Centralizing over-the-counter markets?*

Jason Allen[a] Milena Wittwer[b]

January 6, 2021

CLICK HERE
for the latest version

Abstract

In traditional over-the-counter (OTC) markets investors trade bilaterally through intermediaries, called dealers. An important regulatory question is whether to centralize OTC markets by shifting trades onto centralized platforms. We address this question in the context of the Canadian government bond market, which is liquid and price-transparent. We document that, even in this market, dealers charge substantial markups when trading with investors. We also show that there is a price gap between large investors who have access to a centralized platform and small investors who do not. We specify a model to quantify how much of this price gap is due to platform access, and assess welfare effects. The model predicts that not all investors would use the platform, even if platform access were universal. Nevertheless, the price gap between small and large investors would close by 35-52%. Further, total welfare would increase by 9-30% because the platform better allocates high-valued buyers to low-valued sellers.

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

*The presented views are those of the authors, not necessarily of the Bank of Canada. All 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; in addition to Darrell Duffie, Liran Einav, Matthew Gentzkow, as well as Luis Armona, Markus Baldauf, Timothy Bresnahan, Nina Buchmann, Jose Ignacio Cuesta, Piotr Dworczak, Jakub Kastl, David K. Levine, Negar Matoorian, Paul Milgrom, Monika Piazzesi, Claudia Robles-Garcia, Paulo Somaini, Robert Wilson, Ali Yurukoglu, and all participants of the IO, finance and marco student workshops at Stanford. Correspondence to: [a]Jason Allen - Bank of Canada, Email: jallen@bankofcanada.ca, [b]Milena Wittwer - Stanford University, Email: wittwer@stanford.edu

1 Introduction

Each year, bonds and many other assets (such as mortgage related securities, currencies, commodities and derivatives) worth trillions of dollars are traded in over-the-counter (OTC) markets. Unlike centralized markets such as stock exchanges, OTC markets are considered to be decentralized because buyers have to 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).

Recently, in a series of antitrust lawsuits, dealers have been accused of exerting market power in several markets (including markets for credit default and interest rate swaps, and corporate, as well as government bonds).[1] This, together with dramatic events in bond markets during the recent COVID crisis, has motivated a policy discussion on whether and how to centralize OTC markets.[2] One solution is to shift trading onto centralized electronic platforms. Yet—even though new laws to promote this shift have already been enacted in some markets—it is unclear whether breaking up the market structure that has persisted for over a century would have sizable effects on prices and welfare.[3] On the one hand, a centralized platform can foster competition between dealers, allowing for more efficient trade matches. On the other hand, it may be costly for investors to use the platform, for instance, when trading on the platform leaks private information.

This paper assesses price and welfare effects from centralizing OTC markets in the Canadian government bond market. Government bonds are safe, simple contracts that are traded in highly liquid markets under low price uncertainty. Therefore, government bond markets are considered to be closer to efficient than most other financial markets.[4] With trade-level data, we show that, even in this market, dealers charge substantial markups. We also show that large (institutional) investors pay systematically lower prices than small (retail) investors. One difference between these two investors groups is that institutional

[1]For an overview of antitrust litigations see Schlam Stone and Dolan LLP (2018).

[2]Logan (2020), an executive vice president of the Federal Reserve Bank of New York, summarizes the problematic events in the market for U.S. government bonds during the COVID crisis and mentions possible changes to the market's structure, including increased trading on platforms.

[3]A prominent example is the Dodd-Frank Act. It mandates that some more standarized (OTC) derivatives must be traded on electronic platforms, called swap execution facilities (SEFs). Another example is the European Union's Markets in Financial Instruments Directive II (MiFID II).

[4]Academic research that supports this view dates back to Roll (1970), Fama (1975), Hamburger and Platt (1975), Pesando (1978). Further, recent policy reports that use the newly collected trade-level data of government bond markets conclude that these markets are among the most liquid in the world (e.g., U.S. Treasury report (2017)).

investors have access to a centralized platform, but retail investors 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 would result if platform access were universal.

There are three features that make the Canadian government bond market a particularly attractive setting to assess whether to centralize OTC markets. First, a reporting regulation allows us to observe trade-level data so that we can zoom in on how institutions behave, unlike existing studies on government bond markets. Second, there exists an electronic platform to which some but not all investors have access. Similar to platforms in many other OTC markets—including the largest ones in the United States—the platform fosters competition between dealers but is costly to use because trades are not anonymous. Third, the Canadian government bond market is more liquid and price-transparent than most other OTC markets. Therefore, our findings could be lower bounds of what can be achieved in other markets.

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. We can also distinguish between institutional investors who have access to the platform and retail investors who must trade bilaterally. On the platform, institutional investors see bid and ask quotes that indicate at what prices they can sell and buy, respectively. We collect these quotes, in addition to quotes that are posted on Bloomberg. The latter are visible to all investors and serve as the market value of each bond.

Our trade-level data allow us to document three novel facts: We show that, even in the Canadian government bond market, dealers charge markups over market value. These markups vary across investors and are systematically smaller for institutional than for retail investors. To test whether this is partially because of platform access, we conduct an event study: We show that an investor who loses platform access because she no longer fulfills the necessary legal requirements obtains worse prices. The price drop is large: It is 10 times higher than the bid-ask spread of platform quotes. This raises the possibility that investors could benefit from platform access.

To assess how much and analyze broader welfare effects, as well as equilibrium effects when centralizing the market, we introduce a model. In the model, dealers and investors have different values for realizing trade. Each dealer has its own investor base, and aims at maximizing trade profit. The game has two main periods. In the first period, dealers (imperfectly) compete with one another by simultaneously posting quotes at which they are willing to trade on the platform. In the second period, institutional investors can enter the platform and expect to trade at the posted quotes. To do so, they have to pay a usage cost, which, for instance, reflects concerns about revealing information to more than one dealer when trading on the platform. This is the only option for retail investors. With both types of investors, the dealer discovers the investor's value for realizing the trade and extracts all surplus in a bilateral trade.[5]

We characterize equilibrium conditions that highlight two features. The first is that dealers charge markups on the platform that depend on two elasticities: how easily investors on the platform switch to bilateral trading, and on how easily they switch to the dealer with the best quote. The second feature is that an investor with platform access selects onto the platform whenever she is willing to pay more than other investors. This is because bilateral prices only depend on the investor's value, while platform prices reflect the values of all investors.

We rely on these two equilibrium features to estimate the factors that drive demand and supply for government bonds. In doing so, we face a challenge that is common in demand estimation: Since dealers update quotes in response to investor demand, quotes are likely endogenous. Our solution is to construct a new cost shifter instrument that changes the dealer's costs to sell, but not investor demand. For this, we use primary auction bidding data. In these auctions, 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, she can more cheaply satisfy investor demand. Her cost to sell decreases unexpectedly either because of how others bid in the auction or because the government issued more than the dealer expected. How much more the dealer wins relative to what she expected to win represents an exogenous cost shifter. To construct it, we exploit estimation techniques from the empirical literature on (multi-unit) auctions.

[5]From a theoretic viewpoint, this is a relatively strong assumption. It implies that a dealer does not adjust bilateral prices when she suffers an unexpected shock to her inventory position, which changes her value for realizing trade. We test whether this implication holds in our data and find supporting evidence. Further, we verify that our main findings are robust to when we allow the investor to capture some (imposed) share of the bilateral trade surplus.

We validate the predictions of the model with what we find in the event study of institutional investors who lose platform access. Since the event study information is not used in the model estimation, the fact that our model can replicate this pattern gives us confidence for policy assessment.

With the model and its estimated parameters, we can quantify how much of the price gap between retail and institutional investors is due to platform access, and assess welfare effects from centralizing the market. We do this by means of two counterfactuals. In the first we allow retail investors to enter the platform under the same conditions as institutional investors. In the second, we remove all platform usage costs. In both, we take into account how dealers and investors respond to the policy.

We find that 35-52% of the price gap between retail and institutional investors is due to costly platform access. The gap does not close entirely because only about 60% of retail investors would use the platform. Even when platform access is free, investors who are willing to pay less than most others stay off the platform.

The implied welfare effects are theoretically ambiguous. While retail investors can only gain from obtaining platform access, institutional investors and dealers may gain or lose depending on how dealers adjust their quotes as the composition of investors who trade on the platform changes. Similarly, total welfare (measured by the total expected gains from trade) might increase or decrease depending on whether the platform is sufficiently competitive. Insufficient platform competition, for instance, can cause dealers to post quotes that distort the allocation away from the first best.

Empirically, we find that both institutional and retail investors gain, and that dealers lose in both counterfactuals. The biggest winners are institutional investors who may (each) gain up to C$1.7 million of extra interest rate earnings per year. Overall, welfare increases by 9-30%. The reason is that the centralized platform allows investors to directly access all dealers. This leads to more efficient trade matches, as investors and dealers match who both attach high values to realizing the trade.

This finding highlights a common advantage of (two-sided) centralized markets, which is that centralized markets promote better matches between counterparties. This is not specific to our setting and likely true in other contexts. In addition, our results have valuable policy implications for OTC markets. They emphasize that granting platform access alone does not shift all bilateral trades onto the platform due to platform usage costs.[6] One possible solution—put forward by industry experts—is to allow investors to trade anonymously on the platform. This could reduce privacy concerns and with it make it less costly to use the platform.

Finally, we quantify how efficient the market is today relative to the first-best to bound how much welfare could be potentially be gained from other types of market reforms. We find that the status quo only achieves 60% of the first-best, which suggests potentially large welfare gains. Crucially, in our analysis, welfare gains would come entirely form reallocating who trades with whom because market participants have heterogenous valuations for realizing trade. This is surprising as government bonds are liquid and safe assets whose market value is publicly known.

Interestingly, if investors directly traded with one another on an all-to-all platform without dealers, welfare would even be lower than in the status quo. This is because dealers no longer "make-markets" by absorbing excess demand or supply. While this finding should be taken with cautionas it might change if we allowed new market participants to enter the market, it emphathizes that market-making is important, even in liquid markets.

Contribution. Our main contribution is to empirically assess price and welfare effects when centralizing OTC markets. Relatively few studies touch upon centralization of financial markets, and typically via reduced-form analysis (e.g., Barklay et al. (2006), Fleming et al. (2017), Abudy and Wohl (2018), Biais and Green (2019), Benos et al. (2020), O'Hara and Zhou (2021)). Most closely related to our paper is Hendershott and Madhavan (2015) (HM). The authors build a model in which investors select between bilateral trading and trading on a platform, trading off lower search costs on the platform against the privacy benefits of trading bilaterally. Unlike HM, we highlight the trade-off that dealers face when choosing quotes and structurally estimate our model, allowing us to conduct counterfactual analyses. This complements HM's rich descriptive analysis of how dealers and investors behave on a platform for U.S. corporate bonds.

By creating a data set with trade-level information, we contribute to a steadily growing literature that analyzes trade-level data of financial markets (e.g., Bessembinder et al. (2006), Harris and Piwowar (2006), Edwards et al. (2007), Green et al. (2007b,a), Hender-

[6]This is in line with concerns that have been raised in the industry: Dealers are accused to "long [have] 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).

shott et al. (2011), Lagos et al. (2011), Brancaccio et al. (2017), Iercosan and Jiron (2017), Di Maggio et al. (2019), Hangströmer and Menkveld (2019), Li and Schürhoff (2019), Collin-Dufresne et al. (2020), Riggs et al. (2020), Hennig (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 (Bessembinder et al. (2020)).

By showing that—even in this market—there is evidence for price discrimination, we add to existing evidence of price discrimination in less liquid or more opaque OTC markets (e.g., Green et al. (2007a), Jankowitsch et al. (2011), Hau et al. (2019), Hendershott et al. (2020)).[7] In particular, we support findings by Hau et al. (2019), who compare price discrimination on and off an electronic platform in the OTC market for foreign exchange derivatives via OLS regressions, in conducting an event study.

By estimating the demand and elasticity of demand 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), Fleming (2003), Krishnamurthy and Vissing-Jørgensen (2012, 2015)), and a young literature that estimates demand for financial assets (e.g., Koijen and Yogo (2019, 2020)). For estimation, we exploit techniques used to study (multi-unit) auctions to construct a cost-shifter instrument for prices outside of the auction (e.g., Hortaçsu (2002), Kastl (2011), Hortaçsu and McAdams (2010), Allen et al. (2020)).[8]

Our theory lies in between the theoretic literature on OTC markets (e.g., Duffie et al. (2005), Weill (2007), Lagos and Rocheteau (2009), Zhu (2012), Hugonnier et al. (2018), Carvalho (2020), Kakhbod and Song (2020)), and a large theoretic literature that studies decentralized or fragmented financial markets (e.g., Glosten (1994), Glode and Opp (2016) Li and Song (2019), Chen and Duffie (2020), Colliard et al. (2020), Rostek and Yoon (2020), Wittwer (2020a,b)).[9] Our model differs from most papers in both literatures in that it focuses on the selection of investors into trading venues, similar to few other papers (e.g., Hendershott and Mendelson (2000), Zhu (2014), Liu et al. (2018), Vogel (2019)).[10]

[7] Price discrimination has also been documented in stock markets (e.g., Bernhardt et al. (2007)).

[8] Further, we apply an approach by Bresnahan (1981, 1987), Berry (1994) and Berry et al. (1995) that is commonly used in the literature on demand estimation to infer marginal costs of firms from observable behavior in a trade setting. Here, marginal costs become values for realizing trade.

[9] A non-exhaustive list of contributions on fragmentation of equity markets includes Hamilton (1979), Mendelson (1987), Chowdhry and Nanda (1991), Stoll (2001), O'Hara and Ye (2011), Baldauf and Mollner (2019), Budish et al. (2019).

[10] In assuming that investors either participate in the bilateral trade or on the platform, we relate to contributions that assume exclusive participation per market segment (e.g., Rust and Hall (2003),

Different from these papers, we highlight the importance of benchmark prices, as in Duffie et al. (2017). They show how exogenous benchmarks affect trading incentives in traditional OTC markets without platforms. We let dealers choose the benchmark prices strategically, incentivizing investors to either trade bilaterally or on a platform.

Unlike most papers in the OTC literature, we do not highlight search frictions or 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 what determines 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 (RFQ) auction. This is in line with our empirical findings.

Paper overview. The remainder of 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 which build the basis for the welfare analysis in Section 8. Section 9 concludes.

2 Institutional Environment

Our model and empirical analysis build on our institutional knowledge. It is therefore useful to explain the setting before going into the analysis.

Market structure. Government bond markets around the world divide into a centralized primary market and a (decentralized) OTC market. The latter often makes up a large part of a country's total bond market.[11]

In the primary market, the government issues bonds via regularly held auctions. Most of the auctioned amount is bought by (primary) dealers who actively trade with investors in the OTC market. In Canada, there are ten such dealers. They handle 99% of the traded volume with investors on a median day, and buy 99% of the amount that is issued to dealers and investors in a primary auction.

Miao (2006), Yoon (2018), Lee and Wang (2019)). Dugast et al. (2019) only recently relaxed this assumption.

[11] For example, the OTC market for Canadian government bonds accounts for 70% of the Canadian bond market.

The OTC market for government bonds is structured similarly to other OTC markets. It splits into a segment in which dealers trade with other dealers (or brokers) and one in which dealers trade with investors. Trades are realized either via bilateral negotiation or an electronic platform(s).

Platforms. Electronic platforms are called Alternative Trading Systems (ATS) in the U.S. and Canada, multilateral trading facility (MTF) in Europe, and Dark Pools for equities. They come in different formats, but have in common that trades are matched (more) easily between buyers and sellers.[12]

In this paper, we focus on dealer-to-investor platforms which match investors to dealers but not to other investors. This increases the competition that dealers face relative to when they negotiate terms of trade bilaterally. Given the dealers' strong influence in OTC markets, however, there are reasons to believe that the existing platforms are not designed to maximize investor surplus. Further, most platforms—including the largest ones in the U.S.—are only accessible to institutional investors.[13]

In Canada, until recently, there was only one ATS in the dealer-to-investor segment: CanDeal (described in more detail below). It is owned by the largest dealers in the market, but is otherwise similar to dealer-to-investor platforms in other markets, including the largest platforms (e.g., MarketAxess, Tradeweb, Bloomberg SEF).

Who are the investors? Government bonds are attractive for many different investors because they offer greater safety and liquidity than other securities and are used in various ways, for instance, as collateral, to hedge when trading risky assets, or to meet short-term obligations or regulatory requirements. To get a sense of who investors are, we manually categorize into types 1,459 investors that we can identify by name (see Appendix A.2). We find that asset managers are the largest investor group. They trade the highest volume, followed by banks and pension funds. Thereafter, we have public entities (such as governments, central banks or universities), insurance companies, firms that offer brokerage services and non-financial companies.

[12]See Bech et al. (2016) and Bessembinder et al. (2020) for an overview of electronic platforms.

[13]A non-exhaustive list of ATS includes MarketAxess (the leader in e-trading for global fixed-income), BGC Financial L.P. (which offers more than 200 financial products), BrokerTec Quote (leading in the European Repo market), Tradeweb Institutional (a global operator of electronic marketplaces for rates, credit, equities and money markets)

How do investors trade? If an investor is interested in trading a bond, she typically consults Bloomberg (or another information provider like Thomson Reuters) to check dealer-advertised prices, and prices that reflect the bond's current average market value (as shown in Appendix Figure A1). Next, the investor contacts her dealer, traditionally over the phone or text. The dealer makes a take-it-or-leave-it offer which the investor either accepts or declines. If she declines, she could call up other dealers to seek more bilateral offers. A typical investor, however, does not do so. This is common across OTC markets, where existing studies show that many investors only trade with a few dealers bilaterally (e.g., Hendershott et al. (2020)).[14] In our setting, investors usually have a single dealer with whom they trade bilaterally (see Appendix Figure A2a).

Institutional investors have an alternative to bilateral trading: They can trade on an electronic platform (CanDeal) (see Appendix FigureA2b). On CanDeal, an investor has access to more dealers simultaneously. She observes quotes that dealers post (only) for platform participants, and has three alternatives on how to trade. All have in common that dealers more directly compete with one another than in the bilateral market segment: Investors can either request quotes from multiple dealers by running a request for quote (RFQ) auction, trade on a limit order book, or contact a dealer and ask to execute a trade at the dealer-posted quote (Berger-Soucy et al. (2018)).

According to CanDeal, among trades that are executed by them, more than 95% of the trade volume goes via RFQ auctions. In such an auction an investor sends a request to up to 4 dealers at a time. 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. This means that the investor has to share potentially valuable information with all the dealers that she asks. Knowing how many, but not which dealers participate, dealers respond with a price. The investor chooses the deal that she likes best (typically the best price) and the trade is executed shortly after.

Who is classified as an institutional investor? To become an institutional investor and obtain the right to access the platform, investors have to fulfill certain legal requirements. Fulfillment of these requirements is monitored through occasional audits. In Canada, non-individuals with total securities under administration or management exceeding C$10

[14] Hendershott et al. (2020) examine the network formation of insurers from 2001-2014 in the corporate bond market with more than 400 active broker-dealers. They find that insurers form small but persistent networks. About one-third of insurers trades with a single dealer annually.

million if Canadian and C$15 million if foreign, and corporations with a minimum net worth of C$75 million classify as institutional. Public entities, like federal governments, Canadian provincial governments and larger Canadian cities, as well as specific crown corporations, also classify as institutional investors. In addition, regulated entities (such as institutions that participate in the Canadian Investor Protection Fund (CIPF) which protects investor assets in case of bankruptcy), members of recognized exchanges and associations, or institutions that are in the business of trading in, or advising on, securities are institutional. Anyone who is not an institutional investor is classified as a retail investor.[15]

Platform usage costs. To use platforms, an eligible investor has to pay a small monthly fee, which ranges between C$725 and C$3,035, depending on usage. 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. This is true in our data, as well as in other markets, such as the U.S. investment-grade corporate bond market.[16]

Switching to the platform rather than immediately realizing the trade with the dealer takes some time and effort. Further, investors could be averse to sharing information about the trade with more than one dealer on the platform (Hendershott and Madhavan (2015)).[17] In addition, investors may not want to damage an existing relationship by appearing to trade on a platform (Hendershott et al. (2020), Riggs et al. (2020)).[18]

3 Data

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

[15]For more details, consult IIROC Rule Book (2020).

[16]According to an industry report 19% of trading in the U.S. investment-grade corporate bond market is on electronic platforms (see McPartland (2016)).

[17]Supporting this idea, a powerful trade group in Washington warns: A "key mechanism suppressing buy-side trading on IDB SEFs [i.e., electronic platforms] and perpetuating the current two-tier market structure [of bilateral and platform trading] is the legacy practice of post-trade name disclosure. Even though otherwise eligible buy-side participants [i.e., investors] have access to all SEFs in theory, the loss of anonymity caused by the continuation of post-trade name disclosure deters buy-side access to IDB SEFs in practice" (Managed Funds Association (2015), p.2).

[18]Relationships might matter for different reasons: The investors may be able to post low collateral with the dealer to trade (low margins). She may have better options of realizing a potentially large trade fast and at a good price when she needs it (balance sheet space) or get better terms for other activities such as lending or borrowing over night (haircuts). Finally, an investor might like to be told that she obtains the most favorable price thanks to her loyalty.

Main data source. The main source is the Debt Securities Transaction Reporting System, MTRS2.0. It is collected by the Industry Regulatory Organization of Canada (IIROC) since November 2015.[19] Our sample entails trade-level information on all bond trades of registered brokers or dealers from 2016-2019. 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.[20]

The most unique feature of the data is that each dealer (and broker) carries a unique legal identifier (LEI). Investors either have a LEI or a dealer-specific identifier. The latter is an anonymous dealer-specific account ID. Of all trades with investors, about 25% are with investors that have a LEI and can be identified by name and traced throughout the market.

Similar to the TRACE data set that has been used to study the U.S. corporate bond market, the MTRS2.0 data is self-reported. We do our best to clean out irregularities (see Appendix A.1). The cleaned sample includes almost all (cash) trades of Canadian government bonds. It misses trades between investors, which do not have to be reported but are rare according to market experts.[21] 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 to the full volume which must be reported to the Canadian Depository for Securities. Our data covers approximately 90% of all trades involving Treasury bills.

Additional data sources. We augment our trade-level data with four additional data sources. First, we obtain bidding data of all government bond auctions between 2016 and July 2019 from the Bank of Canada. We 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

[19] This might seem like a short time-period, but is in fact rather long compared to other countries. The U.S., for example, only began recording similar information in TRACE mid-2017.

[20] For a full description of all available variables, see MTRS 2.0 User Guide (2016).

[21] 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 (`https://regportal.azurewebsites.net/hcprod/frbny/20200929/`).

in-house trades, that is, trades between a dealer and one of its subsidiaries.

Third, we obtain data on quotes from the electronic platform (CanDeal). For each security that is traded on the platform, we observe an average of hourly bid and ask quotes (per security) that were displayed on the platform. These quotes are indicative. They are meant to give investors an idea of the prices at which they can sell (bid) and buy (ask). We do not observe the quotes that each dealer posts, nor do we have any other information of trading behavior on the platform. This means, for example, that we do not know whether a trade realized via an RFQ auction or on the limit order book.

Fourth, we download the hourly Bloomberg Generic (BNG) mid-price for each security. It is the average of the BNG bid and ask price that are publicly posted on a Bloomberg screen.[22] We use this price as a proxy for a bond's true market value that is commonly known to everyone. We think that this is a reasonable assumption because the BNG mid-price is very close to the price at which dealers trade with one another (see Appendix Figure A3). The inter-dealer market price, in turn, is often taken as the true value of a security in the related literature.

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 7am and after 5pm) 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 model, we impose some additional restriction so as to construct an instrument for quotes using bidding data of the primary auctions. We focus on primary dealers, and drop trades after July 2019 because we currently do not have auction data for the second half of 2019. Due to data reporting, we exclude one dealer, which leaves us with 9 dealers. Further, we exclude trades that are realized before the outcome of a primary auction was announced. This is 10:30 am for T-Bills and 12:00 am for Bonds. Appendix Table 4 summarizes all sample restrictions.

[22]The BNG bid and ask prices are calculated by Bloomberg using prices that are contributed to Bloomberg (including the CanDeal quotes) and any other information that Bloomberg considers relevant with the goal of providing "consensus" pricing. The concrete methodology is proprietary.

3.1 Key variables and market features

Unit of measurement. Following market convention, we convert each price into a yield to maturity and report our findings in terms of yields rather than prices. This yield is the annualized interest rate that equates the price with the present discount value of the bond. A higher price implies a lower yield and vice versa.[23]

All yields are expressed in basis points (bps). 1 bps is 0.01%. It 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. Explaining all of them is beyond the scope of this paper. Our approach, instead, is to control for factors that are not endogenous in our model (as discussed in Section 5.3): trade size, time-trends and unobservables that are security specific.

