Phi(-z)^{J-2}\phi(-z)]\,\phi(z)\,dz. \tag{13}$$

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

Proof. We derive conditions that are satisfied in a symmetric equilibrium via backward induction. For notational convenience, we drop subscripts whenever possible.

Guess that there is an equilibrium in which a dealer with signal x submits a linear function $\beta(x) = q + x + \sigma s$, where σ is a parameter, and s determines the noisiness of the dealer's signal.

Note that conditional on x, $v_2|x \sim N([\mu_v/h + x/h']/[h+h'], 1/[h+h'])$, where $h = 1/\sigma_v^2, h' = 1/s^2$. As $\sigma_v^2 \to \infty$, $h \to 0$, and at limit $v|x \sim N(x, s^2)$. Hence $[v_2 - x]/s = \omega$ where $\omega \sim N(0,1)$.

The dealer's problem is, given x,

$$\arg\max_b \int_{-\infty}^{\infty} [q + v_2 - b] F(\beta^{-1}(b)|v_2)^{J-1} dF(v_2|x),$$

where $F(\beta^{-1}(b)|v_2)^{J-1}$ is the distribution of maximum of others' signals given v_2, evaluated at $\beta^{-1}(b)$, and $F(v_2|x)$ is the distribution of v_2 given signal x. The first-order condition is

$$0 = \int_{-\infty}^{\infty} [-F(z|v_2)^{J-1} + (q + v_2 - \beta(x))(J-1)F(z|v_2)^{J-2} f(z|v_2)\frac{\beta^{-1}(b)}{db}]\,dF(v_2|x)\,,$$

evaluated at $z = \beta^{-1}(\beta(x)) = x$, so $\frac{d\beta^{-1}(\beta(x))}{db} = 1/\beta'(x)$.

To evaluate $F(x|v_2)$, observe that it is the probability that another bidder's signal $x' < x$ given v_2. Since $x' = v_2 + \omega' s$ this is the probability that $v_2 + \omega' s < x$ or $\omega' < [x - v_2]/s$, which is $F(x|v_2) = \Phi([x - v_2]/s)$. From above, $[x - v_2]/s = -\omega$, so $F(x|v_2) = \Phi(-\omega)$ and $f(x|v_2) = \phi(-\omega)/s$. Lastly, since $[v_2 - x]/s = \omega$, $F(v_2|x) = \Phi(\omega)$ and $f(v_2|x) = \phi(\sigma)/s$. Taken together, the first-order condition becomes

$$0 = \int_{-\infty}^{\infty} [-\Phi(-\omega)^{J-1}\beta'(x) + (q + x + \omega s - \beta(x))(J-1)\Phi(-\omega)^{J-2}\phi(-\omega)/s]\,\phi(\omega)/s\,d\omega\,.$$

Given $\beta(x) = q + x + \sigma s$, the FOC is satisfied when σ solves (13). □

C.3 Model extension: Loyalty benefit

In our model extension, we let investors earn a loyalty benefit from trading with their home dealer, both on and off the platform, to captures the value of keeping a close relationship with one dealer independent of the way in which the bond is traded. Formally, the investor's value is

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

where Ξ_{tij} depends on whether the dealer is the investor's home dealer or not

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

The dealer gives the loyalty benefit to the investor independent of how the trade realizes thanks to some binding agreement. To achieve sufficient statistical power, in the estimation we assume that all home investors of a dealer enjoy the same loyalty benefit:

$$r_{tij} = \begin{cases} 0 & \text{(benchmark model)} \\ r & \text{for all investors } i \text{ and dealers } j \text{ (extended model).} \end{cases}$$

It is straightforward to derive the analogous propositions of Propositions 1 and 2 for this model extension. To save space, we omit repeating the proofs as they follow the same structure.

Proposition 4 (Investors)**.**
(i) A retail investor with shock ν_{ti}^R buys bilaterally from home dealer d at

$$y_{tid}^R = \theta_t + \nu_{ti}^R - \xi_{td}. \tag{14}$$

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

$$y_{tid}^I = \theta_t + \nu_{ti}^I - \xi_{td} \quad \text{if } \psi_{td}(q_t) \leqslant \nu_{ti}^I, \tag{15}$$

$$\text{where } \psi_{td}(q_t) = \mathbb{E}[\max_{k \in \mathcal{J}_t} \tilde{u}_{tik}(\boldsymbol{\epsilon_{tik}})] - \theta_t - c_t - \mathbb{I}(j = d)r \tag{16}$$

$$\text{with } \tilde{u}_{tij}(\boldsymbol{\epsilon_{tij}}) = \xi_{tj} + \mathbb{I}(j = d)r + q_{tj} + \sigma\boldsymbol{\epsilon_{tij}}. \tag{17}$$

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

An investor now obtains the loyalty benefit r whenever she trades with her home dealer. Thanks to a binding agreement, the dealer gives this benefit to the investor independent of how the trade realizes. Therefore, the investor obtains the same yield on the platform as she did in the benchmark model. Without such an agreement, the dealer would offer a different platform yield and extract at least some of the loyalty benefit from the investor.

Proposition 5 (Dealers). *Dealer j posts a quote q_{tj} that satisfies*

$$q_{tj}\Big(1+\frac{1}{\eta^E_{tj}(q_t)}\Big(1-\frac{\partial\pi^D_{tj}(q_t)}{\partial q_{tj}}\Big/S_{tj}(q_t)\Big)\Big)=v^D_{tj}, \tag{18}$$

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 ϵ_{tij} are EV1 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 $S^l_{tj}=s^l_{tj}(q_t)\Pr(\boldsymbol{\nu^I_{ti}}\leqslant\psi_{kj}(q_t))$ 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 $\pi^D_{tj}(q_t)=\mathbb{E}[v^D_{tj}-(\boldsymbol{\nu^I_{ti}}+\theta_t-\xi_{tj})|\psi_{tj}(q_t)\leqslant\boldsymbol{\nu^I_{ti}}]$.

