Learning by Trading:

The Case of the US Market for Municipal Bonds*

Giulia Brancaccio†, Dan Li‡, Norman Schürhoff§

January 2020

Abstract

We study information acquisition as a motive for trade in decentralized markets. The setting is the secondary market for US municipal bonds, a main source of financing for local governments. There is no central exchange for municipal bonds, and regional dealers intermediate trades among retail investors. We provide novel facts suggesting that dealers leverage their trading activity to learn about the investors' demand for municipal bonds. We then design and estimate a dynamic model for this market that focuses on dealers' incentives to trade. Our estimates reveal that dealers are willing to pay 12% of the intermediation spread (7 bp) for the information they acquire by trading with one investor. Dealers' motives for trade interact with policies that increase the availability of public information about trading activity in decentralized markets. Intended to improve liquidity, these policies are the object of an intense regulatory activity in the US and abroad. In our counterfactual analysis we show that these policies can weaken dealers' incentives to acquire information. This can decrease trade activity by up to 8% and reduce investors' welfare by up to 10%. Finally, we exploit a policy improving transparency that was implemented in the market for municipal bonds to validate our approach.

*I am deeply indebted to my advisors, Bo Honoré, Myrto Kalouptsidi, and Jakub Kastl, for their continued guidance and support. For helpful comments and discussions at various stages of this project we thank Panle Jia Barwick, Nick Buchholz, Sylvain Chassang, Alessandro Gavazza, Jean-François Houde, Jan De Loecker, Tommaso Denti, Adam Kapor, Eduardo Morales, Motohiro Yogo, Batchimeg Sambalaibat, Clara Santamaria, Molly Schnell, Theodore Papageorgiou, Claudia Robles, and Pierre-Olivier Weill.

†Cornell University, Department of Economics, 464 Uris Hall, Ithaca NY 14850, USA. Email: giulia.brancaccio@cornell.edu

‡Board of Governors of the Federal Reserve System, Washington, DC 20551, USA. Email: dan.li@frb.gov

§Faculty of Business and Economics and Swiss Finance Institute Université de Lausanne CH-1015 Lausanne, Switzerland. Email: norman.schuerhoff@unil.ch

1 Introduction

In many markets, there are strong incentives to acquire information. When agents can learn from their actions, they may face a trade-off between experimenting to acquire information and choosing (myopically) the most profitable action. This trade-off is particularly relevant in markets with decentralized trade, where there is limited information about prices and trade volume. In such cases, agents often rely on their own trading activity to gather information about market fundamentals: negotiating with others reveals information about one counterpart's valuation of the asset, which, in turn, provides information about the overall state of the market. If information is sufficiently valuable, agents may find it optimal to engage in more informative, but less profitable trades. In this paper, we empirically study information acquisition as a motive for trade in the context of a decentralized market with trades totaling over $3 trillion a year: the secondary market of US municipal bonds.

States and municipalities throughout the United States depend on the municipal bond market to raise funds for investments in schools, highways, and other public projects. There is no central exchange to trade municipal securities. Instead, financial institutions registered as broker-dealers intermediate trades among investors. This paper focuses on dealers' incentives to trade and their implications for liquidity and the well-functioning of the market. We show that information acquisition is an important driver of trading activity in this market. Dealers' information-seeking incentives influence the effectiveness of regulatory interventions in this market. We consider policies improving access to public information about trade activity. Intended to improve liquidity in decentralized financial markets, these policies are the object of a lively debate and intense regulatory activity both in the US and abroad. We show that improving access to public information might weaken financial intermediaries' motives to acquire information and, ultimately, reduce their incentives to provide liquidity. This result highlights an unintended consequence of market transparency that may thwart the benefits of these policies.

Several features of the market for municipal bonds make it a natural candidate to study the interaction between trading and information acquisition. First, the lack of a central exchange for municipal securities implies that learning about trading prices requires participating in a trade directly. Second, a large number and variety of bonds are outstanding at any given time, and each asset includes a variety of complex special provisions. The lack of standardization of the assets complicates pricing and makes information acquisition a first-order issue. Finally, in recent years, the regulatory body for the municipal bond market has taken a number of concrete steps to improve access to information about market activity for market participants.

These changes allow us to directly test the importance of information acquisition in explaining the behavior of financial institutions active in the market.

We leverage a detailed dataset of transactions on the secondary market of municipal bonds to uncover two novel facts that illustrate the relevance of information acquisition for motivating trade. First, we find that the correlation in the pricing behavior of two dealers increases after they have traded with one another. A variety of placebo tests suggest that this result is not spurious and indicates that this change in behavior is likely driven by information acquisition. Trade, therefore, can be a source of valuable information.

Second, we study the impact of a policy that increases access to public information about trading prices in the municipal bond market. To do that, we leverage the idea that improving transparency will have stronger consequences for assets for which incomplete information is more severe, which we proxy using uninsured assets. The results suggest that trade between dealers and investors for uninsured assets falls compared to insured ones. Moreover, the policy impacts the cost of issuing debt for municipalities, as yields for uninsured assets increase compared to insured ones. Information acquisition motives for trade can explain this surprising result. Greater availability of public information about market activity may weaken dealers' incentives to acquire information and reduce their incentives to trade. This, in turn, can lead to a decline in market liquidity.

With these facts in mind, we build a dynamic model of trading in decentralized markets that focuses on the dealers' incentives to trade. We build on an inventory model of dealers' behavior. Forward-looking dealers accumulate a costly inventory of municipal bond holdings by trading with other dealers as well as short-lived retail investors. Investors' valuations for the asset change over time due to a persistent common demand shock. The demand shock is unobserved, and there is no public information about it. This uncertainty about the demand for municipal bonds introduces incentives to acquire information in the dealers' decision to trade. Indeed, trading is costly but trading prices are informative about the unobserved demand shock. When retail investors are involved, each trade has the same information content, but trading with more investors allows the dealer to sample more observations. When other dealers are involved, trading prices are informative about the counterpart's valuation of the asset, which in turn reveals what he knows about the state of the market. Some dealers have better information than others due to their trading history. To capture the decentralized nature of trade, we allow dealers to only observe a summary statistic of the past trading activity of their peers, which we denote as "experience."

We estimate the model using data on municipal bonds trade before the introduction of market transparency. The estimation proceeds in two steps. A dealer's trading decisions depend on their type, which includes their information about the demand shock and their experience, and is unobserved by the econometrician. In the first step of our estimation we recover this unobserved type. To recover the dealers' experience we exploit an implication of the model: more experienced dealers will pay lower prices to buy assets in the inter-dealer market. Our baseline specification compares how prices for trades executed by a specific seller, in a specific month and asset, change depending on the trading history of the buyer. We focus on comparisons for a fixed month and seller to ensure that the estimates are robust with respect to market-wide shocks.

To recover the dealers' information about the demand shock we assume that dealers only acquire information through trade or through public signals observed by everyone, the econometrician included. This assumption allows us to directly trace out the dealers' information set at every point in time. We exploit a Hansen-Sargan test for over-identifying restrictions to test for additional unobserved sources of information available to the dealers. The test suggests that dealers have no information about the demand for an asset in periods in which they do not trade the asset. This confirms that learning activities in the market for municipal bonds are strongly connected to "realized" trade, justifying our approach.

We then estimate the remaining primitives, which include dealers' trading and inventory costs, as well as investors' valuations and their entry costs. We recover dealers' costs from their optimal choice probabilities, via method of simulated moments, following the dynamic discrete choice literature (Rust (1987)). Then, we obtain investors' valuations directly from observed prices.

The estimates show that for the average dealer, the information content of a trade with an investor is worth 12.41% of the average intermediation spread (i.e., the difference between the price at which they buy and sell the asset), corresponding to 7 basis points. We also find that the value of this information is higher for assets with more volatile market prices as well as riskier assets. Finally, we find that the information acquired through trade allows dealers to increase the precision of their estimate of the asset's market value by 30%.

We use the estimated model to explore the impact on dealers' incentives to trade of policies improving the availability of public information about trade activity. We first look at the impact of these policies on market liquidity. This is an interesting outcome variable per se, as historically it has been the target of policies addressing the inefficiencies of decentralized financial markets. We find that improving market transparency can reduce the volume of trade for an asset up to 8%. Two effects are at play. On the one hand, transparency naturally improves trading volume because it lowers uncertainty; on the other, it weakens the dealers' incentives to acquire information, thus reducing trade activity and counteracting, possibly overriding, the positive effect of transparency on liquidity.

The net effect of these two forces varies across different types of assets, depending on the underlying primitives. We leverage this heterogeneity to validate our model. In particular, we compare the models' predictions for the impact of the introduction of market transparency with its observed impact. We find that our model, through its focus on the dealers' incentives to trade, captures 51% of the heterogeneity in the effect of the policy on trade activity across different types of assets. Furthermore, the data is consistent with model's prediction for the impact of the policy on the intermediation spread, trading behavior for different types of dealers, and the differential reaction to the policy of trade with investors and inter-dealer trade.

Finally, we exploit our model to trace out how the predicted changes in trade activity translate into changes in the agents' welfare. Better information allows the dealers both to improve the timing of their trades and to spare the costs involved in trading to acquire information. Both this forces contribute to improving dealers' profits by 2.7%. Instead, investors' welfare can decline by up to 10% following the introduction of market transparency. Indeed, under market transparency the dealers exercise their market power by curbing their trading activity. Therefore, there are investors who would profitably exchange an asset, but are not served in a transparent market. By counteracting this effect, information acquisition motives for trade can benefit investors and reduce the assets' misallocation.

In conclusion, we shed new light on the role of experimentation in decentralized financial markets. We argue that in these markets trade can be a source of valuable information about the market fundamentals. Obtaining this information, therefore, becomes an additional motive for trade. While information acquisition motives for trade are often overlooked, we show that they affect the impact of regulatory intervention in decentralized financial markets. In particular, they generate a channel through which market transparency may have an unintended impact on market liquidity. The policy maker may want to take these repercussions into account to more completely assess the impact of market transparency in opaque OTC markets.

Related literature. This paper is at the intersection of three principal strands of the literature. The basic trade-off between learning and sacrificing immediate payoff is focal in the literature on strategic experimentation. We empirically quantify the strength and implications of this basic trade-off in the context of decentralized financial markets. Therefore, our paper contributes to the literature that studies asset prices and allocations in OTC markets. Finally, we integrate these concepts with ideas from empirical studies of industry dynamics.

Experimentation has long been studied in economics, mostly from a theoretical standpoint (for a survey, see Hörner and Skrzypacz (2016)). Several papers within this literature explicitly share our focus on experimentation as a motive for trade —most notably Aghion et al. (1993), Grossman et al. (1977), Mirman et al. (1993), and Kihlstrom et al. (1984). Our focus remains an empirical one. For this reason, we strip the incentives to experiment to their minimal components. This makes the agents' problem tractable, allowing us to bring the model to the data.

Several papers, such as Leach and Madhavan (1993) and Bloomfield and O'Hara (1999, 2000), have discussed the implications of experimentation for the trading behavior of agents in financial markets. Furthermore, Wolinsky (1990), as well as Golosov et al. (2014), and Blouin and Serrano (2001) explore the linkages between trading and information diffusion in a decentralized market with private information. Their objective is to study under what conditions all relevant information is revealed over time. Despite this interest, direct quantification of the role of experimentation and measurement of its implications for market structure have remained scarce. We contribute to this literature by employing a tractable analytical framework to empirically study the role of incentives to experiment as a motive for trade.

We integrate the literature on experimentation with the literature studying the trading behavior of agents in decentralized markets. In particular, our model draws from a rich tradition of papers using search models to study asset prices and allocations in OTC markets. Early papers include Gehrig (1993), Spulber (1996), and Rust and Hall (2003). Most recent papers build on the framework of Duffie et al. (2005, 2007), to study the implication of search frictions for the functioning of OTC markets. Recent works include Gavazza (2011b,a), Hugonnier et al. (2012), Hugonnier et al. (2014a), Farboodi et al. (2016), Glode and Opp (2016), Hugonnier et al. (2018) and Neklyudov (2019). We focus on a different feature of decentralized markets compared to these works: the lack of public information about trade activity.[1] For this reason, we borrow the structure of these models and enrich it with incomplete information and learning: in our setup the decision to trade depends not only on inventory management and search costs but also on experimentation.[2] Furthermore, traditionally this literature models meeting among agents as random. Instead, in our model we allow dealers to direct their offer to a specific counterpart.

Our paper is also related to the growing empirical literature on OTC markets. Bessembinder et al. (2006), Harris and Piwowar (2006), and Green et al. (2006), study transactions cost and price discovery

[1] Duffie and Manso 2007 have a similar focus, but they focus on information diffusion rather than information acquisition.
[2] Few papers in this literature allow for asymmetric information. A recent example is Lester et al. (2018).

in OTC markets. Moreover, Li and Schürhoff (2019), Di Maggio et al. (2017), and Hollifield et al. (2017) document a number of stylized facts both about the structure of network defined by inter-dealer trades and about how dealers' trading prices depend on their position in this network. While we don't target explicitly this stylized facts, in Appendix F.1, we show that our model can give rise to this facts in equilibrium in response to dealers' incentives to acquire information.[3]

Finally, our paper is related to the literature on industry dynamics (e.g., Hopenhayn (1992), Ericson and Pakes (1995)) which characterizes Markov-perfect equilibria in entry, exit, and investment choices given some uncertainty in the evolution of the states of firms and their competitors. Instead of these choices, we model the agents' problem as a series of trading and pricing decisions. Since agents interact with one another repeatedly, this problem generates a particularly high-dimensional state space. We introduce a number of innovations to mitigate the computational burden. First, to simplify trading decisions, we assume that agents only observe a summary statistic of the trading history of other market participants. This assumption not only permits solving for equilibrium policies of agents and simplifies the agents' inference, but also reflects a more realistic behavioral model for decentralized markets. Moreover, given the large number of dealers in the market, we assume that the distribution of dealers' private states is perfectly forecastable by agents, conditional on the demand shocks they are trying to learn. This approach has precedent in the literature on firm dynamics (Weintraub et al. (2008)). Finally, our empirical methodology borrows from the literature on the estimation of dynamic setups (e.g., Rust (1987), Aguirregabiria and Mira (2007), Bajari et al. (2007), and Pakes et al. (2007)) in exploiting conditional choice probabilities to obtain information on the value functions and, in turn, on the primitive of interest.

2 Institutional setting and data

2.1 The secondary market for municipal bonds

Municipal bonds are debt securities issued by states, cities, and other local governments to fund day-to-day obligations and to finance capital projects. Municipal bonds total $3.7 trillion in principal and $300 billion of municipal debt is issued every year. Their importance cannot be overstated: in 2017 municipal bonds were the main source of funding for 75% of the public investment in infrastructure in the US.

We focus on the secondary market for municipal bonds, where these assets can be traded after issuance. There is no central exchange for municipal securities. Instead, the market is decentralized, and financial

[3] See Babus and Kondor (2018) for an alternative model in which information acquisition gives rise to these stylized facts.

institutions registered with the SEC as dealers intermediate trades among investors. Dealers execute nearly all transactions in a "principal capacity," buying the assets directly and holding them in inventory until they find a buyer.[4] Every year there are more than 2,000 active broker-dealers, and the largest market share is around 10%.

Dealers mainly interact with retail investors. Indeed, to ease credit access for local governments, interest accrued on municipal bonds is exempt from individual income taxes both at the federal and the local level. Due to this obvious tax advantage, the large majority of the municipal debt outstanding is held by private investors either directly or through mutual funds.[5] Consistent with this, trades on the secondary market are small: the median trade is worth $25,000, and 80% of trades have a value of less than $100,000.

Municipal bonds are considered "buy-and-hold" assets, as investors tend buy them at issuance and hold them until maturity. Indeed, municipal bonds are considered to be a relatively safe investment[6], which limits speculation. Instead, investors' participation in the secondary market is mostly driven by liquidity needs –factors such as retirement or alternative investments.

Despite the low default risk, the secondary market for municipal bonds is highly illiquid and volatile. As an example, for the average dealer the purchasing price of a given asset ranges between 97% and 102% of the average sale price (respectively at the 10th and 90th percentile) within a year. This corresponds to fluctuations in the price of the asset corresponding to almost twice the average yearly return for a municipal bonds.

The main driver of the illiquidity of the market for municipal bonds is the lack of depth in the market. First, demand for municipal bonds is segmented geographically due to the tax treatment of these assets. Indeed, most states exempt interest earned from municipal bonds initiated within their borders and tax the earnings from out-of-state municipal bonds. Moreover, municipal bonds are notably heterogeneous. Over our sample period there are 1.5 million different assets outstanding, issued by more than $50,000$ different units of state and local governments.[7] Furthermore, most municipal bonds are equipped with various types of special provisions, ranging from nonstandard interest payment frequencies to embedded derivatives, that further decrease standardization.[8] Naturally, dealers aim to exploit fluctuations in the

[4] As opposed to trades executed in principal capacity, we can classify as purely intermediated trades those trades where a dealer buys and sells a given asset within five minutes. These trades represent less than 5% of total trades in our dataset.

[5] Based on the Federal Reserve Flow of Funds, in 2004 private households owned 53% of the existing municipal securities directly, and 26% through mutual funds. Insurance companies held the remaining 20% of municipal securities.

[6] As an example, for Aa- and A- rated municipal bonds, the 10-year cumulative default rate is .03% compared to 0.8% for corporate bonds.

[7] For comparison, this is 20 times the number of corporate bond types.

[8] As an example, callable bonds are redeemable by the issuer before maturity, while a sinking fund provision requires the

market prices of municipal bonds to "buy low and sell high." To this end, the dealers need to anticipate the liquidity needs of a very selected group of investors interested in holding municipal bonds.[9] In this paper, we study how dealers' incentives to learn what "their firms' clients want" affect their trading behavior.

2.2 Market transparency

Due to the lack of a centralized trading mechanism and the lack of standardization, agents don't have access to direct information about trading activity in the market for municipal bonds. Moreover, while several market indexes are available, they are too coarse to effectively reduce the uncertainty for the pricing of individual bonds.[10] However, access to public information about trade activity has improved in recent years, due to a steady push from the SEC to improve market transparency. In particular, on June 23, 2003, the MSRB started distributing daily summaries about the trading activity in the market during the previous day ("next -day reporting"). Moreover, starting on January 31, 2005, the MSRB mandated that details of all transactions in US municipal bonds be reported on a timely basis and posted online almost immediately.[11] Investors seem to have embraced the new source of information with enthusiasm: on the first day of 15-minute trade reporting, the Bond Market Association reported that the website on which trades were reported averaged about 10,000 visits per minute.[12]

The stated objective behind policies improving market transparency is to increase the assets' liquidity by improving investor participation and trade activity.[13] An asset is considered more liquid "if it is more certainly realizable at short notice without loss." Therefore, liquidity is valuable per se, as long as investors value immediacy. Moreover, a liquid secondary market is a crucial condition to lower the cost of raising capital. As an example, Wang et al. 2008 estimate that municipal bond issuers pay $13 billion a year to compensate investors for the risks implied by the illiquidity of the market. Increasing liquidity in this market, therefore, would translate to huge savings for local governments and municipalities.

issuer to retire a portion of debt each year. Furthermore, nonstandard interest payment frequencies are not uncommon, and most assets have some form of credit enhancement. In fact, Harris and Piwowar 2006 show that 86%of the outstanding assets contain at least one of these special provisions.

[9] As an example, from Feldstein et al. (2008), "Municipal traders must have a sense of the kinds of bonds their firms' clients want, which means they must have a clear and consistent communication with the salesforce, which in turn must maintain open lines with the investor clients."

[10] Since 1995 the Municipal Securities Rulemaking Board (henceforth, MSRB) has published information about the volume of trade and average trading price for assets traded more than four times during the previous day, which covers around 5% of the assets traded. Moreover, several widely watched municipal bond indexes are compiled by "*The Bond Buyer*." These indexes are based either on dealers' estimates for the price of a hypothetical bond or on the activity in the primary market.

[11] Asquith et al. 2013 study the effect of a similar policy intervention in the market for corporate bonds, and find similar results.