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 (buy_{thsij}), a flexible function of trade size, $f(quantity_{thsij})$, an hour-day fixed effect (ζ_{th}) and a security-week fixed effect (ζ_{sw}):

$$yield_{thsij} = \alpha + \beta * buy_{thsij} + f(quantity_{thsij}) + \zeta_{th} + \zeta_{sw} + \epsilon_{thsij}.$$

Here, we choose a third degree polynomial for $f(quantity_{thsij})$, yet our results are not sensitive to how we control for trade size. We then construct the residual yield: $y_{thsij} = \hat{\alpha} + \hat{\epsilon}_{thsij}$ if the investor sells and $y_{thsij} = \hat{\alpha} + \hat{\beta} + \hat{\epsilon}_{thsij}$ if she buys. The remaining variation in y_{thsij} is by construction smaller than in the raw data (see Appendix Figure A4).

In addition, we normalize the Bloomberg yield by subtracting the estimated fixed effects $(\hat{\zeta}_{th}, \hat{\zeta}_{sw})$. We call the residualized trade yield y_{thsij} and the normalized Bloomberg yield θ_{ths}, and use these normalized yields throughout the paper, that is, for the reduced-form evidence, as well as to estimate our model. Our reduced-form evidence is robust to using the yields in the raw data rather than the normalized yields.

[23] To incorporate all market conventions as described in IIAC (2020), we follow the common practice of market participants and use Bloomberg's YAS function to convert prices into yields. To illustrate how the yield relates to the price, consider a simple example: Assume that you buy a bond that is issued today and will mature in two years. At the maturity date you obtain its face value of C$100. Before that, you receive a coupon of C$10 every 6 months. The yield (to maturity) is the annualized interest rate that equates the price with the present discount value of the bond: price $= \sum_{t=1}^{2y} \frac{\text{coupon}}{(1+\frac{\text{yield}}{2})^t} + \frac{\text{face value}}{(1+\frac{\text{yield}}{2})^{2y}}$. Here, paying a price of C$137.53 would give a yield of 1%.

Key market features. The typical trade between one of the 10 dealers and one of the 546,048 investor ids is small (see Appendix Figure A10). It involves a bond that is actively traded because it was issued in a primary auction less than 3 months ago. 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, suggesting that market reforms could affect welfare.

4.1 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 is the yield for the investor, independent of whether she buys or sells.

Figure 1 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 gap implies that a retail investor earns an annual interest that is $4/150 * 100 = 2.67\%$ lower than an institutional investor.

4.2 Better yields thanks to platform access?

Whether market reforms that shift trading onto the platform could have sizable effects on prices or welfare depends on what drives the yield gap. One possibility is that retail investors are willing to pay more and therefore pay higher prices, independent from whether they have access to the platform or not. Another possibility is that institutional investors realize better yields in parts because they have platform access. If this is so, there is scope for market reforms.

The ideal but infeasible experiment to establish a causal link between platform access and yields is to randomly assign the right to access the platform to some investors (the treatment group) but not all investors (the control group). We, instead, leverage the fact that 90 institutions lost the right to access the platform in our sample, because they no longer fulfill certain legal requirements to be classified as an institutional investor.

Why do investors lose platform access? Investors might lose the institutional status for different reasons. The first is that the investor no longer holds sufficient capital or is no longer willing to prove that it does.[24] The second is that the investor terminates her membership with a regulated entity, such as the CIPF.[25] The third is that she stops selling securities, offering investment advice or managing a mutual fund which implies that she no longer is classified as a registrant. Our data does not allow us to disentangle these reasons because switchers are not reported with their legal identifier, which means that they are anonymous to us.

Platform access and yields. To conduct an event study, we define the event of investor i as the first time that we observe this investor as a retail investor. 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

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

$$\text{where } D_{mi} = 1 \; m \text{ months before/after } i \text{ loses access, 0 otherwise}$$

$$\zeta_i, \zeta_{th}, \zeta_s, \zeta_j = \text{investor, hour or the day, security, dealer fixed effects, respectively.}$$

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. To obtain sufficient power, we bucket time into months and pool the buy and sell side of the trade.

[24]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.

[25]The CIPF is by far the largest of such entities. It protects investor assets in case of bankruptcy and publishes a list of current and past members online. During our sample period, 23 institutions canceled or lost membership to CIPF. If these institutions traded Canadian government bonds and do not hold sufficient capital to fulfill the size requirements, they would be switchers in our data.

The size of the hour, security and dealer fixed effects is pinned down by trade information of retail investors who never obtain access in our sample, as these investors are likely more similar to those who lost access than 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 2). The yield drops on average by 1 bps in the first month and decreases further by about 4 bps thereafter. This suggests that platform access matters for yields.

The relationship would be causal if the loss of platform access was an exogenous event. This means that the investor does not choose to lose access, and that there are no systematic changes over time that affect the yield other than the loss of platform access. Among these two conditions, we think of the first as less problematic because there is no obvious reasons, such as fees or costly obligations, for which an investor might want to lose the status of being institutional. This is especially true if institutional investors obtain better yields.

The second condition would be violated if the investor changes the way in which she trades, and, as a result, obtains worse yields. For example, an investor might trade less frequently and become less sophisticated, so that she realizes worse yields in bilateral negotiations. To lessen this concern, we show that we observe no significant change in variables that capture different aspects of trade behavior when the investor loses platform access (see Appendix Figures A5 and A6).

The condition would also fail if retail investors obtain worse yields because they are riskier trading partners than institutional investors. One example would be a firm that suffers a big financial loss which reduces the firm's net worth below the threshold to classify as an institutional investor and increases the probability to default on a trade. If the dealer knew this, she would price in the additional risk. Often, however, the dealer does not have such detailed information. One exception would be if the investor was publicly excluded from the CIPF because she imposed a threat to the other CIPF members. This is a drastic measure, however, that is rarely taken by the CIPF so that it is unlikely to be driving our results.

Summary. We have gathered novel suggestive evidence that platform access matters for trade yields. This implies that there is scope to increase investor surplus by centralizing the market and shifting bilateral trades onto platforms. To quantify such benefits and assess welfare effects more generally, we now set up a model.

5 Model

We first introduce the model and then discuss the main simplifying assumptions and extensions. Beyond the model-free empirical evidence, the model allows us to quantify total gains from trade which measures total welfare. For this, we need to know the valuations for realizing trade. In the data, we only observe the transaction prices, which lie somewhere between the value of the buyer and the value of the seller. Further, we can account that institutional investors select between trading bilaterally and trading on the platform. This selection problem can bias estimates of OLS regressions. Lastly, we can take into account how dealers and investors respond when changing the market rules and centralize the market by allowing all investors to trade on the platform.

In line with the rest of the paper and the estimation, we formulate the model using yields rather than prices. To make this conversion, it helps to keep in mind that the yield is like a negative price: the higher the yield, the lower the price and vice versa.

We consider two separate games: One in which dealers sell to investors and one in which dealers buy from investors. Here we present the setting with buying investors and leave the game with selling investors for Appendix C.2.

5.1 Model overview

Dealers sell a bond to institutional and retail investors. They set quotes to maximize the expected profit from trading with investors. The quotes are posted publicly. They inform the investors about the yields that they may expect to realize when buying on the platform. This is motivated by the empirical fact that the individual trade yields on the platform are on average identical to the quotes that dealers post on the platform. Given the quotes, institutional investors decide whether to buy on the platform or bilaterally with a dealer, while retail investors can only buy bilaterally.

In choosing a quote, a dealer faces a tradeoff: When she decreases the quote, the platform becomes less attractive for institutional investors. As a result, more of them stay off the platform and buy bilaterally. On the other hand, when she increases the quote, more investors who enter the platform buy from her, increasing her platform market share.

An institutional investor also has a tradeoff: 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 she 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.

5.2 Formal details

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.

On a fixed day t, $J_t \geqslant 2$ dealers sell a bond to infinitely many investors, bilaterally or on an electronic platform. Each transaction is a single unit trade. The market value of the bond is $\theta_t \in \mathbb{R}^+$. It is commonly known and exogenously determined by macro-economic factors.

Motivated by the empirical feature that investors tend to trade with a single dealer bilaterally, each investor has a home dealer. We call this dealer d, short for d_i. Each dealer, thus, has a home investor base. It consists of two investor groups, the institutional and the 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 $\kappa^I + \kappa^R = 1$. Among 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.

Dealers. Each dealer seeks to maximize profit from trading with investors. Ex-post, dealer j obtains a profit of $\pi_{tj}(x) = v_{tj}^D - x$, when selling one unit at yield x. $v_{tj}^D \in \mathbb{R}$ is known to all dealers but may be unobserved to investors. One may think of it as a reservation value, which reflects the minimal price at which the dealer would sell to the investor. It would take the form of a continuation value in a dynamic version of the game.

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 to hold in reality—for instance, because it is costly to hold inventory or to be available to trade at any point in time—we refrain from imposing $v_{tj}^D = \theta_t$.

At what yield the dealer sells depends on whether she sells on the platform or bilaterally. In a bilateral offer, the dealer observes how much the investor is willing to pay. In our main specification, she extracts all trade surplus by setting the yield equal to the investor's full willingness to pay, specified below. This assumption captures the feature that dealers are able to extract more (here all) information from investors, and with it surplus, when negotiating in person than on the platform.

In contrast, when selling on the platform the yield hinges on the quote, q_{tj}, that the dealer posts on the platform. The quote is chosen simultaneously with all other dealers at the beginning of the game. It is not targeted to a specific investor, but valid for all institutional investors who have access to the platform. Specifically, a quote signals the yield that an institutional investor expects to obtain when buying from the dealer on the platform. When choosing the quote each dealer maximizes the expected profit from selling to institutional investors. She believes to sell at this quote on the platform, and takes the expectation over how much an institutional investor is willing to pay.

Retail investors. Each retail investor buys from her home dealer d. She obtains a surplus of $y^R_{tid} - v^R_{tid}$ when buying at yield y^R_{tid}, where

$$v^R_{tid} = \theta_t + \nu^R_{ti} - \xi_{td} \text{ with } \nu^R_{ti} \overset{iid}{\sim} \mathcal{F}^R_t \text{ and } \xi_{tj} \sim \mathcal{G}_t$$

is the investor's (true) willingness to pay. It splits into two parts. The first, $\theta_t + \nu^R_{ti}$, is the investor's reservation value for the bond. It is the difference between the value of holding a unit of the bond versus not. We assume that this value is additively separable into the market value of the bond, θ_t, and an individual liquidity shock of the investor, ν^R_{ti}. The shock is drawn iid from a commonly known distribution with a continuous CDF $\mathcal{F}^R_t(\cdot)$ that has a strictly positive density on the support. It is private information to the investor. In reality, the liquidity shock may be driven by individual hedging or trading strategies, balance sheets concerns, or cash needs of an institution.

The second component of the investor's willingness to pay, ξ_{tj}, is the quality of the chosen dealer. For simplicity, it is common to all investors. It is drawn from an arbitrary distribution with CDF $\mathcal{G}_t(\cdot)$. In practice, dealer quality captures anything that makes trading with this dealer attractive. 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 (such as offering investment advise on a broad range of securities including foreign stocks or debt, having automated payment systems, or underwriting debt).

Institutional investors. An institutional investor obtains a similar surplus as a retail investor, $y^I_{tid} - v^I_{tid}$. The difference is that we allow the liquidity shocks and yields to systematically differ across investor groups:

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

Which yield the institutional investor realizes depends on whether she buys bilaterally from her home dealer or on the platform. In a bilateral trade, the yield equals the investor's full willingness to pay v_{tid}^{I} as for retail investors. On the platform, an institutional investor can choose any dealer. The yield is similar to the posted quote of the chosen dealer. How similar depends on how fiercely dealers compete for the investor on the platform (see Appendix B.1 for motivating empirical evidence).

Platform competition. The transaction yield that an investor realizes when buying from a dealer on the platform equals to the posted quote plus a friction:

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

The friction term captures the fact that the platform might not be perfectly competitive. It consists of a known parameter σ that measures the degree of competition on the platform and scales up an idiosyncratic shock ϵ_{tij}. This shock is drawn iid from a commonly known distribution with CDF $\mathcal{H}_t(\cdot)$ and captures anything that prevents an investor from buying from the best dealer with the highest posted quote.

One microfoundation for the friction term is formalized in Appendix D. Here, we model the auction that investors run when entering the platform after they observed the posted quotes. We cannot estimate this microfoundation because we do not observe how investors trade on the platform, but we can derive theoretic insights. For instance, we show that bid shading causes the winning yield to differ from the posted quote. An alternative microfoundation could formalize the fact that a dealer might not respond in an RFQ auction (as in Liu et al. (2018), Riggs et al. (2020)).[26]

The framework nests the two extreme forms of competition: When $\sigma = 0$, all investors buy from the dealer who offers the best quote and quality. It is the one with $\max_j\{\xi_{tj} + q_{tj}\}$. In that case, the platform is perfectly competitive and dealers compete à la Bertrand. If

[26]Such a model would be inspired by Liu et al. (2018). In their model, an investor with value v runs an RFQ auction. She asks $J \geqslant 2$ dealers, who all have the same reservation price q, for a quote. Each dealer responds with some probability that is strictly less than one. Conditional on receiving at least one response by a dealer, the expected price of the RFQ auction is $q + \sigma\epsilon$, where σ is a function of model parameters and $\epsilon = (v - q)$.

$\sigma \to \infty$, 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.

Platform usage cost. To use the platform and choose from any of the dealers, the investor has to pay a cost c_t. It represents any obstacle to access the platform, including privacy concerns or relationship costs (as described in Section 2), and is motivated by additional empirical evidence (see Appendix B.1). For simplicity, in our main specification, the cost is common to all investors and commonly known.

Timing of events. The sequence of the game is inspired by discussions with market experts (as described in in Section 2):

(0) Dealers observes $\theta_t, \xi_t, v_{tj}^D$, and investors observe θ_t.
N_t^G investors of both groups $G \in \{I, R\}$ draw liquidity shocks ν_{ti}^G

(1) Dealers simultaneously post $q_{tj} \in \mathbb{R}^+$ for everyone to see.

(2) Each investor contacts her home dealer, who observes ν_{ti}^G and offers $y_{tid}^G = v_{tid}^G$.

(3) Retail investors accept the offer. An institutional investor can accept the offer or enter the platform. In the latter case, she pays cost c_t, observes the platform shock ϵ_{tij} and decides from which dealer to buy at $q_{tj} + \sigma\epsilon_{tij}$.

Equilibrium. A pure strategy equilibrium can be derived by backwards induction.

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

$$y_{tid}^R = \theta_t + \nu_{ti}^R - \xi_{td}. \tag{3}$$

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

$$y_{tid}^I = \theta_t + \nu_{ti}^I - \xi_{td} \quad \textit{if } \psi_t(q_t) \leqslant \nu_{ti}^I \tag{4}$$

$$\textit{where } \psi_t(q_t) = \mathbb{E}[\max_{k \in \mathcal{J}_t} \tilde{u}_{tik}(\boldsymbol{\epsilon_{tik}})] - \theta_t - c_t \textit{ with } \tilde{u}_{tij}(\boldsymbol{\epsilon_{tij}}) = \xi_{tj} + q_{tj} + \sigma\boldsymbol{\epsilon_{tij}}. \tag{5}$$

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

The proposition characterizes where investors buy and at what yields. By assumption, retail investors always buy at a yield that equals their willingness to pay (statement (i)).

Statement (ii) tells us that an institutional investor trades on the platform if she expects the surplus from buying on the platform minus the platform usage cost, $\mathbb{E}[\max_{k\in\mathcal{J}_t}\tilde{u}_{tik}(\boldsymbol{\epsilon_{tik}}) - (\theta_t + \nu^I_{ti})] - c_t$, to be higher than the zero surplus that she receives in a bilateral trade. This is the case for investors who are more urgent to trade in that they 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 the investor's individual willingness to pay.

Overall, the proposition highlights that yields for the 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, the others buy bilaterally.

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

$$q_{tj}\Big(1 + \frac{1}{\eta^E_{tj}(q_t)}\Big(1 - \frac{\partial \pi^D_{tj}(q_t)}{\partial q_{tj}}\Big/ S_{tj}(q_t)\Big)\Big) = v^D_{tj} \tag{6}$$

where $\eta^E_{tj}(q_t)$ is dealer j's quote elasticity of demand on the platform and $\frac{\partial \pi^D_{tj}(q_t)}{\partial q_{tj}}$ is the marginal profit that the dealer expects from bilateral trades with institutional investors, which is normalized by the size of her platform market share, $S_{tj}(q_t)$.

Formally, $\eta^E_{tj}(q_t) = q_{tj}\frac{\partial S_{tj}(q_t)}{\partial q_{tj}}/S_{tj}(q_t)$ with $S_{tj}(q_t) = \sum_{j\in J_t}\Pr(\boldsymbol{\nu^I_{ti}} \leqslant \psi_t(q_t))\Pr(\tilde{u}_{tki}(\boldsymbol{\epsilon_{tki}}) < \tilde{u}_{tij}(\boldsymbol{\epsilon_{tij}})\ \forall k \neq j)$ where $\tilde{u}_{tij}(\boldsymbol{\epsilon_{tij}})$ and $\psi_t(q_t)$ are defined in (5) of Proposition (1); and $\pi^D_{tj}(q_t) = \mathbb{E}[v^D_{tj} - (\boldsymbol{\nu^I_{ti}} + \theta_t - \xi_{tj})|\psi_t(q_t) \leqslant \boldsymbol{\nu^I_{ti}}]$.

Proposition 2 characterizes the quotes that 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 true valuation, v^D_{tj}. To obtain an intuition for 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 firm(s) who set price(s) 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 dealer is an oligopolist. Her marginal cost is v^D_{tj}. Since she chooses quotes in yields (rather than prices) the markup becomes 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 for investors.

In summary, when choosing the quote the dealer trades off the profit that she can make from selling bilaterally where she extracts a higher trade surplus with the profits that she can gain on the platform when stealing some of the investors of other dealers.

5.3 Discussion

Our model builds on several simplifying assumptions. First, we assume that there is a given number of dealers and investors who actively trade on a day and that no trade between them fails. We do so because—similar to most existing studies with trade-level data—we only observe realized trades. This implies that we have no information to estimate how many investors or dealers chose not to trade.

This assumption is not problematic for the dealers that we consider, which are the primary dealers. The reason is that they have an obligation to actively trade as "market makers" to keep the status of a primary dealer. In line with this, we observe that the least active dealer trades on 98% of dates. For investors, the assumption is more likely to hold for shorter than longer bonds, which make up the majority of trades in our sample. The reason is that short bonds typically serve as overnight collateral, as hedging vehicles, or for cash management on short notices so that investors seek to realize the trade by the end of the day.

Second, our game does not connect multiple days. In particular, we assume that the dealers' and investors' values for the bond are independent of prior trades. This implies that we set aside trading strategies that connect multiple days or involve buying and selling. Dealers and investors can still trade everyday, 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 assume in our main specification that the dealer charges a bilateral price that equals to the investor's full willingness to pay. Here, we rule out that the dealer has to form expectations over how much the investor is willing to pay or that the investor has some bargaining power in the bilateral negotiation. We do this because there is no information in the data that allows us to identify the dealer's beliefs or bargaining power.

From a theoretic viewpoint, this is a relatively strong assumption. It implies that dealers do not adjust the yields that 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 in Appendix B.3, and find supporting evidence.

In addition, we later verify that our findings are robust when we allow the investor to extract some surplus in the bilateral trade. For this we extend the model in Appendix C.3. In this extension, the dealer and investor Nash-bargain in a bilateral trade, and the investor has bargaining power $\phi \in [0, 1]$. The bilateral yield now lies in between the dealer's valuation and the investor's valuation, and the investor captures a fraction ϕ of the bilateral trade surplus. When $\phi = 0$, the extended model is equivalent to our benchmark model.

Fourth, we assume that all trades have the same size (normalized to one). This assumption is mainly motivated by the fact that most trades are small and of similar size in our sample (see Appendix Figure A10 and Appendix B.2 for additional evidence). In the related theory literature, this assumption is common as it makes models more tractable (e.g., Duffie et al. (2005), Weill (2007), Duffie et al. (2017), Lee and Wang (2019)). For us, it implies that we cannot analyze whether and how trade sizes and volume change as we change the market rules. Further, it implies that all of our estimation findings hold per unit of the trade, rather than the total amount of the trade.

Finally, we abstract from order splitting in assuming that investors either trade bilaterally or on the platform but not both. This is motivated by the empirical observation that investors typically maximally trade once per day, and do not split orders (see Appendix Figure A11).

6 Estimation

We described the model for investors who buy. Bringing in the other side of the trade, we have four investor groups, indexed by G: retail and institutional investors who buy (R and I), and retail and institutional investors who sell (R^* and I^*).

For all investor groups, we want to estimate the daily distribution of the liquidity shocks (F_t^G), in addition to the daily quality of the dealers (ξ_{tj}) and the degree of competition on the platform (σ). We allow the cost of using the platform and the dealer's valuation to depend on whether the investor buys (c_t and v_{tj}^{D}) or sells (c_t^* and v_{tj}^{*D}). Details of the estimation procedure are in Appendix E.

6.1 Identifying assumptions

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

Assumption 1. *The liquidity shocks $\boldsymbol{\nu_{ti}^G}$ are conditional on $\{\theta_t, q_t, \xi_t\}$ 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. Given that we observe very few investors who trade several times within the same day, we believe that this concern it not of first order.

Assumption 2. *All investor groups value the quality of a dealer, $\boldsymbol{\xi_{tj}}$, equally. W.l.o.g. we decompose $\boldsymbol{\xi_{tj}}$ into a part that is persistent over time and a part that might vary: $\xi_{tj} = \xi_j + \chi_{tj}$. The time-varying part, χ_{tj}, is drawn iid across dealers within a day.*

This assumption imposes that all investors value the quality of a dealer in the same way, and that the dealer provides the same quality on and off the platform. This is plausible for factors such as the dealer's balance sheet space, the probability of default or deliver, or the types of services that the dealer offers outside of the government bond market.

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, for example, an exogenous cost shifter that changes the dealer's cost of providing liquidity but is independent of the unobservable demand shocks.

A promising candidate for a cost shifter is how much the dealer won in the last primary auction. This is because dealers make profit from buying at auction and selling the bonds at higher prices to investors. In line with this idea, we observe that auction prices are systematically lower than prices in the OTC market. A dealer can therefore more cheaply satisfy investor demand when winning a lot at auction, as she does not have to buy or borrow the bonds at a higher price post-auction.[27]

One problem with using how much a dealer won as a cost shifter is that this amount depends on how the dealer bid in the auction; and when the dealer bids, she likely anticipates investor demand. First, because it is common for investors to place orders with their dealer before the auction, and second because investors who want to buy in the auction must bid via a dealer. This means that the dealer can observe the investor's demand before passing the bid onto the auctioneer (Hortaçsu and Kastl (2012)). Together this implies that the amount won is likely correlated with the unobservable demand shocks, which would bias the results.

Our solution is to subtract the amount that the dealer expected to win when bidding from the amount that she actually won:

$$\begin{aligned} won_{\tilde{t}j} &= \text{amount dealer } j \text{ won at the last auction day } \tilde{t} \\ &\quad - \text{amount she expected to win when placing her bids.} \end{aligned} \tag{7}$$

To compute the expected amount, we model the bidding process in the auction and use recent techniques of the empirical literature on (multi-unit) auctions (see Appendix E.1 for details). In a nutshell, we estimate how much a dealer expects to win in each auction by resampling bids and simulating market clearance: We randomly draw bids with replacement from the other bidders and let the market clear according to the auction rules. This gives one realization of how much the dealer wins. We repeat this many times to obtain the empirical distribution of how much the dealer wins. With it, we compute the amount that the dealer expected to win in the auction.

In our estimation, we then rely on the following exclusion restriction:

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

To better understand whether this assumption is plausible, it helps to think through where the surprise shock and with that the identifying variation of the instrument comes from. It

[27]Notice that there is some tension between our assumption that the value of dealers today is independent of the trades that realized yesterday, and the idea that dealers can more cheaply satisfy investor demand when they win a lot at auction.

In principle, winning a lot could also increase the cost of providing liquidity. This is because a dealer has to carry a large inventory position if she wins a lot. This can be costly as it involves taking on the risk that prices or demand might fall.

comes from how much or or less the government issues in the auction relative to what it announced it would and from how auction participants other than the dealer herself bid.

With this in mind, the biggest threat to identification is the following scenario: One dealer receives a negative shock and bids less. This makes the other dealers win more than expected. If this also increases how much demand the dealers face because investors substitute from the unlucky dealer to the remaining dealers, the exclusion restriction would be violated. However, in our data, we see relatively few substitution 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 think 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 that 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. It implies that the model restricts how investors substitute across different dealers.[28] Since this type of substitution is rare in the data and, hence, not the focus of this paper, we do not think that this simplifying assumption is problematic for what is first-order.