The dealer's markup equation is analogous to equation (7). The main difference is that the platform market share, $S_{tj}(q_t)$, splits into J_t elements given that there are J_t different investor types, each belonging to a different home dealer.

D Details regarding the estimation

D.1 Construction of the supply shock instruments

In modeling the primary auction and estimating the bidders' values, we follow Hortaçsu and Kastl (2012) and Allen et al. (2020). For a detailed discussion of all assumptions and derivation of the equilibrium, we refer to these papers.

Auction model. In the auction, there are two groups of bidders: N_d dealers and N_c customers, who are investors who bid at auction. All of them draw a private signal.

Assumption 5. *Dealers' and customers' private signals s^d_j and s^c_j are for all bidders j independently drawn from common atomless distribution functions F^d and F^c with support $[0,1]^M$ and strictly positive densities f^d and f^c.*

The bidder's group and signal affect how much she values the bond.

Assumption 6. *A bidder j of group $g \in \{d, c\}$ with signal s_j^g values amount q by $v^g(q, s_j^g)$. This value function is nonnegative, measurable, bounded strictly increasing in s_j^d for all q and weakly decreasing in q for all s_j^g.*

Given their values, bidders place bids. Each bid is a step function that characterizes the price the bidder would like to pay for each amount.

Assumption 7. *Each bidder has the following action set:*

$$A = \begin{cases} (b, q, K) : dim\,(b) = dim(q) = K \in \{1, ..., \overline{K}\} \\ b_k \in [0, \infty) \text{ and } q_k \in [0, 1] \\ b_k > b_{k+1} \text{ and } q_k > q_{k+1} \forall k < K. \end{cases}$$

Dealers can submit their bids directly to the auctioneer. Customers have to place their bids with one of the dealers. This might give the dealer additional information. To capture this, we define the information that is available to dealer j before placing her bid by Z_j. We call $\theta_j^d = (s_j^d, Z_j)$ the dealer's type. The type of a customer is her private signal s_j^c.

Definition 2. *A pure strategy is a mapping from the bidder's set of types to the action space: $\Theta_j^g \to A$. It is a bidding function, labeled $b_j^g(\cdot, \theta_j^g)$ for bidder j of group g with type θ_j^g.*

Once all bidders submit their step function, the market clears at the lowest price at which the aggregated submitted demand satisfies the total supply.

The supply is unknown to each bidder when she places her bid because a fraction of it goes to noncompetitive tenders. These are bids that specify only an amount that is won with certainty.

Assumption 8. *Supply $\boldsymbol{Q}$ is a random variable distributed on $[\underline{Q}, \overline{Q}]$ with strictly positive marginal density conditional on s_j^g $\forall i, g = c, d$.*

Bidder j wins amount q_j^c at market clearing, and pays the amount she offered to win for each unit won.

Definition 3. *A Bayesian Nash equilibrium in pure strategies is a collection of functions* $b_j^g(\cdot, \theta_j^g)$ *that for each bidder* j *and almost every type* θ_j^g *maximizes the expected total surplus,* $\mathbb{E}\left[\int_0^{\boldsymbol{q_j^c}}[v(x, s_j^g) - b_j^g(x, \theta_j^g)]dx\right]$.

We focus on type-symmetric BNE in which all dealers and all customers play the same strategy. One can show that in any type-symmetric BNE, every step k in the bid function $b_j^g(\cdot, \theta_j^g)$ has to satisfy

$$v^g(q, s_j^g) = b_k + \frac{\Pr(b_{k+1} \geqslant \boldsymbol{P^c}|\theta_j^g)}{\Pr(b_k > \boldsymbol{P^c} > b_{k+1}|\theta_j^g)}(b_k - b_{k+1}) \tag{19}$$

for all but the last step and $b_k = v^g(\bar{q}(\theta_j^g), s_j^g)$ at the last step, where $\bar{q}(\theta_j^g)$ is the maximal amount the bidder may be allocated in equilibrium. Here P^c denotes the market-clearing price.

Estimation. We estimate how much each dealer expects to win in this equilibrium, at the time at which she places the bids, in three steps.

First, we estimate the distribution of the residual supply curve a dealer faces. This curve is the total supply minus the total demand of all other bidders. For this, 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 submitted 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 already 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 the distribution of how much the dealer won in the auction, q_j^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., that $b_k \geqslant \boldsymbol{P^c} > b_{k+1}$.

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

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

where K_j are the steps in dealer j's bidding function. To obtain our instrument $won_{j\tilde{t}}$

of auction $\tilde{t}$, we subtract $\mathbb{E}$[amount dealer j wins | bids] from the amount that bidder j actually won, which we observe.

D.2 Estimation procedure

D.2.1 Benchmark model

We explain the estimation for buying institutional investors in detail. For buying retail investors, we match the expectation and variance of the bilateral yields via GMM (similar to step 2 below). For selling investors, the estimation is analogous.