[12] See Schultz 2012.

[13] Testimony of Chairman Arthur Levitt Before the House Subcommittee on Finance and Hazardous Materials, Committee on Commerce, Concerning Transparency in the United States Debt Market and Mutual Fund Fees and Expenses.

Changes to the quality of public information about trade activity interact with dealers' incentives to acquire information. Indeed, these policy interventions offer a laboratory to observe information acquisition motives for trade in action, which we leverage in Section 3. Furthermore we use this policy intervention to validate our model: in Section 8 we show that information acquisition motives for trade explain a large fraction in the variation of the effect of market transparency across states. The regulatory push toward market transparency is a widespread trend across different financial markets, both in the US and abroad.[14] This intense regulatory activity concerning transparency in decentralized financial markets motivates us to study, in Section 8, how dealers' information acquisition motives for trade interact with market transparency.

2.3 Data

Our main data source is the proprietary Transaction Reporting System audit trail from the MSRB. In an effort to improve market transparency, the MSRB has required dealers to report all transactions in municipal securities since 1998. The transactions data cover the 6-year period from January 2000 to December 2005. For every transaction involving municipal bonds, our data provide information about the terms of the trade, such as the trading price, date and time of the trade as well as par value (the value at maturity of the asset exchanged, or the volume of the trade) of the asset, and an asset identifier. Significantly, we observe identifiers for the dealer firm intermediating each trade: for customer trades, the data identify the dealer buying and the dealer selling the bond, while for trades among dealers, the data identify the dealers on each side of the trade. In addition to the comprehensive transactions data, we obtained reference information on all municipal bonds, including issuance date, maturity, coupon, taxable status, ratings, call features, issue size, and issuer characteristics from Thomson SDC. Finally, we obtain the time series for market bond indexes, as well as monthly municipal mutual fund flows from Bloomberg.

We filter the transactions to eliminate data errors and ensure data completeness. For a bond to be in our sample, it must have complete descriptive data in the SDC and satisfy a number of trade-specific filters and bond-specific filters (fixed or zero coupon, non-derivative, non-warrant, not puttable, maturity $\geq$1 year, \$5K denomination). Since this paper focuses on the secondary market, we remove all trades during the first 90 days after issuance and less than one year away from maturity.[15]

[14] As an example, in 2002 a similar provision was imposed by the Financial Industry Regulatory Authority (FINRA) in the market for US corporate bonds. After that, agency-backed securities, asset-backed securities, and 144A transactions followed. Moreover, in July 2018 FINRA began requiring its member firms to report US Treasury securities transactions, even though those prices are currently not disseminated to the public. Europe is also in the process of instituting a post-trade transparency regime which will affect a broad range of instruments.

[15] As a result of these filters, we retain 65% of all the transactions included in the initial dataset.

Our final data set involves $20,207,244$ trades on the secondary market between 2000 and 2005, involving $587,224$ unique assets. As shown in Table 1, on average \$34 million worth of assets are bought or purchased by private investors every month. The average price is \$99.45, across sales and purchases, with substantial variation (the overall standard deviation is \$10.68, and the median standard deviation within each month is \$10.44). The difference between the price paid to and from investors within a month (the "intermediation spread") is on average 2%. Consistent with the description of the market in Section 2.1, the trade size is on average \$70,000 (the median is \$25,000) and institutional size trades (above 1 million) happen sporadically (they represent 1% of the total trades).

There are $4,072$ different dealers active in the market over our sample period. The largest dealer intermediates 10% of total trades, while the second largest dealer has less than a 5% market share. We obtain a similar picture if we use a narrower definition of "market" that takes into account the possibility that dealers specialize. For instance, the highest market share by state of issuance is on average 11%.

Interaction on the inter-dealer market is sparse: the inter-dealer trade is one-third that of trade with investors. Finally, as can be seen in the last row of Table 1, on average (across dealers) one of every two transactions on the inter-dealer market involves a new counterpart.[16] This is consistent with dealers trading in the inter-dealer market to acquire information.

Table 1: Summary statistics

	Mean	St. Dev	Median	Min	Max
Trading Price	99.484	10.68	101.52	36.644	116
Intermediation spread	2.1	1.55	1.19	-0.23	6.8
Monthly trade to investors (10^7 USD)	3.45	0.51	3.38	2.35	4.85
Monthly inter-dealer trade (10^7 USD)	1.50	0.25	1.48	1.02	2.25
Trade size (1,000 USD)	72.05	190.92	25	5	2,245
Overall dealers' market share (%)	0.043	0.40	0.00026	2^{-7}	11.6
Share of inter-dealer trades with a new counterparty (%)	0.475	0.301	0.415	0.002	1

Notes: The table provides summary statistics for trading activity on the secondary market for US municipal bonds. Data come from the proprietary Transaction Reporting System audit trail from the MSRB, and cover the universe of transactions in this market between 2000 and 2005.

[16] Given a trade between dealers d and d' involving an asset issued in state s, we classify d' as a new counterpart if we never observed them trading an asset issued in state s. If we classify as a new counterpart a dealer with whom d has never traded before, the average share of trades with new partners becomes 0.28.

3 Motivating evidence

In this section, we present some patterns of trading in the market for municipal bonds that suggest that: (i) dealers acquire information through trade; and (ii) dynamic incentives to acquire information are an important determinant of dealers' incentives to provide liquidity. Incentives to acquire information, therefore, can have important repercussions both on market liquidity and, through that, on the cost of capital for municipalities.

3.1 Learning through inter-dealer trade

In markets where public information about trading activity is limited, agents need to rely on private interactions with other agents to aggregate the information dispersed among market participants. In particular, through the negotiations to trade an asset a dealer can acquire information about his counterpart's valuation of the asset. This, in turn, provides valuable information about the overall state of the market. Here, we use data on inter-dealer trades to argue that dealers acquire information through their trading activity. Furthermore this information is valuable, since it leads dealers to change their trading behavior.

In particular, we look at the pricing behavior of pairs of dealers who have recently traded with one another in the inter-dealer market. Suppose that the bargaining process revealed to the dealers some information about their opponent's assessment of the market. Then, after the trade, the dealers' pricing strategies should change to account for the new information. Consistent with this intuition, we test whether dealers tend to price an asset more similarly after having traded with one another.

Let i denote a generic inter-dealer trade between dealers d_i and $\tilde{d}_i$ for asset a_i. For every such trade we construct the absolute difference in the average prices $\hat{p}_{a_i,d_i,t}$ and $\hat{p}_{a,\tilde{d}_i,t}$ at which each dealer trades asset a_i with investors in period t, where t ranges from five periods before and ten periods after inter-dealer trade i. The average price $\hat{p}_{a,d,t}$ is constructed as

$$\hat{p}_{\mathrm{a},d,t} = \frac{\sum_{j=1}^{N_{j,\mathrm{a},d,t}} p_{j,\mathrm{a},d}\mathrm{par}_{j,\mathrm{a},d}}{\sum_{j=1}^{N_{j,\mathrm{a},d,t}} \mathrm{par}_{j,\mathrm{a},d}},$$

where j ranges across the dealer's trades in period t, while $p_{j,\mathrm{a},d}$ and $\mathrm{par}_{j,\mathrm{a},d}$ are, respectively, the price and the quantity exchanged in trade j.

We stack all the inter-dealer trades and estimate Equation 1 below, where $\mathrm{d2d}_{i,t-k}$ is a dummy indicating whether inter-dealer trade i took place in $t-k$. The coefficient β_k tells us how the difference in dealers' pricing strategies is affected by the inter-dealer trade. We include pair×asset fixed effects, $\alpha_{a_i,d_i,\tilde{d}_i}$, to absorb any initial differences between dealers' pairs, and cluster standard errors at the pair and asset level.

$$\left|\hat{p}_{a_i,d_i,t} - \hat{p}_{a_i,\tilde{d}_i,t}\right| = \alpha_{a_i,d_i,\tilde{d}_i} + \sum_{k=-5}^{10} \beta_k \text{d2d}_{i,t-k} + \epsilon_i \tag{1}$$

Figure 1 plots the coefficients β_k for $k \in \{-5, \ldots, 10\}$. The coefficients are plotted relative to the difference in prices among the two dealers in the period before the trade ($k = -1$), which we normalize to zero. The figure shows that after two dealers trade with one another, the absolute difference in their pricing strategies falls overall by \$0.18, around 2%, compared to the period before the inter-dealer trade.

Figure 1: Change in pricing behavior after inter-dealer trade

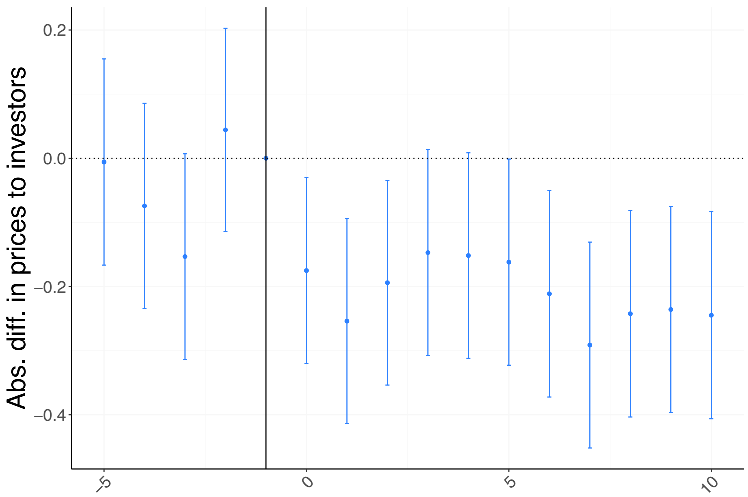

	$\left\|\hat{p}_{a,t,d} - \hat{p}_{a,t,\tilde{d}}\right\|$
after inter-dealer trade$_{a,d,d',t}$	-0.173^{***}
	(0.040)
N	465,053
R^2	0.558

$^{***}p \leq 0.01$, $^{**}p \leq 0.05$, $^{*}p \leq 0.1$

Notes: The figure on the left-hand side plots the weekly regression coefficients and 95% confidence intervals from estimating Equation 1 using the sample of inter-dealer trades between January 2000 and June 2003. Note that this excludes trades executed after the improvement in market transparency described in Section 2.2. The table shows the estimate of the cumulative average effect of an inter-dealer trade on the difference in the price at which the involved dealers trade with investors. Dealers trade asset with one another in period -1. The outcome variable is the difference, among the two dealers, in the weighted average price for trades of asset with investors. The coefficients are plotted relative to the difference in prices to investors among the two dealers in -1, which are normalized to zero. Due to the infrequency of trades, we consider a three-day interval as a period. The results are robust to different fixed effects, and a similar pattern emerges using the logarithm of the difference in prices.

In Appendix A we show the results of several placebo tests. First, we compare the result in Figure 1 with the impact of an inter-dealer trade on dealers' pricing strategies after the introduction of "next day reporting" described in Section 2.1. The introduction of this system increased market transparency and made information about trading prices easily accessible for everyone in the market. In a transparent market, the informative content of a single inter-dealer trade should become negligible. Indeed, we show that the price response in Figure 1 disappears. This suggests, as an example, that the pattern in 1 is not driven by the inventory effect. Next, one might worry that the result in 1 is driven by a release of public information that equalizes the trading behavior across *all* market agents. To verify that this is not the case, for each inter-dealer trade i, completed in period t_i, between dealers d_i and $\tilde{d}_i$, we create a placebo inter-dealer trade in period t_i between d_i and a fictitious trading partner $\tilde{f}_i$[17]. If the results in Figure 1 were driven by a common shock, we should see dealers d_i and $\tilde{f}_i$ pricing the asset more similarly after period t_i, despite not having traded with each other. Instead, the results in Appendix A show that the pattern in Figure 1 disappears if we focus on these fictitious trades.

3.2 The effect of market transparency

For years the SEC has been warning private investors and Congress about the need to improve access to information about trade activity in the market for municipal bonds. As described in Section 2.1, this pressure from the SEC culminated in a series of provisions aimed at improving market transparency. In particular, on June 23, 2003, the MSRB started distributing daily summaries about the trading activity in the market during the previous day.[18]

Proponents of market transparency argue that the lack of public information about trading activity gives dealers an informational advantage vis-à-vis their clients. Market transparency, by leveling the playing field, would increase investors' participation, improve liquidity, and benefit the market at large.[19] This argument, however, ignores *dealers'* incentives to trade. Information acquisition motives for trade, in particular, can substantially erode the positive effects of transparency. Indeed, when public information about market activity is limited, trading with investors allows dealers to acquire valuable information about the market value of the asset. This generates an additional motive for trade that market transparency might weaken.

We explore the effect of the 2003 policy change through a difference-in-difference set-up. We leverage the idea that improving transparency will have stronger consequences for assets for which incomplete information is more severe. A typical example of these assets in the market for municipal bonds is uninsured assets. Issuers that meet certain credit criteria can purchase municipal bond insurance policies from large private insurance companies. The insurance guarantees the payment of principal and interest on

[17] The fictitious partner $\tilde{f}_i$ has traded asset a_i at least once in the previous year and has traded at least once with dealer d_i. The results are robust if we select fictitious partner $\tilde{f}_i$ to also have the same size and specialization as d's observed trading partner.

[18] Asquith et al. 2013 study the effect of a similar policy intervention in the market for corporate bonds and find similar results.

[19] For instance, in the speech before the Bond Market Association SEC commissioner Arthur Levitt remarked, "The undeniable truth is that transparency helps investors make better decisions, and it increases confidence in the fairness of the markets. And, that means more efficient markets, more trading, more market liquidity."

a bond issue if the issuer defaults. Pricing for insured assets, therefore, is more straightforward compared to pricing for the uninsured ones and depends less on unobserved factors. To confirm this intuition, we regress several market outcomes across assets on a dummy variable that equals one if the asset is uninsured. The results, shown in Table 2, confirm that prices for uninsured assets are both more volatile over time and dispersed across trades. Furthermore, uninsured assets trade at a higher price and are issued at a higher yields.

We study the effect of market transparency on trading activity, which we measure as the number of trades per week between dealers and investors. Since most municipal bonds trade infrequently, we use one week as the minimum unit of time. Our main specification is

$$y_{it} = \phi_i + \kappa_t + \gamma post_t + \lambda post_t \times unins_i + \epsilon_{it}, \tag{2}$$

where y_{it} is bond i's outcome in week t; ϕ_i is a vector of asset fixed effects; κ_t is a vector of week fixed effects; and $post_t$ is an indicator for the trade outcomes on weeks after the policy intervention. Finally, $unins_i$ is an indicator that equals one if the asset's principal is uninsured. Since there are repeated observations per bond, in all estimates, the standard errors are clustered by bond. In Equation 2, any pre-existing difference between assets is captured by the fixed effects ϕ_i, while the effects of the policy that accrue to all bonds are absorbed by coefficient γ. The coefficient of interest is λ, which estimates the effect of transparency on trading outcomes for uninsured assets.

Table 3 reports estimates of the parameters in Equation 2 for three different estimation windows covering 2, 4, and 6 months surrounding the policy intervention. Figure 6 in Appendix A.2 shows the event study. The estimate of the effect of transparency on the number of trades per day is negative and significant for all three estimation windows, as the number of trades drops by approximately 0.03, which corresponds to 12% of the average level of trade before dissemination. In AppendixA.2 we show that this result is robust to using alternative measures of trading activity. In particular, we show that also inter-dealer trade falls after the introduction of market transparency.

A liquid secondary market is a crucial condition to lower the cost of raising capital. Decreasing liquidity in this market, therefore, might be very costly for local governments and municipalities. Indeed, in Appendix A.2 we show that the introduction of market transparency increased offering yield for uninsured assets compared to insured ones.

Table 2: Insured vs. uninsured assets

	Est	Standard Error	Mean
Weekly standard deviation purchasing price	0.030**	0.005	0.73
Weekly purchasing price	1.279**	0.031	100.43
Standard deviation of market price	0.047**	0.005	1.65
Offering Price	0.113**	0.006	3.942

$^{**}p \leq 0.05$, $^{*}p \leq 0.1$

Notes: Regressions reflect a regression with the listed variable as the dependent variable and a dummy for whether the asset is uninsured as the independent variable. To produce the table, we focus on the weeks before market transparency. This confirms the intuition that pricing for uninsured assets is more uncertain.

Table 3: Difference-in-difference estimates

	Number of Trades		
	2 Months	3 Months	6 Months
uninsured * $post_t$	-0.024**	-0.031**	-0.026**
	(0.005)	(0.004)	(0.003)
N	2,322,302	3,551,756	7,240,118
Level		Issuer-Week	

$^{**}p \leq 0.05$, $^{*}p \leq 0.1$

Notes: The table presents the output from the difference-in-difference regression that measures the effect of the change in market transparency on trading activity. We use insured assets as control group. Observations are at the asset-week level, and standard errors are clustered at the asset level.

3.3 Why a model?

The results in Section 3.1 suggest that dealers acquire information through their trading activity. Since this information is valuable, acquiring it strengthens the dealers' incentives to trade and to provide liquidity. This observation is particularly relevant given the intense regulatory push to improve market transparency across different financial markets and different countries. Indeed, the results in Section 3.2 suggest that the weakening of information acquisition motives for trade might offset the positive effects of market transparency for liquidity that drive the policymaker.

Importantly, however, the evidence presented in Sections 3.1 and 3.2 falls short along three dimensions. First, a reduced0form approach is not sufficient to assess conclusively the role of dealers' information acquisition as a driver of market liquidity. Indeed, the policy experiment studied in Section 3.2 cannot control for other channels through which the policy affects dealers' and investors' behavior.

Second, a model is necessary to assess the consequences for welfare of policies that improve market transparency. Two opposing forces distort the level of liquidity provision: information acquisition and dealers' market power. While information motives for trade bolster dealers' trading activity, their market power pushes them in the opposite direction. For this reason, a decline in the volume of trade could either be welfare enhancing, if information acquisition is indeed driving dealers to trade more compared to the efficient outcome, or welfare decreasing, in the opposite case. To evaluate the impact of information acquisition motives for trade on a that policy improves market transparency, it is necessary to understand the incentives and the trade-offs driving dealers' motives for trade, and quantify empirically their welfare implications.

Finally, the results in Sections 3.1 and 3.2 are hard to generalize to the new markets at which the regulatory activity is directed. Instead, a model focusing on dealers' motives for trade can help identify the critical features that determine the success of the policy for a specific class of assets.

4 Empirical model

In this section, we introduce an empirical model of trade in decentralized markets. Time t is discrete and the horizon is infinite. The market is populated by two types of risk neutral agents: short-run investors and long-run dealers. Finally, a unique asset is traded.

We build on a standard inventory model of dealers' behavior: dealers accumulate a costly inventory of an asset by buying and selling the asset to investors. Investors' valuations for the asset change over time due to a persistent common demand shock. The dealers' aim is to leverage aggregate fluctuations in the investors' demand for the asset to "buy low and sell high." Therefore, a dealer's fundamental motive for trade is *inventory management*. The common demand shock is unobserved and there is no public information available about it. This introduces a second motive for trade: *information acquisition*. Trading is costly but it reveals information about investors' valuations. Since this information is valuable, acquiring it strengthens the dealers' incentives to trade and to provide liquidity. Finally, every period dealers can trade with one another. The main focus of the model is trade between dealers and investors. However, as we explain in Section 5, modeling dealers' decisions in this context of inter-dealer trade allows us to isolate information acquisition motives for trade.

In each period the timing is as follows. First, the unobserved common demand state is realized, and dealers pay a cost that depends on their accumulated inventory. Next, potential investors observe their private valuation of the asset v and decide whether to enter the market. Dealers, then, start searching for investors to either buy or sell the asset. The prices for these trades will depend on the current value of the common demand shock. Finally, dealers can trade with one another.