Assumptions 1 - 4 summarize the main assumptions that are essential for identification. In addition, 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: $\xi_{t0} = 0$ for all 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 4.[29]

[28] Train (2002) explains the restrictions and implications of this assumption on pp. 54-57 in detail.

[29] 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.

Parametric Assumptions.
(i) Platform frictions $\boldsymbol{\epsilon_{tij}}$ *are Extreme Value Type 1 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.

6.2 Identifying variation

We estimate the model separately for each investor group. Here we only discuss the estimation for institutional investors who buy, and leave the other groups for Appendix E.2.

Key variables. The first set of variables that is used in the estimation are market shares. We define the dealer's market share on the platform (among investors who enter the platform) as

$$s_{tj} = \frac{\text{number of times dealer } j \text{ sold on the platform on day } t}{\text{number of total sales on the platform on day } t} \tag{8}$$

and the dealer's bilateral market share relative to her platform market share as

$$\rho_{tj} = \frac{\text{number of times dealer } j \text{ sold bilaterally on day } t}{\text{number of times dealer } j \text{ sold on day } t}. \tag{9}$$

The second set of variables are the normalized yields of Section 3.1. We approximate the bond's daily market value by the normalized Bloomberg yield, averaging across securities and hours of the day: $\theta_t = mean_{hs}(\theta_{ths})$. Further, we approximate the quote at which dealer j sells on a day by the average yield at which she sells on the platform on that day: $q_{tj} = mean_i(y_{tij} | \text{investor } i \text{ buys from dealer } j \text{ on the platform})$. We think that this is a reasonable approximation because the posted average quote (across dealers) that we observe is very similar to the realized trade yields on the platform.

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

To derive an intuition of how to identify the competition parameter (σ), 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 lowest cost and with it the lowest quote q_{tj}— captures the entire platform market share on day t. As σ increases, this dealer loses more and more of her market share to the other dealers. Which dealer gains how much of the market share depends (besides σ) on the dealers' costs shocks. Therefore, σ is mainly identified from the within-day correlation between the dealers' (daily) platform market shares and their cost shifters which is shown in Figure 3a.

The dealer's quality (ξ_{tj}) is determined by how the dealers split the platform market when posting the same (or very similar) quotes: Dealers with higher ξ_{tj} capture a higher market share. 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 3b).

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 4. 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 valuation (v^D_{tj}) from the markup equation (6) 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 that is used in IO to infer marginal costs of firms from firm behavior following Bresnahan (1981).

7 Estimation results

We report the estimates for a median day, for example $\hat{\mu}^I = median_t(\hat{\mu}^I_t)$. We choose the median rather than the average day because it is more robust to outliers. It could, for instance, be the case that (to us) unobservable events drive abnormal trade behavior on few days. An overview of all estimates can be found in Table 1.

7.1 Parameter estimates

Investor's values. We find that buying retail investors have greater urgency to trade in that they are willing to pay $(\hat{\mu}^R - \hat{\mu}^I)/\hat{\mu}^I * 100\% \approx 226\%$ more above the bond's market value than institutional investors. When selling, the difference is smaller, with about 109%. This could be because retail investors who also sell and not only buy differ from retail investors who only buy in that they are actively trading. Taken together this suggests that the yield gap between retail and institutional investors is not entirely driven by platform access, but that differences in the willingness to pay account for some of it. Below we quantify how much.

Yield elasticity of demand. The yield elasticity of demand on the platform (which is mainly driven by the degree of platform competition) is about 174-179. This means that the demand of an institutional investor is relatively inelastic with respect to changes in the quotes on the platform. For example, if a dealer increases the quote at which she sells by 1 bps, demand for this dealer increases by 1.74%. This increase is relatively small compared to a quote change that is 1.7 times higher than the (median) bid-ask spread that is posted on Bloomberg. This says that even if the dealer was willing to sell at a price that is lower than the price at which she usually buys, she would only sell 1.74% more.

The elasticity of demand of an individual investor is similar, even though 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.

Dealer's values and quality. Typically, a dealer values the bonds similar to an average institutional investor on both sides of the trade. For instance, when the dealer sells, he values the bond like an average institutional investor who sells. This is plausible given that the average institutional investor is similar to a dealer in that it is a large financial institution who frequently trades.

While different dealers attach very similar (median) values to the bonds, they systematically differ in quality (see Figure 5). This suggests that there might be welfare gains in matching investors to dealers of higher quality. We quantify these gains below.

Platform usage costs. In line with concerns that have been raised by industry experts, we find that high costs prevent investors from using the platform. With about 3.5 bps, the median cost is much bigger than the actual fee to trade on the platform, which lies between C$725 and C$3,035. For a typical institutional investor this fee is equivalent to a yield loss of about 0.04-0.17 bps per trade. Intuitively, the cost is high because only 30% of platform-eligible investors use it on an average day, even though the expected gain from using the platform is relatively large since dealers differ in quality (see Figure 5a).

7.2 Model fit

Before assessing price and welfare effects from centralizing the market, we validate whether our parsimonious model can replicate the event study in Section 4.2. 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 (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.

8 Counterfactual exercises

We use the model to quantify how much of the gap between retail and institutional investors' yields is due to platform access, and assess changes to investor surplus, dealer profits, and total welfare when centralizing the market by shifting trades onto the platform.

We do this by conducting two counterfactuals. We let retail and institutional investors have access to the platform on which dealers post quotes that are valid for any investor who uses the platform. In the first counterfactual, all investors can use the platform but have to pay the estimated usage costs. In the second, we eliminate all entry barriers by setting usage costs to 0. This removes any type of friction that prevents investors from using the platform.

We think of the zero-cost case as a benchmark that, depending on what drives the platform usage costs, might not be achievable in reality. For example, if the costs reflect that investors value keeping a close relationship with a dealer, it is not obvious whether and how one would remove them. If, instead, the costs are mostly driven by privacy concerns, mandating anonymous trades would remove the costs.

In all counterfactuals we take into account how dealers and investors respond to the changes in hypothetical 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 this equilibrium, all investors select onto the platform as in Proposition 1 and dealers set quotes similar to Proposition 2. Different from the status quo, however, dealers now behave as if there was a "representative" investor who draws liquidity shocks from a normal distribution with mean $\mu_t = \sum_G \kappa^G \mu_t^G$ and standard deviation $\sigma_t = \sum_G \kappa^G \sigma_t^G$, where κ^G is the fraction of trades of investors in group $G \in \{R, I\}$ on an average day. For retail investors $\kappa^R = 0.1$, for institutional investors $\kappa^I = 0.9$.[30]

Throughout, we take the ex-ante perspective, which means that we take the expectation over how many retail versus institutional investors seek to trade and how much they are willing to pay. Further, we keep the number of trades fixed because our data does not allow us to estimate how likely it is that a trade is realized. For the welfare assessment, for example, this means that we focus on the question of who trades with whom and abstract from any gains or losses that may arise because more or less investors trade as market rules change.

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 32-35% when the investor is buying. When the investor is selling, the percentage change is larger with 51-52% because the yield gap in the status quo is smaller.[31]

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.[32] These investors obtain a worse yield than institutional investors because they are willing to accept them in bilateral trades.

Platform participation is weak in part because it is costly. If we eliminate these costs,

[30]In Appendix D.4 we show that this holds numerically, given our parameters. In theory, this is not always the case. The reason is that dealers do not take into account that retail investors may more strongly select onto the platform than institutional investors when setting quotes as if there were a representative investor.

[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]The fraction of retail investors who would use the platform is a bit higher than the fraction of institutional investors, because retail investors are on average willing to pay more. Their liquidity shocks are higher. Hence, more of them select onto the platform (see Proposition 1).

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

We present welfare effects for investors who buy, but our findings generalize to selling investors since the buy- and sell-side of the market are close to symmetric. Further, we show in Appendix F that the main welfare findings are not sensitive to the assumption that the dealer extracts all trade surplus in the bilateral trade, but instead only captures a share of the trade surplus.

Who wins and who loses? To assess who wins and loses, we compute by how much the expected investor surplus and dealer profit change when going from the status quo to the counterfactual market rules.

Definition 1.

(i) The expected surplus of an investor i in group G who buys on t is

$$EU_t^G = \sum_j S_{tj}^G(q_t)\mathbb{E}[(q_{tj} + \sigma\boldsymbol{\epsilon_{tij}}) - (\theta_t + \boldsymbol{\nu_{ti}^G}) + \xi_{tj} - c_t | \textit{buy on the platform from dealer } j].$$

(ii) The expected profit of dealer j from selling to group G on t is

$$PI_{tj}^G = \mathbb{E}[v_{tj}^D - \boldsymbol{y_{tij}^G} | \textit{buy bilaterally}] + S_{tj}^G(q_t)\mathbb{E}[v_{tj}^D - (q_{tj} + \sigma\boldsymbol{\epsilon_{tij}}) | \textit{buy on the platform from dealer } j].$$

In both cases, Proposition 1 specifies where the investor buys and at what yield. $S_{tj}^G(q_t)$ *is the probability that an investor in* $G \in \{I, R\}$ *buys from dealer j on the platform. It is defined analogous to* $S_{tj}(q_t)$ *in Proposition 2 if the investor group has platform access, and is 0 otherwise.*

In theory, retail investors cannot lose surplus when obtaining platform access, but it is unclear whether institutional investors or dealers benefit from the change. This depends on how dealers adjust their quotes. If they set quotes that are more favorable to investors, dealers lose and institutional investors win.

Whether this is the case depends on how the two elasticities that govern the quotes change as the composition of investors on the platform changes. A lower elasticity of demand on the platform toughens platform competition and therefore leads to better quotes.

A lower cross-market elasticity has the opposite effect: The less easily investors switch onto the platform, the lower is the incentive for dealers to post unattractive quotes to prevent investors from using the platform. Which of the two effects dominates is an empirical question.

We find that the second effect slightly dominates, so that quotes become better for investors as platform access becomes universal. When allowing retail investors to enter the platform at estimated costs, the quotes change very little. The reason is that quotes are more strongly targeted to institutional investors who make up 90% of the market, and institutional investors already have platform access in the status quo. However, as we eliminate platform usage costs, the average quote increases by about 0.05 bps, which is roughly $1/10^{th}$ of the median bid-ask spread.

As a result, free platform access brings higher gains to investors and larger losses to dealers than costly platform access (see Figure 6): Retail investors gain about 4 bps, institutional investors about 1 bps, and dealers lose about 1 bps (per unit), when access if free.

For an average retail investor, who trades about C$86 million (units) per year, this is equivalent to earning about 4bps*C$86 = C$34 thousand more interest in a year. For an average institutional investor the monetary gain is larger because institutional investors trade larger amounts and more often in a year than retail investors: C$1.7 million. This gain is significant. It would double the revenue that an average institutional investor makes from interest on any type of investment in a year. Dealers, who trade the most, expect a monetary loss of about C$27 million per year. This is sizable reduction in her yearly revenue from interest on any type of investment of 15% (see Appendix A.3 for details).

Welfare. We measure welfare as the expected gains from trade between a dealer and an investor, excluding the platform usage cost.

Definition 2. *The expected welfare is* $W_t = \sum_G \kappa_G W_t^G$ *where*

$$W_t^G = \sum_j \mathbb{E}[v_{tj}^D - v_{tj}^G(\boldsymbol{\nu_{ti}^G})|\textit{investor } i \in g \textit{ buys from dealer } j] \tag{10}$$

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

Whether welfare increases as more investors enter the platform depends on who matches with whom. To see this, we compute the change in welfare when going from the status quo to the counterfactual world:

$$\Delta W_t = \sum_G \kappa_G \sum_j \Delta\gamma_{ij}^G * (v_{tj}^D + \xi_{tj}). \tag{11}$$

Here $\Delta\gamma_{tj}^G$ abbreviates the change in the probability that an investor in group G buys from dealer j on day t. Welfare increases, $\Delta W_t > 0$, as investors become more likely to buy from "more efficient" dealers, i.e., dealers with higher $(v_{tj}^D + \xi_{tj})$.[33]

We find that welfare increases 9% when platform access becomes universal and by about 30% when access is free (see Figure 7a). This translates into a sizable monetary gain of C$123-411 million per year, or roughly 0.008-0.026% of GDP. The reason is that investors are more likely to match with more efficient dealers. For instance, a retail investor is 18% more likely to buy from the most efficient dealer (see Appendix Figure A7a).

To better understand where the welfare gain comes from, we decompose it into how much value is generated because dealers with higher values, v_{tj}^D, versus dealers with better quality, ξ_{tj}, are more likely to sell. We find that almost the entire welfare gain comes from matches to dealers with higher values (see Figure 7c).

This finding suggests that dealers cannot freely sell and buy as much they please. 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. Such constraints triggered dramatic events in the U.S. market for government bonds during the recent COVID-19 crisis: When dealers failed to absorb enough bonds onto their balance sheets to meet the extraordinary supply of investors, the Federal Reserve System (FED) purchased $1 trillion of U.S. government bonds in March 2020 to rescue the market (Duffie (2020)). In addition, the FED temporarily relaxed balance sheet constraints for government debt.

[33] In theory, it is ambiguous whether welfare increases as more investors trade on the platform. On the platform, each investor i is free to choose among all of the dealers. She picks the dealer with the highest $(q_{tj} + \sigma\epsilon_{tij}) + \xi_{tj}$. This dealer must not be more efficient than the dealer who was chosen in the status quo because the platform is not perfectly competitive. Dealers sell at quotes that differ from their true values v_{tj}^D. They strategically set these quotes in response to investor behavior, and dealers with higher v_{tj}^D must not necessarily post higher quotes. Further, this negative effect is amplified when too many investors use the platform relative to what would be optimal if the platform was frictionless (see Appendix Figure A7b).

How high are potential welfare gains? In this paper, we focus on market reforms that shift bilateral trading onto platforms on which investors run (RFQ) auctions with dealers. We view such a shift as a feasible first step into the right direction, but there are other reforms that could affect trading. To determine the size of any potential gains, we quantify how efficient the market is today relative to the first-best.

In the first best, a single dealer—the one with the highest $v_{tj}^D + \xi_{tj}$—sells to all investors on day t. This is because, in our model, over the corse of a day dealers have no capacity constraints and their values are constant. In reality, these assumptions might fail, for instance, because a single dealer cannot absorb the daily excess supply of all investors onto her own balance sheet. This is especially critical during times of financial distress. During our sample period, the market conditions were stable with no extraordinary imbalances. We therefore think that the assumptions are reasonable.

We compare four different market settings to the first best: the status quo, the two counterfactuals in which all investors have access to the platform, and an additional counterfactual, where we remove the dealers and let investors directly trade with one another.

This last counterfactual approximates an environment in which trades between investors (but not dealers) realize via an efficient market mechanism, such as an efficient batch auction as suggested by Budish et al. (2015): Each day, the market clears at the price that equates expected investor demand with supply. All investors who seek to buy (sell) are willing to pay a price that lies above (below) the clearing price buy (sell). Importantly, we assume that dealers no longer smooth out investor demand and supply over time. One reason could be that they no longer earn sufficient profits when the market clears via an efficient mechanism that eliminates all markups on trades.

Our findings are shown in Figure 8: The status quo achieves roughly 60% efficiency, which suggests that there are potentially large welfare gains from market reforms. Our first counterfactual–allowing all investors platform access at the estimated costs–does very little. The second counterfactual, where we eliminate all costs leafs to a large increase in welfare and we achieve 80% efficiency. Finally, letting investors trade directly with one another would lead to lower welfare than the status quo. This is because dealers no longer absorb excess supply or demand of investors on days on which demand and supply does not balanced perfectly. Crucially, this is not an endogenous outcome of running an efficient mechanism but an assumption. Therefore, this finding highlights how important dealers are in providing liquidity, and should not be taken as an argument against efficient mechanisms.

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. In our analysis, these gains come entirely form reallocating who trades with whom, because market participants have heterogenous valuations for realizing trade. This is surprising as government bonds are liquid and safe assets whose market value is publicly known.

In addition, our results highlight that it is easier to form more efficient matches between counterparties when the market is centralized than when it is decentralized. This is a common advantage of (two-sided) centralized markets, which is not specific to over-the-counter markets and likely true in many other contexts.

8.3 Robustness

We conduct several tests to verify the robustness of our findings in Appendix F. 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 that a dealer won as instrument. We also verify that measurement errors in the dealers' qualities (ξ_{tj}) do not significantly bias the distribution of the liquidity shocks. In addition, we allow for dealer specific platform usage costs (c_{tj}), and restrict the sample to exclude occasionally large trades.

Second, we verify that our estimates and welfare findings are robust when we allow the investor to capture some trade surplus in the bilateral trade. For this, we rely on the extended model in Appendix C.3, in which the investors captures a trade surplus of ϕ in a bilateral trade. While we cannot identify this parameter with our data, we can test how the model estimates and counterfactual findings change as we increase ϕ from zero (as in the benchmark model) to positive values.

Taken together, our robustness tests confirm our expectations and suggest that our main findings are qualitatively robust.

9 Conclusion

In this paper, we use trade-level data of the Canadian government bond market to study whether to centralize OTC markets by shifting bilateral trades onto centralized platforms. We show that platform access can lead to better prices for investors. This brings large gains to investors and losses to dealers. Overall, welfare increases because high-valued buyers are more easily matched with low-valued sellers. This highlights an advantage of (two-sided) centralized platforms beyond our specific setting.

Based on our findings, policy reforms should shift more trades onto centralized platforms. To achieve this goal, it is not sufficient to grant platform access to more investors or to open up platforms in markets in which there are none. Instead, one would have to find ways to reduce costs that prevent investors from using the platform. Future research with data to analyze what drives these costs could helps us understand what type of policy reform works best.

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. Bank of Canada working paper 2020-29.

Allen, J. and Thompson, J. R. (2019). Variable pay: Is it for the worker or the firm? *Journal of Corporate Finance*, 58:551–566. Bank of Canada Staff Working Paper.

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

Baldauf, M. and Mollner, J. (2019). Trading in fragmented markets. *Journal of Financial and Quantitative Analysis*, pages 1–29.

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.

Bech, M., Illes, A., Lewrick, U., and Schrimpf, A. (2016). Hanging up the phone - electronic trading in fixed income markets and its implications. *BIS Quarterly Review*, pages 79–94.

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

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

Bernhardt, D., Dvoracek, V., Hughson, E., and Werner, I. M. (2007). Why do larger orders receive discounts on the london stock exchange? *The Review of Financial Studies*, 18(4):1343–1368.

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

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 25(2):242–262.

Bessembinder, H., Maxwell, W., and Venkataraman, K. (2006). Market transparency, liquidity externalities, and institutional trading costs in corporate bonds. *Journal of Financial Economics*, 82(2):251–288.

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.

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.

Bresnahan, T. F. (1987). Competition and collusion in the American automobile industry: The 1955 price war. *The Journal of Industrial Economics*, 35(4):457–482.

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

Budish, E., Lee, R. S., and Shim, J. J. (2019). Will the market fix the market? A theory of stock exchange competition and innovation. NBER Working Paper No. 25855.

Carvalho, A. (2020). Corporate debt standardization and the rise of electronic bond trading. Working paper.

Chen, D. and Duffie, D. (2020). Market fragmentation. NBER working paper No. 26828.

Chowdhry, B. and Nanda, V. (1991). Multimarket trading and market liquidity. *The Review of Financial Studies*, 4:483–511.

Colliard, J.-E., Foucault, T., and Hoffmann, P. (2020). Inventory management, dealers' connections, and prices in OTC markets. *Journal of Finance*, forthcoming. HEC Paris Research Papers No. FIV-2018-1286.

Collin-Dufresne, P., Junge, B., and Trolle, A. B. (2020). Market structure and transaction costs of index CDSs. *Journal of Finance*, 75(5).

Di Maggio, M., Egan, M. L., and Franzoni, F. (2019). The value of intermediation in the stock market. NBER working paper 26147.

Duffie, D. (2020). Still the world's safe haven? Redesigning the U.S. Treasury market after the COVID-19 crisis. Hutchins Center Working Paper #62.

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. NBER working paper 25887.

Edwards, A. K., Harris, L. E., and Piwowar, M. S. (2007). Corporate bond market transaction costs and transparency. *Journal of Finance*, 62(3):1421–1451.

Fama, E. F. (1975). Short-term interest rates as predictors of inflaction. *The American Economic Review*, 65(3):268–282.

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

Fleming, M. J. (2003). Measuring Treasury market liquidity. *Economic Policy Review*, 9(3):83–108.

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.

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. (2016). Asymmetric information and intermediation chains. *American Economic Review*, 106(9):2699–2721.

Glosten, L. (1994). Is the electronic open-limit order book inevitable? *Journal of Finance*, 49:1127–1161.

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

Green, R. C., Hollifield, B., and Schürhoff, N. (2007b). Financial intermediation and the costs of trading in an opaque market. *The Review of Financial Studies*, 20(2).

Hamburger, M. J. and Platt, E. N. (1975). The expectations hypothesis and the efficiency of the treasury bill market. *The Review of Economics and Statistics*, 57(2):190–199.

Hamilton, J. L. (1979). Marketplace fragmentation, competition, and the efficiency of the stock exchange. *Journal of Finance*, 34:171–87.

Hangströmer, B. and Menkveld, A. J. (2019). Information revelation in decentralized markets. *The Journal of Finance*, 74(6):2751–2787.

Harris, L. E. and Piwowar, M. S. (2006). Secondary trading costs in the municipal bond market. *Journal of Finance*, 61(3):1361–1397.

Hau, H., Hoffmann, P., Langfield, S., and Timmer, Y. (2019). Discriminatory pricing of over-the-counter derivatives. IMF Working paper.

Hendershott, T., Jones, C. M., and Menkveld, A. J. (2011). Does algorithmic trading improve liquidity? *Journal of Finance*, 66(1):1–33.

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

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.

Hendershott, T. and Mendelson, H. (2000). Crossing networks and dealer markets: Competition and performance. *Journal of Finance*, 55:2071–2115.

Hennig, T. (2020). Estimating a model of decentralized trade with asymmetric information. Job market paper.

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

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

Hortaçsu, A. 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.

Hugonnier, J., Lester, B., and Weill, P.-O. (2018). Frictional intermediation in over-the-counter markets. NBER Working Paper No. 24956.

Iercosan, D. and Jiron, A. (2017). The value of trading relationships and networks in the CDS market. Working paper.

IIAC (2020). Canadian conventions in fixed income markets. Accessible at: `http://iiac.ca/wp-content/uploads/Canadian-Conventions-in-Fixed-Income-Markets-1.1.pdf`.

IIROC Rule Book (2020). Rule 1. Accessible at: `https://www.iiroc.ca/industry/rulebook/Pages/default.aspx`.

Jankowitsch, R., Nashikkar, A., and Subrahmanyam, M. G. (2011). Price dispersion in OTC markets: A new measure of liquidity. *Journal of Banking and Finance*, 35:343–357.

Kakhbod, A. and Song, F. (2020). Liquidity,welfare and transparency in dynamic otc markets. Working paper.

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.

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.

Lagos, R. and Rocheteau, G. (2009). Liquidity in asset markets with search frictions. *Econometrica*, 77(2):403–426.

Lagos, R., Rocheteau, G., and Weill, P.-O. (2011). Crises and liquidity in over-the-counter markets. *Journal of Economic Theory*, 146:2169–2205.

Lee, T. and Wang, C. (2019). Why trade over the counter? when investors want price discrimination. Working paper.

Li, D. and Schürhoff, N. (2019). Dealer networks. *Journal of Finance*, 74(1):91–144.

Li, W. and Song, Z. (2019). Dealers as information intermediaries in over-the-counter markets. 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. Accessible at: `https://www.newyorkfed.org/newsevents/speeches/2020/log201023`.

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, accessible at: `https://www.greenwich.com/fixed-income-fx-cmds/corporate-bond-electronic-trading-continues-growth-trend`.

Mendelson, H. (1987). Consolidation, fragmentation, and market performance. *Journal of Financial and Quantitative Analysis*, 22:189–207.

Miao, J. (2006). A search model of centralized and decentralized trade. *Review of Economic Dynamics*, pages 68–92.

MTRS 2.0 User Guide (2016). Debt securities transaction reporting. Accessible at: `https://www.iiroc.ca/industry/marketmonitoringanalysis/Documents/DebtSecuritiesTransactionReportingMTRS2_UserGuide_en.pdf`.

O'Hara, M. and Ye, M. (2011). Is market fragmentation harming market quality? *Journal of Financial Economics*, 100:459–474.

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

Pesando, J. E. (1978). On the efficiency of the bond market: Some Canadian evidence. *Journal of Political Economy*, 86(6):1057–1076.

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. Working paper.

Roll, R. (1970). *The Behavior of Interest Rates: The Application of the Efficient Market Model to U.S. Treasury Bills*. New York: Basic Books.