Step 1. Denote dealer j's market share on the platform by $s_{tj}(q_t, \xi_t, \sigma)$. When ϵ_{tij} are extreme value type 1 distributed:

$$s_{tj}(q_t, \xi_t, \sigma) = \frac{\exp(\delta_{tj})}{\sum_{k \in \mathcal{J}_t} \exp(\delta_{tk})} \text{ with } \delta_{tk} = \frac{1}{\sigma}(\xi_{tk} + q_{tk}) \text{ for all } k \in \mathcal{J}_t. \tag{20}$$

Abbreviate $s_{tj}(q_t, \xi_t, \sigma)$ by s_{tj} 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 Assumption 2 plus Normalization 1 to obtain

$$\log(s_{tj}/s_{t0}) = \zeta_j + \zeta_t + \frac{1}{\sigma}\tilde{q}_{tj} + rest_{tj}, \tag{21}$$

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 3, which implies $\mathbb{E}[\boldsymbol{rest_{tj}}|won_{\tilde{t}j}, \zeta_t, \zeta_j] = 0$, we can estimate σ in a linear IV regression in which we instrument $\tilde{q}_{tj}$ by $won_{\tilde{t}j}$ and include dealer ζ_j and date ζ_t fixed effects.

With this, we compute $\hat{\xi}_{tj}$ for all $j \neq 0$ and the cutoff that determines whether an investor buys bilaterally from home dealer d or on the platform:

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

Step 2. For each day t, we estimate the remaining parameters via GMM by matching the expectation and variance of the bilateral yield of a buying institutional investor, and the probability that such an investor buys bilaterally. To compute the predicted moments, we rely on $\boldsymbol{\nu^I_{ti}} \sim N(\mu^I_t, \sigma^I_t)$ and Proposition 1.

D.2.2 Extended model

Additional assumptions. First, we need to assign a home dealer to each investor. To do this, we restrict our sample. We take investors that we can observe with their LEI, and define the dealer with whom an investor trades the most as her home dealer. In addition, to conduct the same counterfactuals as with the benchmark model, we take all retail investors who are rarely reported with a LEI. We assume that retail investors trade with a single dealer and that they earn the same loyalty benefit as institutional investors. Second, to obtain sufficient power, we pool trades within a week and define the period t as one week rather than a single day. This implies that we can no longer account for unobservable movements of prices and demand/supply within the same week.

Changes to the estimation procedure. Dealer j's total market share on the platform now splits into one part that comes from the dealer's home investors and another part that comes from the home investors of other dealers. Conditional on entering the platform, the probability that a home investor buys from dealer j is

$$s_{tj}^{j}(q_t, \xi_t, \sigma, r) = \frac{\exp(\frac{1}{\sigma}(q_{tj} + \xi_{tj} + r)}{\sum_j \exp(\frac{1}{\sigma}(q_{tj} + \xi_{tj} + \mathbb{I}(j = d)r)}.$$

Conditional on entering the platform, the probability that a home investor of dealer k buys from dealer j is

$$s_{tj}^{k}(q_t, \xi_t, \sigma, r) = \frac{\exp(\frac{1}{\sigma}(q_{tj} + \xi_{tj})}{\sum_j \exp(\frac{1}{\sigma}(q_{tj} + \xi_{tj} + \mathbb{I}(j = k)r))}.$$

Abbreviating the market shares as above, normalizing all shares by the shares of the benchmark dealer 0, and taking logs, we obtain the following:

$$\log(s_{tj}^{k}/s_{t0}^{k}) = \zeta_j + \mathbb{I}(k = j)\frac{r}{\sigma} + \zeta_t + \frac{1}{\sigma}\tilde{q}_{tj} + rest_{tj}, \tag{21'}$$

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

From here, we can estimate σ and loyalty benefit r. The rest of the estimation is analogous to the benchmark model.

D.3 Formal details of the counterfactual

Let $\bar{\boldsymbol{\nu}}_{\boldsymbol{ti}}$ be drawn 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 $\kappa^R = 0.1$ and $\kappa^I = 0.9$.

Benchmark model. An investor in $G \in \{I, R\}$ with shock ν_{ti}^G buys bilaterally from home dealer d at $y_{tid}^G = \theta_t + \nu_{ti}^G - \xi_{td}$ if $\psi_t(q_t) \leqslant \nu_{ti}^G$, where $\psi_t(q_t)$ is defined in (6) of Proposition 2. Otherwise, the investor enters the platform, where she observes ϵ_{tij} and buys from the dealer with the maximal $\tilde{u}_{tij}(\epsilon_{tij})$ at $q_{tj} + \sigma\epsilon_{tij}$, with $\tilde{u}_{tij}(\epsilon_{tij})$ as in (6). Dealer j posts a quote q_{tj} that satisfies equation (7) of Proposition 2, but with $S_{tj}(q_t) = \sum_{j \in J_t} \Pr(\bar{\boldsymbol{\nu}}_{\boldsymbol{ti}} \leqslant \psi_t(q_t)) \Pr(\tilde{u}_{tki}(\boldsymbol{\epsilon_{tki}}) < \tilde{u}_{tij}(\boldsymbol{\epsilon_{tij}})\ \forall k \neq j)$, where $\pi_{tj}^D(q_t) = \mathbb{E}[v_{tj}^D - (\bar{\boldsymbol{\nu}}_{\boldsymbol{ti}} + \theta_t - \xi_{tj}) | \psi_t(q_t) \leqslant \bar{\boldsymbol{\nu}}_{\boldsymbol{ti}}]$.