4.1 Environment

Inventory model for dealers. Each dealer faces a pool of potential buyers and sellers i who, at the beginning of each period, decide whether to enter the market after observing their private valuation for the asset v_{it}. Whether they succeed in trading or not, investors exit the market at the end of the period. Valuations vary across investors to reflect heterogeneity in liquidity needs, portfolio holdings, and tax advantages. We assume that v_{it} depends on an idiosyncratic component and a common demand shock $\theta_t \in \Theta$. The demand shock θ_t represents common factors that affect investors' willingness to pay for the asset, such as the profitability of alternative investments. We denote by $F_{\text{inv}}(v|\theta_t, \text{b})$ and $F_{\text{inv}}(v|\theta_t, \text{s})$ the distributions of valuations among the buyers and sellers active in the market. The common shock θ_t is the only source of aggregate uncertainty in the market and evolves over time according to a discrete Markov chain with transition matrix $\mathbb{P}_\theta$. While each investor knows his own valuation, they do not observe the common shock θ_t.[20] Dealers, too, do not directly observe the realizations of the demand shock θ_t.

Dealers are forward-looking players with time preferences determined by a constant discount rate $\beta > 0$. In every period, each dealer decides whether to search for investors interested in buying the asset, to search for investors interested in selling the asset, or to pass: $a \in \{\text{b}, \text{s}, \text{pass}\}$. As discussed in Section 4.3, the focus on inter-temporal, rather than cross-sectional, intermediation is natural in the market for municipal bonds. If a dealer has decided to search, he draws search costs $c \sim F_a(c)$ sequentially. After every draw, the dealer can either pay the cost and contact an investor drawn at random from those active in period t, or he can decide not to pay the cost and pass. Regardless of this choice, the dealer has the opportunity to draw another cost (and possibly meet another investor) with probability $1-\gamma$, while with probability γ he moves to the second stage.

We assume that, when meeting an investor, the dealer immediately learns the investor's valuation. Then, dealer and investor enter an alternating offer bargaining game à la Binmore et al. (1986) to determine the terms of trade for u_a units of the asset, where the quantity u_a is an exogenous constant. The bargaining starts with an offer from the investor to the dealer.[21] If the dealer accepts, the parties trade the asset at the proposed price. If the dealer rejects the offer, the negotiation can either break down exogenously or continue with a dealer's counteroffer to the investor. Dealers and investors alternate offers in this fashion

[20] In other words, investors don't know how their own valuation correlates with that of other investors.
[21] The outcome of the bargaining is the same regardless who makes the first offer.

until either one offer is accepted or the bargaining breaks down.

Note that the bargaining between dealer and investor is a one-sided incomplete-information game: the dealer knows the investor's value for the asset, but the investor does not know the dealer's. As we discuss in Section 4.3, this assumption allows us to specify the bargaining in a parsimonious way without substantially distorting the dealer's dynamic incentives to acquire information through trading.[22] It is also consistent with the evidence that the dealers have an informational advantage vis-à-vis investors.[23] Finally, note that the dealer does not commit to trading the asset when calling an investor. Instead, we allow for the possibility that the dealer will decide not to buy the asset after observing the investor's valuation. However, in Section 4.2 we argue that this never happens in equilibrium.

By meeting with the investors and observing their valuation of the asset, dealers acquire information about the current realization of the demand state, since investors' valuations are noisy signals about θ_t. This information is valuable, since it allows the dealer to anticipate changes in the resale value of the asset and, therefore, to improve the future timing of his trading decisions. In Section 4.2 we describe in more detail dealers' belief updating after trading with investors.

Assets bought and not sold accumulate over time and form a dealer's inventory. Carrying inventory is costly: in every period dealer d pays a cost $\kappa(\cdot)$ that captures frictions that prevent dealers from increasing their balance-sheet size, such as the cost of capital (usually inventory is levered) or limits to exposure to risk.

Inter-dealer trade. After trading with investors, dealers can trade with one another. Inter-dealer trade proceeds as follows. A constant share α of the dealers is randomly selected to be "potential sellers," while the remaining are labeled as "potential buyers."[24] Each potential seller d has the option to make an offer to sell u_{d2d} units of the asset to a potential buyer, where u_{d2d} is an exogenous constant.

Contacts between buyer and seller are not random. Instead, we assume that each dealer is characterized by a publicly observed summary statistic of his history of trades, which we call *experience* $e \in \{1, \ldots, E\}$. A dealers' experience captures a his reputation for being active and informed about the market for a certain class of municipal bonds.[25] If a potential seller decides to make an offer, he chooses the experience

[22] Note that the dealer does not commit to trade the asset when calling an investor. Therefore, we allow for the possibility that the dealer will decide not to buy the asset after observing the investor's valuation.

[23] See, as an example, Green et al. (2010).

[24] The role of α is similar to that of the number of potential entrants in standard entry games. Indeed, α sets an upper bound on the total volume of trade in the market, while the total quantity traded in the inter-dealer market remains endogenous since dealers can decide not to engage in trade, and buyers can reject the offer received.

[25] As an example, from Feldstein et al. (2008): "A good regional firm should keep a robust inventory of the region's bonds always available so that customers across the country will think of its firm when looking to buy and sell local bonds."

level $\tilde{e} \in \{1, \ldots, E\}$ of the counterpart, subject to logit shocks. Then he proceeds to make a take-it-or-leave-it offer $q \geq 0$ to a random potential buyer with the chosen experience level.[26] Potential buyers decide whether to accept the offer they receive or not. In case of trade the seller pays a trading cost $c_{\tilde{e}}^{\text{d2d}}$ that depends on the counterparts' experience. This cost captures the time and effort put into finalizing the trade as well as external motives to inter-dealer trade.

Beyond experience, inter-dealer trade is "*anonymous*": dealers do not keep track of the identity of their trading counterparts. Therefore, when interacting with their peers dealers only know their counterpart's experience, while they don't know their counterpart's inventory level or beliefs. As we discuss in detail in Section 4.3, a public summary statistic like experience allows us to model the dealers' decision about whom to trade with in the inter-dealer market, without sacrificing model tractability.

By trading with one another, dealers acquire information about each other's information regarding θ_t. Several aspects of this interaction convey information to the dealers, e.g., the agreed-upon price, the negotiation process itself, the casual interaction between the parties, etc. To keep the model tractable, we model this information exchange by assuming that, after trading, dealers communicate their posterior belief to each other.[27] This assumption allows us to leave in the background the dealers' strategic decision about what information to reveal. However, as we discuss in Section 4.3, this provides a reasonable assumption given the structure of the municipal bond market.

Dealers also accumulate experience through their trading activity. A dealer with experience e achieves experience e' with probability $r^{\text{d2i}}(e'|e)$ after *trading* with an investor, and with probability $r^{\text{d2d}}(e'|e, \tilde{e})$ after *trading* with a dealer with experience $\tilde{e}$. The transition matrices r^{d2i} and r^{d2d} are increasing (with respect to first order stochastic dominance) in their arguments, capturing the idea that experience grows as dealers trade and that it grows faster if they trade with more experienced counterparts. Furthermore, experience depreciates over time: we denote by $r^{\text{dep}}(e'|e)$ the probability that a dealer with experience e at the end of the period will have experience e' at the beginning of the next.[28]

While a dealer's experience is not directly payoff-relevant it affects its standing on the inter-dealer market. More experienced dealers have better information about the common shock θ_t, on average, since experience is increasing in past trading activity and dealers acquire information by trading. Thus, trading with experienced dealers conveys valuable information about the common demand shock θ_t. This implies that experienced dealers will receive more and higher offers on the inter-dealer market. Note that this

[26] Following the literature, we assume that if a potential buyer receives multiple offers, he only observes one of them, chosen at random.

[27] Duffie and Manso (2007) adopts a similar model for how information is exchanged.

[28] Given dealers' history, experiences are drawn independently across dealers.

also means that, for potential sellers, the decision on whom to make an offer to involves a trade-off: while trading with experienced dealers is more valuable, making an offer to an experienced counterpart is also riskier and costlier.[29]

Finally, we make two important assumptions to simplify the inference problem faced by dealers. First, borrowing from the literature on social learning,[30] we assume that each dealer behaves as if the information received from any other dealer is *independent* of what he already knows, conditional on the realization of θ_t and the counterpart's experience. This is a reasonable assumption in the context of a large market where dealers share a common history of trades with very low probability. Second, we assume that potential buyers and sellers only update their beliefs based on *realized trade*. As we discuss in Section 4.3 this assumption is driven by empirical concerns, as our data do not show offers to sell the asset that were rejected by the buyer.

4.2 Behavior

Each dealer is characterized by his type, which consists of his inventory $x \in \{0, 1, \ldots, \bar{x}\}$, his experience $e \in \{1, \ldots, E\}$, as well as his current beliefs $\pi \in \Delta(\Theta)$ about the unobserved demand state. A dealer's inventory and his beliefs are private information; instead his experience is publicly observed.

In this paper, we focus on a steady state of the model such that: (i) the fraction of dealers with a given inventory, belief, and experience depends on θ but not time; and (ii) the fraction of dealers with a given experience is constant over θ and time.

We first derive the optimal behavior of dealers vis-à-vis investor, then we describe their behavior on the inter-dealer market. Where it does not generate confusion, we drop the d subscript and use a "tilde" to denote state variables of dealer d's trading counterpart. For instance we use $\tilde{e}$ instead of $e_{\tilde{d}}$ to denote the experience of dealer d's trading counterpart $\tilde{d}$.

Trade with investors. Consider a dealer with type $\omega = (\pi, x, e)$. Denote by $V_0(\omega)$ the value for the dealer at the beginning of the period, by $V_1(\omega, a)$ the value of the dealer who has decided to search

[29] In Section 7.1, we show that this trade-off holds in the data.

[30] It is standard in the social learning literature to assume that agents learn through DeGroot rules-of-thumb models, which often involve double-counting information. Most notably, Ellison and Fudenberg 1993, 1995 present benchmarks for the rule-of-thumb learning models. Moreover Chandrasekhar et al. 2012 exploit an experimental setup to argue that a DeGroot rule-of-thumb model of learning might provide a better description of agents learning on a network, compared to standard Bayesian updating.

investors of type $a \in \{\mathrm{b}, \mathrm{s}\}$, and by $W(\omega)$ the dealer's value from trading with other dealers. We have:

$$V_0(\omega) = -\kappa(x) + \mathbb{E}\left[\max\left\{W(\omega) + \epsilon_{\text{pass}}, V_1(\omega, \mathrm{b}) + \epsilon_{\mathrm{b}}, V_1(\omega, \mathrm{s}) + \epsilon_{\mathrm{s}}\right\}\right].$$

The dealer pays inventory cost $\kappa(x)$, observes action-specific shocks $\epsilon \in \mathbb{R}^3$ drawn from a type I extreme value (Gumbel) distribution with standard deviation σ_ϵ, and decides whether (i) to avoid meeting investors altogether and move on to inter-dealer trade obtaining value $W(\omega)$, (ii) to search for buyers and obtain value $V_1(\omega, \mathrm{b})$, or (iii) to search for sellers and obtain value $V_1(\omega, \mathrm{s})$.

Let $p(\omega, v, a)$ be the trading price between an investor of type a with value v and the dealer. A dealer who has decided to search for investors of type $a \in \{\mathrm{b}, \mathrm{s}\}$ has value function

$$V_1(\omega, a) = \begin{cases} \mathbb{E}\left[\max\left\{-c + \mathbb{E}\left[-p(\omega, v, a) + V_2(\omega'(v, a), a)\right], V_2(\omega, a)\right\}\right] & \text{if } a = \mathrm{b} \\ \mathbb{E}\left[\max\left\{-c + \mathbb{E}\left[p(\omega, v, a) + V_2(\omega'(v, a), a)\right], V_2(\omega, a)\right\}\right] & \text{if } a = \mathrm{s} \end{cases}, \tag{3}$$

where

$$V_2(\omega, a) = \gamma V_1(\omega, a) + (1 - \gamma) W(\omega).$$

Indeed, the dealer draws search cost $c \sim F_a(\cdot)$ and decides whether to contact an investor. If he does, he draws at random an investor with valuation $v \sim F_{\text{inv}}(\cdot|\theta_t, a)$, and pays (or receives) price p. As we discuss below, if a dealer decides to meet an investor, trade happens with probability one. Note that, not knowing θ_t, the dealer forms expectations about the valuation of the investor that he will meet based on his belief π. After trading with an investor, the dealer's type changes to $\omega'(v, a)$: the dealer updates his beliefs as he learns that the investor, of type a, has valuation v for the asset, and his inventory and experience evolve to account for the trade. Finally, after the dealer concludes the negotiations with the investors, or if he decides not to pay the search cost, he obtains value $V_2(\omega, a)$: with probability γ he can draw a new search cost and restart with valuation $V_1(\omega, a)$, while with probability $1 - \gamma$ he moves on to inter-dealer trade where he obtains value $W(\omega)$.

Let $\mathbb{P}(a|\omega)$ denote the probability that a dealer of type ω chooses action $a \in \{\mathrm{b}, \mathrm{s}\}$. We have:

$$\mathbb{P}(a|\omega) = \frac{\exp\left(\sigma_\epsilon^{-1} V_1(\omega, a)\right)}{\sum_{a'} \exp\left(\sigma_\epsilon^{-1} V_1(\omega, a')\right) + \exp\left(\sigma_\epsilon^{-1} W(\omega)\right)}. \tag{4}$$

Similarly, let $\mathbb{P}(n+1|\omega, a, n)$ be the probability that a dealer of type ω trades at least one additional unit, having traded n units with investors of type $a \in \{\mathrm{b}, \mathrm{s}\}$. We have

$$\mathbb{P}\left(n+1|\omega, \mathrm{b}, n\right) = \frac{\gamma F_{\mathrm{b}}\Big(\mathbb{E}\left[-p\left(\omega, v, \mathrm{b}\right) + V_2\left(\omega'\left(v, \mathrm{b}\right), \mathrm{b}\right)\right] - V_2\left(\omega, \mathrm{b}\right)\Big)}{1 - \gamma F_{\mathrm{b}}\Big(\mathbb{E}\left[-p\left(\omega, v, \mathrm{b}\right) + V_2\left(\omega'\left(v, \mathrm{b}\right), \mathrm{b}\right)\right] - V_2\left(\omega, \mathrm{b}\right)\Big)} \tag{5}$$

and

$$\mathbb{P}\left(n+1|\omega, \mathrm{s}, n\right) = \frac{\gamma F_{\mathrm{s}}\Big(\mathbb{E}\left[p\left(\omega, v, \mathrm{s}\right) + V_2\left(\omega'\left(v, \mathrm{s}\right), \mathrm{s}\right)\right] - V_2\left(\omega, \mathrm{s}\right)\Big)}{1 - \gamma F_{\mathrm{s}}\Big(\mathbb{E}\left[p\left(\omega, v, \mathrm{s}\right) + V_2\left(\omega'\left(v, \mathrm{s}\right), \mathrm{s}\right)\right] - V_2\left(\omega, \mathrm{s}\right)\Big)} \tag{6}$$

Note that the probability $\mathbb{P}\left(n+1|\omega, a, n\right)$ does not depend on n.[31]

Bargaining. When a dealer and an investor meet, the terms of trade are determined by an alternating offer bargaining game à la Binmore et al. (1986). According to this bargaining protocol, the risk that the negotiation breaks down provides the key incentive for the two parties to reach an agreement. We assume that the probability of a breakdown equals $1 - \exp\left(-\left(1-\rho\right)\Delta\right)$ after the dealer rejects an offer and $1 - \exp\left(-\rho\Delta\right)$ in the opposite case, where $\rho, \Delta \in (0,1)$ are fixed constants. The difference in the breakdown probability after a dealer's or an investor's offer, summarized by the parameter ρ, captures the different bargaining power of the two players. For higher values of ρ, the probability of a breakdown after the dealer's rejection is lower; this strengthens his bargaining position and induces him to reject better offers. Following the literature, we consider the limit outcome of the bargaining game as Δ converges to zero.

For a dealer with type $\omega = (\pi, x, e)$ who meets an investor with valuation v, the value of the asset is

$$\Delta V_2\left(\omega, v, a\right) = \begin{cases} V_2\left(\left(\pi'\left(v, \mathrm{b}\right), x, e\right), \mathrm{b}\right) - \mathbb{E}V_2\left(\left(\pi'\left(v, \mathrm{b}\right), x - u_{\mathrm{b}}, e'\right), \mathrm{b}\right) & \text{if } a = \mathrm{b} \\ \mathbb{E}V_2\left(\left(\pi'\left(v, \mathrm{s}\right), x + u_{\mathrm{s}}, e'\right), \mathrm{s}\right) - V_2\left(\left(\pi'\left(v, \mathrm{s}\right), x, e\right), \mathrm{s}\right) & \text{if } a = \mathrm{s} \end{cases}. \tag{7}$$

Whether the dealer concludes the bargaining successfully or not, he obtains value V_2 and updates his beliefs to $\pi'\left(v, a\right)$ after observing the investor's valuation v. Moreover, if the dealer completes the trade his inventory and his experience change to account for the trade.[32]

The bargaining protocol closely follows the model analyzed by Menzio (2005), who extends the Coase conjecture building on Grossman and Perry (1986) and Gul and Sonnenschein (1988). In particular, Menzio (2005) shows that, under the standard restrictions of stationarity of equilibrium strategies and a

[31]See Section C for a derivation.

[32]The expectation is Equation 7 captures the stochastic change in experience after the dealer has concluded the trade. Note that experience only evolves if the trade is actually completed.

monotonicity requirement on beliefs, every sequential equilibrium of the bargaining game under one-sided private information implies immediate agreement at the same terms. The agreed upon price between a dealer of type ω and a buyer with valuation v is

$$p(\omega, v, a) = \begin{cases} \rho v + (1-\rho) \max_{\omega'} \Delta V_2(\omega', v, \mathrm{b}) & \text{if } a = \mathrm{b} \\ (1-\rho) \min_{\omega'} \Delta V_2(\omega', v, \mathrm{s}) + \rho v & \text{if } a = \mathrm{s} \end{cases}. \tag{8}$$

Moreover, the dealer sells the asset if and only if $v - \max_{\omega'} \Delta V_2(\omega', v, \mathrm{b}) \geq 0$, and the dealer buys the asset if and only if $\min_{\omega'} \Delta V_2(\omega', v, \mathrm{s}) - v \geq 0$.

The intuition for the result is as follows: the investor aims to screen dealers who attach a high value to the asset. A take-it-or-leave-it offer would be the most effective way to do that; however, as $\Delta \to 0$, the investor's ability to commit to making a unique offer evaporates and the investor loses his ability to screen different types of dealers. For this reason, the outcome of the bargaining is equivalent to that of a full information version of the game where the investor faces the worst possible dealer's type. This inability of the investor to screen is reflected in the fact that the trading price is independent of the dealer's type. To ease notation we will drop the dependence of the trading price on ω.

Investors entry decision. At the beginning of the period, potential investorsobserve their private valuation of the asset v and decide whether to enter the market. If an investor is contacted by a dealer, he has to incur cost ϕ_a to conduct the negotiations. If the negotiations are successful, the investor either pays or receives the agreed-upon price $p(v, a)$. Therefore, investors enter the market if an only if the surplus from negotiating with a dealer is positive.[33]

Substituting for the price expression in Equation 8, it is easy to see that the surplus for an investor of type a with valuation v for the asset is

$$\begin{cases} (1-\rho)(v - \max_{\omega'} \Delta V_2(\omega', v, \mathrm{b})) - \phi_{\mathrm{b}} & a = \mathrm{b} \\ (1-\rho)(\min_{\omega'} \Delta V_2(\omega', v, \mathrm{s}) - v) - \phi_{\mathrm{s}} & a = \mathrm{s} \end{cases}. \tag{9}$$

Given Equation 9, if an investors' valuation is large enough, he will enter as a buyer. Vice versa, if an investors' valuation is small enough, he will enter as a seller. Crucially, Equation 9 implies that all meetings between a dealer and an investor result in a trade. Especially, while we allow for the possibility

[33] Note that the cost ϕ_a is only incurred if the investor is indeed contacted by a dealer.

that dealers meet investors only to learn their valuation, without the intention of trading the asset, this result ensures that this never happens in equilibrium.