Rostek, M. J. and Yoon, J. H. (2020). Exchange design and efficiency. Working paper.

Rust, J. and Hall, G. (2003). Middlemen versus market makers: A theory of competitive exchange. *Journal of Political Economy*, 111(2):353–403.

Schlam Stone and Dolan LLP (2018). The manipulation monitor: A guide to financial market manipulation antitrust litigation. Accessible at: https://www.schlamstone.com/welcome-to-schlam-stone-and-dolan-llps-newest-blog-the-manipulation-monitor-a-guide-to-financial-market-manipulation-antitrust-litigation/.

Stoll, H. R. (2001). Market fragmentation. *Financial Analysts Journal*, 57:16–20.

Train, K. (2002). *Discrete Choice Methods with Simulation*. Cambridge University Press.

U.S. Treasury report (2017). A financial system that creates economic opportunities, capital markets. Accessible at: `https://www.treasury.gov/press-center/press-releases/Documents/A-Financial-System-Capital-Markets-FINAL-FINAL.pdf`.

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

Weill, P.-O. (2007). Leaning against the wind. *Review of Economic Studies*, 74:1329–1354.

Wittwer, M. (2020a). Connecting disconnected financial markets? *American Economic Journal: Microeconomics*, forthcoming.

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

Yoon, J. H. (2018). Endogeneous market structure: Over-the-counter versus exchange trading. Working paper.

Zhu, H. (2012). Finding a good price in opaque over-the-counter markets. *Review of Financial Studies*, 25(4):1255–1285.

Zhu, H. (2014). Do dark pools harm price discovery. *The Review of Financial Studies*, 27(3):747–789.

Figure 1: Yield gap

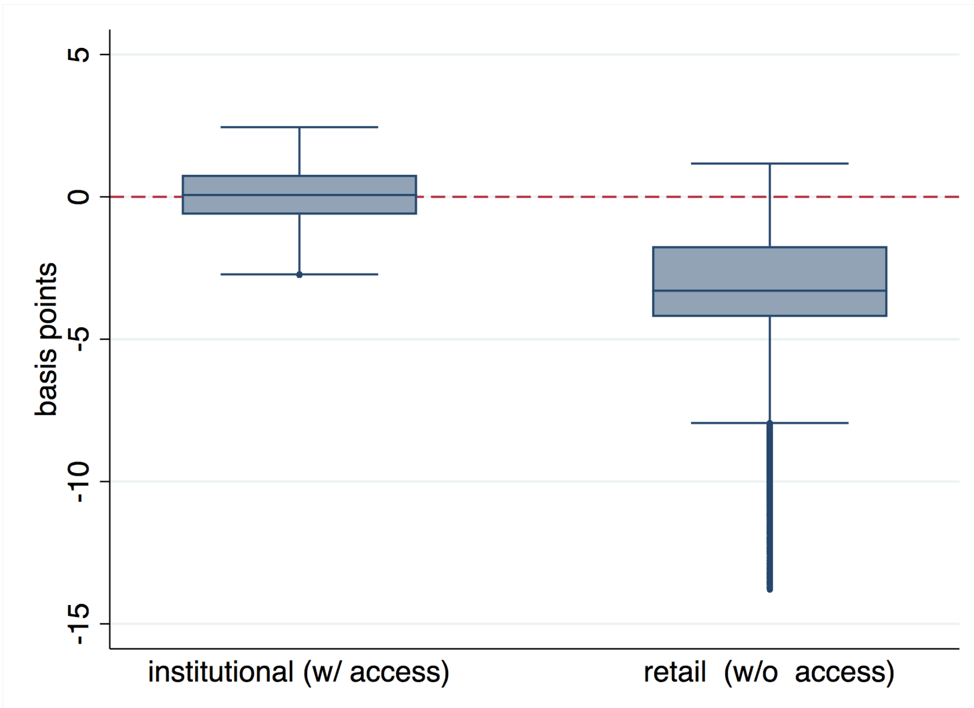

The graph shows a box plot of markups for institutional and retail investors, excluding the upper and lower 5% of the distribution. To construct these markups, we regress $(y_{thsij} - \theta_{ths})^{+}$ as defined in (1) on indicator variables that distinguish retail from institutional investors ($retail_{thsij}$)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

Figure 2: Event study: Yield drop when losing platform access

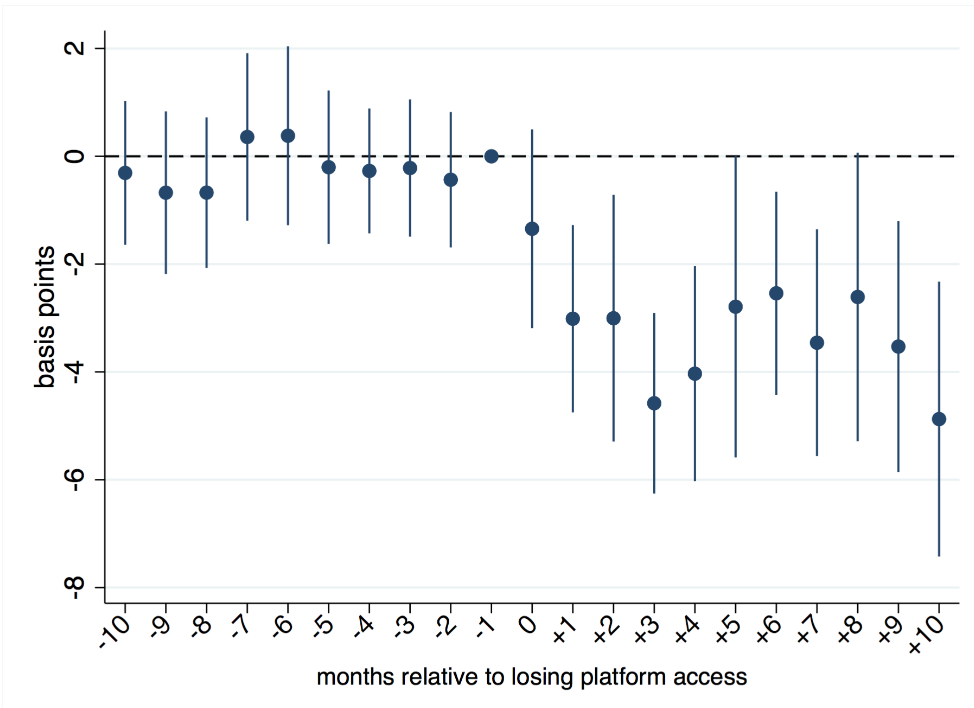

The graph shows the β_m estimates and the 95% confident intervals of the event study regression (2), $(y_{thsij} - \theta_{ths})^+ = \zeta_i + \sum_{m=Mi^-}^{Mi^+} \beta_m D_{mi} + \zeta_{th} + \zeta_s + \zeta_j + \epsilon_{thsij}$, for 10 months before and after an investor i switched from having access 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 one month prior to it. Standard errors are clustered at the investor-level.

Figure 3: Identifying variation for σ and ξ_{tj}

(a) Dealers' daily platform market shares and their cost shifters

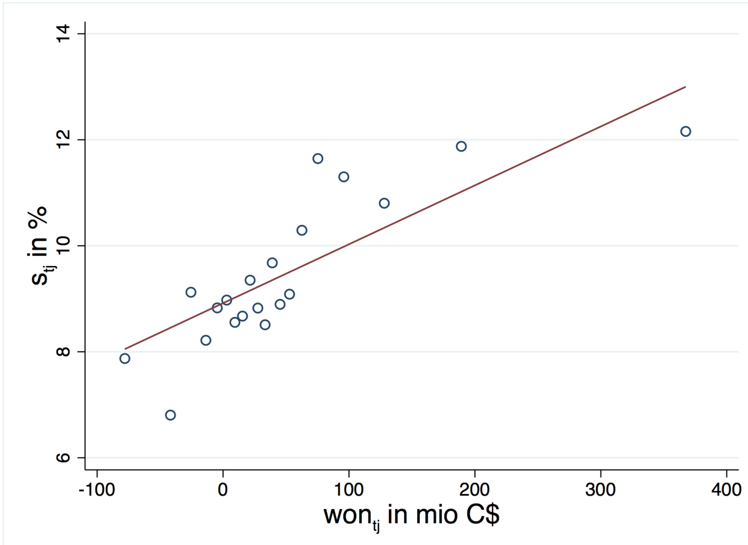

The graph shows a binned scatter plot to visualize the correlation between the the dealers' daily market shares on the platform (s_{tj}) and the amount that each dealer unexpectedly won in the most recent primary auction ($won_{\tilde{t}j}$) when partialling out day fixed effects. The steepness of the slope identifies the degree of competition on the platform. The steeper it is, the more competitive is the platform.

(b) Dealers' daily platform market shares when posting the best quote

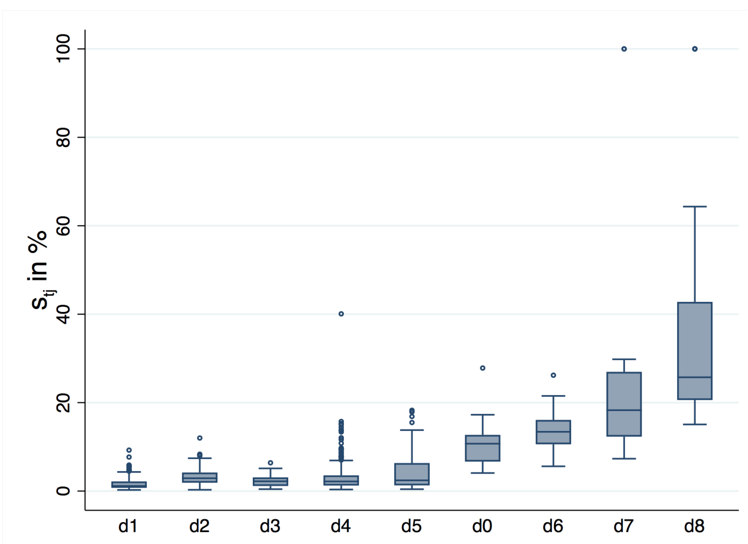

The graph displays a box plot for each dealer. Each shows the distribution of how much of the total platform market this dealer captures on days on which she posts the best quote relative to the other dealers. The market share is measured in percent. Dealers are labeled from d1 to d8. d0 is benchmark dealer whose quality is normalized to 0.

Figure 4: Identifying variation for c_t, μ_t^I, σ_t^I

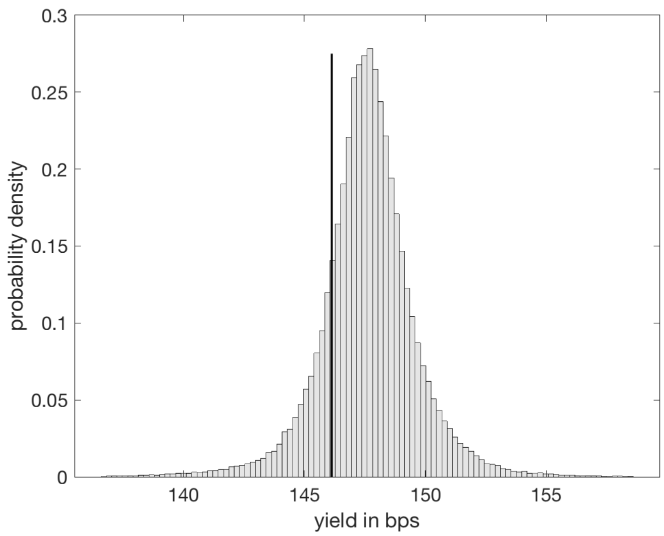

This graph illustrates which moments in the data pin down the platform usage cost (c_t) and the distribution of the investors' liquidity shocks (μ_t^I, σ_t^I). It shows a probability density histogram of the yields (in bps) that institutional buyers realize, excluding the upper and lower 0.1 percentile of the distribution, in addition to a black line. This line is the average cutoff that determines whether an institutional investor buys bilaterally or on the platform according to Proposition 1, i.e. $mean_{tj}(\hat{\psi}(q_t) + \theta_t - \hat{\xi}_{tj})$. The shape of the distribution above the line identifies the distribution of the investor's liquidity shocks, conditional on trading bilaterally. The position of the line, and with it the platform usage costs, is determined by the share of institutional buyers who trade bilaterally versus on the platform.

Figure 5: Dealer's quality and values

(a) Box plots of $\hat{\xi}_{tj}$ for each j

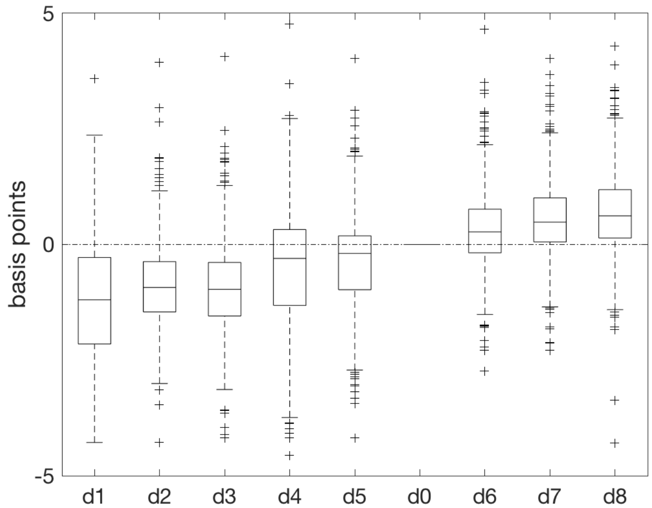

(b) Box plots of $(\hat{v}alue^{D}_{tj} - \theta_t)$ for each j

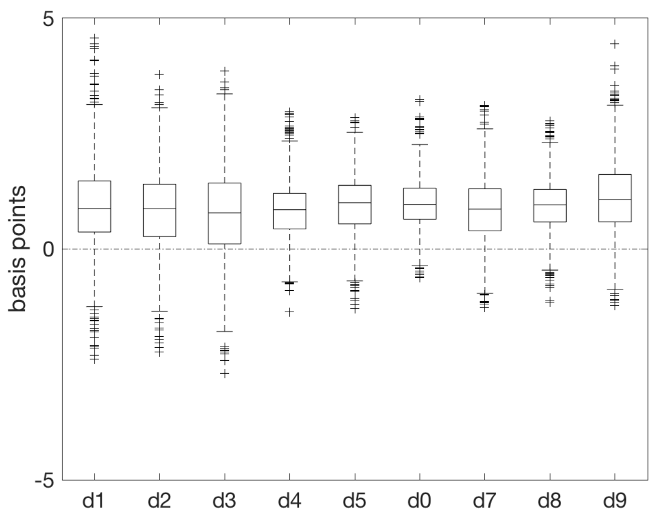

Graph 5a shows a box plot of the estimated qualities $(\hat{\xi}_{tj})$ for each dealer j, labeled d0 up to d8. The size of each $\hat{\xi}_{tj}$ is relative to the quality of the benchmark dealer d0. Graph 5b displays box plots of the estimated daily values that a dealer has when selling to an investor $(\hat{v}alue^{D}_{tj})$ net of the bonds market value (θ_t). All box plots exclude the upper and lower 1% of the respective distribution. All variables are measured in bps.

Figure 6: Who wins how much when platform access is universal?

(a) Universal access: Benchmark model

(b) Universal access: Extended model

(c) Free universal access: Benchmark model

(d) Free universal access: Extended model

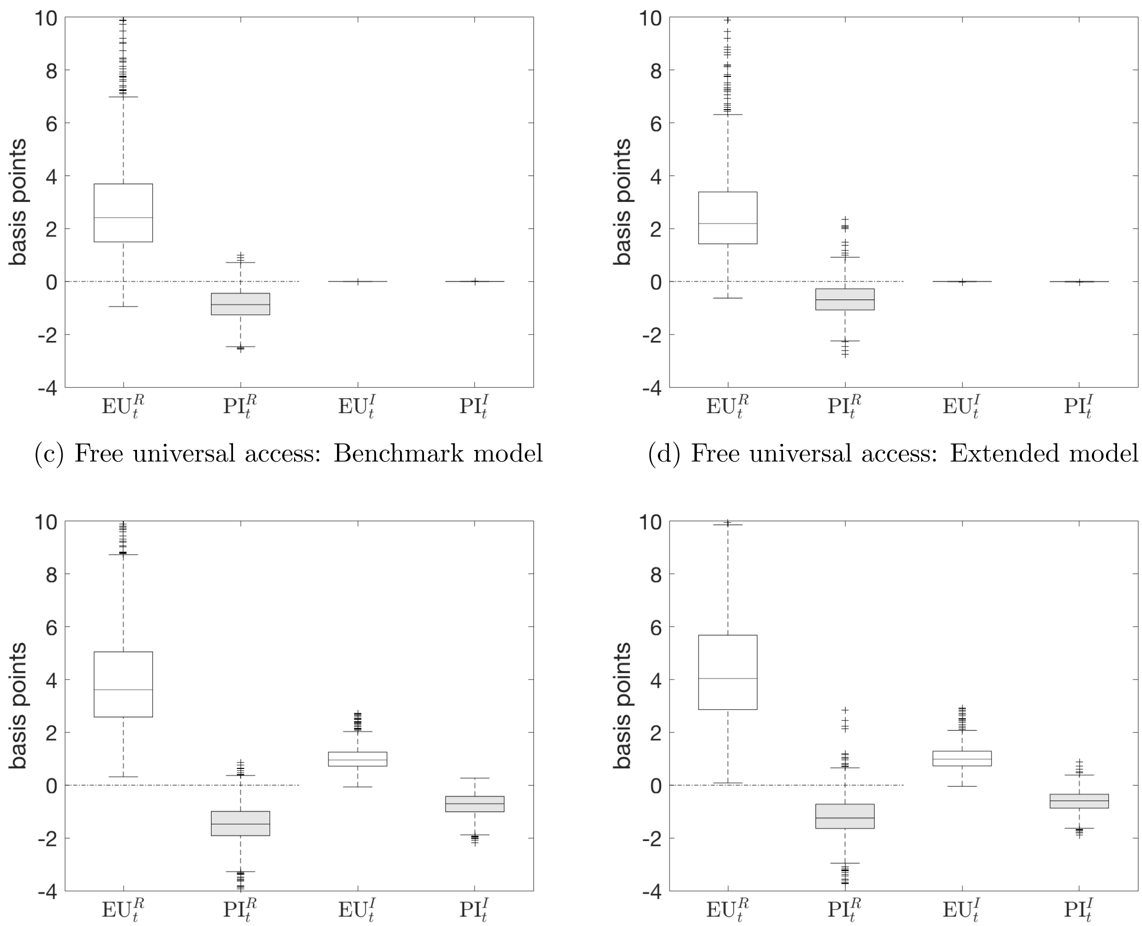

The graphs on the RHS show box plots of the change in investor surplus (EU_t^G) and dealer profit (PI_t^G) of Definition 1 for the two investor groups (G), retail investors R and institutional investors I, when allowing retail investors to access the platform at estimated costs, and when eliminating all usage costs. The graphs on the LHS, show the analogous box plots but for the extended model model of Appendix C.3 in which we allow the investor to extract 10% of the trade surplus in a bilateral trade. The change is measured in bps. The upper and lower 1% of the respective distribution is excluded in all box plots.

Figure 7: Welfare change and decomposition

(a) Increase in welfare: Benchmark model

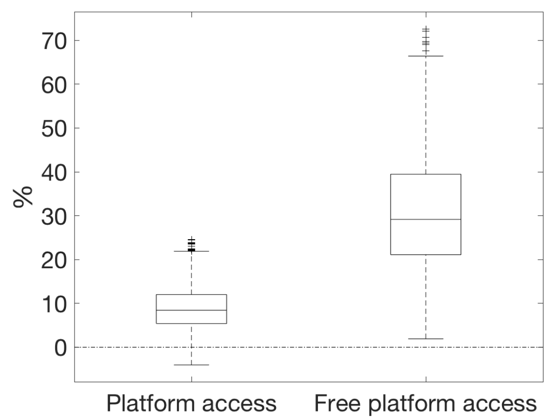

(b) Increase in welfare: Extended model

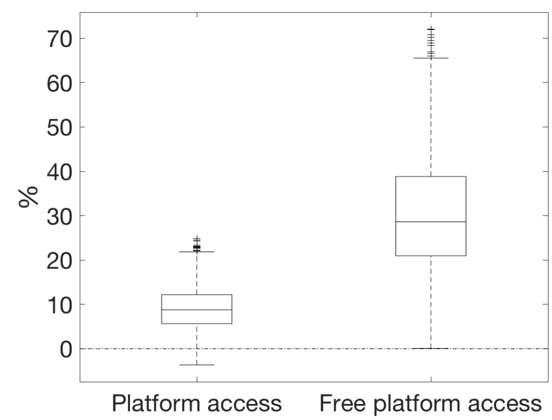

(c) Welfare decomposition: Benchmark model

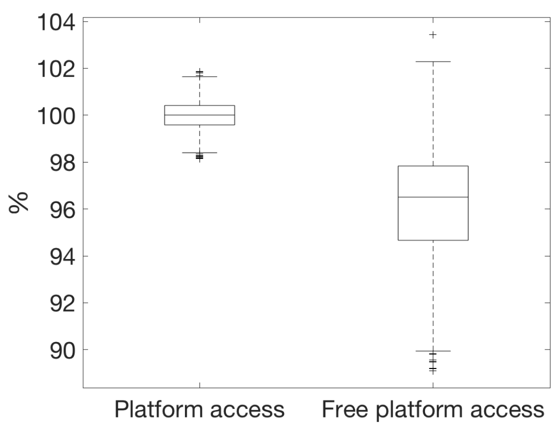

(d) Welfare decomposition: Extended model

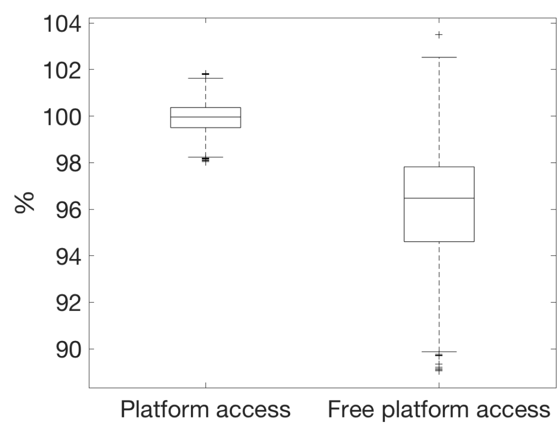

Graph 7a illustrates by how much welfare increases when making platform access universal or free. In both cases, it shows the distribution of the percentage change in welfare, $\Delta W_t / W_t * 100\%$, over days with W_t as in (10) and ΔW_t as in (11). Graph 7c shows how much of the welfare gain is due to more matches to dealers with higher v_{tj}^D rather than dealers with higher quality ξ_{tj} as percent of the total change in welfare. It plots the distribution of $\sum_G \kappa_G \sum_j \Delta\gamma_{ij}^G v_{tj}^D / \Delta W_t * 100\%$ over days for both counterfactuals. Graphs 7b and 7d show the analogous box plots of the extended model of Appendix C.3 in which the investor extracts 10% of the trade surplus in a bilateral trade. In all box plots the upper and lower 1% of the respective distribution is excluded.

Figure 8: How efficient is the market?

(a) Benchmark model

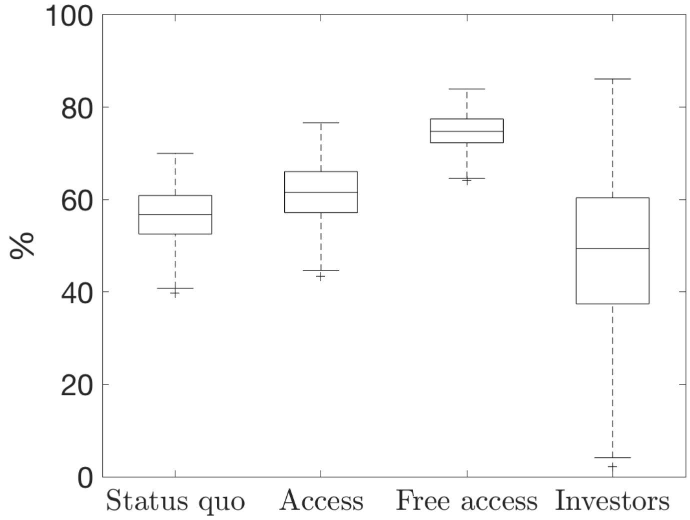

(b) Extended model

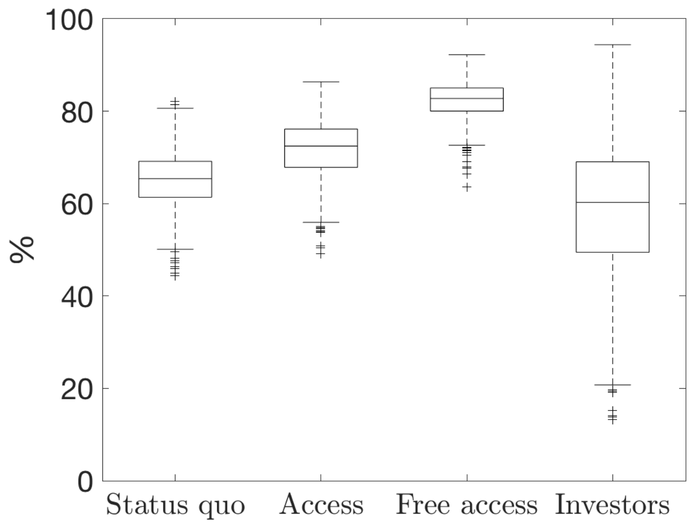

These graphs show the total expected gains from trade (welfare) as percent of what could be achieved in the first best (per day) in four different market settings. Graph 8a displays the results of the benchmark model in which investors have no bargaining power ($\phi = 0$), and Graph 8b of the extended model of Appendix C.3 when investors and dealers have the same bargaining power ($\phi = 0.5$). Each box plots shows the distribution over days for one of the four settings: the status quo, the counterfactuals in which all investors have platform access at the estimated platform usage costs and for free, as well as the counterfactual in which investors trade with one another. The upper and lower 1% of the respective distribution is excluded in all box plots.