Extended model with loyalty benefit. An investor in $G \in \{I, R\}$ with shock ν_{ti}^G buys bilaterally from home dealer d at $y_{tid}^G = \theta_t + \nu_{ti}^G - \xi_{td}$ if $\psi_t(q_t) \leqslant \nu_{ti}^G$, where $\psi_t(q_t)$ is defined in (16) of Proposition 5. Else, the investor enters the platform, where she observes ϵ_{tij} and buys from the dealer with the maximal $\tilde{u}_{tij}(\epsilon_{tij})$ at $q_{tj} + \sigma\epsilon_{tij}$, with $\tilde{u}_{tij}(\epsilon_{tij})$ as in (17). Dealer j posts a quote q_{tj} that satisfies equation (7) of Proposition 5 but with $S_{tj}^l = s_{tj}^l(q_t) \Pr(\bar{\boldsymbol{\nu}}_{\boldsymbol{ti}} \leqslant \psi_{kj}(q_t))$, $\pi_{tj}^D(q_t) = \mathbb{E}[v_{tj}^D - (\bar{\boldsymbol{\nu}}_{\boldsymbol{ti}} + \theta_t - \xi_{tj}) | \psi_t(q_t) \leqslant \bar{\boldsymbol{\nu}}_{\boldsymbol{ti}}]$.

E Robustness analysis

Instrument. We conduct several tests regarding our instrument (see Appendix Table A11). First, we check whether $\hat{\sigma}$, which governs the yield elasticity of demand on the platform, is biased in the expected direction when we do not instrument the quotes by $won_{\tilde{t}j}$ and replace Assumption 3 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.

Next, we use a different instrument for the quote; namely, the amount a dealer won on the most recent auction day rather than the amount she won unexpectedly. The advantage of this instrument is that it is model-free because we can read it off the data. The big disadvantage is that it does not address the concern that dealers might anticipate investor demand and bid accordingly in the auction.

We obtain similar estimates with both instruments when allowing for dealers to systematically differ in quality. However, the instrument is weak. We therefore check by how much the estimates change when dropping the dealer fixed effect by imposing $\xi_j = 0 \, \forall j$. In both specifications, this increases the correlation between our instrument and the platform quotes. The instrument becomes stronger, but we no longer control for unobservable differences between dealers that might drive differences in the quotes. $\hat{\sigma}$ decreases from 0.77 to 0.49 when using $won_{\tilde{t}j}$ and from 0.68 to 0.62 when using the amount the dealer actually won as instrument.

Trade size. We verify that our estimates are not driven by occasionally large trade sizes (see Appendix Table A12). We do so because our model abstracts from trade size since most trades are small and similar in size. However, investors occasionally trade large amounts, and if they do, it is more likely that they trade bilaterally than on the platform (see Appendix Table A6). To test that our estimates are robust to these rare occasions, we re-estimate the model on a subsample of trades, excluding the 5% largest trades of an investor who trades more than one time. Investors who trade a single time do not trade large amounts.

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 θ_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 A3). 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 dealer's 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 θ_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 A13 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. However, the decrease is small. For instance, when we allow all investors to trade on the platform at the estimated cost, the benchmark model predicts a welfare gain of C$ 127 million per year when $x = 0\%$ and C$ 120 million when $x = 16\%$.

Appendix Figure A1: Investor types and where they trade

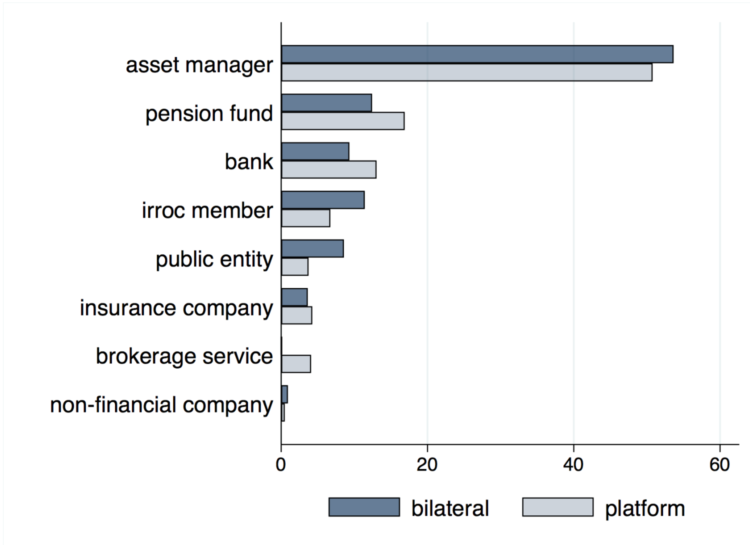

Appendix Figure A1 shows how much each investor type trades on the platform versus bilaterally as a percentage of the total amount investors trade.

Appendix Figure A2: CanDeal quotes versus average platform trade yields

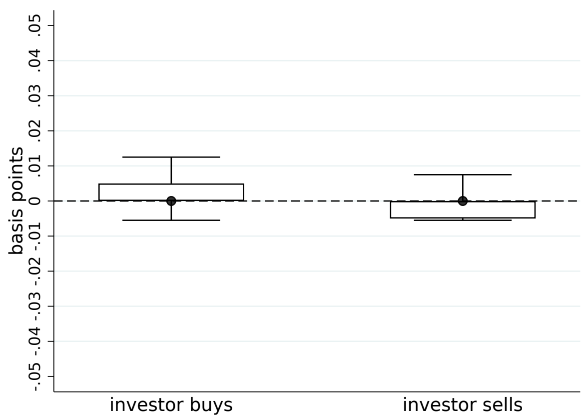

The left-handed box plot of Appendix Figure A2 depicts the distribution of the difference between the hourly average yield at which a security was bought on CanDeal and the security's ask quote posted on CanDeal (averaged per minute). The right-handed box plot shows the analogous for when the investor sells with the bid quote replacing the ask quote. We see that the yield at which an investor buys is extremely close to the ask quote and the yield at which she sells is extremely close to the bid quote. The median quote and execution yield are identical on both sides of the trade.