Inter-dealer trade. On the inter-dealer market, "potential sellers," make take-it-or-leave-it offers to "potential buyers" to a dealer of chosen type to exchange the asset. Consider the situation of a potential buyer with type (π, x, e), who receives an offer to buy a unit of the asset at price $\tilde{q}$ from a dealer with experience $\tilde{e}$. The dealer decides whether to accept the offer by comparing the value from purchasing the asset at price $\tilde{q}$ and from rejecting the offer:

$$W_{\mathrm{b}}\left(\left(\pi, x, e\right), \tilde{q}, \tilde{e}\right) = \max\left\{-\tilde{q} + \mathbb{E}\left[\beta V_0\left(\pi'\left(\tilde{\pi}, \tilde{e}, \tilde{q}\right), x + u_{\mathrm{d2d}}, e'\right) | \tilde{e}, \tilde{q}\right], \beta V_0\left(\pi'\left(\tilde{q}, \tilde{e}\right), x, e\right)\right\}. \tag{10}$$

If the dealer accepts the offer, he pays price $\tilde{q}$ and his inventory and experience evolve to account for the trade. Moreover, the dealer updates his beliefs to $\pi'\left(\tilde{\pi}, \tilde{e}, \tilde{q}\right)$ in order to account for the new information he observed, namely (i) the offer $\tilde{q}$ he received, (ii) his counterpart's experience $\tilde{e}$, and (iii) the posterior belief $\tilde{\pi}$ that his counterpart communicates. When deciding whether to accept the offer, the dealer computes the expectation over his counterpart's posterior belief $\tilde{\pi}$, conditional on the offer $\tilde{q}$ received as well as on the counterpart's experience $\tilde{e}$. Importantly, the offer $\tilde{q}$ depends both on the counterpart's inventory and on his beliefs. For this reason, a no-trade theorem does not apply. If the dealer rejects the offer, he only learns that an offer at price $\tilde{q}$ was made by a dealer with experience level $\tilde{e}$ and his belief evolves to $\pi'\left(\tilde{q}, \tilde{e}\right)$. Whether the offer is accepted or not, the dealer moves on to the next period, where he will obtain value βV_0.

Next, consider the situation of a potential seller with type (π, x, e) who has decided to make an offer to a potential buyer with experience $\tilde{e}$. We have:

$$W_{\mathrm{s}}\left(\left(\pi, x, e\right), \tilde{e}\right) = \max_{q \geq 0} \mathbb{E}\left[\mathbb{I}\left\{\text{acc. } q\right\}\left(q + \beta V_0\left(\pi'\left(\tilde{\pi}, \tilde{e}, \text{acc. } q\right), x - u_{\mathrm{d2d}}, e'\right) + c_{\tilde{e}}^{\mathrm{d2d}}\right) + \mathbb{I}\left\{\text{rej. } q\right\}\beta V_0\left(\pi'\left(\tilde{e}, \text{rej. } q\right), x, e\right)\right] \tag{11}$$

If the offer is accepted, the seller pays cost $c_{\tilde{e}}^{d2d}$ and price q, and his inventory and experience change to account for the trade. Moreover, he observes his counterpart's posterior $\tilde{\pi}$ and his beliefs evolve to $\pi'\left(\tilde{\pi}, \tilde{e}, \text{acc. } q\right)$ to account for the fact that (i) offer q was accepted by a dealer with experience $\tilde{e}$, and (ii) his counterpart communicated posterior belief $\tilde{\pi}$. If the offer is rejected, instead he updates his information to account for the rejection. Whether the offer is accepted or not, the dealer moves on to the next period, where he will obtain value V_0.

Finally, the value of a potential seller deciding to whom to make an offer is

$$W_s(\omega) = \mathbb{E}\left(\max\left\{\max_{\tilde{e}\in\{1,\dots,E\}}\left\{W_s(\omega,\tilde{e}) + \xi^{\tilde{e}}\right\}, \beta V_0(\omega) + \xi^0\right\}\right). \tag{12}$$

Indeed, he can either wait for the next period, when he restarts with value βV_1, or make an offer to a dealer with experience $\tilde{e}$, in which case he pays cost $c_{\text{d2d}}^{\tilde{e}} + \xi^{\tilde{e}}$, and decides what to offer as captured by $W_s(\omega,\tilde{e})$. We assume that shocks $\xi \in \mathbb{R}^{E+1}$ are draws from a double exponential distribution with standard deviation σ_ξ. This implies that the probability that the dealer makes an offer to a dealer with experience $\tilde{e}$ can be written as

$$\mathbb{P}_{\text{d2d}}(\tilde{e}|\omega) = \frac{\exp\left(\frac{-c_{\tilde{e}} + W_s(\omega,\tilde{e})}{\sigma_\xi}\right)}{\exp\left(\frac{\beta V_0(\omega)}{\sigma_\xi}\right) + \sum_{\tilde{e}}\exp\left(\frac{-c_{\tilde{e}} + W_s(\omega,\tilde{e})}{\sigma_\xi}\right)}. \tag{13}$$

Equation 13 can be interpreted as an approximation to the discrete choice problem faced by the dealer. Indeed, as the cost shock variance σ_ξ converges to zero, the dealer chooses with certainty the action associated with the highest utility.

Equilibrium. Dealers' policy functions depend on their current type $\omega = (\pi, x, e)$, as well as on their beliefs about the policies of competitors. Other dealers' beliefs and inventory are unobservable. Dealers, therefore, do not observe the valuation of the asset and the policy functions of their peers, but rather have beliefs about them. Beliefs over other dealers' policies and valuations determine dealers' behavior in the context of inter-dealer trade. Similarly to Weintraub et al. 2008, we assume that dealers' conjectures about their peers' inventory x and beliefs π are anchored to their long-run distribution. To allow for learning in the context of inter-dealer trade, however, these conjectures depend on the long-run distribution of dealers' inventory x and beliefs π *conditional* on the unobserved common demand shock θ_t.

Definition. An equilibrium for the market described in the previous section is a distribution $F^*(\omega|\theta)$ of dealers' type $\omega = (\pi, x, e)$ and conditional on demand shock θ, together with distributions $F_{\text{inv}}(\cdot|\theta, a)$ of investors' valuations such that:

E1 The distribution $F_{\text{inv}}(v|\theta, a)$ of investors' valuations is consistent with (9);

E2 Trading decisions $\mathbb{P}(\tilde{e}|\omega)$, $\mathbb{P}(a|\omega)$, and $\mathbb{P}(n+1|\omega, n, a)$ are defined in (13), (4), (5), and (6);

E3 Offers and replies in the inter-dealer market achieve the optimum in (10) and (11);

E4 Conjectures in (10), (11), and (3) are correct given $F^*(\omega|\theta)$ and $F_{\text{inv}}(v|\theta, a)$;

E5 $F^*(\omega|\theta)$ is stationary. In particular, let $TF^*(\omega|\theta)$ be the end-of-period distribution of dealers' types ω implied by transitions r^{d2d},r^{d2i},r^{dep}, choice probabilities (13), (4), (5), (6), and initial distribution $F^*(\omega|\theta)$, conditional on demand shock θ. Moreover, let $\mathbb{P}^*_\theta(\theta', \theta)$ be the (forward) probability $P(\theta_{t-1} = \theta'|\theta_t = \theta)$.[34] Then,

$$F^*(\omega|\theta) = \sum_{\theta'} \mathbb{P}^*{}_\theta\left(\theta', \theta\right) TF^*\left(\omega|\theta'\right). \tag{14}$$

E6 The distribution of experience in the population does not depend on θ.

It is worth discussing the stationarity condition **E5** more in detail. The distribution $F^*(\omega|\theta)$ describes the distribution of types ω among dealers at the beginning of the period, conditional on the demand shock taking value θ. As the agents trade during the period according to their equilibrium trategies, their type evolves. We denote by $TF^*(\omega|\theta)$ the distribution of dealers' types ω at the end of the period. Condition **E5** requires that the distributions $F^*(\omega|\theta)$ and $TF^*(\omega|\theta)$ are consistent with each other. In particular, the right hand side of Equation 14 computes the average distribution of dealers types at the end of the previous period across realizations of the demand shock θ', conditional on this period's demand shock being θ. Condition **E5** requires that this average coincide with $F^*(\omega|\theta)$.

4.3 Discussion

We close this section with a discussion of several of our assumptions and some caveats.

We begin our discussion with the **bargaining model between dealers and investors**. We model this interaction as a one-sided incomplete information, as we assume that the dealer learns an investor's valuation of the asset upon meeting him. This allows us to sidestep the intricacies of a bargaining game with two-sided incomplete information and common values. Nonetheless, this assumption is consistent with the perception of the dealers' informational advantage vis-à-vis the investors, a perception that has shaped government intervention in this market (see Section 2.2). Furthermore, this assumption does not substantially distort the dealer's dynamic incentives to acquire information through trading, which is the

[34]Denoting by $\mathbb{P}_\theta(\theta)$ the stationary distribution associated to transition matrix $\mathbb{P}_\theta(\theta', \theta)$, we can write

$$\begin{aligned}\mathbb{P}^*_\theta\left(\theta, \theta'\right) &= P\left(\theta_{t-1} = \theta, \theta_t = \theta'\right)\mathbb{P}_\theta\left(\theta'\right) \\ &= \frac{\mathbb{P}_\theta\left(\theta', \theta\right)\mathbb{P}_\theta\left(\theta'\right)}{\mathbb{P}_\theta\left(\theta\right)}\end{aligned}$$

focus of our paper. Even if the dealer didn't learn the investor's valuation of the asset during a trade, he would observe the realized price, which, in turn, would help him to forecast future prices. This would generate the same information acquisition motive for trade that is present in our model. Finally, note that the dealer does not commit to trading the asset when calling an investor. Therefore, we allow for the possibility that the dealer will decide not to buy the asset after observing the investor's valuation.

Next, it is important to emphasize our focus on **inter-temporal**, rather than cross-sectional, **intermediation**. This focus motivates the assumption that dealers cannot both buy and sell the asset to investors within a period. The assumption is natural in the market of municipal bonds where we observe a "comparable" amount of sales and purchases for only 7% of dealer-month-issuer pairs.[35] While this is atypical for OTC financial markets, it is not surprising for municipal bonds. Indeed, in the market for municipal bonds, a dealer can dramatically improve his revenue by waiting to sell an asset at the appropriate time. To roughly quantify this improvement, we compare the highest markup a dealer would obtain by selling an asset within one month of its purchase at the average market price, with the markup he would obtain by selling the asset within one quarter of its purchase. In the latter case, the average dealer would improve his intermediation spread from 1% to 4%.

Next, we turn to the model of **learning during inter-dealer trade**. Through inter-dealer trade, dealers acquire information about one another's assessment of the market. Several aspects of this interaction convey information to the dealers, e.g., the agreed-upon price, the negotiation process itself, the casual interaction between the parties, etc. To keep the model tractable, we capture these different ways of learning by assuming that, *after trading*, dealers communicate their posterior belief to each other in a non-strategic way. We believe this is a reasonable approximation of what happens in reality. In our model, the dealers' communication happens *after* trade, and therefore, it only affects the strategic interactions two dealers might have in the future. The market for municipal bonds is large and the same two dealers interact infrequently. Moreover, evidence suggests that dealers do not compete directly for the same investors.[36]

We assume that in the context of inter-dealer trade, dealers only **update their beliefs based on realized trade**. In other words, in our estimation, we disregard the information that dealers obtain when they learn that an offer was rejected. This assumption is driven mainly by empirical concerns. Our data do not show offers that were rejected by the buyer. Therefore, we cannot identify changes in dealers'

[35] That is, for 7% of triplets composed by a dealer d, month t and issuer a, it is $\frac{\max\{\text{sales}_{a,d,t},\text{purhases}_{a,d,t}\}}{\text{sales}_{a,d,t}+\text{purhases}_{a,d,t}} < 0.75$.

[36] See, as an example, Green et al. (2010).

beliefs that derive from offers to trade that were rejected. However, this assumption is consistent with anecdotal evidence that suggests that there are strong reputation concerns involved in soliciting quotes only for their informational content, without the actual intention to buy or sell the asset. Moreover, in Appendix B we use Hansen-Sargan test for over-identifying restrictions[37] to show the results of a test suggesting that learning activities in the market for municipal bonds are strongly connected to "realized" trade. This suggests that the empirical bite of this assumption is limited.

Finally, we discuss the **role of dealers' experience** in our model. Dealers are ex-ante homogeneous. Differences in inventory and information emerge over time because of differences in trading history. Each dealer's trading history is his private information, but it is relevant for his peers, as they try to anticipate his behavior. As an example, the offer made by a dealer depends on his assessment of the likelihood that the offer will be accepted. This, in turn, depends on the dealer's assessment of the counterpart's inventory. Fully modeling dealers' expectations about their peers' trading histories is cumbersome, since they are highly dimensional objects. The theoretical literature on trade in OTC markets has largely sidestepped this issue, assuming that (i) dealers' types are observable and (ii) meetings between dealers are random. This approach is not appropriate for our setup since, as discussed in Section 5, the dealers' decisions about which dealer to trade with is an important source of identification. A public summary statistic like experience is a parsimonious solution to these issues since it allows us to model this decision without sacrificing model tractability.

It is also worth remarking that, given this setup, the concept of dealers' experience is related to that of dealers' centrality in the inter-dealer network.[38] For this reason, our model speaks to the growing empirical literature describing the features of the inter-dealer network, and how dealers' trading prices and allocations depend on their position in the network.[39] While we don't target explicitly these stylized facts, in Appendix F.1, we show that they arise in equilibrium in response to dealers' incentives to acquire information.

5 Identification

There are two sets of unknown we need to recover. First, we discuss how we identify the equilibrium distribution of the dealers' types, which consists on their experience and their beliefs about the unobserved demand state. Then, we discuss how we identify the model primitives, i.e., dealers' trading and inventory

[37] We use the specification test suggested in Dickstein and Morales (2015).

[38] Two agents are considered connected in the network if they have traded within a pre specified time interval.

[39] Most notably, Li and Schürhoff (2019).

costs, as well as the distribution of investors' valuations and their entry costs. Appendix E discusses how we recover the sequence $(\theta_t)_{t \geq 1}$ of realizations of the demand shock as well as their transition matrix $\mathbb{P}_\theta$. Here, we treat these objects as known by the econometrician.

Dealers' types: experience. As a reminder, we define a dealer's experience as a publicly observable summary statistics of his history of trades. In particular, denote by $e(h_t)$ the experience of a dealer with trading history h_t. Recovering the dealers' experience boils down to recovering function $e(\cdot)$, since the history h_t is known.

To recover the dealers' experience we use a key implication of the model described in Section 4. Trading with more experienced counterparts allows dealers to acquire better information about the state of the demand shock θ_t. For this reason, the sellers will offer lower prices to a dealer with higher experience. Following this insight, we exploit the trading prices in the inter-dealer market to identify the dealers' experience. In particular, we look at systematic differences in prices across trades with different buyers for a given seller, for a given month and asset. The correlation of these prices variation with the buyers' trading histories h_t allows us to pin down the shape of the function $e(\cdot)$.

Dealers' types: beliefs. What a dealer knows about the demand shock is summarized by his beliefs π. As part of our estimation, we recover the equilibrium distribution $F^*(\pi|e,\theta)$ of dealers' beliefs π conditional on the demand shock θ and the dealer's own experience e.

To recover the distribution $F^*(\pi|e,\theta)$ we leverage the two important assumptions of our model. First, the equilibrium conditions described in Section 4 ensures that $F^*(\pi|e,\theta)$ satisfies the fixed point

$$F^*(\pi|e,\theta) = \sum_{\theta'} \mathbb{P}^*{}_\theta\left(\theta',\theta\right) T F^*\left(\pi|e,\theta'\right), \tag{15}$$

where T maps $F^*(\pi|e,\theta)$ to the distribution of dealers' beliefs at the end of the period, after the dealers have traded according to their equilibrium strategies and updated their beliefs based their trading activity. Second, as we argue below, the operator $T(\cdot)$ only depends on the dealers' observed trading behavior and the distribution $F^*(\pi|e,\theta)$. Combining these two observations, we can identify $F^*(\pi|e,\theta)$ as the solution to Equation 15, after substituting the operator T with its empirical counterpart.

The operator $T(\cdot)$ captures dealers updating their beliefs as they gather new information about the demand shock by (i) trading with investors and (ii) trading with other dealers.[40] Consider first a dealer

[40] In the estimation, we also assume that dealers observe a public signal correlated with the last period's trading activity.

with prior π who meets an investor of type a with valuation v. Denote by $\pi'(v,a)$ the dealer's posterior belief after this interaction and p the trading price. As discussed in Section 4.2, every meeting between a dealer and an investor is concluded with a trade; moreover, the trading price is a deterministic function of the investor's valuation. For this reason we can write the dealer's posterior beliefs $\pi'(v,a)$ as a function of the trading price p and its distribution, which we denote by $F_{\mathrm{p}}(p|\theta,a)$. We have, by Bayes rule,

$$\pi'(v,a) = \frac{dF_{\mathrm{p}}(p|\theta,a)\,\pi(\theta)}{\sum_{\theta'} dF_{\mathrm{p}}(p|\theta',a)\,\pi(\theta')}. \tag{16}$$

The left hand side of Equation 16 is a known function of the dealers' prior π. Indeed, for every trade between a dealer and an investor that we observe in our data, the trading price p is observed. Moreover, we can treat the sequence of realized demand shocks $\{\theta_t\}_{t=0}^{T}$ as known from our estimation of the demand process. Finally, the distribution of prices $F_{\mathrm{p}}(p|\theta,a)$ can be estimated through a simple kernel estimator.

Then, consider a dealer with prior beliefs π and experience e who receives an offer $\tilde{q}$ from a dealer with experience $\tilde{e}$ and prior belief $\tilde{\pi}$. As described in detail in Section 4.2, the dealer updates his beliefs to account for the new information he gathers through this interaction: namely, (i) the fact that a dealer with experience $\tilde{e}$ made offer $\tilde{q}$ to a dealer with experience e, and (ii) the the fact that a dealer with experience $\tilde{e}$ communicated posterior belief $\tilde{\pi}$ after the trade.[41] By Bayes rule, we can write the dealer's posterior belief as

$$\pi'(\tilde{\pi},\tilde{e},\tilde{q}) = \frac{\mathbb{P}(\{\text{offer } \tilde{q} \text{ from } \tilde{e}\} \cap \{\text{observes } \tilde{\pi} \text{ from } \tilde{e}\}|\theta)\,\pi(\theta)}{\sum_{\theta'} \mathbb{P}(\{\text{offer } \tilde{q} \text{ from } \tilde{e}\} \cap \{\text{observes } \tilde{\pi} \text{ from } \tilde{e}\}|\theta')\,\pi(\theta')}. \tag{17}$$

To simplify Expression 17, it is worth emphasizing that the dealer's update will not depend on the action played by his counterpart ((i) above). Indeed, while the offer received by the dealer conveys information about what his counterpart knows about θ_t, the beliefs $\tilde{\pi}$ communicated by the counterpart are a sufficient statistic for this information. In light of this, we can write the dealer's posterior beliefs as

$$\begin{aligned} \pi'(\tilde{\pi},\tilde{e}) &= \frac{\mathbb{P}(\{\text{observes } \tilde{\pi} \text{ from } \tilde{e}\}|\theta)\,\pi(\theta)}{\sum_{\theta'} \mathbb{P}(\{\text{observes } \tilde{\pi} \text{ from } \tilde{e}\}|\theta')\,\pi(\theta')}. \\ &= \frac{dF^*(\tilde{\pi}|\tilde{e},\theta)\,\pi(\theta)}{\sum_{\theta'} dF^*(\tilde{\pi}|\tilde{e},\theta)\,\pi(\theta')}. \end{aligned} \tag{18}$$

To construct this signal, we regress the average market price of municipal bonds in each state on (i) the interest rate for 1, 5, and 10 year Treasuries, (ii) public indexes about the performance of the municipal bond market, and (iii) total new municipal debt issued.