Table 1: Estimates (median across days)

buys	$\hat{\mu}^I$	$\hat{\mu}^R$	$\hat{\sigma}^I$	$\hat{\sigma}^R$	$\hat{c}$	$1/\hat{\sigma}$	$\hat{\eta}$
	-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	$\hat{\mu}^{I^*}$	$\hat{\mu}^{S^*}$	$\hat{\sigma}^{I^*}$	$\hat{\sigma}^{S^*}$	$\hat{c}_t^*$	$1/\hat{\sigma}$	$\hat{\eta}^*$
	+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)	

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

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

This table 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 (see Figure 2) but collapsing 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 yield that an investors expects (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.

Appendix

A Additional information on data

A.1 Data cleaning and sample restrictions

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

We clean out trades that exhibit yields that are extreme relative to the public Bloomberg price on day t in the same hour h of the same security s (see Appendix Figure A8). Specifically, we drop the upper and lower 1% of the distribution of the markup $(y_{thsij} - \theta_{ths})^+$ which is defined in (1) for each investor group. These outliers might be reporting errors as it is difficult to rationalize why anyone would be willing to pay a markup that is higher than a couple of percents of the bond's value. They could also reflect an extreme urgency to trade of an institution that faces an internal problem. Alternatively, prices could be extreme because the trade is part of a larger investment package, which we do not observe.

We correct 100 cases in which subsidiaries of reporting dealers or brokers are labeled as retail investors, and drop 20 observations that were reported without retail/institutional indicator. Among investors who switch from retail to institutional or vice versa, we drop investors who do not permanently switch. This excludes trades with investors that are in a gray area. For example, CIBC Investor Services Inc. (a subsidiary of CIBC) classifies as retail investor according to the rules, even though CIBC (the Canadian Imperial Bank of Commerce) is one of 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.

To avoid double counting inter-dealer trades (which should be reported on both sides of the trade) we eliminate trades that are reported twice. We omit a more detailed description since this paper focusses on dealer-to-investor trades.

We focus on CanDeal or bilateral trades only, which means that we ignore 0.41% of the observations because the reported trading venues are incorrect. In these rare cases, the dealer makes a mistake and reports something other than these the existing trading venues, often the ID of her counterparty.

Finally, when a Bloomberg quote for a security is missing in an hour of the day (6.69% of observations), we use the daily average Bloomberg quote for this security.

A.2 Classification of investors

Investors are classified at the subsidiary-level as follows:

Dealer	Primary dealer, e.g. BMO, CIBC, TD, Caisse Desjardins, HSBC, etc. Government Distributor IIROC dealer member who is required to report trades acc. to MTRS guidelines[34]
Bank	Financial institution that accepts retail deposits Non retail deposit-taking financial institution Credit union Credit union central (e.g., Central 1 is a pooling facility for Ontario credit unions) Trust company as listed as on OSFI website[35] Firms that originate mortgages and finance on wholesale or securitization market Custodian bank
Other IRROC member	Dealer member of IIROC (according to the IIROC-counterparty type variable) who is not a primary dealer. These are firms that are registered with IIROC but are not a reporting dealer or broker. Many of them act as asset managers. Two examples are Dubeau Capital & Cie (a wealth manager) and Liquidnet (which is an investment network that connects asset managers). These institutions are not to be confused with what we call "dealers" in this paper.[36]
IDB	The three inter-dealer-brokers (IDB) in Canada (Freedom, Shorcan, Tullet Prebon)
Other broker	Any firm that describes itself or is described as a "brokerage firm" (financial institutions that acts as middle man between buyers and sellers)
Asset manager	Asset/investment/wealth management company Subsidiary of dealers that focus on asset management (e.g., CIBC or BMO Asset Management, CIBC Investor Services) Investment management fund including hedge funds Investment fund (e.g., mutual funds, bond funds, ETFs)
Insurance	Insurance company Reinsurance company (provides reinsurance to other insurance companies)
Public entity	Crown corporation (owned by the government, like CMHC), Municipality, Education institution, Foreign government agency, Province, Central bank, IIROC
Non-financial company	Biotech company, Energy company, Food processing company, Manufacturing company, Mining company, Retail company, Technology company, Utility company, Transportation, Real estate company, Charity
Other	Transfer agency (keeps records of who owns a firm's assets and how) All-to-all ATS, dealer-to-dealer platform Exchange (IMX, CDCC, or ICE), CDS clearing or central counterparty

[34] View the MTRS 2.0 User Guide (2016)

[35] See `http://www.osfi-bsif.gc.ca/eng/wt-ow/Pages/wwr-er.aspx?sc=1&gc=2`)

[36] For more information, see the list of IIROC regulated institutions (which are referred to as "dealer" by IIROC) under `https://www.iiroc.ca/industry/Pages/Dealers-We-Regulate.aspx`.

A.3 Yearly revenue from interest

To get a sense of the magnitudes of changes in the investor surplus and dealer profit, we look at the latest available data on regulatory financial reports (Form 1 of the IIROC Rule Book).[37] These data used to be collected by the IIROC and include balance sheet information of dealers, institutional and retail investors. For us relevant, is the yearly (net) interest revenue that an institution earned from any type of security. A detailed description of the entire data set can be found in Allen and Thompson (2019).

Unfortunately, the data collection stopped in the year 2010. Therefore, our calculations should be taken with a grain of salt, as they rely on the assumption that the average interest rate revenue of an institution in 2016-2019 is similar to the one in 2010. The fact that we do not observe any drastic changes in the previous years, gives us confidence that this is not too far fetched.

The average dealer made about C$178 million in interest rate revenue of any type of asset, including government bonds in 2010. Therefore, a loss of C$27 million would imply a reduction in revenue of roughly 15%. An average institutional investor earned about C$1.6 million, and an average retail investor about C$3 million.

B Additional reduced-form evidence

B.1 Usage costs and platform competition

In this section, we provide reduced-form evidence that motivate two elements of the model, the competition parameter (σ) and the cost to use the platform (c_t).

We show that yields are better on than off the platform (see Appendix Table 5). This could be due to higher competition on the platform, or due to selection onto the platform whenever yields there are better. The magnitude of this effect drops when including an investor fixed effect. This is reassuring. We would expect that investors who trade more than once would not willingly leave large rents to their dealers in bilateral trades but instead select onto the platform whenever it is opportune. Yet, even these investors are willing to accept worse prices when trading with their dealer one-to-one than on the platform.

Further, we document, that despite better platform yields, most investors trade bilaterally (see Appendix Figure A9). On a typical day, only 35-40% of the institutional investors realize trades on the platform. This suggest that something prevents them from using the platform, and further motivates the platform usage cost.

B.2 Trade size

In our model, we abstract from trade size. Here we provide evidence to justify this assumption and discuss its implications.

Distribution of trade sizes. The vast majority of trades is small in our sample, and trades are very similar in size, both of bilateral trades and on the platform (see Appendix Figure A10a). In comparison, dealers and brokers trade much larger amounts (see Appendix Figure A10c).

[37] Available at `https://www.iiroc.ca/Rulebook/MemberRules/Form1_en.pdf`

This motivates our assumption to normalize trade size to one. It is also the reason for which our model abstracts from the price impact that an order might have on the market price, or incentives to split the order into smaller trades or across trading venues.

Further, the fact that for investors no trade is larger than C\$25 million is important because the quotes that dealers post are not valid for larger amounts. If an investor seeks to trade a larger amount, he has to contact a dealer who tells him quotes for these large amounts depending on current market conditions.

Yields and trade size. One might be concerned the yield that an investor realizes mostly depends on the amount he seeks to trade. In that case, our model would miss a key driver of what determines the trade yield. To reduce this concern, we show that trade size is not highly correlated with the yield that an investor realizes, once we control for time trends and other factors that drive the trade yield. The estimations findings in Appendix Table 8 confirm that the effect of trade-size is not statistically significant in explaining the yield.

Venue choice and trade size. Lastly, we analyze whether investors trade different amounts on versus off the platform (see Appendix Table 9). We find that investors trade slightly smaller amounts on the platform. This relationship is mostly driven by outliers (see Appendix Figure A10). When excluding these outliers from the sample, the correlation between the trade size and the trading venue disappears.

Our interpretation is that investors occasionally trade large amounts, for instance, because they urgently seek liquidity. In these rare cases, they 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. It could also be that the investor is more reluctant to share information with more than her dealer if she seeks to trade a larger amount than usual.

In future research, we plan to investigate this hypothesis further. In this paper, we focus on the regular small trades. In the model, any reason for which an investor trades bilaterally with a dealer is absorbed by the the cost to use the platform.

B.3 Bilateral price equals willingness to pay

Here we gather evidence to support the assumption that the dealer sets a bilateral price equal to the investor's full willingness to pay, which splits into the market value of the bond and and an investor-specific liquidity need. If this is so, the bilateral yield would be independent of the dealer's costs. Otherwise, if the investor had some bargaining power, the bilateral yield would depend on how costly it is for the dealer to realize the trade.

We test this by regressing the bilateral yield, y_{thsij}, on the market value, θ_{ths}, the cost-shifter, $won_{\tilde{t}j}$, a dealer, day and investor fixed effect:

$$y_{thsij} = \alpha + \beta\theta_{ths} + \gamma won_{\tilde{t}j} + \zeta_j + \zeta_i + \zeta_t + \epsilon_{thsij}. \tag{12}$$

As a robustness check, we run a second specification in which we use the bilateral markup as defined in (1) as outcome variable:

$$(y_{thsij} - \theta_{ths})^+ = \alpha + \gamma won_{\tilde{t}j} + \zeta_j + \zeta_i + \zeta_t + \epsilon_{thsij}. \tag{13}$$

We find that both γ estimates are close to 0 and insignificant (see Appendix Table 6).[38] The findings of regression (12) with unclustered standard errors imply that when $won_{\tilde{t}j}$ increases by one standard deviation, the change in the yield lies in (-0.0170 bps, +0.0124 bps) for buying investors and in (-0.0072 bps, +0.0187 bps) for selling investors. When clustering standard errors, either at the dealer- or investor-level, or both, these intervals become even tighter.

As comparison, we run the same regression but with platform quotes rather than bilateral yields to see whether dealers adjust platform quotes when winning more at auction. Our findings in Appendix Table 7 confirm our expectations: When a dealer wins more, she increases the ask quote (because it is cheaper for her to sell) and increases the bid quote (because she already holds a lot of inventory). All point estimates are now significant and are larger than in Appendix Table 6.

Taken together, this is in line with our assumption that a dealer charges a price that equals to the investor's willingness to pay in bilateral trades, because it suggests that dealers do not (systematically) adjust the yield that they charge in bilateral trades when their cost shifts unexpectedly.

C Mathematical Appendix

C.1 Proof of Propositions 1 and 2

On any fixed day, the game has three periods. Its equilibrium conditions can be derived by backwards induction. Proposition 1 summarizes the investor's decision in the third and second stage of the game, and Proposition 2 characterizes the quotes that dealers post in the first stage. For notational convenience, we drop the t subscript, as well as superscript I (that stands for institutional buyer) throughout the proof.

Proposition 1. Statement (i) holds by assumption. To derive statement (ii) begin in the last stage: Conditional on entering the platform, investor i buys from dealer j if

$$u_{ij} > u_{ki} \forall k \neq j \text{ where } \tilde{u}_{tij}(\epsilon_{tij}) = \xi_{tj} + q_{tj} + \sigma\epsilon_{tij}$$
$$\Leftrightarrow \xi_j + q_j + \sigma\epsilon_{ji} > \max_k\{\xi_k + q_k + \sigma\epsilon_{ki}\}.$$

Given that ϵ_{ki} are iid Type 1 Extreme Value distributed, dealer j's market share on the platform (of investors who enter the platform) is

$$s_j(q) = \Pr(\boldsymbol{\epsilon_{ki}} < \boldsymbol{\epsilon_{ji}} + V_{ij} - V_{ki}\ \forall k \neq j) \text{ with } V_{ij} = \frac{1}{\sigma}(\xi_j + q_j) \tag{14}$$
$$\Leftrightarrow s_j(q) = \frac{\exp(\delta_j)}{\sum_k \exp(\delta_k)} \text{ with } \delta_j = \frac{1}{\sigma}(\xi_j + q_j). \tag{15}$$

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 what

[38] The direction of the point estimate suggests that the dealer charges a higher bilateral markup if she wins more in the auction. However, the confidence intervals are almost symmetrical around 0, so that we would refrain from taken the point estimates at face value.

is on the RHS of the following expression when entering the platform. She decides to buy bilaterally if

$$\begin{aligned} &0 \geqslant -(\theta+\nu_i) + \mathbb{E}[\max_{k\in\mathcal{J}} \tilde{u}_{ki}(\boldsymbol{\epsilon_{ki}})] - c \\ \Leftrightarrow\ &\psi(q) \leqslant \nu_i \text{ with } \psi(q) = \mathbb{E}[\max_{k\in\mathcal{J}} \tilde{u}_{ki}(\boldsymbol{\epsilon_{ki}})] - \theta - c \text{ where } \tilde{u}_{ji}(\boldsymbol{\epsilon_{ji}}) = \xi_j + q_j + \sigma\boldsymbol{\epsilon_{ji}} \end{aligned} \tag{16}$$

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. By (16) and the assumption that $\boldsymbol{v_i}$ are iid across i, the dealer knows that the probability that an investor chooses to trade bilaterally is

$$\rho_d(q) = \Pr(\boldsymbol{v_i} \geqslant \psi(q)) = 1 - F(\psi(q)). \tag{17}$$

Since the dealer caters to a unit mass of buyers, this probability is also the dealer's bilateral market share. Similarly, the probability that an investor enters the platform and chooses the dealer, which is

$$S_d(q) = \sum_j F(\psi(q)) * s_d(q) \text{ with } s_d(q) \text{ defined in (15),} \tag{18}$$

is the dealer's market share on the platform. Anticipating this, the dealer chooses q_d to

$$\max_{q_d} \pi_d(q) = \max_{q_d}\{\pi_d^D(q) + \pi_d^E(q)\} \tag{19}$$

where $\pi_d^D(q)$ is the expected profit from bilateral trades and $\pi_d^E(q)$ from platform trades:

$$\pi_d^D(q) = \int_{\psi(q)}^{\infty} (value_d - (\theta+\nu-\xi_d))f(\nu)d\nu \text{ given that } y_d = \theta+\nu-\xi_d \tag{20}$$

$$\pi_d^E(q) = S_d(q)(value_d - q_d). \tag{21}$$

Next we take the partial derivative w.r.t. q_d. Doing so, we abbreviate $\frac{\partial\pi(q)}{\partial q_d} = \pi_d'(q)$ and similarly for any other function, with slight abuse of notation.

$$0 = \pi_d'(q) \Leftrightarrow 0 = (value_d - (\theta+\psi(q)-\xi_d))\rho_d'(q) + S_d'(q)(value_d - q_d) - S_d(q) \tag{22}$$

where we have used that $\rho_d'(q) = (-1)f(\psi(q))\psi'(q)$ by (17). Rearranging gives

$$\Leftrightarrow q_d\left(1 + \frac{S_d(q)}{S_d'(q)}\frac{1}{q_d} - (value_d - (\theta+\psi(q)-\xi_d))\frac{\rho_d'(q)}{S_d'(q)}\frac{1}{q_d}\right) = value_d. \tag{23}$$

To obtain the equation in the proposition, substitute out for terms that equal elasticities:

$$\eta_d^D(q) = \frac{\rho_d'(q)}{\rho_d(q)}q_d \text{ and } \eta_d^E(q) = \frac{S_d'(q)}{S_d(q)}q_d \Rightarrow \frac{\rho_d'(q)}{S_d'(q)} = \frac{\eta_d^D(q)}{\eta_d^E(q)}\frac{\rho_d(q)}{S_d(q)} \tag{24}$$

and the expression for the marginal profit earned in the bilateral segment of the market

$$\pi_d^{D'}(q) = (value_d - (\theta+\psi(q)-\xi_d))\eta_d^D(q)\frac{\rho_d(q)}{q_d} \text{ by (20) and (24).}$$

Equation (23) turns into the markup equation when relabelling $\frac{\pi_d^D(q)}{\partial q_d} = \frac{\pi_d^{D'}(q)}{S_d(q)}$. □

C.2 Model with selling investors

When the dealer sells to the investor, the objective functions change. A dealer generates the following profit ex-post: $\pi_{tj}(x) = x - v_{tj}^{*D}$ when buying at yield x. An investor in group $G^* \in \{R^*, I^*\}$ obtains a surplus of $v_{ti}^{G^*} - y_{ti}^{G^*}$ when selling bilaterally, where

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

The payoff that determines if an institutional investor enters the platform is

$$\max_j \{v_{tij}^{I^*} - q_{tj}^* + \sigma\epsilon_{tij}\} - c_t^* \text{ with } \epsilon_{tij} \overset{iid}{\sim} \mathcal{H}_t \text{ and } c_t^* \in \mathbb{R}, \sigma \in \mathbb{R}^+.$$

Proposition 3 (Selling investors)**.**

(i) A retail investor with shock $\nu_{ti}^{R^*}$ sells bilaterally to home dealer d at

$$y_{tid}^{R^*} = \theta_t + \nu_{ti}^{R^*} + \xi_{td}. \tag{25}$$

(ii) An institutional investor with shock ν_{ti}^{I*} sells bilaterally to home dealer d at

$$y_{tid}^{I*} = \theta_t + \nu_{ti}^{I^*} + \xi_{td} \quad \text{if } \psi_t^*(q_t^*) \geqslant \nu_{ti}^{I^*} \tag{26}$$

$$\text{where } \psi_t^*(q_t^*) = -\mathbb{E}[\max_{k\in\mathcal{J}_t} \tilde{u}_{tik}^*(\boldsymbol{\epsilon_{tik}})] + c_t^* - \theta_t \text{ with } \tilde{u}_{tij}^*(\boldsymbol{\epsilon_{tij}}) = \xi_{tj} - q_{tj}^* + \sigma\boldsymbol{\epsilon_{tij}}. \tag{27}$$

Otherwise, she enters the platform. On the platform, she observes ϵ_{ti} and sells to the dealer with the maximal $\tilde{u}_{tij}^*(\epsilon_{tij})$ at q_{tj}^*.

The proof is analogous to the proof of Proposition 1. □

Proposition 4 (Dealers for selling investors)**.** *Dealer j posts a quote q_{tj}^* that satisfies*

$$q_{tj}^*\Big(1 + \frac{1}{\eta_{tj}^{E^*}(q_t^*)}\Big(1 - \frac{\partial \pi_{tj}^{D^*}(q_t^*)}{\partial q_{tj}^*}\Big/ S_{tj}^*(q_t^*)\Big)\Big) = v_{tj}^{*D} \tag{28}$$

where $\eta_{tj}^{E^*}(q_t)$ is the elasticity of supply on the platform and $\frac{\partial \pi_{tj}^{D^*}(q_t^*)}{\partial q_{tj}^*}$ is the marginal profit that the dealer expects to pay in bilateral trades with institutional investors normalized by the share of investors who buy from her on the platform, $S_{tj}^*(q_t^*)$.
Formally, $\eta_{tj}^{E^*}(q_t^*) = q_{tj}^* \frac{\partial S_{tj}^*(q_t^*)}{\partial q_{tj}^*} / S_{tj}^*(q_t^*)$ with platform market share $S_{tj}^*(q_t) = \sum_{j\in J_t} \Pr(\psi_t^*(q_t^*) \leqslant \boldsymbol{\nu_{ti}^{I*}}) \Pr(\tilde{u}_{tki}^*(\boldsymbol{\epsilon_{tki}}) < \tilde{u}_{tij}^*(\boldsymbol{\epsilon_{tij}})\ \forall k \neq j)$ where $\tilde{u}_{tij}^*(\boldsymbol{\epsilon_{tij}})$ and $\psi_t^*(q_t^*)$ are defined in (27) of Proposition (3); and $\pi_{tj}^{D^*}(q_t^*) = \mathbb{E}[\boldsymbol{\nu_{ti}^{I*}} + \theta_t + \xi_{tj} - v_{tj}^{D^*} | \boldsymbol{\nu_{ti}^{I*}} \leqslant \psi_t^*(q_t^*)]$.

The proof is analogous to the proof of Proposition 2. □

C.3 Model extension: Bargaining power

In this extension, we allow investors to extract some trade surplus in the bilateral trade. We focus on the buy-side only to avoid redundancy. The propositions for the sell-side are analogous.

Let the dealer and investor bargain under complete information, and denote the investor's bargaining power by $\phi \in [0,1)$. The dealer and an investor of group G agree on the following bilateral yield:

$$y_{tid} = \phi\ v_{tj}^D + (1-\phi)v_{tij}^G \tag{29}$$

where v_{tj}^D and v_{tij}^G are the dealer's and investor's value.

The equilibrium characterization can be derived similarly to before.

Proposition 5 (Buying investors)**.**

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

$$y_{tid}^R = \phi\ v_{td}^D + (1-\phi)(\theta_t + \nu_{ti}^R - \xi_{td}). \tag{30}$$

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

$$y_{tid}^I = \phi\ v_{td}^D + (1-\phi)(\theta_t + \nu_{ti}^I - \xi_{td}) \quad \text{if } \psi_{td}(q_t) \leqslant \nu_{ti}^I \tag{31}$$

$$\text{where } \psi_{td}(q_t) = (\mathbb{E}[\max_{k \in \mathcal{J}_t} \tilde{u}_{tik}(\boldsymbol{\epsilon_{tik}})] - c_t - \phi(v_{td}^D + \xi_{td}))/(1-\phi) - \theta_t \tag{32}$$

$$\text{with } \tilde{u}_{tij}(\boldsymbol{\epsilon_{tij}}) = \xi_{tj} + q_{tj} + \sigma\boldsymbol{\epsilon_{tij}}. \tag{33}$$

Otherwise, he enters the platform. On the platform, he observes ϵ_{ti} and buys from the dealer with the maximal $\tilde{u}_{tij}(\epsilon_{tij})$ at q_{tj}.

The proof is analogous to the proof of Proposition 1. $\square$

The proposition is the analogue of Proposition 1. It characterizes where investors buy and at what yields in the extended model. By assumption, retail investors always buy at a yield that leaves them with a fraction ϕ of the total trade surplus (statement (*i*)). Institutional investors buy on the platform if she expects the surplus from a bilateral trade, $\phi(v_{td}^D - (\theta_t + \nu_{ti}^R - \xi_{td}))$, to be lower than the expected surplus from buying on the platform, $\mathbb{E}[\max_{k\in\mathcal{J}_t} \tilde{u}_{tik}(\boldsymbol{\epsilon_{tik}}) - (\theta_t + \nu_{ti}^I)]$, minus the cost she needs to pay to use the platform. As in the benchmark model, this is the case for investors who draw low liquidity shocks.