Appendix Figure A3: Distribution of trade sizes

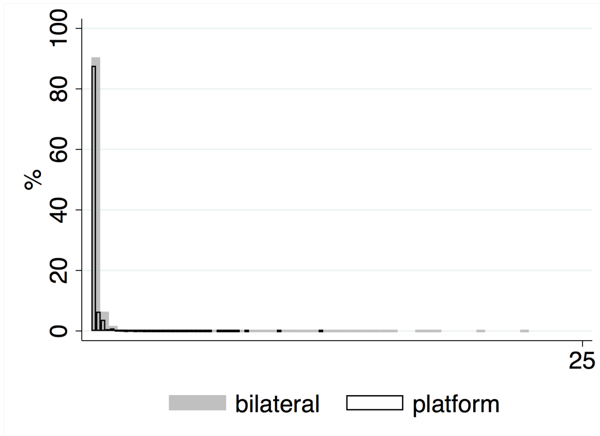

Appendix Figure A3 shows a probability density histogram of trade sizes in bilateral and platform trades. Trade size is measured in million C$.

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

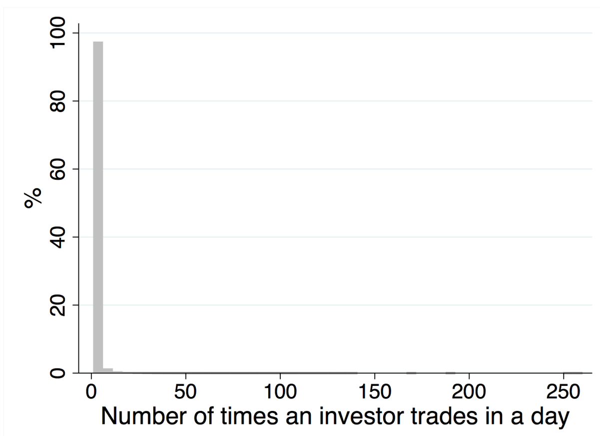

Appendix Figure A4 shows a probability density histogram of the number of times an investor trades—i.e., either sells or buys—in a day. The graph is similar when counting how many trading venues (bilateral vs. platform) an investor uses in a day. We see that in more than 95% of the cases, the investor only trades one time, and that on the rare occasions when the investor trades multiple times, she typically trades bilaterally and on the platform. Our model abstracts from these rare events.

Appendix Figure A5: Event study—observable trade behavior

(a) Trade size

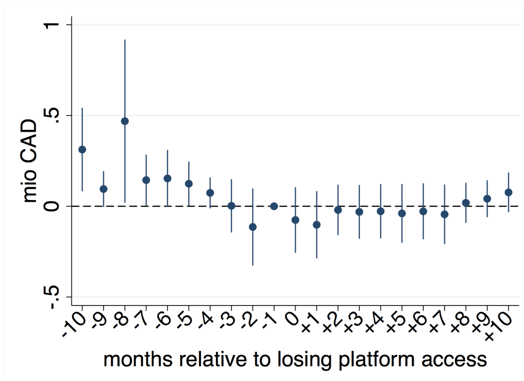

(b) Weeks-to-maturity of traded bonds

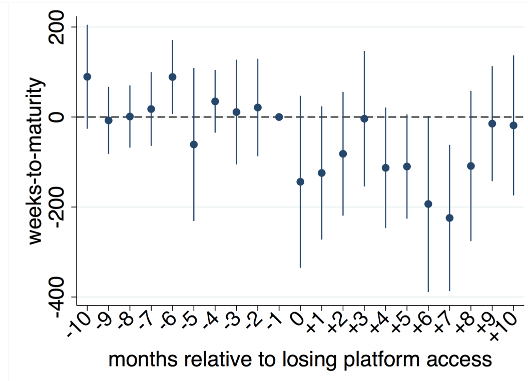

(c) Duration of traded bonds

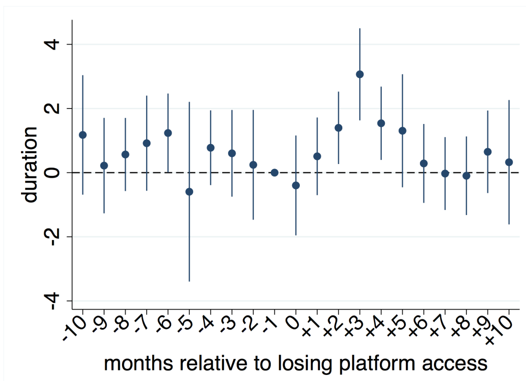

(d) Convexity of traded bonds

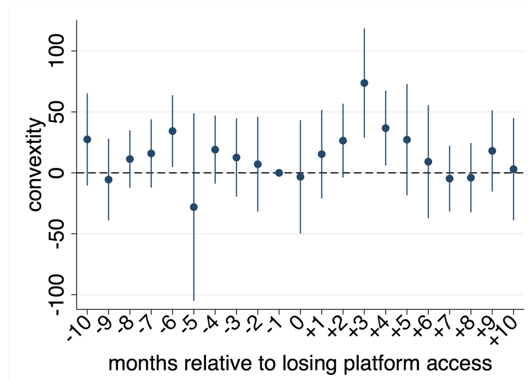

These figures visualize changes in observable trade behavior when the investor loses platform access. They show the β_m estimates and the 95% confidence intervals of regressions that are similar to the event study regression (2), but with outcome variables that capture trade behavior. For Figure A5a we regress $quantity_{thsij} = \zeta_i + \sum_{m=Mi^-}^{Mi^+} \beta_m D_{mi} + \zeta_{th} + \zeta_s + \zeta_j + \epsilon_{thsij}$ 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 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 primary dealer due to data reporting.