[41] Similarly, consider a dealer with prior beliefs π whose offer q is accepted by a dealer with experience $\tilde{e}$ and posterior $\tilde{\pi}$. The dealer will update his beliefs based on (i) the fact that a dealer with experience $\tilde{e}$ accepted offer q, and (ii) the posterior belief $\tilde{\pi}$ that his counterpart communicates.

As above, if the distribution $F^*(\pi|e,\theta)$ was known, for every inter-dealer trade that we observe in our data the left hand side of Equation 16 would be a known function of the dealers' prior π.

Primitive parameters: dealers' costs. On the dealer's side there are four sets of primitive parameters in the model, which we allow to vary across assets issued by different states: (1) the inventory cost κ, (2) the distribution of search costs faced by dealers $F_c(\cdot|a)$, (3) the search costs for inter-dealer trade $c^{\text{d2d}} = \{c_{\tilde{e}}\}_{\tilde{e}=1}^{E}$ and (4) the standard deviations $(\sigma_\epsilon, \sigma_\xi)$ associated with the preference shocks. These parameters determine dealers' optimal behavior, given prices, through their value functions. Therefore, we can exploit variation in the dealers' observed behavior allows to identify them. While the data are used to jointly identify the parameters, we can consider what variation in the data allows for the identification of each parameter.

First, the estimate for the value of information hinge crucially on the estimates for the standard deviations $(\sigma_\epsilon, \sigma_\xi)$ associated with the cost shocks affecting the dealers' trading decisions. The intuition is as follows: dealers value information because it allows them buy or sell the asset at the right time. If standard deviations $(\sigma_\epsilon, \sigma_\xi)$ are large, the dealer's trading decisions depend mostly on the unobserved cost shocks, rather than on his inventory or information. In this case, information acquisition becomes irrelevant.

The identification of the standard deviations $(\sigma_\epsilon, \sigma_\xi)$ is driven by differences in dealers' trading decisions for different prior beliefs π, conditional on the observed trading prices. Consider, as an example, the identification of the standard deviation σ_ξ of the cost shock associated with inter-dealer trade. When deciding with whom to trade on the inter-dealer market, dealers face the following trade-off: selling an asset to a more experienced counterpart is less remunerative, but it conveys more information about the common demand shock. Columns (I) and (II) in Table 4 confirm that more experienced buyers pay lower prices in the inter-dealer market, after controlling for a rich set of fixed effects. How a dealer solves this trade-off as a function of his prior belief π pins down σ_ξ. In particular, if a dealer is more uncertain about the demand shock, then information is weakly more valuable.[42] Therefore, if information is valuable, dealers with more uncertain prior beliefs will tend to solve this trade-off in favor of trading with more experienced dealers. The strength of this substitution pattern pins down the magnitude of σ_ξ. Columns (III) and (IV) of Table 4 confirm that this substitution pattern is present in the data. Indeed, when a dealer has a more uncertain prior, as measured by Shannon's entropy,[43] he tends to substitute into the

[42] This is a general property of the value of information.

[43] In particular, Entropy$(\pi) = \sum_\theta \pi(\theta) \log(\pi(\theta))$

more informative and expensive trade with experienced dealers.

Furthermore, identification of the cost of inventory κ is driven by differences in the dealers' trading decisions as a function of their inventory x, conditional on observed trading prices: as κ increases, the dealer will tend to liquidate the assets as his inventory increases, holding everything else fixed. Similarly, the sheer trading costs $\{c^{\mathrm{b}}, c^{\mathrm{s}}\}$ are pinned down by the overall level of trade. Identification of trade costs in inter-dealer trade $\{c_{\tilde{e}}\}_{\tilde{e}=1}^{E}$ depends in a similar fashion on the level of inter-dealer trade for dealers with different experience.

Table 4: Identification of the value of information

	Inter-dealer price (log) (I)	(II)	Counterpart's experience (log) (III)	(IV)
Buyer's experience (log)	-0.001**	-0.002**		
	(0.0001)	(0.0002)		
Prior uncertainty (log)			0.015***	0.013***
			(0.002)	(0.002)
FE	Seller × Month × Issuer		Seller	
	-	Buyer	-	Month
Controls	Buyer's inventory and experience		Dealer's inventory and experience	
N	1,535,608	1,535,608	147,108	147,108

$^{**}p \leq 0.05$, $^{*}p \leq 0.1$

Notes: In columns (I) and (II) of the table, we regress the logarithm of the price in inter-dealer trades on the experience of the buyer of the asset. We include a seller×month×asset fixed effect to absorb market-wide shocks to prices. Finally, we control for the buyer's experience and his inventory. The second specification additionally includes a fixed effect for the buyer. For columns (III) and (IV) of the table, the dependent variable is the logarithm of the experience of a dealer' counterpart in the inter-dealer market. The independent variable is the dealer's prior uncertainty measured using Shannon's enthropy. Additionally, we control for the dealer's experience and his inventory

6 Dealers' types: estimation and results

The estimation is based on a comprehensive dataset on trading activity in the market for municipal bonds described in Section 2. We focus on the period between January 2000 and between June 2003. Indeed, after June 2003 the policy described in Section 2.2 muted dealers' incentives to privately acquire information about market fundamental in the market for municipal bonds. In Section 8.1, we exploit the data between June 2003 and December 2005 to validate our model. We group the assets based on the state of issuance and estimate the model separately for ten different states.[44] Finally, we assume that θ_t can take three

[44]We divide the states into deciles based on the total outstanding of municipal bonds and estimate the model for the median state within each bracket, namely: Vermont, Arizona, Kansas, Nebraska, Mississippi, Indiana, Wisconsin, New Jersey, New

values: θ_L, θ_M, and θ_H.

6.1 Dealers' experience

A dealer's experience is a publicly observable summary statistics $e(h_t)$ of his history of trades h_t. We begin specifying the function $e(\cdot)$. Following Benkard (2000, 2004), we assume that a dealer's experience evolves as a first-order deterministic process. At time t a dealer has experience e_t. The dealer then trades with n_t^{inv} investors and with n_t^{dealer} dealers with experiences $\left(e_{1,t}, \ldots, e_{n_t^{\text{dealer}},t}\right)$. Between periods t and $t+1$, the dealer's existing stock of experience depreciates by a factor δ, while he acquires a new stock of experience that depends on n_t^{inv} and $\left(e_{1,t}, \ldots, e_{n_t^{\text{dealer}},t}\right)$. This process is approximated by the following equation:

$$e_{t+1} = \delta e_t + n_t^{\text{inv}} + \alpha \frac{1}{n_t^{\text{dealer}}} \sum_{\tilde{d}=1}^{n_t^{\text{dealer}}} e_{\tilde{d}t}, \tag{19}$$

where α is a reduced-form parameter capturing frictions that hinder communication during trade with other dealers.[45]

The intuition for this specification is as follows. A dealer's experience offers a proxy for the precision of the information that a dealer has been able to gather about the unobserved demand shock θ_t through his trading activity. Hence, a dealer's stock of experience is constantly being eroded as information becomes less relevant over *time*. Trading allows a dealer to acquire new information about the demand state: the informational content of trade with investors depends on the number of trades; instead, the information content of an inter-dealer trade depends on the counterparts' experience.[46]

Given the functional form assumption in Equation 19, recovering dealers' experience boils down to estimating the parameters α and δ. To recover these parameters, we leverage trading prices p_i in the inter-dealer market and estimate the baseline specification

$$\log(p_i) = \alpha_{s_i \times \mathrm{m}_i \times a_i} + \alpha_{b_i} + \psi_0 \log\left(e_{b_i,\mathrm{m}_i}(\delta, \alpha)\right) + \psi_1 x_{b_i,\mathrm{m}_i} + u_i, \tag{20}$$

York, and Michigan. This approach also allows us to exclude states without individual income tax. Indeed, due to the different tax treatment, the market for bonds issued by these states is integrated at the national level.

[45] For simplicity, the model in Section 4 assumes that, within a month, each dealer trades only once in the inter-dealer market. However, in the data we observe multiple inter-dealer trade per dealer. To bridge the gap between the data and the model we assume that the average experience of the trading partners is relevant for the evolution of a dealer's experience. However, we experiment with numerous formulations for experience and find similar results. Among others, we try: weighting the contribution of the trading partners' experience by the size of trades; using the total volume of trade, rather than the number of trades $n_{d,t}$; and using the log-sum of the trading partners' experience.

[46] We show in Section 6.2 that the experience we recover is indeed positively correlated with the precision of a dealer's information about the state of demand.

where s_i, b_i, a_i and m_i denote, respectively, the seller, the buyer, the asset involved in the trade, and the month in which trade happens.[47] The parameter ψ_0 measures the discount that a seller is willing to grant to an experienced buyer.[48] We estimate the parameters in Equation 20 using non-linear least squares.

Identification of the parameters in Equation 20 relies on comparisons of inter-dealer prices in trades for specific asset a_i, seller s_i and month m_i. Especially, Equation 20 attributes systematic differences in price across trades executed by seller s_i in month m_i to differences in experience level $e_{b_i,\mathrm{m}_i}(\delta, \alpha)$ of the buyers involved in the transactions. The fixed effect $\alpha_{s_i \times \mathrm{m}_i \times a_i}$ absorbs market-wide shocks to prices, as well as the seller's persistent heterogeneity that might affect prices, while the fixed effect α_{b_i} absorbs the buyers' persistent heterogeneity. Finally, we control for the buyers' inventory in x_{b_i,m_i}.

Results. The first panel of Table 5 summarizes the distribution of the estimates for parameters α and δ, while the estimates are plotted in Figure 14 in Appendix A

Consistent with the volatile nature of the market, information appears to be short-lived. The value of δ in the average state is 53%; hence, only 2% of the experience that dealers accumulate lives through six months. Unsurprisingly, the estimates show that there is substantial heterogeneity in the persistence of experience across states. Indeed, differences in the estimates for δ reflect the different features of the market for municipal bonds across states. As an example, the first row in the second panel of Table 5 shows that information depreciates less slowly (higher δ) in states where the demand process is more persistent, based on our demand estimates. As described in Section 2, illiquidity is the key driver of demand volatility in the market for municipal bonds. To test directly the connection between experience persistence and demand volatility, in Table 5 we correlate our estimate for δ with (i) market depth, as measured by the states' total outstanding municipal debt, as well as (ii) a standard measure for market illiquidity proposed by Amihud (2002). In particular, we measure asset j's liquidity as

$$\mathrm{Amihud}_j = \frac{1}{n_j} \sum_{t=1}^{n_j} \frac{|r_{jt}|}{Q_{jt}}, \tag{21}$$

where r_{jt} is asset j's monthly return in month t,[49] n_j is the number of months for which returns r_{jt} can be computed, and Q_j is the total trading volume in millions \$. Amihud's illiquidity measures can be

[47] Note that, given parameters α and δ, Equation 19 pins down dealers' experience given an initial condition. In particular, we impose that dealers' experience at the beginning of the sample is zero. Moreover, we drop the first year of data to estimate the parameters.

[48] This parameter doesn't have a structural interpretation and we don't use it anywhere else in the estimation.

[49] Following Acharya et al. (2013), we measure the monthly returns for an asset as the percentage change in the average price in a month compared to the average price in the previous month.

interpreted as the "price response associated with one dollar of trading volume, thus serving as a rough measure of price impact." A state's illiquidity consists of the average illiquidity of its assets. As shown in Table 5, we find that information depreciates more slowly in states where the market has more depth and is more liquid.

Finally, the last row of Table 5 compares the contribution to experience of trade with dealers and investors. In particular, for every state we compute the number of investors with whom a dealer would need to trade to acquire a stock of experience comparable to that acquired by trading with the average dealer, which we denote by $MRS_{\text{dealer, inv}}$. Unsurprisingly, the contribution of inter-dealer trade to experience is more relevant than direct trade with investors: in the average state trading the average dealer has the same impact on a dealer's experience as trading with 5 investors.

Table 5: Estimates of the experience process

	Average	Standard Deviation
δ	0.531	0.231
α	0.352	0.269
$MRS_{\text{dealer, inv}}$	5.55	9.84

	Correlation with δ
Persistence	40.61%
Outstanding muni. debt	18.62%
Amihud illiquidity	−28.48%
Share of Aaa-rated assets	57.57%

Notes: The table summarizes the estimates of the experience process defined in Equation 19. We cluster the assets based on that state of issuance and estimate the experience process independently across groups. The first table summarizes the distribution for the estimates of δ, measuring the persistence of experience over time, and α, which captures frictions that hinder communication during trade with dealers. The second table shows how the estimates for δ correlate with various measures of illiquidity.

We consider a coarse discretization of experience, and we run a k-means algorithm with three centers to identify the clusters. The experience process described in 19 determines not only dealers' experience, but also the transition of a dealer's experience over time. To minimize the effects of the discretization on the value and policy functions, we approximate the process in 19 making the evolution of experience stochastic. In particular, given the discretized experience process we estimate the transition matrices r^{d2i}, r^{d2d}, and r^{dep} separately using a multinomial logit.

6.2 Dealers' beliefs about the demand shock

Based on the arguments in Section 5 we estimate the distribution $F^*(\pi|e,\theta)$ as the solution to

$$F^*(\pi|e,\theta) = \sum_{\theta'} \mathbb{P}^*{}_{\theta}\left(\theta',\theta\right)\hat{T}F^*\left(\pi|e,\theta'\right), \tag{22}$$

where $\mathbb{P}^*{}_\theta(\theta',\theta)$ is the (forward) probability $P(\theta_{t-1}=\theta'|\theta_t=\theta)$ and $\hat{T}$ maps $F^*(\pi|e,\theta)$ to the distribution of dealers' beliefs at the end of the period, given the dealers' observed trading behavior.

We exploit the following fixed-point algorithm. Each iteration of the algorithm starts with a guess for the dealers' beliefs $\pi^{(m)} = \left(\pi_{dt}^{(m)}\right)_{d\in D, t\in\{0,\ldots,T\}}$ for each dealer d and period t in our dataset. We first use a kernel estimator to recover the distribution $F^{(m)}(\pi|e,\theta)$ of the beliefs $\pi^{(m)}$ across dealers, conditional on their experience e and the state θ.[50] This corresponds to a guess for the equilibrium distribution $F^*(\pi|e,\theta)$. Next, we update this guess. In particular, we use Equations 16, 18, together with the dealers' observed trading behavior, to recover each dealer's posteriors beliefs $\pi'^{(m)}_{d,t}$ at the end of period t. This step depends on $F^{(m)}(\pi|e,\theta)$ through Equation 18. For every period t, we set a new guess for each dealer's beliefs equal to

$$\pi_{d0}^{(m+1)}(\theta) = \mathbb{P}^*_\theta(\theta)$$
$$\pi_{dt}^{(m+1)}(\theta) = \sum_{\theta'} \mathbb{P}_\theta\left(\theta,\theta'\right)\pi'^{(m)}_{d,t-1}(\theta),$$

where $\mathbb{P}^*_\theta$ is the stationary distribution associated to transition matrix $\mathbb{P}_\theta$.[51] Once the new guess for the dealers' beliefs is computed, the algorithm moves to the following iteration. The procedure continues until the distribution $F^{(m)}(\pi|e,\theta)$ converges.

Results. To illustrate the results, we look at the uncertainty in dealers' forecasts of trading prices for dealers with different experience levels. We measure this uncertainty using the precision of dealers' forecast for the average selling price

$$RMSE(\pi) = \sqrt{\mathbb{E}\left(\left[\mathbb{E}_\pi\left(p\left(v_{it},\mathrm{s}\right)\right) - \mathbb{E}^*\left(p\left(v_{it},\mathrm{s}\right)|\theta_t\right)\right]^2\right)}, \tag{23}$$

where $\mathbb{E}_\pi(p(v_{it},\mathrm{s}))$ is the expected market selling price for an asset for a dealer with belief π, and $\mathbb{E}^*(p(v_{it},\mathrm{s})|\theta_t)$ is the average selling price for the asset, given the realization of demand shock θ_t.

In Table 6 we show the average RMSE for different types of dealers across states, based on the estimated distribution $\hat{F}^*(\pi|e,\theta)$. The first column reports the upper bound for the RMSE when the dealers only observe a public signal about the state. The last three columns show the RMSE attained in equilibrium for dealers with different experience levels. Experimentation allows dealers to improve the precision of

[50] Note that at this stage, each dealer's experience is known.

[51] We initialize $\pi^{(0)}$ to the dealers' posterior if they only updated based on trades with investors. In particular, we set $\pi^{(0)}_{d,0}$ equal to the stationary distribution of the demand shock θ_t, and use Equation (16) to compute initial guess $\pi^{(0)}$.

their estimate on average by 16%. As expected, experienced dealers have better information than the rest of the market, as they are able to improve the precision of their prediction by 31%.

Table 6: Precision of information

	Uninformed Dealers	**Market Average**	**Inexperienced Dealers**	**Experienced Dealers**
RMSE	1.588	1.344	1.422	1.112
Percentage	100	84.5%	89.4%	69.9%

Notes: The table reports the average RMSE for different classes of dealers. For each state we compute compute the average RMSE based on the equilibrium distribution of dealers' types, the table resports the average of this measure across states. The first column reports the measure for players with access only to public information. The second column reports the average across all players, and the last two columns distinguish among experienced and inexperienced players.

7 Estimation of the model's primitives

We now turn to the estimation of the model's primitives. On the dealer's side there are four sets of primitive parameters in the model, which we allow to vary across assets issued by different states: (1) the inventory cost κ, (2) the distribution of search costs faced by dealers $F_c(\cdot|a)$, (3) the search costs for inter-dealer trade $c^{\text{d2d}} = \{c_{\tilde{e}}\}_{\tilde{e}=1}^{E}$ and (4) the standard deviations $(\sigma_\epsilon, \sigma_\xi)$ associated with the logit shocks. As we explain in detail below, the estimation amounts to matching the observed trading decisions of dealers to the model's predicted choice probabilities.

On the investor side we recover the distribution of investors' valuations F_v, and their entry cost ϕ. In this case, the estimation amounts to matching the observed trading prices to the prices outlined in Equation 8.

We calibrate the discount factor to $\beta = 0.999$, the probability that a dealer can search for multiple investors to $\gamma = 0.95$, and the dealers' bargaining power to $\rho = 0.75$.[52] Finally, we set the size of different types of trades u_{s}, u_{b}, and u_{d2d} equal to their sample median.[53]

7.1 Dealer's costs

The key aim of our exercise is to separate dealers' information and inventory motives for trade. To ensure that the recovered inventory motives for trade are not shaped by our functional form assumptions, we

[52] The results are similar for $\rho = 0.5$ and $\rho = 0.9$.
[53] On average, u_{s} equals \$23,888, u_{b} equals \$21,667, and u_{d2d} equals \$41,389.

assume a flexible parameterization of the inventory costs $\kappa(x)$. In particular, inventory costs are a third-order polynomial of a dealer's inventory x[54] with parameters $(\kappa_j)_{j=1}^{3}$. Moreover, we assume that the search costs a dealer faces when trading with investors have the form $c = c^a + \sigma_c \epsilon_a$ for $a \in \{\text{b}, \text{s}\}$, where ϵ_a is a random draw from the logistic distribution.

The model does not yield a closed-form solution to dealers' behavior given the vector of parameter values. Hence, we turn to simulation methods to estimate the parameter $\tau = \left\{(\kappa_j)_{j=0}^{3}, c^{\text{d2d}}, c^{\text{b}}, c^{\text{s}}, \sigma_c, \sigma_\xi, \sigma_\epsilon\right\}$. In particular, we exploit the model's predictions for dealers' behavior, described in Section 4.1, to build a moment-based procedure. We use a nested fixed point algorithm to solve for the dealers' value functions at every guess of the parameter values; then, given the value functions, we construct and match a rich set of moments, discussed below, that are pertinent to dealers' optimal trade choice probabilities.