Proposition 6 (Dealers for buying investors)**.** *Dealer j posts a quote q_{tj} that satisfies the markup equation (6) of Proposition 2. The dealer j's quote elasticity of demand on the platform, $\eta_{tj}^E(q_t)$, and the profit that the dealer expects from bilateral trades with institutional investors, $\pi_{tj}^D(q_t)$, now differ: $\eta_{tj}^E(q_t) = q_{tj}\frac{\partial S_{tj}(q_t)}{\partial q_{tj}}/S_{tj}(q_t)$ with $S_{tj}(q_t) = \sum_{j\in J_t} \Pr(\boldsymbol{\nu_{ti}^I} \leqslant \psi_{tj}(q_t)) \Pr(\tilde{u}_{tki}(\boldsymbol{\epsilon_{tki}}) < \tilde{u}_{tij}(\boldsymbol{\epsilon_{tij}})\ \forall k \neq j)$ where $\tilde{u}_{tij}(\boldsymbol{\epsilon_{tij}})$ and $\psi_{tj}(q_t)$ are defined in (32) of Proposition (5); and $\pi_{tj}^D(q_t) = (1-\phi)\mathbb{E}[v_{tj}^D - (\boldsymbol{\nu_{ti}^I} + \theta_t - \xi_{tj})|\psi_{tj}(q_t) \leqslant \boldsymbol{\nu_{ti}^I}]$.*

The proof is analogous to the proof of Proposition 2. $\square$

Relative to the benchmark model, the dealer extracts a lower trade surplus in the bilateral trades. Therefore, her bilateral trade profit is lower. This implies that the incentive to set low platform quotes and with that deter investors from entering the platform is lower than in the benchmark model.

In the counterfactual exercise, we compute the expected investor surplus and dealer profit as follows:

Definition 3.

(i) The expected surplus of an investor i in group G who buys on t is

$$EU_t^G = \phi\mathbb{E}[\boldsymbol{y_{tij}^G} - v_{tij}^G | \textit{buy bilaterally}] \\ + \sum_j S_{tj}^G(q_t)\mathbb{E}[(q_{tj} + \sigma\boldsymbol{\epsilon_{tij}}) - (\theta_t + \boldsymbol{\nu_{ti}^G}) + \xi_{tj} | \textit{buy on the platform from dealer } j].$$

(ii) The expected profit of dealer j from selling to group G on t is

$$PI_{tj}^G = (1-\phi)\mathbb{E}[v_{tj}^D - \boldsymbol{y_{tij}^G} | \textit{buy bilaterally}] \\ + S_{tj}^G(q_t)\mathbb{E}[v_{tj}^D - (q_{tj} + \sigma\boldsymbol{\epsilon_{tij}}) | \textit{buy on the platform from dealer } j].$$

In both cases, Proposition 5 specifies where the investor buys and at what yield. In particular, $S_{tj}^G(q_t)$ *is the probability that an investor in* $G \in \{I, R\}$ *buys from dealer* j *on the platform. It is defined as in Proposition 2 if the investor group has platform access, and 0 otherwise.*

Welfare is computed as in Definition 2 but with Proposition 5 specifying from which dealer the investor buys.

D Micro-foundation of the platform friction

The goal of this section is to microfound the platform friction term, $\sigma\epsilon_{tij}$, in the structural model. For simplicity, let dealers be ex-ante identical, i.e., abstract from quality ξ_j. Later on, we explain how one could extend the theoretic model to include the quality term. Further, we restrict attention to institutional investors who have access to the platform. Including retail investors with no access is straightforward, and analogous to the structural model.

D.1 Theoretic model

Overview. As in the structural model, the theoretic model features two trade-offs. First, a dealer faces a trade-off when posting a quote on the platform. A low quote makes the platform relatively less attractive than trading bilaterally. A high quote, however, means that the dealer captures more market share on the platform.

Because the platform is not frictionless, the actual transaction yield on the platform equals to the posted quote plus a friction. This friction may cause the investor to choose different dealers than under perfect competition. Here, we microfound where the friction might come from by modeling the auctions that investors run on the platform.

The second trade-off is the one that an investor faces when deciding whether to trade on the platform or bilaterally: To use the platform, the investor has to pay a cost, but she can extract a higher share of the trade surplus than in a bilateral trade. The reason is that, on the platform, dealers compete with one another in a first-price auction.

Before running the auction, each dealer receives a (possibly correlated) signal about how much she wants to realize the trade with a particular investor. The signal determines by how much each dealer is willing to markup or discount her posted quote. In equilibrium, the dealer with the highest signal wins the auction and the investor pays a price that equals to the posted quote of the winning dealer plus a stochastic term. This term scales up by a parameter, sigma, which is called the competition parameter, σ, in the structural model.

Formal details. $J \geqslant 2$ dealers sell a bond to investors. Each transaction is a single unit trade. Each investor has a home dealer, called d (short for d_i), and each dealer has a home investor base. It consists of a unit mass of investors.

Dealers. A dealer aims at maximizing profit. Ex-post, the dealer obtains a profit of $v - x$, when selling one unit at yield x, and valuing the bond by v. The dealer's value splits into a part that is commonly known to all dealers, v_1, and a part that is unknown v_2. It is drawn iid form a commonly known normal distribution:

$$v = v_1 + v_2 \text{ with } v_1 \in \mathbb{R} \text{ and } v_2 \overset{iid}{\sim} N(\mu_v, \sigma_v^2).$$

Investors. All investors can either buy bilaterally from their home dealer or use the platform. In making this decision, each investor i aims at maximizing surplus. She obtains a surplus of $x - \nu_i$ when buying form a dealer at yield x, and valuing the bond by

$$\nu_i \overset{iid}{\sim} N(\theta + \mu_\nu, \sigma_\nu^2) \text{ where } \theta \in \mathbb{R}.$$

Both θ and the normal distribution are commonly known.

Timing of events. The game has two stages. In the first stage, each dealer j posts a quote, q_j, on the platform, simultaneously with all other dealers. The quote signals the benchmark yield that an investor can obtain when buying from this dealer on the platform. In the second stage, each investor observes her private value, ν_i, and decides to buy bilaterally from her home dealer, or to buy on the platform. In a bilateral trade, the home dealer observes ν_i and charges $y_i = \nu_i$. On the platform, the investor runs a first-price auction with all dealers. To do so, she has to pay a cost $c \in \mathbb{R}^+$.

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

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

Given these signals, each dealer submits her bid. A bid determines the markup (or discount) that the investor receives relative to her posted quote. The dealer who offers the best markup wins and and the investor obtains a yield that equals to the posted quote of the winning dealer plus the markup.

Equilibrium. In a symmetric equilibrium, dealers and investors behave in a similar way to how they behave in the structural model. Here, we discuss differences and similarities. All proofs are in Appendix D.3.

Proposition 7 (Investor's choice and the platform auction)**.**

(i) In a symmetric equilibrium, in which all dealers post the same quote q^*, an investor with shock ν_i buys bilaterally from her home dealer at $y_i = \nu_i$ if

$$\nu_i > \psi(q^*) \text{ where } \psi(q^*) = q^* + \sigma s + \mathbb{E}[\max_j \{\boldsymbol{x_{ij}}\}] - c. \tag{34}$$

Otherwise she enters the platform, and runs an auction.
(ii) The dealer with the highest signal, x_{ij}, wins the auction, and the investor obtains the following yield on the platform:

$$y_{ij}^E = q^* + \sigma(x_{ij}/\sigma + s) \text{ where} \tag{35}$$

$$\sigma \text{ solves } 0 = \int_{-\infty}^{\infty} [-\Phi(-z)^{J-1} + (z-\sigma)(J-1)\Phi(-z)^{J-2}\phi(-z)]\,\phi(z)\,dz. \tag{36}$$

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

The first statement of the proposition characterizes which investors buy bilaterally and which on the platform. This is determined by a cutoff that is of the same form as in our structural model. As before, investors with less urgency to trade (in that they aim for a high yield due to a high liquidity shock) stay off the platform. For them it is better to buy bilaterally because the platform quote is targeted to an investor with an average willingness to pay, rather than the investor's individual willingness to pay.

The second statement tells us that the investor receives a yield that equals to the posted quote plus an additional term, like in the structural model. Here, the additional term depends on a parameter, σ, the dealer's signal, x_{ij}, and its standard deviation, s.[39]

The competition parameter σ is negative and depends on the number of dealers in the auction. Intuitively, it should increase as the number of dealers increases as the competition becomes tougher. For example, when $J = 2$, $\sigma \approx -1.7$, when $J = 3$, $\sigma \approx -1.5$.
When posting quotes, dealers anticipate the behavior in the second stage of the game. In equilibrium, quotes fulfill a similar equation as in the structural model.

Proposition 8 (Dealer's quotes)**.**
In a symmetric equilibrium, dealer j posts a quote that satisfies the following equation:

$$q_j\left(1 + \frac{1}{\eta^E(q)}\left(1 - \frac{\partial \pi^D(q)}{\partial q_j}\right)/S(q)\right) = v_1 - \sigma s. \tag{37}$$

Here, $\eta^E(q)$ is a dealer's quote elasticity of demand on the platform and $\frac{\partial \pi^D(q)}{\partial q_j}$ is the marginal profit that the dealer expects from bilateral trades with investors. It is normalized by the size of the dealer's platform market share, $S(q)$.

All dealers post the same quote $q_j = q^*$ since they are ex-ante identical.

Equation (37) takes the same form as equation (6) in our structural model with one difference: On the right-hand side, we no longer have the dealer's (known) value v_1 but her (known) value minus a term that determines the size of the markup that the dealer charges

[39] A similar results holds for the special case, in which dealers draw independent private signals before running the auction (i.e., $v_2 = 0$ and $x_{ij} \sim N(\mu_x, s^2)$). The difference is that the stochastic term no longer separates into a scaler that multiplies a shock term.

in the auction on the platform. The reason is that the dealer now no longer obtains the yield q on the platform, but instead q plus the auction markup. Anticipating this, she adjusts the quote that she posts.

D.2 Discussion: Theoretic vs. structural model

The theoretic model gives rise to similar, but not identical equilibrium conditions as the structural model. Here, we discuss the main differences and how to overcome them.

In the theoretic model, the ex-ante symmetrical dealers post the same quote on the platform, and the platform yield is

$$y_{ij}^E = q^* + \sigma(x_{ij}/\sigma + s) \text{ where } x_{ij} = v_2 + s\zeta'_{ij} \text{ with } v_2 \overset{iid}{\sim} N(\mu_v, \sigma_2^2) \text{ and } \zeta'_{ij} \sim N(0,1), s \in \mathbb{R}.$$

In the structural model, dealers differ in quality, ξ_j, and therefore post different quotes. The investor obtains the following yield on the platform:

$$y_{ij}^E = q_j + \sigma\epsilon_{ij} \text{ where } \epsilon_{ij} \text{ is iid.}$$

One could extend the theoretic model and include a quality term ξ_j in the investor's surplus. Then dealers would post different quotes depending on their quality. To keep the model tractable, one would have to assume that the auction determines the markup over the posted quote and that the dealer with the best markup wins, independent of the posted quote. Otherwise, the auction would become asymmetric.

This extended theory model would differ from the structural model in that the platform shocks may be correlated. Instead, for the structural estimation we assume that the shocks are iid. Such an assumption is common in the literature on demand estimation, but, as the microfoundation highlights, imposes a restriction. To achieve independence in the theory model, one would have to assume that the dealers draw independent signals, conditional on their value v_1.

D.3 Proof of Proposition 8 and 7

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

In stage 3 each dealer observes her signal x and participates in a fist-price auction that pins down the markup (or discount) that the investor pays over the posted quote. The size of the markup is determined by v_2 and the dealer's signal about it. If dealers knew v_2 and were truthful in the auction, the markup would equal to v_2.

Guess that there is an equilibrium in which a dealer with signal x submits a linear function $\beta(x) = x + \sigma s$, 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 $\sigma_v^2 \uparrow \infty$, $h \downarrow 0$, and at limit $v|x \sim N(x, s^2)$. Hence $[v_2 - x]/s = \zeta$ where $\zeta \sim N(0,1)$.

The dealers problem is, given x,

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

where $F(z|v_2)^{J-1}$ is distribution of maximum z of others' signals given v_2, and $F(v_2|x)$ is distribution of v given signal x. The first-order condition is

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

evaluated at $z = \beta^{-1}(\beta(x)) = x$, so $\frac{d\beta^{-1}(\beta(x))}{db} = 1/\beta'(x)$. Thus the first-order condition is

$$0 = \int_{-\infty}^{\infty} [-F(x|v_2)^{J-1} \beta'(x) + (v_2 - \beta(x))(J-1) F(x|v_2)^{J-2} f(x|v_2)] \, dF(v_2|x) \,.$$

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

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

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

$$0 = \int_{-\infty}^{\infty} [-\Phi(-\zeta)^{J-1} + (\zeta - \sigma)(J-1)\Phi(-\zeta)^{J-2}\phi(-\zeta)] \, \phi(\zeta) \, d\zeta$$

This is just an equation for the parameter σ of the bid function $\beta(\cdot)$, independent of x and s.

In this equilibrium, the dealer with the highest signal wins the auction with the investor. The investor buys from the dealer with the highest markup. Assume this dealer is dealer j. Then the investor receives $q_j + x_j + \sigma s$ on the platform.

In stage 2 an investor with shock ν enters the platform if

$$0 < \mathbb{E}[\max_j \{q_j + \boldsymbol{x_j}\}] + \sigma s - c - \nu \Leftrightarrow \nu < \psi(q)$$
$$\text{where } \psi(q) = \mathbb{E}[\max_j \{q_j + \boldsymbol{x_j}\}] + \sigma s - c.$$

In the equilibrium, where $q_j = q^*$ for all j, the cutoff simplifies to the cutoff of the proposition.

In stage 1 a dealer maximizes

$$\max_{q_j} \pi_j(q) = \max_{q_j} \{\pi_j^D(q) + \pi_j^E(q)\}$$

where $\pi_j^D(q)$ is the expected profit from bilateral trades and $\pi_j^E(q)$ from platform trades:

$$\pi_j^D(q) = \int_{\psi(q)}^{\infty} (v_1 + \mu_v - \nu) f(\nu) d\nu \text{ given } v_2 \sim N(\mu_v, \sigma_v^2) \text{ and iid from } \nu.$$

$$\pi_j^E(q) = \sum_j F(\psi(q)) \int \mathbb{I}(q_j + x_j > q_k + x_k \forall k \neq j)(v_1 + x_j - (q_j + x_j + \sigma s)) f_x(x) dx.$$

Here the integral is a multidimensional integral over all x and $f_x(\cdot)$ the multidimensional density function. The platform profit simplifies to

$$\pi_j^E(q) = S_j(q)(v_1 - (q_j + \sigma s)) \text{ where } S_j(q) = JF(\psi(q)) \int \mathbb{I}(q_j + x_j > q_k + x_k \forall k \neq j) f_x(x) dx$$

is the probability that some investor enters the platform and buys from dealer j. In equilibrium,

$$\pi_j'(q) = 0 \Leftrightarrow 0 = (v_1 + \mu - \psi(q))\rho_j'(q) + S_j'(q)(v_1 - (q_j + \sigma s)) - S_j(q) \text{ with } \rho_j'(q) = (-1) f(\psi(q))\psi'(q).$$

This condition rearranges to

$$q_j + \sigma s + \frac{S_j(q)}{S_j'(q)} - (v_1 + \mu - \psi(q)) \frac{\rho_j'(q)}{S_j'(q)} = v_1 \Leftrightarrow q_j \left(1 + \frac{1}{\eta_j^E(q)} \left(1 - \frac{\partial \pi^D(q)}{\partial q_j}\right) / S_j(q)\right) = v_1 - \sigma s$$

Since the condition is identical for all dealers, they post the same quote in equilibrium: $q_j = q^*$. □

D.4 Counterfactual: How dealers set quotes

Here we only derive the optimality condition for the buy-side. The sell-side is analogous. Retail investors now behave like institutional investors in the status quo, i.e., according to Proposition 1. This implies that

$$\rho_d^G(q) = \Pr(\boldsymbol{v_i^G} \geqslant \psi(q)) = 1 - F^G(\psi(q)) \tag{17'}$$

is the probability that an investor in group $G \in \{I, R\}$ buys bilaterally from home dealer d, and

$$S_d^G(q) = \sum_j F^G(\psi(q)) * s_d(q) \text{ with } s_d(q) \text{ defined in (15)} \tag{18'}$$

is probability that an investor in group G enters the platform and chooses this dealer.

Anticipating investor behavior, the dealer chooses q_d to

$$\max_{q_d} \pi_d(q) = \max_{q_d} \left\{ \sum_G \kappa_G \left(\pi_d^{GD}(q) + \pi_d^{GE}(q) \right) \right\} \tag{38}$$

where $\pi_d^{GD}(q)$ is the expected profit from bilateral trades with investors in group G and $\pi_d^{GE}(q)$ from platform trades:

$$\pi_d^{GD}(q) = \int_{\psi(q)}^{\infty} (value_d - (\theta + \nu^G - \xi_d)) f^G(\nu) d\nu^G \text{ given that } y_d^G = \theta + \nu^G - \xi_d \tag{20'}$$

$$\pi_d^E(q) = S_d^G(q)(value_d - q_d). \tag{21'}$$

Taking the partial derivative w.r.t. q_d, and simplifying gives the following condition

$$\begin{aligned}&\sum_G \kappa^G \{(value_d - (\theta_t + \psi(q) - \xi_d)) f^G(\psi(q))\} + \sum_G \kappa^G \{ \sum_d F^G(\psi(q))\} \\ &= \sum_G \kappa^G \Big[\sum_d f^G(\psi(q)) s_d(q) + \sum_d F^G(\psi(q)) \frac{1}{\sigma}(1 - s_d(q)) \Big] (value_d - q_d)\end{aligned} \tag{39}$$

Simplifying further, one can derive a markup equation that is similar to Proposition 2. This condition is more complicated than the markup equation of Proposition 2 when investors draw liquidity shocks from a normal distribution with mean $\mu_t = \sum_G \kappa^G \mu_t^G$ and standard deviation $\sigma_t = \sum_G \kappa^G \sigma_t^G$. This is because the two investor groups may select with different intensity onto the platform.

However, we show numerically, that given our estimated parameters, the quotes that fulfill (39) are very close to the quotes that the dealer choses if she behaves as if she trades with a representative investor whose liquidity shock is drawn from $N(\mu_t, \sigma_t)$. Appendix Figure A12 shows a histogram of the difference between the quotes.

E Details regarding the estimation

E.1 Construction of the cost-shifter instrument

In modeling the bidding process we follow Hortaçsu and Kastl (2012) and Allen et al. (2020) who model the Canadian government bond auctions, and recover how much bidders are willing to pay. For a detailed discussion of all assumptions, we refer to these papers.

E.1.1 Auction model

In the auction, there are two types of bidders, N_d dealers ($g = d$) and N_c customers ($g = c$) who are investors that bid at auction. All bidders are indexed by j. Each of them is characterized by what information she has (by Assumption 5) and by how much she values different amounts of the bond (by Assumption 6).

Assumption 5. *Dealers' and customers' private signals s_j^d and s_j^c 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.*

Assumption 6. *A bidder j in 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.*

In auction $\tilde{t}$, bidders place bids for amounts of the bond that they seek to purchase at different prices. More specifically, a bid is a step function that characterizes the price that 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. Since the dealer can observe the customer's bid, she might be able to extract valuable information. We allow for this, and define the information that is available to dealer j before placing her (final) 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 4. *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 G with type θ_j^g.*

Once all bidders submitted 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 parts of the amount that the Bank of Canada announced to issue a week prior to the auction goes to non-competitive 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 how much she offered to win for each unit that she won.

Definition 5. *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, $\mathbb{E}[\int_0^{q_j^c} [v(x, s_j^g) - b_j^g(x, \theta_j^g)]dx]$.*

We focus on type-symmetric BNE in which all dealers and all customers play the same strategy.

E.1.2 Estimating the amount that a dealer expects to win

If we assume that bidders behave as in the auction model, we can estimate how much a bidder expects to win in equilibrium, at the time at which she places the bids. In a nutshell, the idea is to fix a bidder and randomly draw from the set of submitted bids. This allows one to simulate market clearing. Repeating this many times, one can recover the distribution of how much the bidder wins, and with it construct the expectation.

More specifically, we first estimate the distribution of the residual supply curve for each dealer. This curve is the total supply minus the 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 by 0. We then find the dealer(s) who observed each of the customer bids, and draw their bids. In rare cases, in which 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 that we already draw 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 by intersecting each of the residual supply curves with the bidding function that the dealer submitted. This gives a distribution of how much the dealer won in the auction. It also specifies for each step of the dealer's bidding function, how likely it is that the market clears at that step. With that, we can compute how much the dealer expected to win:

$$\begin{aligned}&\text{amount dealer } j \text{ expected to win when bidding}\\ &= \sum_{k}^{K_j} \hat{\Pr}(\text{market clears at step } k) * \hat{\mathbb{E}}[\boldsymbol{q_k^c}|\text{market clears at step } k]\end{aligned}$$

where K_j are the steps in dealer j's the bidding function.

E.2 Main estimation procedure

We explain the estimation for buying investors. The procedure is analogous for selling investors with the difference that it builds on Propositions 3 and 4 rather than 1 and 2.

E.2.1 Institutional investors

The estimation follows three steps.

Step 1. Dealer j's market share on the platform, denoted $s_{tj}(q_t, \xi_t, \sigma)$, is a function of (q_t, ξ_t, σ). When ϵ_{tij} are Extreme Value Type 1 distributed, $s_{tj}(q_t, \xi_t, \sigma)$ has a closed form solution:

$$s_{tj}(q_t, \xi_t, \sigma) = \frac{\exp(\delta_{tj})}{\sum_{j=1}^{J_t} \exp(\delta_{tj})} \text{ with } \delta_{tj} = \frac{1}{\sigma}(\xi_{tj} + q_{tj}) \text{ for all } j \in \mathcal{J}_t. \tag{40}$$

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 ($j = 0$), and take logs to obtain

$$\log(s_{tj}/s_{t0}) = \delta_{tj} - \delta_{t0} = \frac{1}{\sigma}(\tilde{\xi}_{tj} + \tilde{q}_{tj})$$

where $\tilde{\xi}_{tj} = \xi_{tj} - \xi_{t0} = \xi_{tj}$ by Normalization 1, and $\tilde{q}_{tj} = q_{tj} - q_{t0}$. By Assumption 2, this is equivalent to

$$\log(s_{tj}/s_{t0}) = \zeta_j + \zeta_t + \frac{1}{\sigma}\tilde{q}_{tj} + rest_{tj} \tag{41}$$

with $\zeta_j = \frac{1}{\sigma}\xi_j, \zeta_t = avg_j\left(\frac{1}{\sigma}\chi_{tj}\right), rest_{tj} = \frac{1}{\sigma}\chi_{tj} - \zeta_t$.

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 the estimated $\hat{\sigma}$, fixed effects and residual of this regression we can compute $\hat{\xi}_{tj}$ for all $j \neq 0$:

$$\hat{\xi}_{tj} = \hat{\sigma}\left(\hat{\zeta}_j + \hat{\zeta}_t + r\hat{est}_{tj}\right) = \hat{\sigma}(\frac{1}{\hat{\sigma}}\hat{\xi}_j + \hat{\zeta}_t + \frac{1}{\hat{\sigma}}\hat{\chi}_{tj} - \hat{\zeta}_t) = \hat{\xi}_j + \hat{\chi}_{tj} = \hat{\xi}_{tj} \text{ by construction.}$$

Step 2. Next, we compute the cutoff that determines whether an investor buys bilaterally form home dealer d or on the platform:

$$\mathbb{E}[\max_{k \in \mathcal{J}_t} \tilde{u}_{tik}(\boldsymbol{\epsilon_{tik}})] \text{ with } \tilde{u}_{tij}(\boldsymbol{\epsilon_{tij}}) = \hat{\xi}_{tj} + q_{tj} + \hat{\sigma}\boldsymbol{\epsilon_{tij}}.$$

Since $\boldsymbol{\epsilon_{tij}}$ are drawn iid from the Extreme Value Type 1 distribution, the expectation has a closed form solution:

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

Step 3. For each day t, we estimate the remaining parameters via GMM. To see this, fix day t and denote $\beta_t = \{\mu_t^I, \sigma_t^I, c_t\}$. Call the number of bilateral trades between institutional investors and dealer j N_{tj} and the total number of platform trades E_t. Further, let y_{tij}^I be the normalized yield of a bilateral trade between investor i and dealer j on day t, refer to $\boldsymbol{y_{tij}^I}$ as the yield that is predicted by the model: $\boldsymbol{y_{tij}^I} = \theta_t + \boldsymbol{\nu_{ti}^I} - \hat{\xi}_{td}$. We match the the following three moments:

$$\frac{1}{J_t}\sum_{j=1}^{J_t}\mathbb{E}[\boldsymbol{y_{tij}^I}|i \in I \text{ investor buys bilaterally}, \beta_t] = \frac{1}{J_t}\sum_{j=1}^{J_t}\frac{1}{N_{tj}}\sum_{i=1}^{N_{tj}} y_{tij}^I$$

$$\frac{1}{J_t}\sum_{j=1}^{J_t}Var(\boldsymbol{y_{tij}^I}|i \in I \text{ buys bilaterally}, \beta_t) = \frac{1}{J_t}\sum_{j=1}^{J_t}\frac{1}{N_{tj}}\sum_{i=1}^{N_{tj}}\left(y_{tij}^I - \mathbb{E}[\boldsymbol{y_{tij}^I}|i \in I \text{ buys bilaterally}, \beta_t]\right)^2$$

$$\Pr(i \in I \text{ buy bilaterally}|\beta_t) = \frac{\sum_j N_{tj}}{\sum_j N_{tj} + E_t}$$

To compute the predicted moments (on the LHS), we rely on the assumption that $\boldsymbol{\nu_{ti}^I} \sim N(\mu_t^I, \sigma_t^I)$ and Proposition 1.