Appendix Table A2: Yields are better on the platform

	(1)		(2)	
platform	0.282	(0.0331)	0.0795	(0.0310)
constant	-0.296	(0.00934)	-0.281	(0.00690)
investor fixed effect (ζ_i)	–		✓	
Observations	1,193,999		806,473	
Adjusted R^2	0.169		0.523	

Appendix Table A2 shows that yields on the platform are better than those off the platform. For this, we regress the markup $(y_{thsij} - \theta_{ths})^+$, as defined in (1), on an indicator variable that assumes value 1 if the trade realizes on the platform ($platform_{thsij}$), hour-day (ζ_{th}), security (ζ_s), and dealer (ζ_j) fixed effects. In column (2) we add investor (ζ_i) fixed effects. Standard errors are in parentheses and clustered at the investor level in column (2).

Appendix Table A3: Balancing positions in the inter-dealer market

	(1)	(2)	(3)	(4)
$net\_supply^C_{tsj}$	-0.161	-0.161	-0.162	-0.162
	(0.00233)	(0.0457)	(0.0458)	(0.0454)
dealer fixed effect (ζ_j)	–	✓	✓	✓
date fixed effect (ζ_t)	–	–	✓	✓
security fixed effect (ζ_s)	–	–	–	✓
Constant	0.259	0.259	0.260	0.261
	(0.255)	(0.105)	(0.105)	(0.112)
Observations	177165	177165	177164	177157
Adjusted R^2	0.026	0.027	0.022	0.021

Appendix Table A3 provides evidence that dealers do not sizably balance their position by trading with other dealers or brokers. We regress a dealer's net supply for a security (in million C$), $net\_supply^C_{tsj}$, in the inter-dealer market segment on the dealer's excess demand in the dealer-investor market, $net\_supply^D_{tsj}$. In both cases, the excess demand is defined as the total amount the dealer sold of a security minus the total amount this dealer bought of the security in the respective market segment. Column (1) shows an OLS regression: $net\_supply^D_{tsj} = net\_supply^C_{tis} + \epsilon_{tis}$. 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 A4: Yields and the supply shocks

(a) Bilateral yields

	(1) buys		(1) sells	
θ_{ths}	0.668	(0.00249)	0.749	(0.00195)
$won_{\tilde{t}j}$	-0.0000199	(0.0000588)	-0.0000567	(0.0000458)
constant	47.73	(0.369)	37.82	(0.288)
Observations	192,342		192,004	
Adjusted R^2	0.724		0.681	

(b) Quotes

	(1) buys		(1) sells	
θ_{ths}	0.0224	(0.000990)	0.0351	(0.000963)
$won_{\tilde{t}j}$	0.000489	(0.0000234)	0.000291	(0.0000226)
constant	144.4	(0.147)	142.7	(0.142)
Observations	192,342		192,004	
Adjusted R^2	0.272		0.237	

Here we provide evidence that dealers do not adjust bilateral yields when hit by an unexpected supply shock, defined in (8). For Appendix Table A4 (a), we regress the bilateral yield of a trade between investor i and dealer j in hour h of date t with security s on the bond's market value, and the dealers' supply shocks, as well as dealer, investor, and date fixed effects: $y_{thsij} = \alpha + \beta\theta_{ths} + \gamma won_{\tilde{t}j} + \zeta_j + \zeta_i + \zeta_t + \epsilon_{thsij}$ for buying investors in column (1) and selling investors in column (2). Appendix Table A4 (b) shows the analogous results when replacing the bilateral yield in both regressions with the platform quote q_{tj}. All yields are in basis points and $won_{\tilde{t}j}$ is in million C\$. Standard errors are in parentheses.

We find that both $won_{\tilde{t}j}$ coefficients in Appendix Table A4 (a) are close to 0 and insignificant. The estimates imply that when $won_{\tilde{t}j}$ increases by one standard deviation, the change in the yield lies in (-0.0145 bps, $+0.0102$ bps) for buying investors and in (-0.0159 bps, $+0.0036$ bps) for selling investors. When clustering standard errors, either at the dealer or investor level or both, these intervals become even tighter. In contrast, the dealer adjusts her quotes according to Appendix Table A4 (b). Both findings are in line with our model assumptions.

Appendix Table A5: Effect of trade size on yields

buy	1.313	(0.135)
θ	0.988	(0.000690)
quantity	0.0246	(0.0388)
quantity^2	-0.0177	(0.0122)
quantity^3	0.00120	(0.000737)
constant	0.881	(0.102)
Observations	806,564	
Adjusted R^2	0.998	

Appendix Table A5 shows the estimation results when regressing the trade yield ($yield_{thsij}$) on an indicator variable that shows the size of the trade (buy_{thsij}), the market value (θ_{ths}), and a function of trade size in C$ 100,000, $\sum_{p=1}^{3} \delta_p (quantity_{thsij})^p$, in addition to hour-day (ζ_{th}), security (ζ_s), dealer (ζ_j), and investor (ζ_i) fixed effects. The findings suggest that trade size is not driving the yield, since all of the coefficient multiplying quantity are statistically insignificant. Standard errors are in parentheses and clustered at the investor level.