The first set of moments we match concerns the dealer's behavior vis-à-vis investors. Namely, we match the probability that a dealer with type ω is observed either buying an asset, $\mathbb{P}(\text{buy}, \text{n} \geq 1|\omega))$, selling an asset, $\mathbb{P}(\text{sell}, \text{n} \geq 1|\omega)$, or not trading, $\mathbb{P}(\text{n} = 0|\omega)$. Moreover, we match the probability $\mathbb{P}(n+1|\omega, a, n)$ that a dealer with type ω trades one additional unit, having already traded n units with investors of type $a \in \{\text{b}, \text{s}\}$. The second set of moments concerns the dealer's behavior on the inter-dealer market. In particular, we match the probability $\mathbb{P}(\text{sells to } \tilde{e}|\omega)$ of observing a dealer with type ω selling an asset to a dealer with experience level $\tilde{e}$.

We estimate the probabilities concerning dealers' behavior vis-à-vis investors and other dealers using separate multinomial logit sieves. Then, we stack these observed choice probabilities for a grid of values for ω and choose the parameter τ that minimizes

$$\hat{\tau} = \arg\min \left(\hat{\mathbb{P}} - \Psi(\tau)\right)' \Sigma \left(\hat{\mathbb{P}} - \Psi(\tau)\right),$$

where $\hat{\mathbb{P}}$ is the stacked vector of the observed choice probabilities, $\Psi(\tau)$ is the vector of respective predicted choice probabilities, and Σ is a weighting matrix.[55]

Results. Table 7 reports the average baseline estimates for the dealer's cost parameters across states, while the full set of estimates is reported in Table 15 in Appendix D.[56] Consistent with industry narratives, search costs are large: our estimates imply that a dealer faces average search costs that range from 12%

[54] As a robustness we run the estimation with a fourth- and fifth-order polynomial and find very similar results.

[55] As a weighting matrix we use the covariance matrix of the estimated choice probabilities $\hat{\mathbb{P}}$.

[56] The standard errors are computed from 100 bootstrap samples with the resampling done at the dealers' level. We combine these bootstrap samples with those from the dealers' experience and beliefs to incorporate the error from the estimation of the dealers' types.

to 16% of the trade size. This corresponds to an average of \$2,200 to find an investor interested in selling an asset and \$4,700 to find an investor interested in buying it.[57] These estimates reflect the difficulty of trading in the market for municipal bonds. Indeed, due to the complexity of these assets, negotiations are long and often require the dealer to engage in a lengthy description of the characteristics of the asset before being able to conclude a trade. The estimates also indicate that finding buyers is more difficult than finding sellers: this is consistent with the industry saying that "municipal bonds are sold, not bought."[58] Finally, the estimates for the standard deviations of the preference shocks, σ_ϵ and σ_ξ, equal respectively 5% and 6% of the average trade price, suggesting that preference shocks do not account for a large part of the trading decision for dealers.

Based on the estimates, adding a dollar of municipal bonds to the inventory of the average dealer increases inventory costs by \$0.002. If we interpret the inventory cost in terms of leverage, the estimates imply that around 10% of the inventory is leveraged, assuming that the dealers can borrow at the deposit funds rate. Unsurprisingly, there is substantial heterogeneity in the shape and curvature of the inventory across assets issued by different states. The estimates suggest that this heterogeneity reflects differences in the rating of the assets traded. In particular, doubling the share of assets with a Aaa rating decreases the cost of inventory by 40% across states, for the average dealer.[59] This is consistent with, for example, dealers targeting a certain value-at-risk level when managing inventory.

Table 7: Baseline cost estimates

	Inventory Costs			**Trade with Investors**		**Interdealer Trade**			**Variance**		
	κ_1	κ_2	κ_3	c_{b}	c_{s}	$c_{\tilde{e}_L}$	$c_{\tilde{e}_M}$	$c_{\tilde{e}_H}$	σ_ϵ	σ_c	σ_ξ
Average	1.28e-03	-4.79e-06	3.83e-09	-0.154	-0.126	0.026	-0.020	-0.011	2.100	0.960	0.495
Standard deviation	2.70e-03	7.06e-06	7.26e-09	0.108	0.027	0.024	0.021	0.01	1.636	0.689	0.121

Notes: The table summarizes the estimates of the trading costs that dealers face. We cluster the assets based on state of issuance and estimate these costs independently across groups. The table reports the average and standard deviation of the estimates across states. All the estimates are in 1,000 USD.

To quantify dealers' incentives to experiment, we explicitly study the value that a dealer assigns to the information conveyed by a trade with an investor. In particular, for a dealer with type $\omega = (\pi, x, e)$, we

[57] Since the dealer will only trade for low realizations of the cost shock, the average search cost actually payed by the dealers is substantially lower than the estimates for c_a: the average search cost payed by the dealers ranges from 1.3% to 2.2% of the trade size.

[58] See, as an example Feldstein et al. (2008).

[59] Rating explains 35% of the variation in inventory cost across classes.

compute

$$VI(\omega) = \mathbb{E}\left[V\left(\omega'(v)\right)\right] - V(\omega), \tag{24}$$

where $\omega'(v) = (\pi'(v), x, e)$ and $\pi'(v)$ denotes the dealer's updated belief after having observed the investor's valuation v. The value VI can be interpreted as the highest price that a dealer of type ω is willing to pay to observe an investor's valuation, without actually trading the asset. To fix magnitudes, we normalize VI by the average intermediation spread.[60]

The estimated value of information is positive and substantial. As reported in the left-hand panel of Figure 2, for the average dealer the information conveyed by a trade with an investor is worth 12% of the intermediation spread, which corresponds to 7 basis points.

To illustrate the properties of VI, the right-hand panel of Figure 2 plots the value of information as a function of a dealer's type for Nebraska, which is the median state in terms of municipal debt outstanding for our sample. Overall, a piece of information is more valuable if it has the potential to affect the dealer's future trading decisions. As an example, a dealer values new information more if he is more uncertain about the current realization of the demand state. As an example, in the extreme case in which a dealer is certain about θ_t, no new information can change his assessment of the market and affect his future trading decisions. The value of information, therefore, is zero. Vice versa, the more uncertain a dealer is about the state of demand, the stronger is the impact of a new piece of information on his beliefs and future trading decisions. This, in turn, increases the value of information for the dealer.

Another key factor that determines a dealer's ability to modify his trading decisions in accord with a new piece of information is the impediments that he faces when trading, namely, inventory and search costs. As an example, when facing a large marginal cost of inventory, a dealer will want to offload his inventory, regardless of what he knows about the demand state. Consistent with this intuition, the right-hand panel of Figure 2 shows that information is more valuable for a lower level of inventory, which is associated with a lower marginal cost of inventory.

[60] In each state, we compute the average value of $VI(\omega)$ based on the equilibrium distribution of dealers' types.

Figure 2: Value of information

	Value of Information	
	Average	Standard Deviation
% of intermediation spread	12.41%	12.61%
Basis Points	7.72	5.64

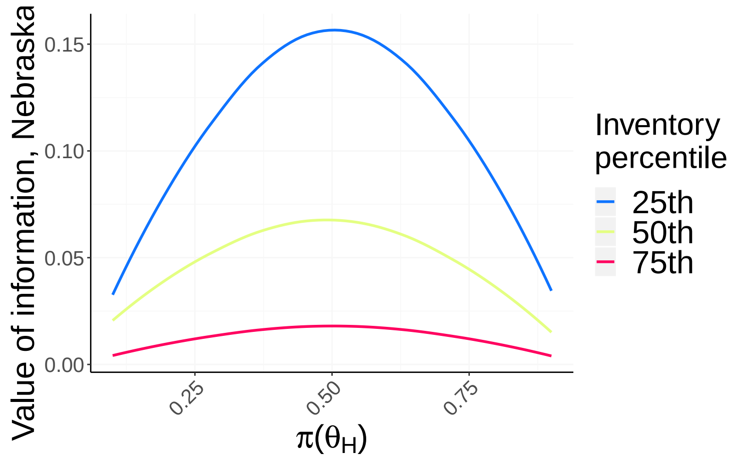

Notes: The right panel of the figure plots the value that a dealer assigns to the information conveyed by trading with an investor as a function of his prior beliefs π and inventory x for Nebraska, which is the median state in terms of municipal debt outstanding in our sample. We normalize the value of information by the average intermediation spread in the market. On the horizontal axis we show the prior probability assigned to the high demand state θ_H. Instead, we set the prior probability for the middle state to zero. Different lines correspond to different levels of inventory. The left panel of the figure reports the average value of information for dealers across states and its standard deviation.

7.2 Investor's Valuations

We use data on trading prices between investors and dealers to recover the investors' valuations point-wise. Consider as an example observing a trade at price $\hat{p}$ between a dealer and a buyer with valuation $\hat{v}$[61]

$$\hat{p} = \rho\hat{v} + (1-\rho)\max_{\omega} \Delta V_2(\omega, \hat{v}, \mathrm{b}), \tag{25}$$

where $\Delta V_2(\omega, \hat{v}, \mathrm{b})$, defined in Equation 7, denotes the valuation for the trade of a dealer in state ω who learned that the buyer has valuation $\hat{v}$.

Note that the dealer's value $\Delta V_2(\omega, \hat{v}, \mathrm{b})$ depends on the investor's valuation $\hat{v}$ only through his own posterior beliefs $\pi'(\hat{v}, a)$. Moreover, as we explain in detail in Section 6.2, we can rewrite the dealer's posterior beliefs as a function of the trading price, rather than the investor's valuation. Therefore, we can rewrite Equation 25 as

$$\hat{p} = \rho\hat{v} + (1-\rho)\max_{\omega} \Delta V_2(\omega, \hat{p}, \mathrm{b}), \tag{26}$$

where the term $\max_{\omega} \Delta V_2(\omega, \hat{p}, \mathrm{b})$ can be easily recovered, given the estimates in Section 7.1.

We calibrate the bargaining coefficient ρ, and leverage Equation 26 to recover the valuation of every investor involved in a trade. Next, we exploit the investors' entry conditions 9 to estimate the investors'

[61]The approach is analogous to the case in which the investor is a seller.

bargaining costs, ϕ_b and ϕ_s. Finally, to recover the distribution of investors' valuations before the entry decision, we assume that among potential investors of type a, the valuations are normally distributed with mean $\mu_v(\theta, a)$ and standard deviation $\sigma_v(\theta, a)$. These parameters are then recovered through maximum likelihood. The full set of estimates is reported in Table 15 in Appendix D.

8 Market transparency

Traditionally, assets in decentralized markets are traded in an opaque environment, with limited or no public information about market activity. However, in recent years, financial authorities in the US and abroad have been steadily improving access to information about trade activity in financial decentralized markets with the objective of increasing liquidity in these markets. This push toward greater transparency is still ongoing, both in the US and abroad.

The argument for the effectiveness of the policy goes as follows: dealers have an informational advantage vis-à-vis investors. By eroding this informational advantage, access to public information should encourage investors' participation in the market and increase market liquidity. This argument, however, ignores a key driving force of market liquidity: the dealers' incentives to trade.

In this section, we exploit our model to quantify the effect of an increase in market transparency *on the dealers' incentives to trade.* To approximate a transparent market, we simulate the model assuming that the terms of trade of all transactions become public at the end of each period. Once public, information about trade activity can be observed, free of charge, by everyone.

The left-hand panel of Figure 3 shows the results when implementing a policy that increases market transparency. For six out of ten states, the policy weakens dealers' incentives to trade, dampening the potentially positive effect of transparency on liquidity. Indeed, on average the volume of trade falls by 3%. There is substantial heterogeneity in the effect of the policy across states, as the change in average trade ranges from a decline in trading activity of −8% for assets issued in New Jersey to an increase of 4% for assets issued in Michigan.

Two effects are at play. First, transparency weakens the incentives to experiment: when information on trading activity is made public, uncertainty about common shock θ_t is drastically reduced. Therefore, the value of additional information conveyed by trade becomes irrelevant. This makes each trade less valuable for the dealers and implies that, conditional on type (π, x, e), dealers are willing to trade more sporadically. Second, improving public information reduces uncertainty about the realization of the demand shock θ_t.

Lower uncertainty implies that dealers are more willing to trade larger quantities of the asset, partially offsetting the first effect.

The balance between these two forces varies substantially across states. First, the model predicts that market transparency will favor states where demand uncertainty is lower. Indeed, demand uncertainty creates the scope to acquire information through trade. Figure 3 shows that the decline in trading activity is strongest in those states where demand is more volatile, such as Arizona, Vermont, and Mississippi. States where demand is more persistent, such as Michigan and Kansas, experience an increase in overall trading activity.

Importantly, demand uncertainty explains only a small portion of the effect of market transparency predicted by the model. Indeed, the first column of the right-hand panel of Figure 3, shows that differences in the persistence of the demand process explain only 9% of the differences in the predicted impact of market transparency across states. As an example, while demand for municipal bonds is relatively stable in New Jersey, this state sees the largest decline in market liquidity following the introduction of market transparency. Based on our estimates, dealers trading bonds issued in New Jersey face a low marginal cost of inventory and low search costs. Following the argument described in Section 7, this implies that information is particularly valuable here. This has a twofold implication for market transparency. First, information acquisition motives for trade are strong in New Jersey. Market transparency, therefore, will have a disproportionately negative effect on trading activity. Second, since incentives to acquire information are stronger, dealers are able to effectively reduce the uncertainty in the demand state through their trading activity. This, in turn, dampens the positive effect of market transparency described above.

We generalize this example by regressing the predicted change in trading activity on (i) the estimated persistence of the demand process, (ii) the estimated marginal cost of inventory and (iii) the estimated search costs for the dealer. The results are shown in the second column in the right-hand panel of Figure 3. The regression confirms the intuition that states where the marginal cost of inventory is larger see a sharper decline in market liquidity following the introduction of market transparency. By a similar token, in states where dealers face higher search costs the volume of trade falls more sharply.

Finally, we exploit our model to trace out how the predicted changes in trade activity translate into changes in agents' welfare. Unsurprisingly, improving the availability of public information about trade activity dramatically benefits dealers. Indeed, better information allows dealers both to improve the timing of their trades and to spare the costs involved in trading to acquire information. Both of these forces contribute to improving the dealers' profits by 2.7% on average.

On the other hand, the weakening of dealers' incentives to provide liquidity is harmful for investors. Under market transparency, the dealers exercise their market power curbing their trading activity. This results in substantial misallocation: there are investors who would profitably exchange an asset, but they are not served in a transparent market, due to the dealers' market power. By counteracting this effect, information acquisition motives for trade benefit investors and reduce the assets' misallocation in an opaque market. Indeed, in the states most affected by market transparency, New Jersey and Arizona, investors' welfare declines by 8.6% and 10.7%. Investors' welfare, instead, increases in states like Michigan and Kansas, where the total volume of trade increases. Overall investors' welfare decreases by 2%.

It is useful to conclude our analysis with some caveat in regards to the interpretation of these results. Our model focuses exclusively on dealers' incentives to trade. Consistent with this, our counterfactual analysis can only capture how market transparency affects dealers' trading behavior, holding investors' behavior fixed. As an example, we don't model how the investors' searching behavior and information change in reaction to market transparency. Naturally, a richer model would be necessary to derive clear cut policy recommendations. Nevertheless, our approach allows us to highlight an overlooked channel through which market transparency may have an unintended impact on market liquidity. The policy maker may want to take these repercussions into account to more completely assess the impact of market transparency in opaque OTC markets.

Figure 3: Impact of market transparency on trading activity

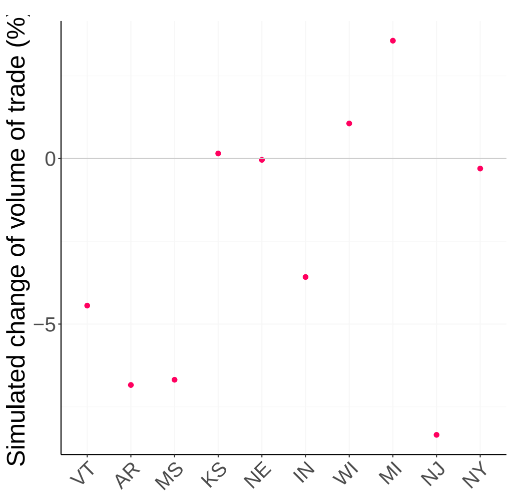

	Predicted change in liquidity	
	(I)	(II)
Demand persistence	2.17	9.613
	(2.079)	(2.520)
Marginal cost of inventory		25.459
		(7.790)
Average search cost		1.046
		(0.353)
R-squared	0.09	0.621

Notes: The left hand panel of the figure above displays the total change in trading volume associated to an increase in market transparency across different classes of assets. The right hand panel of the figure regresses the predicted outcome of the policy across states with the model's estimated parameters.

8.1 Model validation

As described in Section 2.2, on June 23, 2003, the MSRB started distributing daily summaries about the trading activity in the market during the previous day ("next -day reporting"), substantially improving access to public information about trade activity for market transparency. While this policy is implemented during our sample period, we only leverage data between January 2000 and June 2003 to estimate in our model. To validate our model, in this section we compare the observed outcome of the policy across assets issued by different states with the model's predictions.

We begin by estimating the observed effect of market transparency on across states following the same approach outlined in Section 3.2. In particular, we estimate Equation 2 separately for assets issued by different states, and we compare the estimated effect of market transparency with the our model prediction.

The left-hand panel of Figure 4 correlates the estimates of the impact of market transparency on total trading volume across states. Our model captures 51% of the heterogeneity in the policy's effect on trade activity across states. Moreover, the model correctly predicts that Michigan would benefit the most from market transparency, while New Jersey would suffer the most. Furthermore, as detailed in the right-hand panel of Figure 4, the model also correctly predicts the observed impact of the policy on the intermediation spread as well as on dealers' trading behavior in the inter-dealer market.

Figure 4: Predicted and observed impact of market transparency

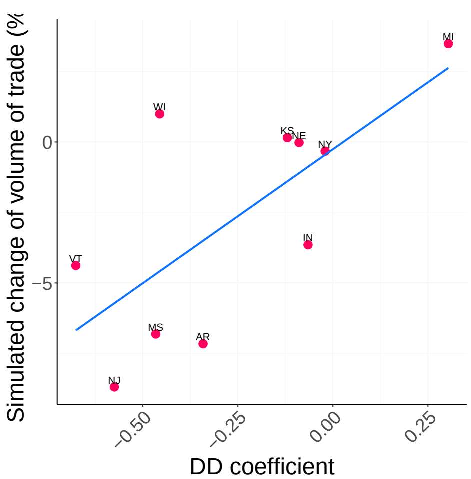

	Correlation between predicted and observed changes	R squared
Volume of trade btw dealers and investors	65%	50%
Dealers' purchases from investors	51%	24%
Inter-dealer trade	47%	23%
Intermediation spread	46%	21%

Notes: The figure shows the results of a validation exercise for our model. We exploit the estimated model to simulate the outcome of a policy intervention that increased market transparency in the market for muncipal bonds. The policy experiment is included in our sample period, but not in the sample we use to estimate the model. The left hand side of the figure compares the simulated effect of the policy on volume of trade with the diff-in-diff estimate of this effect. The right hand panel of the figure describes the correlation between the simulated and observed effect of the policy on intermediation spreads, dealers' purchases of the asset, inter-dealer trade, and total volume of trade between dealers and investors.

9 Conclusion

In this paper, we shed new light on the role of experimentation in decentralized opaque markets. These markets are common in wholesale trade markets and markets for investment goods. We argue that in these markets trade can be a source of valuable information about the market fundamentals. Obtaining this information, therefore, becomes an additional motive for trade.

To characterize incentives to experiment, we use a detailed dataset of transactions on the secondary market of municipal bonds. We first use the dataset to provide reduced form evidence suggesting that incentives to experiment are a first-order motive for trade in the market. To rationalize these facts we build a dynamic model of trade in decentralized markets where agents are uncertain about the underlying value of the asset traded. Finally we argue that information acquisition motives for trade shape the effectiveness of policies targeting market liquidity.

References

Acharya, Viral V, Yakov Amihud, and Sreedhar T Bharath, "Liquidity risk of corporate bond returns: conditional approach," *Journal of Financial Economics*, 2013, *110* (2), 358–386.

Admati, Anat R. and Paul Pfleiderer, "A theory of intraday patterns: Volume and price variability," *The Review of Financial Studies*, 1988, *1* (1), 3–40.