E.2.2 Retail investors

For each day t, we estimate the μ_t^R and σ_t^R by GMM, matching the average and standard derivation of the trade yield:

$$\frac{1}{J_t}\sum_{j=1}^{J_t}\mathbb{E}[\boldsymbol{y^R_{tij}}|i \in R \text{ buys}, \{\mu_t^R, \sigma_t^R\})] = \frac{1}{J_t}\sum_{j=1}^{J_t}\frac{1}{R_{tj}}\sum_{i=1}^{R_{tj}} y^R_{tij}$$

$$\frac{1}{J_t}\sum_{j=1}^{J_t} Var(\boldsymbol{y^R_{tij}}|i \in R \text{ buys}, \{\mu_t^R, \sigma_t^R\}) = \frac{1}{J_t}\sum_{j=1}^{J_t}\frac{1}{R_{tj}}\sum_{i=1}^{R_{tj}}\Big(y^R_{tij} - \mathbb{E}[\boldsymbol{y^R_{tij}}|i \in R \text{ buys}, \{\mu_t^R, \sigma_t^R\}]\Big)^2$$

where R_{tj} is the number of retail investors who buy from dealer j on day t, y^R_{tij} is the observed yield and $\boldsymbol{y^R_{tij}} = \theta_t + \boldsymbol{\nu^R_{ti}} - \hat{\xi}_{tj}$ is yield that is predicted by the model.

F Robustness analysis

The first set of robustness checks mainly regards the estimate of the competition parameter, σ, and the dealer's qualities, ξ_{tj} (see Appendix Table 10). First, we check whether $\hat{\sigma}$, which governs the quote 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 by $\mathbb{E}[q_{tj}\chi_{tj}|\zeta_t, \xi_j] = 0$. 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 downwards 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.

When including dealer fixed effects, that is, when allowing for dealers to systematically differ in quality, we find similar estimates. However, the instrument is very weak. We therefore check by how much the estimates change when dropping the dealer fixed effect by imposing $\xi_j = 0 \forall j$. This increases the correlation between our instrument (in both specifications) and the platform quotes, so that the instrument becomes stronger. At the same time, we no longer control for unobservable differences between dealers that might drive differences in the quotes. The implied $\hat{\sigma}$ decreases from 0.77 to 0.49 when using our preferred instrument and from 0.68 to 0.62 when using the amount the dealer won as instrument. For more details, see Appendix Table 10.

The second set of robustness tests regards the estimates of the usage costs and the distribution of the liquidity shocks. To verify that the distribution of the liquidity shock is not much biased by measurement errors in the quality of the dealers ξ_{tj}, we estimate the model under the assumption that the bilateral yield equals to the market value plus the liquidity shock, i.e., $y^G_{ti} = \theta_t + \nu^G_{ti}$ rather than $y^G_{tid} = \theta_t + \nu^G_{ti} - \xi_{td}$. We find that estimates are similar (see Appendix Table 11). For our policy analysis we rely on our main specification in which investors obtain 0 surplus when trading bilaterally with their home dealer. This implies that there is no obvious reason for them to switch home dealers. That is no longer the case when $y^G_{ti} = \theta_t + \nu^G_{ti}$. Now the investor gets a surplus of ξ_{td} from a bilateral trade and would have incentives to switch home dealers if other home dealers provide higher quality.

Further, we check by how much the estimates change when we allow the platform usage cost to be dealer specific, c_{tj}. This cost is identified from how many trades each dealer j realizes on vs. off the platform in period t rather than how many trades are realizes on vs. off the platform in total. To obtain sufficient power for each dealer, we pool five trading days, and let the period t be a business week, rather than a single day. This implies that we cannot directly compare the point estimates in Appendix Table 12 to the estimates of our main specification in Appendix Table 1. For instance, the average of the liquidity shocks becomes a little bit smaller and the variance increases. This is the case even when the platform usage cost is common to all dealers but we let the period be a business week and not a day. Allowing the platform cost to be dealer specific, thus, does not change the other parameters significantly. But it reveals that the platform usage cost differs across dealers. One explanation for this could be that investors attach different values for maintaining a close business relationship with different dealers.

In addition, we verify that our estimates are not driven by occasionally large trade sizes (see Appendix Table 13). We do so, because our model abstracts from trade size as most trades are small and similar in size. However, occasionally, investors trade large amounts, and if they do, it is more likely that they trade bilaterally than on the platform (see Appendix Table 9). To test that our estimates are robust to these rare occasions, we re-estimate the model on a subsample of trades. More specifically, we exclude the 5% largest trades of an investor who trades more than one time. Investors who trade a single time, do not trade large amounts.

Finally, we test how sensitive our estimates are to the assumption that the dealer extracts all rents in the bilateral negotiation and sets a price that equals the investor's willingness to pay. For this, we rely on the extended model in Appendix C.3. We first re-estimate all model parameters imposing a positive bargaining power ϕ of 0.1. Then, with the parameter estimates of the extended model, we repeat the counterfactual exercise to test robustness of our main welfare findings.

To estimate the extended model, we extend the estimation procedure of the benchmark model which is explained in Appendix E.2. The key difference is that we can no longer back out the dealer's valuation for the bond from the markup equation once we have estimated all other model parameters. Instead, the estimation procedure finds the implied valuation for the dealer for each constellation of parameters until it finds the set of parameters for which all moments of Steps 2 and 3 in the estimation are matched.

We find that the most sensitive parameter to the choice of the investor's bargaining power ϕ is the investor's average willingness to pay for the bond, i.e., the average liquidity shock (see Appendix Table 14). We would expect that a buying investor who has some bargaining power would pay a price that is lower than her true willingness to pay, and vice versa for the selling investor. This is true for the point estimates of all investor groups but for selling retail investors. For this investor group standard errors are relatively large because retail investors do not sell as often in a day.

Similarly, our welfare findings go in the expected direction: Investor surplus increases slightly as we increase ϕ and dealer profits decrease (see Figure 6). The effects on total welfare and the welfare decomposition remain essentially unchanged (see Figure 7).

Appendix Figure A1: Bloomberg screen

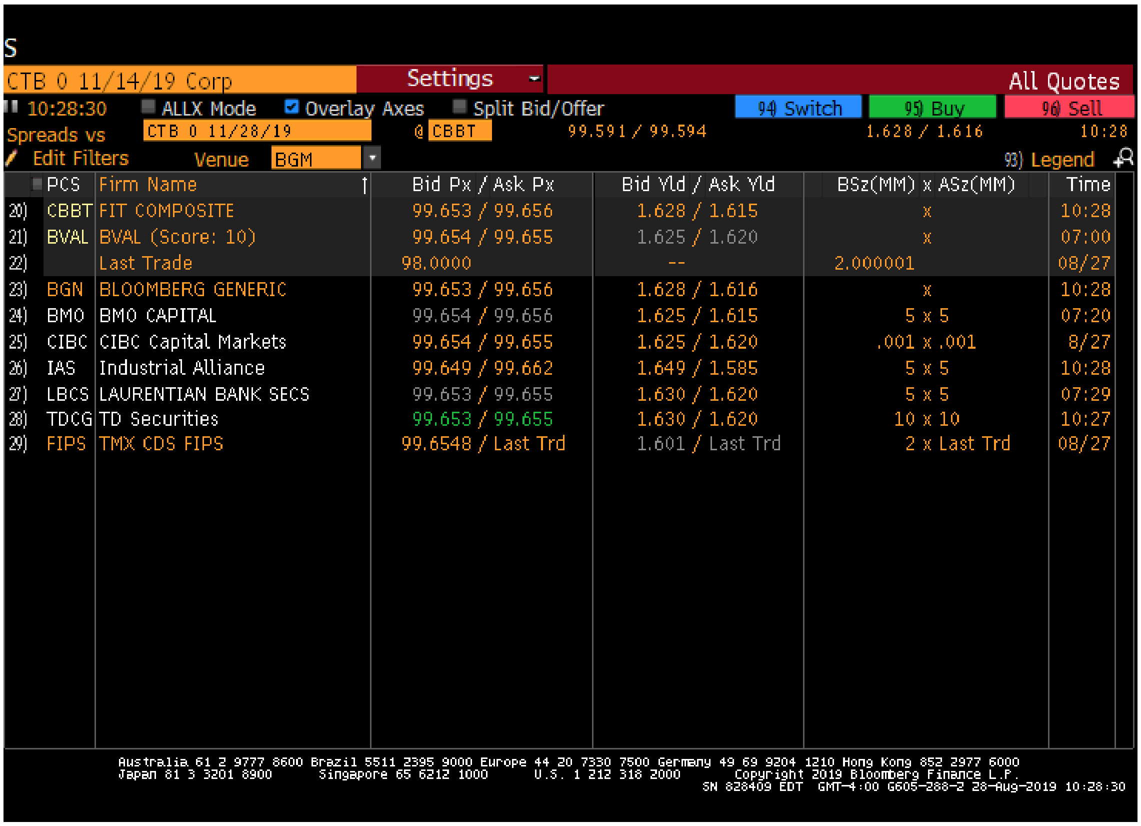

The graph depicts a screen shot of a Bloomberg Terminal that was taken on 08/28/2019. It shows the quotes for a T-Bill that matures on 11/14/19. This is shown in the upper right corner in yellow: "CTB 0 11/14/19 Corp". In the second column, we see the names of the firms that provide quotes. These quotes are shown in third and fourth column in form of prices and yields, respectively. There are two types of quotes. First, there are average quotes that serve as a benchmark for the bill's value. An example is the quote by Bloomberg Generic. Second, there are dealer-specific quotes. For instance, at 7:29 am Laurentian Bank SECS posted quotes at which it was willing to buy (bid) and sell (ask). The time is specified in the last column.

Appendix Figure A2: How investors trade

(a) With how many dealers do investors trade?

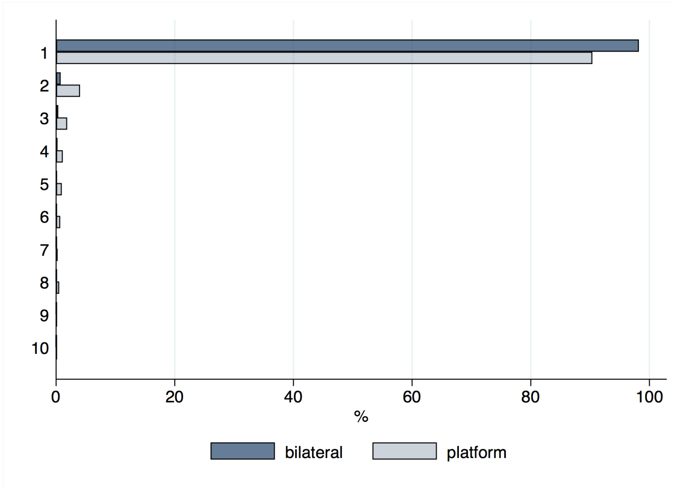

The graph shows how many investor ids trade with 1,. . .,10 dealers (bilaterally or on the platform) as a fraction of all ids. Almost all trade with a single dealer. This is partially due to data limitations as some of the ids are dealer-account numbers, and with that dealer-specific. When restricting the sample to investors with LEIs, the fraction of ids that trade with a single dealer decreases to about 70%. This is still very high, considering that these investors are very active traders in this market.

(b) Which investors trade bilaterally vs. on the platform?

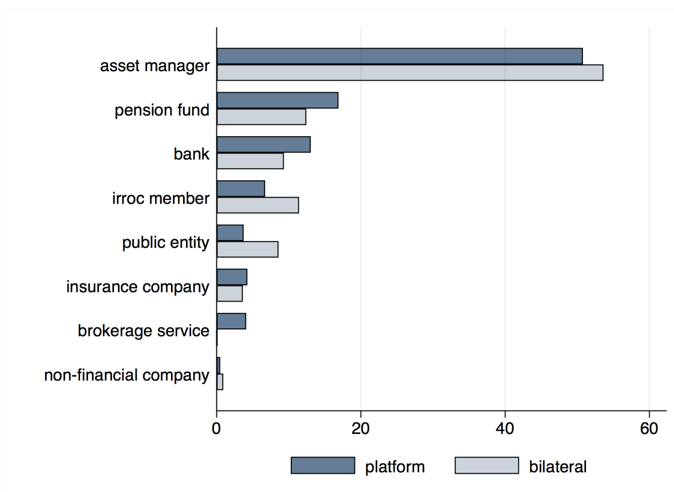

The graph shows how much each investor group (as classified in Appendix A.2) trades on the platform versus bilaterally as percentage of the total amount that institutional investors trade in the sample. The types of investors are similar across trading venues.

Appendix Figure A3: Bloomberg price ≈ inter-dealer market price

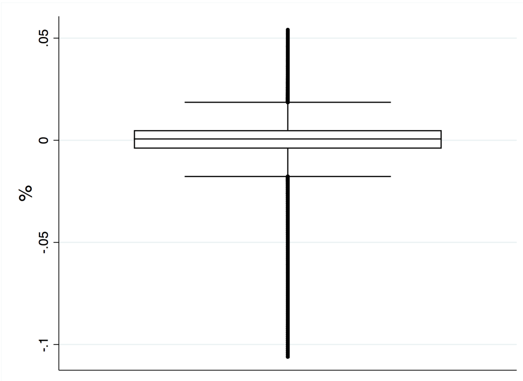

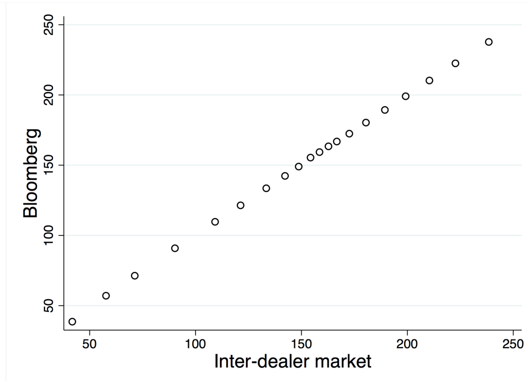

The graph on the LHS shows a box plot of the difference between the Bloomberg yield for security s in hour h of a day t and the yield that two dealers charge when trading the same security in that hour, excluding the upper and lower 1% of the distribution. The graph on the RHS shows a binned scatter plot of the Bloomberg yield and the inter-dealer market yield. The correlation is essentially one.

Appendix Figure A4: Normalization of trade yields

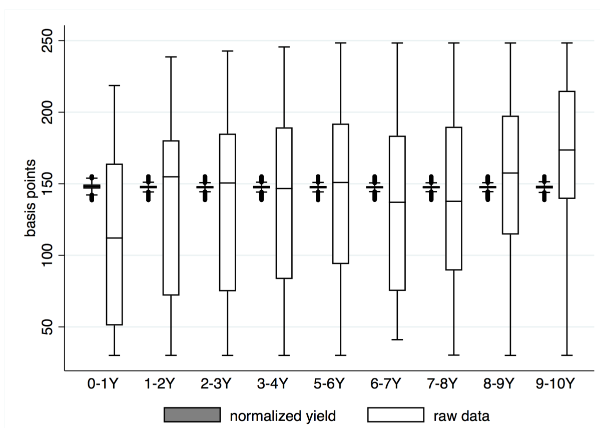

The graph shows box plots of the normalized and observed yields (of Section 3.1) in the raw data per maturity class, excluding the upper and lower 1% of each distribution. Each class pools bonds that have 0-1, 1-2, ... 9-10 years left to maturity. We see that the box plots of the normalized yields have a much smaller range than the box plots of the yields in the raw data. This is mostly because we have partialed out time trends. Further, the medians lie on a straight line because we have taken out the term-structure.

Appendix Figure A5: Event study—Observable trade behavior

(a) Trade size

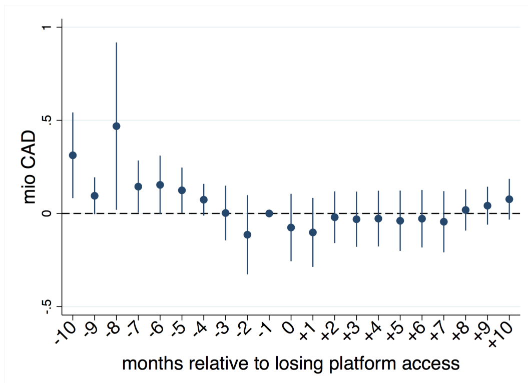

(b) Weeks-to-maturity of traded bonds

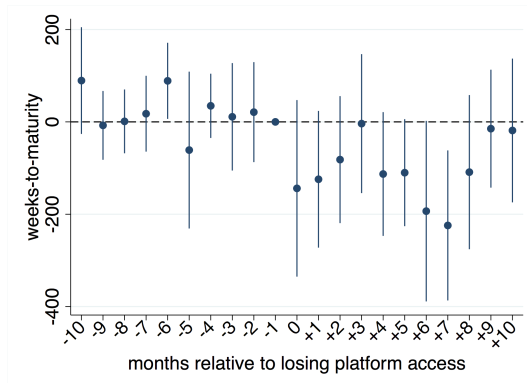

(c) Duration of traded bonds

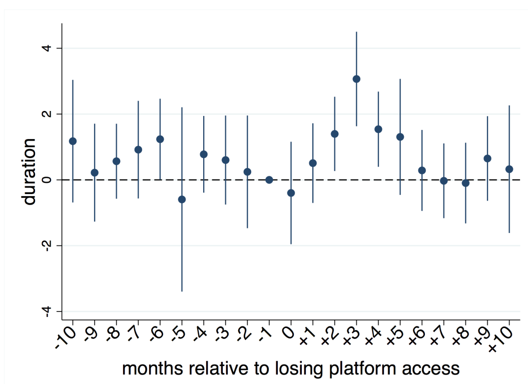

(d) Convexity of traded bonds

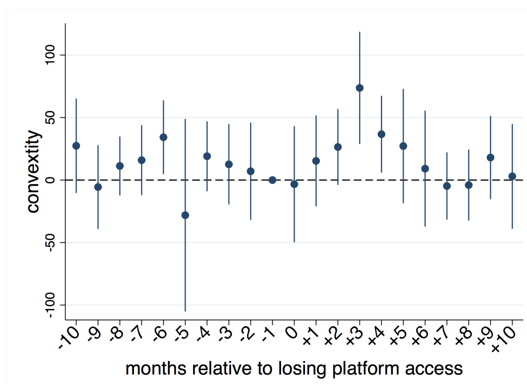

These graphs visualize changes in observable trade behavior when the investor loses platform access. They show the β_m estimates and the 95% confident intervals of regressions that take a similar form as the event study regression (2) but with outcome variables that capture trade behavior. For graph 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 if the amount traded changes. Graphs A5b - A5d illustrate if the investor trades bonds with different characteristics, namely: the length to maturity, the duration (which approximate 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 as it would absorb any characteristic of the bond. All standard errors are clustered at the investor-level.

Appendix Figure A6: Event study: Observable trade behavior per month

(a) Number of dealers per month

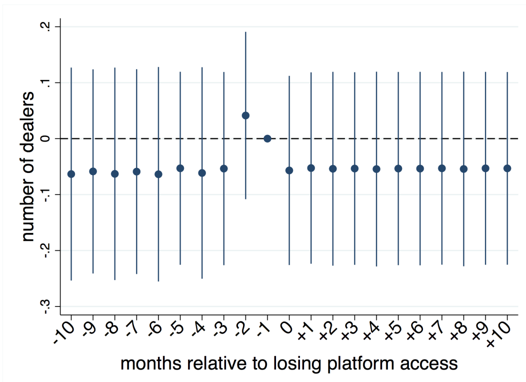

(b) Number of trades per month

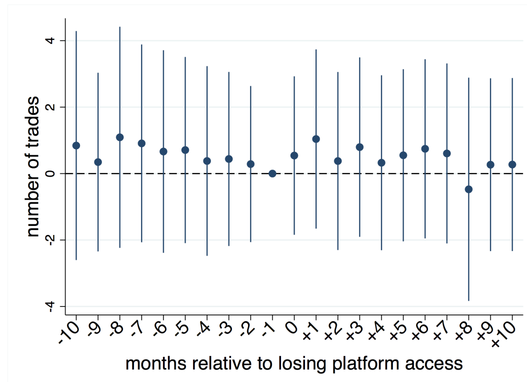

(c) Total amount traded per month

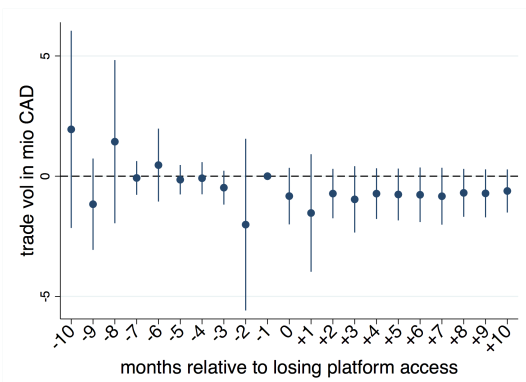

(d) Net demand per month

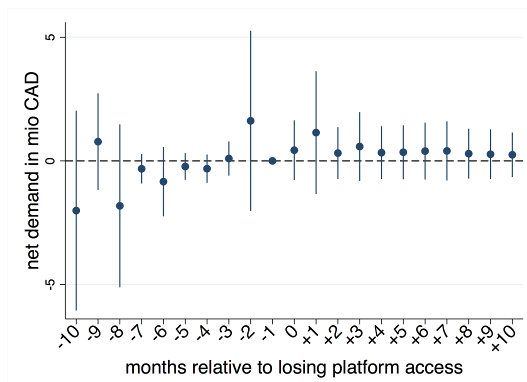

These graphs visualize whether there is a systematic change in observable trade behavior per month when the investor loses platform access. All graphs show the β_m estimates and the 95% confident intervals of a regression that takes a similar form as the event study regression (2) but aggregating observations by month m and investor i. Specifically, we regress: $outcome_{mi} = \xi_i + \sum_{m=Mi^-}^{Mi^+} \beta_m D_{mi} + \epsilon_{mi}$. In Graph A6a, the outcome variable is the number of dealers with whom investor i trades in month m, in A6b it is the number of times the investor trades in the month, in A6c it is the total amount that the investor traded in the month, and in A6d it is the investor's total net demand (which is the total amount she bought minus the total amount she sold) in the month. All quantities are measured in million C$. Standard errors are clustered at the investor-level.

Appendix Figure A7: Changes in choice probabilities in the counterfactual

(a) Probability to trade with efficient dealer (b) Platform entry due to frictions

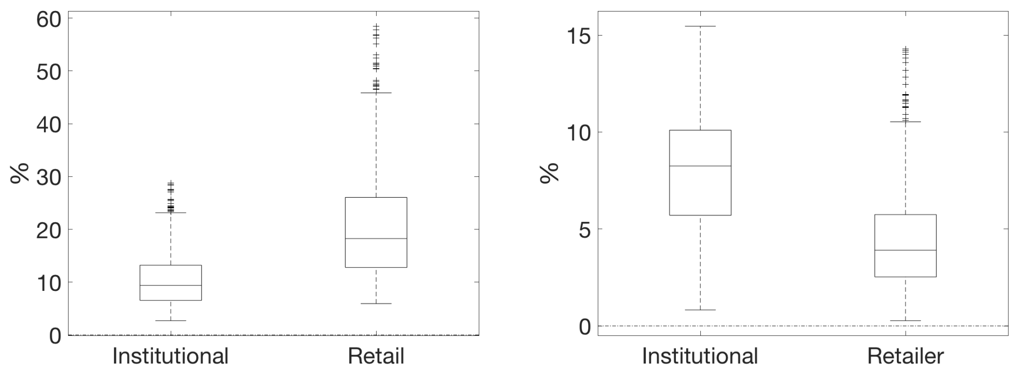

Graph A7a shows how much more likely it is that a randomly drawn institutional/retail investor buys from the efficient dealer $j_t^* \in \text{argmax}_j\{v_{tj}^D + \xi_{tj}\}$ when allowing retail investors to use the platform for free relative to the status quo. The probability is measured in percent, and the distribution is taken over all trading days. This is also true for Graph A7b which shows how much more likely it is for an investor to buy on the existing, imperfect platform when access is free relative to when access is free and the platform is frictionless ($\sigma = 0$). In all box plots the upper and lower 1% of the respective distribution is excluded.