Appendix Table A6: Trade size and venue choice

	(1)	(2)	(3)
quantity	-0.101	-0.0373	-0.0213
	(0.00104)	(0.00911)	(0.0163)
constant	0.353	0.330	0.331
	(0.000452)	(0.00157)	(0.00207)
investor fixed effect (ζ_i)	–	✓	✓
Observations	1,234,945	784,809	738,344
Adjusted R^2	0.008	0.462	0.468

Appendix Table A6 shows whether institutional investors trade different amounts on versus off the platform. Column (1) gives the estimation results when regressing an indicator for whether trade realizes on or off the platform on the trade size (quantity). In column (2), we add an investor fixed effect (ζ_i). In column (3), we exclude the 5% largest trades of an investor to show that the statistically significant negative correlation between platform and quantity is driven by occasional large trades. Our interpretation is that in these rare cases, investors prefer to trade bilaterally with their dealer. This could be because their dealer works as an insurance for rainy days and offers a better deal than other dealers with whom the investor does not maintain a close business relationship. In this paper, we focus on regular small trades. Standard errors are in parentheses, and clustered at the investor level in columns (2) and (3).

Appendix Table A7: Estimates of the extended model

(a) With estimated loyalty benefit

	$\hat{\mu}^I$	$\hat{\sigma}^I$	$\hat{c}$	$1/\hat{\sigma}$	$\hat{r}$
Buys	-1.01	$+2.69$	$+2.03$	$+1.02$	$+2.08$
	(0.15)	(0.07)	(0.35)	(0.32)	(0.05)
Sells	$\hat{\mu}^{I^*}$	$\hat{\sigma}^{I^*}$	$\hat{c}_t^*$	$1/\hat{\sigma}$	$\hat{r}$
	$+1.08$	$+2.71$	$+2.14$	$+1.02$	$+2.08$
	(0.16)	(0.08)	(0.35)	(0.32)	(0.05)

(b) With zero loyalty benefit

	$\hat{\mu}^I$	$\hat{\sigma}^I$	$\hat{c}$	$1/\hat{\sigma}$	r (set)
Buys	-1.01	$+2.73$	$+3.53$	$+1.02$	0
	(0.10)	(0.08)	(0.18)	(0.32)	
Sells	$\hat{\mu}^{I^*}$	$\hat{\sigma}^{I^*}$	$\hat{c}_t^*$	$1/\hat{\sigma}$	r (set)
	$+1.08$	$+2.75$	$+3.66$	$+1.02$	0
	(0.10)	(0.09)	(0.19)	(0.32)	

Appendix Tables A7 (a) and (b) are similar to Table 1 but for the extended model with loyalty benefit, estimated using trades with investors that have a LEI and with a period being a week. In (a) we use the estimated loyalty benefit, in (b) we set the loyalty benefit to 0. The relationship cost is defined as the difference between the cost estimate in (b) and the cost estimate in (a). Each table shows the median over all weeks of the point estimates (in bps) for institutional investors. The corresponding medians of the standard errors are in parentheses. All estimates but the cost estimate are similar in size to those of Table 1 which is estimated on the full sample of investors.[a]

[a]Keep in mind, however, that only the sum of the liquidity shock and dealer base qualities can be compared across models due to different normalizations. For illustration, consider a buy-side investor as example. She is willing to accept a yield that is $\hat{\mu}^I - median_{j,t}(\hat{\xi}_{tj}) = -1.00 - 0.38 \approx -0.62$ different from the market value when buying from a dealer of median base quality. In the benchmark model, she is willing to accept a yield that is $\hat{\mu}^I - median_{j,t}(\hat{\xi}_{tj}) = -0.82 + 0.28 \approx -0.54$ different. These numbers are identical in both models for the buy-side and similar on the sell-side.

Appendix Table A8: Order splitting in primary auctions affects trade yields

	sell (1)	sell (2)	buy (3)	buy (4)
N_dealers	0.459	0.626	0.0836	-0.527
	(0.167)	(0.222)	(0.221)	(0.245)
bid_size	-0.00221	-0.00706	0.00152	0.00401
	(0.000856)	(0.00191)	(0.00260)	(0.00273)
Constant	146.8	128.8	149.4	127.4
	(0.00910)	(0.000396)	(0.0143)	(0.00236)
Observations	632914	3011	607060	1804
Adjusted R^2	0.922	1.000	0.185	1.000

Appendix Table A8 provides evidence that sharing information with dealers in Treasury auctions harms future trade yields. In columns (1)-(2) we regress the trade yields selling investors realize after a primary auction ($yield_{thsji}$) on the number of dealers with whom an investor shared information in this auction ($N\_dealers_{ai}$), the size of her bid in that auction (bid_size_{ai}), hour-security (ζ_{ths}), dealer (ζ_j) and similar trade size (ζ_q) fixed effects. Columns (3)-(4) show the estimates for investors who buy. In (1) and (3), we include all trades. In (2) and (4), we restrict the sample to bilateral trades of investors with LEIs who trade with a single dealer in our sample. Standard errors are in parentheses and clustered at the security level.

Appendix Table A9: Order splitting in primary auctions affects trade yields

	sell (1)	sell (2)	buy (3)	buy (4)
N_dealers=1	0.962	1.008	1.329	-1.296
	(0.362)	(0.0976)	(1.042)	(0.470)
N_dealers=2	0.902	4.164	1.272	-0.782
	(0.406)	(0.792)	(1.193)	(0.492)
N_dealers=3	1.379		-1.232	
	(0.470)		(1.692)	
N_dealers=4	1.506		2.008	
	(0.659)		(1.713)	
N_dealers=5	0.946		2.677	
	(0.437)		(2.721)	
N_dealers=6	0.565		0.190	
	(0.662)		(0.854)	
bid_size	-0.00243	-0.0199	-0.0000763	0.00653
	(0.000903)	(0.00390)	(0.00123)	(0.00261)
Constant	146.8	128.8	149.4	127.4
	(0.0131)	(0.000911)	(0.0279)	(0.00256)
Observations	632914	3011	607060	1804
Adjusted R^2	0.922	1.000	0.184	1.000