Aghion, Philippe, Maria Paz Espinosa, and Bruno Jullien, "Dynamic duopoly with learning through market experimentation," *Economic Theory*, 1993, *3* (3), 517–539.

Aguirregabiria, Victor and Pedro Mira, "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, 2007, *75* (1), 1–53.

Amihud, Yakov, "Illiquidity and stock returns: cross-section and time-series effects," *Journal of financial markets*, 2002, *1* (5), 31–56.

Arellano, Manuel and Stephen Bond, "Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations," *The Review of Economic Studies*, 1991, *58* (2), 277–297.

Asquith, Paul, Thom Covert, and Parag Pathak, "The effects of mandatory transparency in financial market design: Evidence from the corporate bond market," *National Bureau of Economic Research*, 2013.

Babus, Ana and Péter Kondor, "Trading and information diffusion in over-the-counter markets," *Econometrica*, 2018, *86* (5), 1727–1769.

Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin, "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 2007, *75* (5), 1331–1370.

Benkard, C. Lanier, "Learning and Forgetting: The Dynamics of Aircraft Production.," *American Economic Review*, 2000, *90* (4), 1034–1054.

_ , "A dynamic analysis of the market for wide-bodied commercial aircraft," *The Review of Economic Studies*, 2004, *7* (31), 581–611.

Bessembinder, Hendrik, William Maxwell, and Kumar Venkataraman, "Market transparency, liquidity externalities, and institutional trading costs in corporate bonds," *Journal of Financial Economics*, 2006, *82* (2), 251–288.

Binmore, Ken, Ariel Rubinstein, and Asher Wolinsky, "The Nash bargaining solution in economic modelling," *The RAND Journal of Economics*, 1986, pp. 176–188.

Bloomfield, Robert and Maureen O'Hara, "Market transparency: Who wins and who loses?," *The Review of Financial Studies*, 1999, *12* (1), 5–35.

_ **and** _ , "Can transparent markets survive?," *Journal of Financial Economics*, 2000, *55* (3), 425–459.

Blouin, Max R. and Roberto Serrano, "A decentralized market with common values uncertainty: Non-steady states," *The Review of Economic Studies*, 2001, *68* (2), 323–346.

Chade, Hector and Edward Schlee, "Another look at the Radner–Stiglitz nonconcavity in the value of information," *Journal of Economic Theory*, 2002, *107* (2), 421–452.

Chandrasekhar, Arun G., Horacio Larreguy, and Juan Pablo Xandri, "Testing models of social learning on networks: Evidence from a framed field experiment," *Work. Pap., Mass. Inst. Technol., Cambridge, MA*, 2012.

Dickstein, Michael J. and Eduardo Morales, "What do exporters know?," *National Bureau of Economic Research*, 2015.

Duffie, Darrell and Gustavo Manso, "Information percolation in large markets," *The American economic review*, 2007, *97* (2), 203–209.

_ **, Nicolae Gârleanu, and Lasse Heje Pedersen**, "Over-the-Counter Markets," *Econometrica*, 2005, *73* (6), 1815–1847.

_ **,** _ **, and** _ , "Valuation in over-the-counter markets," *The Review of Financial Studies*, 2007, *20* (6), 1865–1900.

Ellison, Glenn and Drew Fudenberg, "Rules of thumb for social learning," *Journal of Political Economy*, 1993, *101* (4), 612–643.

— **and** — , "Word-of-mouth communication and social learning," *The Quarterly Journal of Economics*, 1995, *110* (1), 93–125.

Ericson, Richard and Ariel Pakes, "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *The Review of Economic Studies*, 1995, *62* (1), 53–82.

Farboodi, Maryam, Gregor Jarosch, and Robert Shimer, "Meeting technologies in decentralized asset markets," *Working Paper*, 2016.

Feldstein, Sylvan G, Frank J Fabozzi et al., *The handbook of municipal bonds*, Vol. 155, John Wiley & Sons, 2008.

Gavazza, Alessandro, "Leasing and secondary markets: Theory and evidence from commercial aircraft," *Journal of Political Economy*, 2011, *119* (2), 325–377.

— , "The Role of Trading Frictions in Real Asset Markets," *American Economic Review*, 2011, *101* (4), 1106–43.

Gehrig, Thomas, "Intermediation in search markets," *Journal of Economics & Management Strategy*, 1993, *2* (1), 97–120.

Glode, Vincent and Christian Opp, "Asymmetric information and intermediation chains," *American Economic Review*, 2016, *106* (9), 2699–2721.

Golosov, Mikhail, Guido Lorenzoni, and Aleh Tsyvinski, "Decentralized trading with private information," *Econometrica*, 2014, *82* (3), 1055–1091.

Gourieroux, Christian, Alain Monfort, and Eric Renault, "Indirect inference," *Journal of Applied Econometrics*, 1993, *8* (S1).

Green, Richard C, Burton Hollifield, and Norman Schürhoff, "Financial intermediation and the costs of trading in an opaque market," *The Review of Financial Studies*, 2006, *20* (2), 275–314.

Green, Richard C., Dan Li, and Norman Schürhoff, "Price discovery in illiquid markets: Do financial asset prices rise faster than they fall?," *The Journal of Finance*, 2010, *65* (5), 1669–1702.

Grossman, Sanford J and Motty Perry, "Perfect sequential equilibrium," *Journal of economic theory*, 1986, *39* (1), 97–119.

Grossman, Sanford J., Richard E. Kihlstrom, and Leonard J. Mirman, "A Bayesian approach to the production of information and learning by doing," *The Review of Economic Studies*, 1977, pp. 533–547.

Gul, Faruk and Hugo Sonnenschein, "On delay in bargaining with one-sided uncertainty," *Econometrica: Journal of the Econometric Society*, 1988, pp. 601–611.

Harris, Lawrence E and Michael S Piwowar, "Secondary trading costs in the municipal bond market," *The Journal of Finance*, 2006, *61* (3), 1361–1397.

Hollifield, Burton, Artem Neklyudov, and Chester Spatt, "Bid-ask spreads, trading networks and the pricing of securitizations: 144a vs," Technical Report, registered securitization. Working paper, CMU and HEC Lausanne 2014.

_ , _ , and _ , "Bid-ask spreads, trading networks, and the pricing of securitizations," *The Review of Financial Studies*, 2017, *30* (9), 3048–3085.

Hopenhayn, Hugo A., "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 1992, *60* (5), 1127–1150.

Hörner, Johannes and Andrzej Skrzypacz, "Learning, Experimentation and Information Design," 2016.

Hugonnier, Julien, Benjamin Lester, and Pierre-Olivier Weill, "Heterogeneity in decentralized asset markets," *National Bureau of Economic Research*, 2014.

_ , _ , and _ , "Heterogeneity in decentralized asset markets," Technical Report, National Bureau of Economic Research 2014.

_ , _ , and _ , "Frictional intermediation in over-the-counter markets," Technical Report, National Bureau of Economic Research 2018.

_ , Semyon Malamud, and Eugene Trubowitz, "Endogenous completeness of diffusion driven equilibrium markets," *Econometrica*, 2012, *80* (3), 1249–1270.

Jovanovic, Boyan, "Learning By Doing and the Choice of Technology.," *Econometrica*, 1996, *64* (6), 1299–1310.

Kalouptsidi, Myrto, "Detection and Impact of Industrial Subsidies: The Case of World Shipbuilding," *American Economic Review*, 2014, *104* (2), 564–608.

Keppo, Jussi, Giuseppe Moscarini, and Lones Smith, "The demand for information: More heat than light," *Journal of Economic Theory*, 2008, *138* (1), 21–50.

Kihlstrom, Richard, L. Mirman, and Andrew Postlewaite, "Experimental Consumption and the 'Rothschild Effect'," *Bayesian Models of Economic Theory*, 1984.

Kyle, Albert S., "Continuous auctions and insider trading," *Econometrica*, 1985, pp. 1315–1335.

_ , "Informed speculation with imperfect competition," *The Review of Economic Studies*, 1989, *56* (3), 317–355.

Leach, J. Chris and Ananth N. Madhavan, "Price experimentation and security market structure," *The Review of Financial Studies*, 1993, *6* (2), 375–404.

Lerman, Kristina, Rumi Ghosh, and Jeon Hyung Kang, "Centrality metric for dynamic networks," *Proceedings of the Eighth Workshop on Mining and Learning with Graphs*, 2010, (70–77).

Lester, Benjamin, Ali Shourideh, Venky Venkateswaran, and Ariel Zetlin-Jones, "Market-making with search and information frictions," Technical Report, National Bureau of Economic Research 2018.

Li, Dan and Norman Schürhoff, "Dealer networks," *The Journal of Finance*, 2019, *74* (1), 91–144.

Maggio, Marco Di, Amir Kermani, and Zhaogang Song, "The value of trading relations in turbulent times," *Journal of Financial Economics*, 2017, *124* (2), 266–284.

Menzio, Guido, "A cheap-talk theory of random and directed search," *Manscript, Univ. Pennsylvania*, 2005.

Miao, Jianjun, "A search model of centralized and decentralized trade," *Review of Economic Dynamics*, 2006, *9* (1), 68–92.

Mirman, Leonard J., Larry Samuelson, and Amparo Urbano, "Monopoly experimentation," *International Economic Review*, 1993, pp. 549–563.

Neklyudov, Artem, "Bid-ask spreads and the over-the-counter interdealer markets: Core and peripheral dealers," *Review of Economic Dynamics*, 2019, *33*, 57–84.

Pagano, Marco, "Trading volume and asset liquidity," *The Quarterly Journal of Economics*, 1989, *104* (2), 255–274.

Pakes, Ariel, Michael Ostrovsky, and Steven T. Berry, "Simple Estimators for the Parameters of Discrete Dynamics Games," *The RAND Journal of Economics*, 2007, *38* (2), 373–399.

Radner, Roy and Joseph Stiglitz, "A Nonconcavity in the Value of Information," *Bayesian models in Economic Theory*, 1984, *5*, 33–52.

Rust, John, "Optimal replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 1987, *55* (5), 999–1033.

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

Ryan, Stephen P., "The Costs of Environmental Regulation in a Concentrated Industry," *Econometrica*, 2012, *80* (3), 1019–1062.

Schultz, Paul, "The market for new issues of municipal bonds: The roles of transparency and limited access to retail investors," *Journal of Financial Economics*, 2012, *106* (3), 492–512.

Spulber, Daniel F, "Market microstructure and intermediation," *Journal of Economic perspectives*, 1996, *10* (3), 135–152.

Su, Che-Lin and Kenneth L Judd, "Constrained optimization approaches to estimation of structural models," *Econometrica*, 2012, *80* (5), 2213–2230.

Wang, Junbo, Chunchi Wu, and Frank X. Zhang, "Liquidity, default, taxes, and yields on municipal bonds," *Journal of Banking & Finance*, 2008, *32* (6), 1133–1149.

Weintraub, Gabriel Y., C. Lanier Benkard, and Benjamin Van Roy, "Markov perfect industry dynamics with many firms," *Econometrica*, 2008, *76* (6), 1375–1411.

Wolinsky, Asher, "Information revelation in a market with pairwise meetings," *Econometrica*, 1990, pp. 1–23.

Appendix

A Additional tables and figures

A.1 Learning through inter-dealer trade

A natural way to potentially falsify the result in Figure 1, is to compare it with the impact of an inter-dealer trade on dealers' pricing strategies after the introduction of "next day reporting" described in Section 2.1. The introduction of this system increased market transparency and made information about trading prices easily accessible for everyone in the market. Under this condition, the informative content of a single inter-dealer trade should become negligible and the price response in Figure 1 should disappear. To test this, we estimate Equation 27 below using inter-dealer trades taking place before and after the policy intervention. The variable opaque_i is a dummy variable that equals one if inter-dealer trade i takes place before the policy intervention, while $\text{post}_{i,t}$ equals one in period t following inter-dealer trade i.

$$\left|\hat{p}_{a_i,d_i,t} - \hat{p}_{a_i,\tilde{d}_i,t}\right| = \alpha_{a_i,d_i,d'_i} + \delta_1 \text{post}_{i,t} + \gamma \text{post}_{i,t} \times \text{opaque}_i + \epsilon_{i,t}. \tag{27}$$

As shown in Column II in Table 8 the coefficient γ is large and significantly different from zero. Figure 5 shows the corresponding event study. This confirms that in the opaque regime inter-dealer trades have a substantially stronger effect on dealers' pricing strategies compared to trades in the transparent regime.

This supports the interpretation that the effect displayed in Figure 1 is indeed driven by information acquisition.

Finally, one might worry that the result in 1 is driven by a release of public information that equalizes the trading behavior across *all* market agents. To verify whether this is the case, we construct a set of placebo trades. For each inter-dealer trade i, completed in period t_i, between dealers d_i and $\tilde{d}_i$, we create a placebo inter-dealer trade in period t_i between d_i and a fictitious trading partner $\tilde{f}_i$[62]. If the results in Figure 1 were driven by a common shock, we should see dealers d_i and $\tilde{f}_i$ pricing the asset more similarly after period t_i, despite not having traded with each other. We stack all the observed and placebo trades and estimate Equation 28 where the variable obs_i is a dummy variable that equals one if trade i is not fictitious, while $\text{post}_{i,t}$ equals one in the periods following the (observed or fictitious) inter-dealer trade.

$$\left|\hat{p}_{a_i,d_i,t} - \hat{p}_{a_i,\tilde{d}_i,t}\right| = \alpha_{a_i,d_i,\tilde{d}_i} + \delta_1 \text{post}_{i,t} + \gamma \text{post}_{i,t} \times \text{obs}_i + \epsilon_{i,t}. \tag{28}$$

Once again, Column III in Table 8 shows that the coefficient γ remains significantly different from zero even after controlling for market-level movements in prices, captured by δ_1. Furthermore the parameter δ_1 is not significant, suggesting that the difference in trading price doesn't fall for dealers who haven't traded with each other.

A.2 The effect of market transparency

A liquid secondary market is a crucial condition to lower the cost of raising capital. As an example, Wang et al. (2008) estimate that municipal bond issuers pay 13 billion a year to compensate investors for the risks implied by the illiquidity of the market. Decreasing liquidity in this market, therefore, might be very costly for local governments and municipalities. To test whether this is the case, we quantify the impact of information dissemination on the cost of capital for municipalities, as measured by yield on new issues. We use specification

$$y_{it} = \kappa_t + \gamma_0 post_t + \gamma_1 unins_i + \lambda post_t \times unins_i + \epsilon_{it}, \tag{29}$$

where y_{it} is the offering yield for issuance i at time t; κ_t is a vector of month fixed effects; $post_i$ is an indicator for the issuances on weeks after the policy intervention; and $unins_i$ is an indicator that equals

[62]The fictitious partner $\tilde{f}_i$ has traded asset a_i at least once in the previous year and has traded at least once with dealer d_i. The result are robust if we select fictitious partner $\tilde{f}_i$ to also have the same size and specialization as d's observed trading partner.

Table 8: Placebo tests

	I	II	III
		$\left\|\hat{p}_{a,t,d} - \hat{p}_{a,t,\tilde{d}}\right\|$	
$\text{post}_{a,d,d',t}$	-0.173^{***}	0.012	-0.042
	(0.040)	(0.054)	(0.027)
$\text{post}_{a,d,d',t} \times \text{opaque}_{a,d,d'}$		-0.186^{***}	
		(0.067)	
$\text{post}_{a,d,d',t} \times \text{obs}_{a,d,d'}$			-0.130^{***}
			(0.042)
N	465,053	687,453	1,184,240
R^2	0.558	0.564	0.599
Level	Pair × Asset × Period		

$^{***}p \leq 0.01$, $^{**}p \leq 0.05$, $^{*}p \leq 0.1$

Notes: The first column in the table shows the estimate of the average effect of an inter-dealer trade on the difference in the price at which the involved dealers trade with investors. To confirm that the change in behavior is indeed driven by information acquisition, Column II shows that the pattern disappears after a policy intervention improving market transparency. To test whether the result is driven by asset-specific demand shocks, in Column III we show that the same pattern is not present for dealers who haven't traded with one another. Due to the infrequency of trades, we consider a three-day interval as a period. In all three regressions we consider a window of five periods before the trade and ten periods after.

one if the asset issued are uninsured.[63] Table 12 shows that the offering yield for uninsured assets increases by 0.09 percentage points (pp) compared to insured assets, corresponding to an increase of 2.3% of the cost of issuing uninsured assets compared to insured ones.

Furthermore, an important implication of our model is that inter-dealer trade should also fall as a consequence of the introduction of market transparency. To test whether this is the case, we quantify the impact of information dissemination on inter-dealer trade using specification Our main specification is

$$y_{it} = \phi_i + \kappa_t + \gamma post_t + \lambda post_t \times unins_i + \epsilon_{it}, \tag{30}$$

where y_{it} is bond i's outcome in week t; ϕ_i is a vector of asset fixed effects; κ_t is a vector of week fixed effects; and $post_t$ is an indicator for the trade outcomes on weeks after the policy intervention. Finally, $unins_i$ is an indicator that equals one if the asset's principal is uninsured. Since there are repeated observations per

[63]The result are robust if we control for the issuance's characteristics such as rating, maturity, or state of issuance.

Figure 5: Change in pricing behavior after inter-dealer trade under transparency

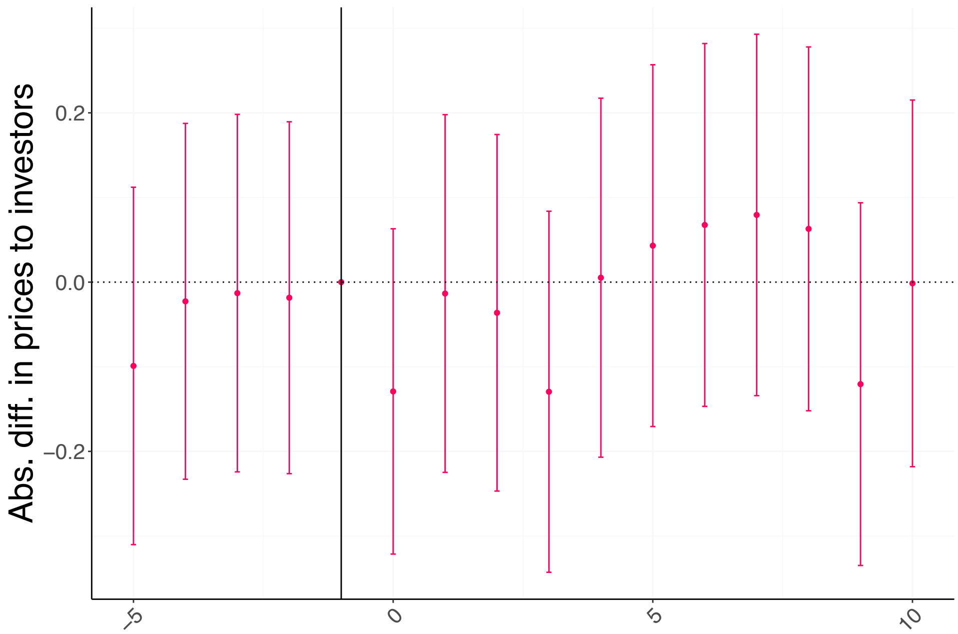

Notes: The figure plots the weekly regression coefficients from estimating Equation 1 after the introduction of market transparency. Dealers trade asset with one another in period -1. The outcome variable is the difference, among the two dealers, in the weighted average price for trades of asset with investors. The coefficients are plotted relative to the difference in prices to investors among the two dealers in -1, which are normalized to zero. Due to the infrequency of trades, we consider a three-day interval as a period. The results are robust to different fixed effects, and a similar pattern emerges using the logarithm of the difference in prices.

Table 9: Effect of market transparency on the cost of capital

	Offering Yield		
	2 Months	3 Months	6 Months
uninsured * $post_t$	0.090**	0.091**	0.046**
	(0.002)	(0.001)	(0.002)
N	19,730	27,165	50,691
Level		Issue	

$^{**}p \leq 0.05$, $^{*}p \leq 0.1$

Notes: The table presents the output from the difference-in-difference regression that measures the effect of the change in market transparency on offering yields (last three columns). We use insured assets as control group. Observations are at the issuance level, and standard errors are clustered at the issuer level.