Appendix Figure A8: Distribution of yields

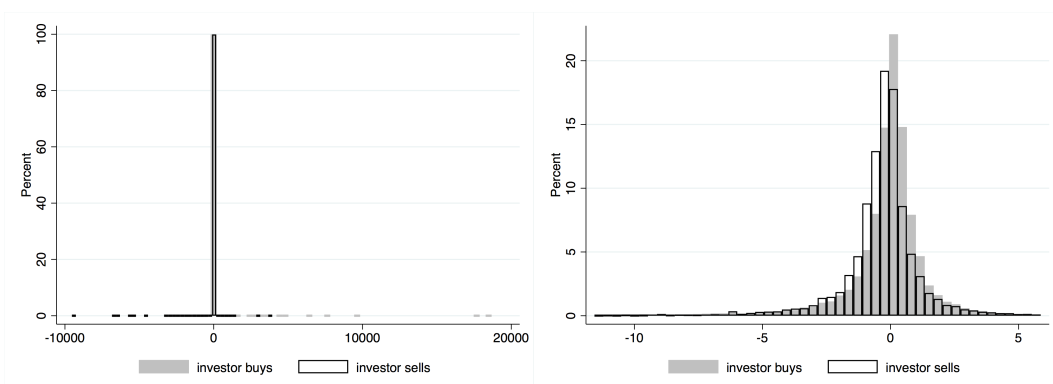

The graphs show histograms of the markup $(y_{thsij} - \theta_{ths})^+$ as defined in (1) for institutional investors. It is measured in bps. The left-sided graph includes the entire sample. The right-sided graph excludes the upper and lower 1% of the markup.

Appendix Figure A9: Market share of the bilateral market segment

(a) By number of trades

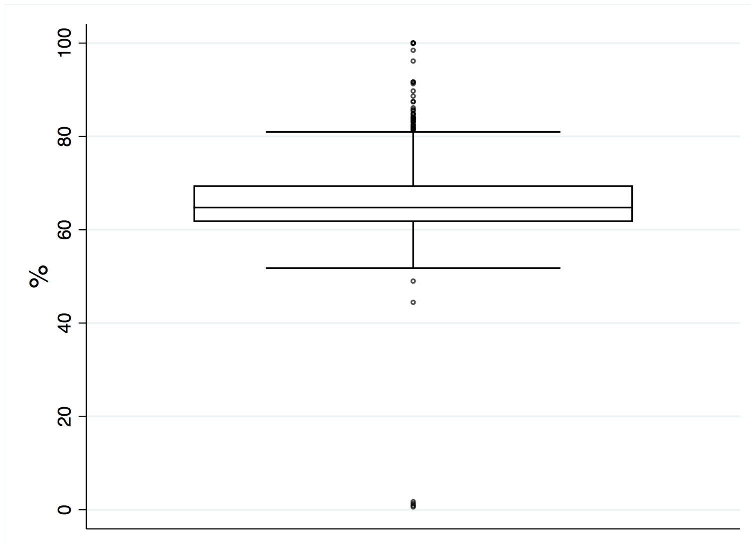

(b) By trade volume

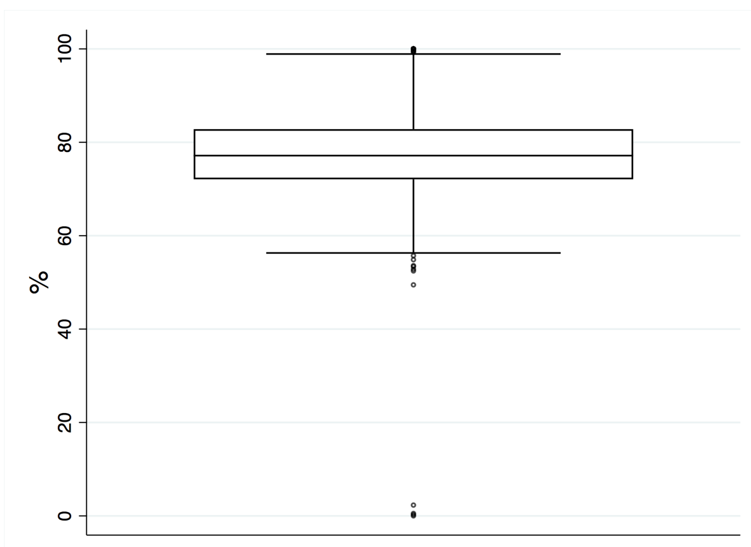

The graph on the RHS shows a box plot of the number of trades that are executed bilaterally as percentage of the total number of trades in a day. The graph on the LHS shows a box plot of the of volume that is traded bilaterally in percentage of the total volume traded on that day.

Appendix Figure A10: Distribution of trade sizes

(a) Dealer-to-investor (full distribution)

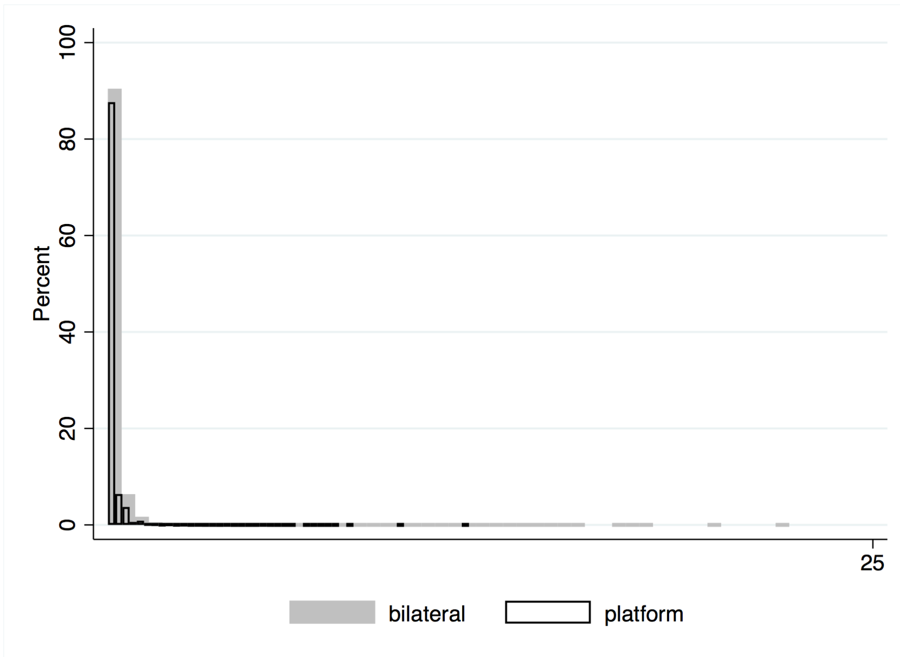

(b) Dealer-to-investor (excluding outliers)

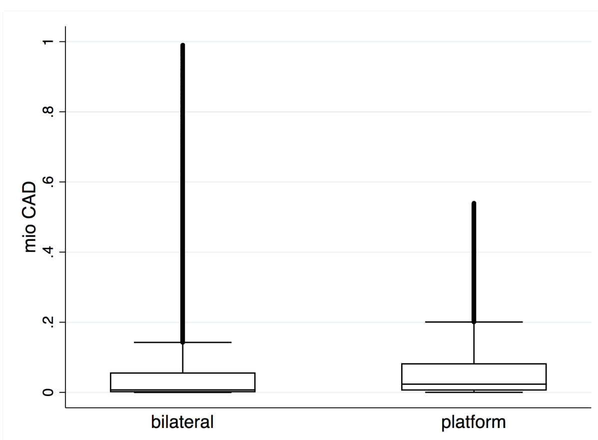

(c) Dealer-to-dealer (excluding outliers)

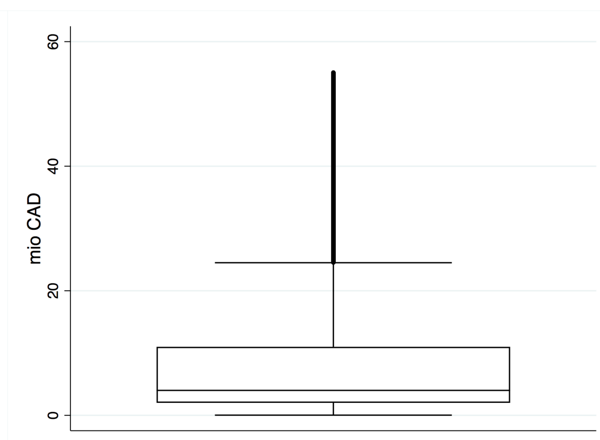

Graph A10a shows a probability density histogram of trade sizes in bilateral trades and platform trades. It illustrates that all trades are small and relatively similar in size. Graph A10b box plots of the trade size distribution which distinguish between bilateral and platform trades, and exclude the upper and lower 3% of the distribution. Graph A10c shows a box plot of the distribution of trade sizes of inter-dealer trades, also excluding outliers. We see that most sizes of dealer-to-investor trades are relatively similar across trade venues. However, in rare occasions, investors trade large amounts, and if they do they tend to trade bilaterally as the long tail in the distribution of bilateral trade sizes suggests. This paper, focusses on the typical trade which is small. All trade size is measured in million C$.

Appendix Figure A11: Do investors split orders in a day?

(a) Number of times an investor trades in a day

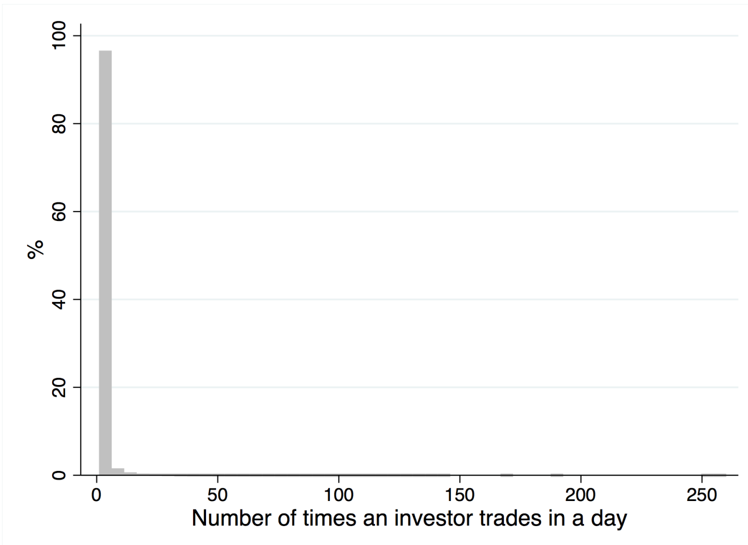

(b) Number of trading venues (bilateral vs. platform) per day

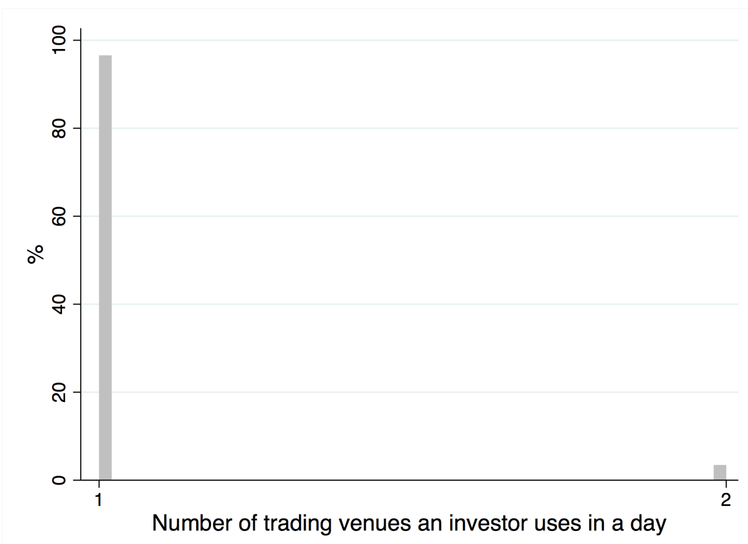

Graph A11a shows a probability density histogram of the number of times that an investor trades, i.e., either sells or buys, in a day. Graph A11b shows a probability density histogram of the number of trading venues (bilateral vs. platform) that an investor uses in a day. We see that in more than 95% of cases, the investor only trades one time, and that in the rare occasions in which the investor trades multiple times, she typically trades bilaterally and on the platform. Our model abstracts from these rare events.

Appendix Figure A12: Difference in counterfactual quotes

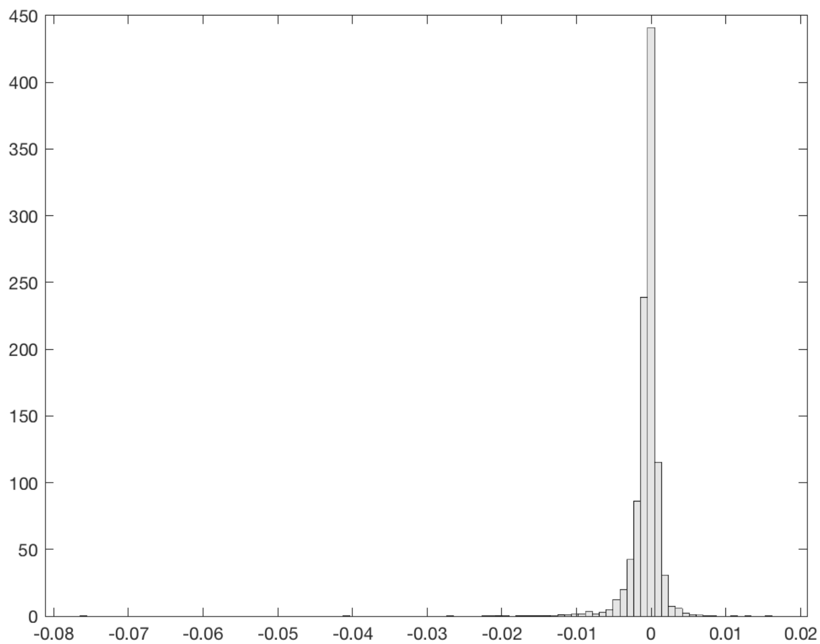

The graph shows a probability density histogram of the difference between two sets of quotes (measured in bps) that are posted for buying investors when platform access is free. The first set of quotes are the quotes that dealers choose when they are sophisticated, as characterized in (39). The second set of quotes fulfill the markup condition of Proposition 2 when liquidity shocks are drawn from $N(\mu_t, \sigma_t)$ where $\mu_t = \sum_G \kappa^G \mu_t^G$ and $\sigma_t = \sum_G \kappa^G \sigma_t^G$. The histograms look similar when platform access is universal but costly, as well as when the investor sells.

Appendix Table 4: 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	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%

This table summarizes how we restrict the raw data. We clean out extreme yields and trades by institutions that are likely reported as institutional investors but are retail or vice versa (as described in Appendix A.1). Further, we exclude in-house trades, trades that are not realized on CanDeal or bilaterally, trades that occur out of (CanDeal) business hours. For the structural estimation, we focus on trades with primary dealers only. We excludes trades after July 2019 because our auction data currently does 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 5: 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	1193999		806473	
Adjusted R^2	0.169		0.523	

To show that yields on the platform are better than off the platform, 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 (ζ_{th}), security (ζ_s), dealer (ζ_j). In column (2) we in addition add investor (ζ_i) fixed effects. Standard errors are in parentheses.

Appendix Table 6: Bilateral yields and the cost shifters

Regression	(12) buys	(12) sells	(13) buys	(13) sells
θ_{ths}	0.656	0.739		
	(0.00262)	(0.00220)		
$won_{\tilde{t}j}$	-0.00000407	0.00000969	-0.00000838	-0.00000518
	(0.0000132)	(0.0000112)	(0.0000140)	(0.0000117)
Constant	49.48	39.29	-1.478	-0.766
	(0.389)	(0.324)	(0.00770)	(0.00665)
Observations	152871	158310	152871	158310
Adjusted R^2	0.728	0.686	0.724	0.568

Appendix Table 7: Quotes and the cost shifters

Regression	(12') buys	(12') sells	(13') buys	(13') sells
θ_{ths}	0.0267	0.0368		
	(0.00118)	(0.00106)		
$won_{\tilde{t}j}$	0.0000982	0.0000424	0.0000955	0.0000522
	(0.00000593)	(0.00000541)	(0.00000505)	(0.00000466)
Constant	143.8	142.5	147.7	147.9
	(0.175)	(0.157)	(0.00276)	(0.00262)
Observations	152871	158310	199247	206783
Adjusted R^2	0.290	0.279	0.283	0.268

Table 6 shows the estimates of regression (12) and (13) for investors who buy (second and fourth column) and investors who sell (third and fifth column). Table 7 shows the analogous results when replacing the bilateral yield in both regressions with the (proxied) platform quote q_{tj}. All yields are in basis points and $won_{\tilde{t}j}$ is in million C\$. Standard errors are in parentheses.

Appendix Table 8: Effect of trade size on yields

buy	0.936	(0.126)
θ	0.823	(0.00650)
quantity	-0.0378	(0.0329)
quantity^2	-0.00605	(0.00956)
quantity^3	0.000658	(0.000542)
Constant	24.58	(0.904)
Observations	806197	
Adjusted R^2	0.999	

This table shows the estimation results when regressing the trade yield ($yield_{thsij}$) on an indicator that shows the size of the trade (buy_{thsij}), the market value (θ_{ths}), a function of trade size, $\sum_{p=1}^{3} \delta_p (quantity_{thsij})^p$, in addition to hour ($\zeta_{th}$), security ($\zeta_s$), dealer ($\zeta_j$), and investor ($\zeta_i$) fixed effects. The findings suggest that trade size is not driving the yield as all of the coefficient multiplying the functions of quantity are statistically insignificant. Standard errors are in parenthesis. They are clustered at the investor-level.

Appendix Table 9: 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	1234945	784809	738344
Adjusted R^2	0.008	0.462	0.468

Here we test whether investors trade different amounts on versus off the platform. The first column shows the estimation results when regressing an indicator for whether trade realizes on or off the platform (platform) on the trade size (quantity). In the second column, we add an investor fixed effect (ζ_i). In the last column, we exclude the 5% largest trades of an investor to show that the negative correlation between platform and quantity is driven by such occasional large trades. The coefficient becomes statistically insignificant with a p-value of 0.19. Standard errors are in parenthesis. They are clustered at the investor-level in column (2) and (3).

Appendix Table 10: 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)	✓	✓	✓	–	–	–
N	8492	8492	8492	8492	8492	8492
Adj. R^2	0.804	0.805	0.805	0.040	0.040	0.040

This table shows how the σ parameter changes depending on the instrument we use. Specifically, it gives the point estimate of regression (41) in Appendix E.2 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 that the dealer won on the most recent auction day. The fourth and seventh columns show the estimate using the simulated in (7) instrument (as reported in the text). Standard errors are in parenthesis.

Appendix Table 11: Estimates (median across days) when $y_{ti}^G = \theta_t + \nu_{ti}^G$

buys	$\hat{\mu}^I$	$\hat{\mu}^R$	$\hat{\sigma}^I$	$\hat{\sigma}^R$	$\hat{c}$	$1/\hat{\sigma}$	$\hat{\eta}$
	-0.87	-2.97	2.56	5.05	3.34	1.29	$+174.13$
	(0.13)	(0.73)	(0.12)	(0.94)	(0.18)	(0.94)	
sells	$\hat{\mu}^{I^*}$	$\hat{\mu}^{S^*}$	$\hat{\sigma}^{I^*}$	$\hat{\sigma}^{S^*}$	$\hat{c}_t^*$	$1/\hat{\sigma}$	$\hat{\eta}^*$
	$+0.97$	$+2.04$	2.62	4.50	3.46	1.29	-174.23
	(0.14)	(0.64)	(0.11)	(0.98)	(0.18)	(0.94)	

This table is the analogue to Table 1, but here we assume that $y_{ti}^G = \theta_t + \nu_{ti}^G$. It shows the median over all days of all point estimates per investor group G, in addition to the implied elasticity of demand ($\hat{\eta}$) and of supply ($\hat{\eta}^*$) on the platform, averaged across days and dealers. The corresponding median in the standard errors are in parentheses. All estimates are measured in bps.

Appendix Table 12: Estimates when the platform usage cost is dealer specific

buys	$\hat{\mu}^I$	$\hat{\sigma}^I$	$\hat{c}_1$	$\hat{c}_2$	$\hat{c}_3$	$\hat{c}_4$	$\hat{c}_5$	$\hat{c}_6$	$\hat{c}_7$	$\hat{c}_8$	$\hat{c}_9$	$\hat{\eta}$
	-0.77	2.67	3.04	3.49	2.71	4.06	3.74	2.95	1.67	3.23	3.56	$+173.82$
sells	$\hat{\mu}^{I*}$	$\hat{\sigma}^{I*}$	$\hat{c}_1^*$	$\hat{c}_2^*$	$\hat{c}_3^*$	$\hat{c}_4^*$	$\hat{c}_5^*$	$\hat{c}_6^*$	$\hat{c}_7^*$	$\hat{c}_8^*$	$\hat{c}_9^*$	$\hat{\eta}^*$
	$+0.89$	2.76	3.42	3.72	2.82	4.00	3.66	3.34	1.97	1.97	3.42	-173.95

This table is similar to Table 1, but here we assume that the platform usage cost is dealer specific and count business week rather than a day as a period t. The table shows the median over all weeks of all point estimates for institutional investors I, in addition to the implied elasticity of demand ($\hat{\eta}$) and of supply ($\hat{\eta}^*$) on the platform, averaged across days and dealers. All estimates are measured in bps.

Appendix Table 13: Estimates (median across days)

buys	$\hat{\mu}^I$	$\hat{\mu}^R$	$\hat{\sigma}^I$	$\hat{\sigma}^R$	$\hat{c}$	$1/\hat{\sigma}$	$\hat{\eta}$
	-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	$\hat{\mu}^{I^*}$	$\hat{\mu}^{S^*}$	$\hat{\sigma}^{I^*}$	$\hat{\sigma}^{S^*}$	$\hat{c}_t^*$	$1/\hat{\sigma}$	$\hat{\eta}^*$
	$+0.91$	$+1.99$	2.86	4.48	-3.46	1.36	-184.51
	(0.13)	(0.67)	(0.10)	(0.94)	(0.17)	(0.25)	

This table 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 ones. 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 ($\hat{\eta}$) and of supply ($\hat{\eta}^*$) on the platform, averaged across days and dealers. All estimates are measured in bps.

Appendix Table 14: Estimates of the extended model ($\phi = 0.1$)

buys	$\hat{\mu}^I$	$\hat{\mu}^R$	$\hat{\sigma}^I$	$\hat{\sigma}^R$	$\hat{c}$	$1/\hat{\sigma}$	$\hat{\phi}$
	-0.96	-3.31	3.08	5.61	3.38	1.29	$+174.57$
sells	$\hat{\mu}^{I^*}$	$\hat{\mu}^{S^*}$	$\hat{\sigma}^{I^*}$	$\hat{\sigma}^{S^*}$	$\hat{c}_t^*$	$1/\hat{\sigma}$	$\hat{\phi}^*$
	$+0.44$	$+2.24$	2.37	5.00	3.69	1.29	-179.59

Appendix Table 15: Estimates of the extended model ($\phi = 0.25$)

buys	$\hat{\mu}^I$	$\hat{\mu}^R$	$\hat{\sigma}^I$	$\hat{\sigma}^R$	$\hat{c}$	$1/\hat{\sigma}$	$\hat{\phi}$
	-1.23	-4.17	3.59	6.72	3.27	1.29	$+174.42$
sells	$\hat{\mu}^{I^*}$	$\hat{\mu}^{S^*}$	$\hat{\sigma}^{I^*}$	$\hat{\sigma}^{S^*}$	$\hat{c}_t^*$	$1/\hat{\sigma}$	$\hat{\phi}^*$
	$+0.71$	$+2.90$	2.79	5.98	3.69	1.29	-178.67

Appendix Table 16: Estimates of the Benchmark Model ($\phi = 0.5$)

buys	$\hat{\mu}^I$	$\hat{\mu}^R$	$\hat{\sigma}^I$	$\hat{\sigma}^R$	$\hat{c}$	$1/\hat{\sigma}$	$\hat{\phi}$
	-2.01	-6.59	5.13	10.07	3.06	1.29	$+174.44$
sells	$\hat{\mu}^{I^*}$	$\hat{\mu}^{S^*}$	$\hat{\sigma}^{I^*}$	$\hat{\sigma}^{S^*}$	$\hat{c}_t^*$	$1/\hat{\sigma}$	$\hat{\phi}^*$
	$+1.53$	$+4.81$	4.07	8.94	3.62	1.29	-177.11

These tables are similar to Table 1 but of the extended model in which the investor has a bargaining power of $\phi = 0.1, \phi = 0.25$ or $\phi = 0.5$. The tables shows the median over all days of all point estimates per investor group G, in addition to the implied elasticity of demand ($\hat{\eta}$) and of supply ($\hat{\eta}^*$) on the platform, averaged across days and dealers. All estimates are measured in bps.