Appendix Table A9 provides evidence that sharing information with dealers in primary auctions harms future trade yields. In columns (1)-(2) we regress the trade yields selling investors realize after a primary auction ($yield_{thsji}$) on indicator variables marking the number of dealers with whom an investor shared information in this auction ($N\_dealers_{ai} = k$ for $k = 1$ until $k = 7$ in theory, but until $k < 7$ in the data), the size of her bid in that auction (bid_size_{ai}), hour-security (ζ_{ths}), dealer (ζ_j) and similar trade size (ζ_q) fixed effects. Columns (3)-(4) show the estimates for investors who buy. In (1) and (3), we include all trades. In (2) and (4), we restrict the sample to bilateral trades of investors with LEIs who trade with a single dealer in our sample. Standard errors are in parentheses and clustered at the security level.

Appendix Table A10: Order splitting and bid sizes in Treasury auctions

	(1)		(2)		(3)	
bid_size	2.155	(0.0782)	2.586	(0.149)	2.950	(1.227)
Constant	1.076	(0.0208)	1.005	(0.0245)	0.944	(0.202)
Observations	4193		4152		4180	
Adjusted R^2	0.153		0.173		0.510	

Appendix Table A10 shows that investors who demand larger amounts in Treasury auctions place their bids with more dealers. In column (1), we regress the bid size of an investor i in an auction a (in billion C\$) on the number of dealers this investor places her bids with: $bid\_size_{ai} = \alpha + \beta N\_dealers_{ai} + \epsilon_{ai}$. In (2), we add an auction fixed effect and use variation within an auction across investors to identify β. In column (3), we use variation across auctions within the same investor, instead and use an investor fixed effect. Standard errors are clustered at the security level in (2) and the investor level in (3).

Appendix Table A11: Robustness of $1/\sigma$ w.r.t. the instrument

Specification	OLS	IV1	IV2	OLS	IV1	IV2
quote coefficient ($1/\sigma$)	0.014	1.467	1.287	-0.093	1.612	2.050
	(0.005)	(0.259)	(0.246)	(0.011)	(0.175)	(0.236)
dealer fixed effect (ζ_j)	✓	✓	✓	–	–	–
Observations	8,492	8,492	8,492	8,492	8,492	8,492
Adjusted R^2	0.804	0.805	0.805	0.040	0.040	0.040

Appendix Table A11 shows how the σ parameter changes depending on the instrument we use. Specifically, it gives the point estimate of regression (21) in Appendix D.2.1, which is the inverse of σ. The second and fifth columns show the OLS estimates, first including a dealer fixed effect and then excluding it. In the third and sixth columns, we instrument the relative quote with the amount the dealer won on the most recent auction day. The fourth and seventh columns show the estimate using the unexpected supply shocks as instruments (as reported in the text). Standard errors are in parentheses.

Appendix Table A12: Estimates (median across days)

Buys	$\hat{\mu}^I$	$\hat{\mu}^R$	$\hat{\sigma}^I$	$\hat{\sigma}^R$	$\hat{c}$	$1/\hat{\sigma}$	$\hat{\eta}^E$
	-0.79	-2.95	2.79	5.14	-3.34	1.36	$+184.51$
	(0.13)	(0.79)	(0.10)	(0.92)	(0.16)	(0.25)	
Sells	$\hat{\mu}^{I^*}$	$\hat{\mu}^{S^*}$	$\hat{\sigma}^{I^*}$	$\hat{\sigma}^{S^*}$	$\hat{c}_t^*$	$1/\hat{\sigma}$	$\hat{\eta}^{*E}$
	$+0.91$	$+1.99$	2.86	4.48	-3.46	1.36	-184.51
	(0.13)	(0.67)	(0.10)	(0.94)	(0.17)	(0.25)	

Appendix Table A12 shows the estimation results when restricting the sample to trades of regular trade sizes, excluding the 5% largest trades of investors who trade more than once. Analogous to Table 1, it shows the median over all days of all point estimates per investor group G, in addition to the implied elasticity of demand ($\hat{\eta}^E$) and of supply ($\hat{\eta}^{*E}$) on the platform, averaged across days and dealers. All estimates are in bps.

Appendix Table A13: Average welfare gain of all counterfactuals, shrunk dealer values

Benchmark model	**Access**	**No fee**	**No fee +**	**No cost**
Imperfect competition	10.30% ($C120)	18.44% (C$215)	22.66% (C$264)	33.91% (C$396)
Perfect competition	38.61% (C$450)	50.33% (C$587)	55.74% (C$650)	68.01% (C$793)
Extended model	**Access**	**No fee**	**No fee +**	**No cost**
Imperfect competition	4.75% (C$85)	8.85% (C$159)	10.74% (C$193)	12.22% (C$219)
Perfect competition	18.14% (C$326)	22.16% (C$398)	23.87% (C$429)	25.19% (C$452)

Appendix Table A13 is the analogous to Table 3 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, $mean_t(\Delta W_t)$ for each counterfactual—universal platform access at estimated cost (Access), at no platform fee (No fee), at no platform fee and no information cost (No fee+), at no cost (No cost) when platform competition is imperfect and perfect—and model specification. The welfare gain is shown in % of welfare achieved in the status quo and as annual monetary gain in million C$. To obtain the latter, we re-weight the welfare gain that our model in which all trades are normalized to one by the actual trade volume that dealers and investors exchange in an average year.