Figure 6: Event study for the number of trades

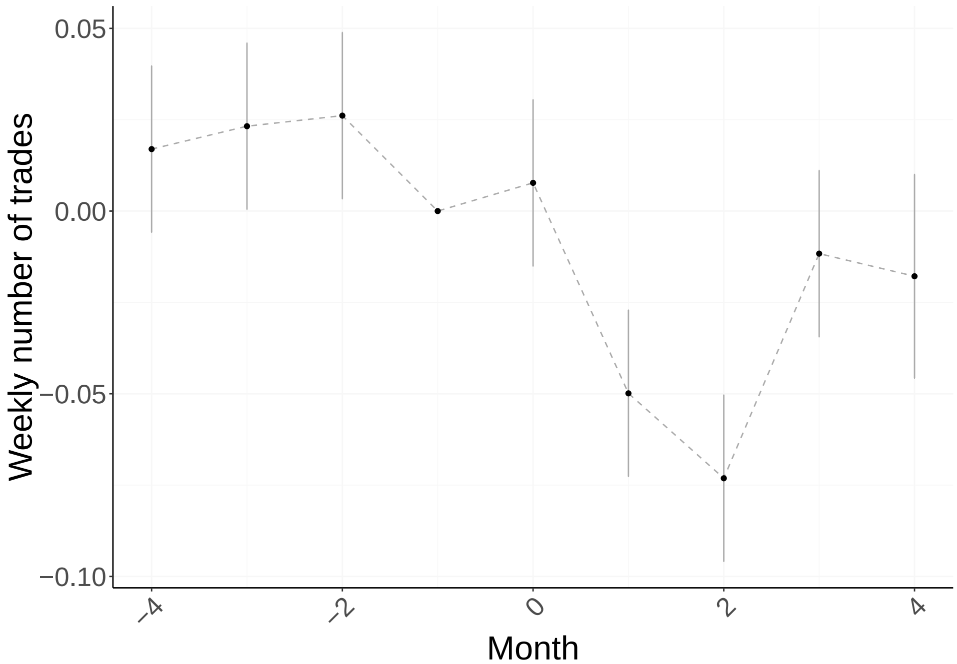

Notes: The figure plots the monthly regression coefficients and 95% confidence intervals from estimating equation 1 for a four-month window around the introduction of "next day reporting." The outcome variable is the number of trades between dealers and investors. The coefficients are plotted relative to the number of trades in $k = -1$, which are normalized to zero. The results are robust to different fixed effects, and a similar pattern emerges using the logarithm of the difference in prices.

bond, in all estimates, the standard errors are clustered by bond. Table 10 shows that the estimate of the effect of transparency on the number of trades per day is negative and significant for all three estimation windows, as the number of trades drops by approximately 0.013, which corresponds to 15% of the average level of trade before dissemination. Finally, Table 11 shows the effect of market transparency on volume of trade between dealers and investors using alternative measures of trading activity: the log number of trades and the log volume of trade.

An alternative measure of market liquidity is the intermediation spread, which measures the difference between the price at which dealers buy and sell the assets. The standard approach to measure the intermediation spread is to compare average ask and bid prices for a given asset within a certain time period. This approach is not robust for municipal bonds, due to the infrequency of trades. Instead, to obtain an estimate of the effect of market transparency, we rely on a triple-difference approach. In particular, denoting p_{ajt} the trading price in trade j for asset a in period t, we estimate the following

Table 10: Difference-in-difference estimates of transparency on inter-dealer trade

	Number of Trades		
	2 Months	3 Months	6 Months
uninsured * $post_t$	-0.008**	-0.013**	-0.013**
	(0.002)	(0.002)	(0.002)
N	1,327,819	2,030,782	2,655,638
Level		Issue-Week	

$^{**}p \leq 0.05$, $^{*}p \leq 0.1$

Notes: The table presents the output from the difference-in-difference regression that measures the effect of the change in market transparency on inter-dealer trading activity. We use insured assets as control group. Observations are at the asset-week level, and standard errors are clustered at the asset level.

regression

$$p_{jta} = \kappa_t + \alpha_a + \gamma_0 post_t + \gamma_1 unins_a + \gamma_2 sale_j \\ + \gamma_3 post_t \times unins_a + \gamma_4 post_t \times sale_j + \gamma_5 sale_j \times unins_a \\ + \lambda post_t \times unins_a \times sale_j + \epsilon_{it},$$

where $post_t$ is an indicator for the trades completed after the policy intervention, $unins_a$ is an indicator that equals one if the asset issued is uninsured, and $sale_j$ is an indication that equals one if the trade was a sale from a dealer to an investor. Finally, κ_t is a week fixed effect, and α_j is an asset fixed effect. The parameter of interest is λ, which measures the disproportionate effect of the policy on selling prices for uninsured assets. The results are summarized in Table 12, that shows that the difference between the price at which dealers sell and buy uninsured asset increase compared to the same quantity for insured assets.

Table 11: Difference-in-difference estimates

	Number of Trades (log)			Volume of Trade (log)		
	2 Months	4 Months	6 Months	2 Months	4 Months	6 Months
uninsured * $post_t$	-0.007**	-0.015**	-0.019**	0.007	-0.037**	-0.055**
	(0.003)	(0.003)	(0.003)	(0.011)	(0.009)	(0.008)
N	479,332	958,664	1,466,192	479,332	958,664	1,466,192
Level	Issue-Week					

$^{**}p \leq 0.05$, $^{*}p \leq 0.1$

Notes: The table presents the output from the difference-in-difference regression that measures the effect of the change in market transparency on trading activity measured by the number of trades (first three columns) and volume of trade (last three columns). We use insured assets as control group. Observations are at the asset-week level, and standard errors are clustered at the asset level.

Table 12: Effect of market transparency on intermediation spread

	Trading Price		
	2 Months	3 Months	6 Months
uninsured * $post_t$ * sale	0.001**	0.001**	0.001***
	(0.0002)	(0.0002)	(0.0002)
N	548,215	820,100	1,102,682
Level	Trade		

$^{**}p \leq 0.05$, $^{*}p \leq 0.1$

Notes: The above table presents the output from the difference-in-difference regression that measures the effect of the change in market transparency on the intermediation spread. We use insured assets as control group. Observations are at the trade level, and standard errors are clustered at the issuer level.

B What do dealers know?

We use the specification test suggested in Dickstein and Morales (2015) to test the assumption that dealers have no information about the market value of the asset in months where they don't participate in trade.

The intuition behind the test is the following: let $y_{d,t}$ be an outcome variable that depends on a decision of dealer d in period t, such as the quantity traded in a certain asset, or the price charged to investors. Let $\mathcal{I}_{d,t}$ also denote dealer d's information set at the beginning of period t. Dealer d's decision about $y_{d,t}$ will depend on dealer d's expectation of the market value for the asset $\mathbb{E}\left(\theta_t|\mathcal{I}_{d,t}\right)$, conditional on what he knows about past realizations of θ_t. Under this scenario, if a variable Z_t belongs to d's information set $\mathcal{I}_{d,t}$, then it must be orthogonal to his forecast error:

$$\mathbb{E}\left[\left(\theta_t - \mathbb{E}\left(\theta_t|\mathcal{I}_{d,t}\right)\right) Z_t\right] = 0.$$

In this case, therefore, Z_t would be an instrument for $\mathbb{E}\left(\theta_t|\mathcal{I}_{d,t}\right)$ in the regression

$$\begin{aligned} y_{d,t} &= \alpha + \beta\theta_t + \beta\left(\mathbb{E}\left(\theta_t|\mathcal{I}_{d,t}\right) - \theta_t\right) \\ &= \alpha + \beta\theta_t + \epsilon_{d,t} \end{aligned}$$

We use this idea to test whether the dealer knows the average market price for an asset that he does not trade in a given month. Table 13 reports the result of this test for different outcome variables $y_{d,t}$ and instruments Z_t. The first two rows test whether the dealer knows the average trading price of an asset in periods in which he does not trade. In all four of the combinations the p-value is zero, suggesting that the average price for the asset, $\theta_{t,a}$, or for assets from the same state $\theta_{t,s}$, does not belong to the dealer's information set when he does not trade. On the contrary, for periods in which the dealers did participate in trade the test cannot reject the null, confirming that dealer d acquires information through trade.

	$y_{d,t} = par_{d,t}^{buy}$	$y_{d,t} = par_{d,t}^{sell}$
$\left(1 - \mathbb{I}\left\{\text{trade in } t-1\right\}\right)\theta_{t-1,a}$	0.00	0.00
$\left(1 - \mathbb{I}\left\{\text{trade in } t-1\right\}\right)\theta_{t-1,s}$	0.00	0.00
$\mathbb{I}\left\{\text{trade in } t-1\right\}\theta_{t-1,a}$	0.02	0.75
$\mathbb{I}\left\{\text{trade in } t-1\right\}\theta_{t-1,s}$	0.10	0.15

Table 13: p-values for Hansen-Sargan test

C Derivation of choice probabilities

First, consider the probability $\mathbb{P}(\text{no trade}|\omega, \text{b}, n)$ that a dealer with type a will not trade after having traded with n investors. First remember that the dealer will contact the investor only when the search cost c drawn satisfies

$$c \leq \mathbb{E}\left[p(\omega, v, \text{s}) + V_2\left(\omega'(v, \text{s}), \text{s}\right)\right] - V_2(\omega, \text{s}) \equiv K_s(\omega, \text{s}).$$

Then, we have

$$\begin{aligned}
\mathbb{P}(\text{no trade}|\omega, \text{s}, n) &= \sum_{k\geq 0} \mathbb{P}(\text{can contact } k \text{ investors})\left(1 - F_s\left(K_s(\omega, \text{s})\right)\right)^k \\
&= \sum_{k\geq 0} (1-\gamma)\gamma^k \left(1 - F_s\left(K_s(\omega, \text{s})\right)\right)^k \qquad (31)\\
&= \frac{(1-\gamma)}{1 - \gamma\Big(1 - F_s\left(K_s(\omega, \text{s})\right)\Big)}. \qquad (32)
\end{aligned}$$

It follows, that the probability that a dealer will trade at least once more, after having traded with n investors satisfies

$$\mathbb{P}(n+1|\omega, \text{b}, n) = \frac{\gamma F_s\left(K_s(\omega, \text{s})\right)}{1 - \gamma\Big(1 - F_s\left(K_s(\omega, \text{s})\right)\Big)}.$$

D Estimates

Table 14: Dealers' cost estimates

	α	δ
Vermont	0.014	0.575
	(0.105)	(0.264)
Arizona	0.451	0.534
	(0.308)	(0.205)
Nebraska	0.493	0.326
	(0.163)	(0.097)
Kansas	0.015	0.911
	(0.012)	(0.041)
Mississippi	0.473	0.175
	(0.46)	(0.47)
Indiana	0.370	0.451
	(0.131)	(0.093)
Wisconsin	0.316	0.397
	(0.105)	(0.264)
New Jersey	0.059	0.900
	(0.017)	(0.033)
New York	0.047	0.586
	(0.042)	(0.099)
Michigan	0.387	0.461
	(0.340)	(0.297)

Notes: The table shows the estimates of the experience process defined in Equation 19. We cluster the assets in our sample based on the state of issuance and estimate the experience process independently across groups. The parameters are estimated through non-linear least squares, and the estimation leverages Equation 20. Identification of the parameters relies on the correlation of the prices dealers pay on the inter-dealer market with their experience, after controlling for seller-month-asset fixed effects.

Table 15: Dealers' cost estimates

	Vermont	Arizona	Nebraska	Kansas	Mississippi	Indiana	Wisconsin	New Jersey	New York	Michigan
κ_0	8.7e-03	0.004	2.1e-04	0.003	-6.4e-04	1.3e-04	-6.69e-04	0.006	-0.003	3.3e-04
	(2.6e-03)	(1.5e-03)	(2.14e-03)	(3.9e-05)	(2.7e-04)	(3.7e-05)	(2.7e-04)	(1.1e-03)	(3.1e-04)	(1.3e-04)
κ_1	-5.7e-05	-2.0e-05	-6.8e-06	-5.1e-06	-4.2e-06	-6.9e-06	-2.84e-06	-1.0e-06	2.4e-06	4.3e-06
	(9.2e-06)	(3.9e-06)	(2.03e-06)	(1.2e-06)	(1.0e-06)	(1.1e-06)	(9.8e-07)	(2.7e-07)	(5.8e-07)	(1.2e-06)
κ_2	1.1e-07	2.1e-08	5.4e-09	1.8e-09	8.3e-10	6.1e-09	2.58e-09	-3.2e-09	-5.3e-10	-7.3e-10
	(2.0e-08)	(4.7e-10)	(1.55e-09)	(1.2e-10)	(2.8e-10)	(1.1e-10)	(1.2e-10)	(7.8e-10)	(2.0e-10)	(3.2e-10)
c_{b}	-0.189	-0.145	-0.091	-0.539	-0.013	-0.006	-0.047	-0.184	0.012	-0.101
	(0.077)	(0.092)	(0.016)	(0.003)	(0.010)	(0.002)	(0.070)	(0.034)	(0.004)	(0.031)
c_{s}	-0.148	-0.139	-0.117	-0.416	-0.188	-0.129	-0.145	-0.150	-0.026	-0.135
	(0.069)	(0.091)	(0.034)	(0.010)	(0.075)	(0.008)	(0.099)	(0.032)	(0.006)	(0.020)
$c_{\tilde{e}_L}$	0	-0.062	-0.003	-0.048	4.75e-04	-0.023	-0.038	-0.007	-0.01	6.89e-05
	(1.5e-04)	(0.013)	(0.009)	(0.009)	(1.9e-04)	(0.007)	(0.006)	(0.002)	(0.0003)	(2.7e-05)
$c_{\tilde{e}_M}$	0.01	-0.034	-0.008	-0.014	-0.004	-0.029	0.018	-0.020	-0.011	0.037
	(3.4e-04)	(0.012)	(0.003)	(0.011)	(0.002)	(0.008)	(0.006)	(0.006)	(0.004)	(0.011)
$c_{\tilde{e}_H}$	-0.012	0.002	-0.006	0.004	-0.018	0.002	0.032	0.001	-0.01	-4.19e-04
	(0.004)	(0.001)	(0.001)	(0.0003)	(6.2e-03)	(0.001)	(0.006)	(0.001)	(0.0001)	(0.0001)
σ_ϵ	1.161	2.216	0.367	4.144	1.501	0.890	1.821	4.586	0.813	0.628
	(0.868)	(0.973)	(0.419)	(0.139)	(1.072)	(0.139)	(0.729)	(0.354)	(0.568)	(0.320)
σ_c	0.959	0.949	1.177	2.482	0.786	0.559	1.349	1.718	0.544	0.590
	(0.428)	(0.609)	(0.093)	(0.037)	(0.318)	(0.037)	(0.675)	(0.302)	(0.050)	(0.211)
σ_ξ	0.590	0.327	0.207	0.691	0.510	0.419	0.431	0.555	0.663	0.455
	(0.176)	(0.118)	(0.161)	(0.045)	(0.513)	(0.045)	(0.122)	(0.064)	(0.331)	(0.096)

Notes: The table summarizes the estimates of dealers' cost parameters. We cluster the assets based on that state of issuance and estimate the parameters independently across groups. To ease the comparison across states, all the fixed costs of trading are normalized by the average trade size. All other estimates are expressed in 1,000 USD.

Table 16: Investors' cost estimates

	ϕ_{b}	ϕ_{s}	$\mathbb{E}_{\text{b}}(v\|\theta_H)$	$\mathbb{E}_{\text{b}}(v\|\theta_M)$	$\mathbb{E}_{\text{b}}(v\|\theta_L)$	$\mathbb{E}_{\text{s}}(v\|\theta_H)$	$\mathbb{E}_{\text{s}}(v\|\theta_M)$	$\mathbb{E}_{\text{s}}(v\|\theta_L)$
VT	-0.025	-0.013	1.063	1.048	1.023	1.045	1.029	1.006
	(0.006)	(0.006)	(2.42e-04)	(2.11e-04)	(2.31e-04)	(2.38e-04)	(2.28e-04)	(3.09e-04)
AR	-0.019	-0.008	1.030	1.021	1.005	1.022	1.013	0.996
	(0.006)	(0.005)	(2.77e-04)	(2.89e-04)	(2.70e-04)	(3.31e-04)	(2.48e-04)	(3.33e-04)
NE	-0.099	0.074	1.062	1.046	1.031	1.043	1.031	1.017
	(0.004)	(0.004)	(2.13e-04)	(2.07e-04)	(1.59e-04)	(9.49e-04)	(3.50e-04)	(2.11e-04)
KS	-0.015	-0.022	1.053	1.028	1.013	1.034	1.011	0.997
	(0.001)	(0.001)	(3.24e-04)	(3.37e-04)	(4.23e-04)	(2.92e-04)	(3.54e-04)	(3.74e-04)
MS	-0.050	0.023	1.022	1.006	0.985	1.009	0.995	0.975
	(0.010)	(0.012)	(2.4e-04)	(2.3e-04)	(2.5e-04)	(3.2e-04)	(2.7e-04)	(2.7e-04)
IN	-0.035	0.014	1.006	0.990	0.966	0.999	0.983	0.959
	(0.002)	(0.003)	(9.53e-05)	(1.28e-04)	(1.24e-04)	(1.69e-04)	(7.59e-04)	(2.80e-04)
WI	-0.001	-0.030	1.009	0.994	0.970	1.002	0.992	0.966
	(0.005)	(0.008)	(3.17e-04)	(2.77e-04)	(2.57e-04)	(3.48e-04)	(2.31e-04)	(2.90e-04)
NJ	0.001	-0.038	1.049	1.026	1.010	1.039	1.017	1.000
	(0.002)	(0.002)	(6.9e-04)	(3.1e-04)	(3.9e-04)	(3.0e-04)	(2.9e-04)	(3.0e-04)
NY	-0.002	-0.019	1.036	1.022	0.999	1.029	1.017	0.995
	(0.001)	(0.001)	(2.76e-04)	(2.54e-04)	(2.42e-04)	(2.64e-04)	(2.44e-04)	(3.03e-04)
MI	-0.017	-0.028	1.044	1.029	1.003	1.036	1.023	0.995
	(0.003)	(0.004)	(2.5e-04)	(2.7e-04)	(2.2e-04)	(2.7e-04)	(2.4e-04)	(2.8e-04)

Notes: The table summarizes the estimates of investors' costs and valuations. We cluster the assets based on that state of issuance and estimate the parameters independently across groups. To ease the comparison across states, both valuation and entry costs are expressed as a fraction of the trade size.

E Demand state transition

We specify the evolution of the demand state variable θ_t as a first-order Markov process with transition $\mathbb{P}_\theta$ and initial distribution $\mathbb{P}_0$. In each period, θ_t takes values in $\{\theta_L, \theta_M, \theta_H\}$.

To recover the process $(\mathbb{P}_\theta, \mathbb{P}_0)$ we leverage a large sample approximation of the distribution of the average market price in trades between dealers and investors. Let $\hat{p}_t(a)$ denote the average price at which dealers trade with investors of type $a \in \{\text{b}, \text{s}\}$ in period t. Given Equation 8, we have:

$$\hat{p}_t(\text{b}) = \frac{1}{N_t(\text{b})} \sum_i \left(\rho v_{it} + (1-\rho) \max_\omega \Delta V_2(\omega, v_{it}, \text{b}) \right),$$

and

$$\hat{p}_t(\text{s}) = \frac{1}{N_t(\text{s})} \sum_i \left((1-\rho) \min_\omega \Delta V_2(\omega, v_{it}, \text{s}) + \rho v_{it} \right)$$

where $N_t(a)$ is the total number of trades between dealers and investors of type a observed in period t.

Note that, for large numbers of trades $N_t(a)$, the average price $\hat{p}_t(a)$ converges in probability to its expected value, which we denote by $\mu(\theta_t, a)$. Similarly, for large numbers of trades $N_t(a)$ the distribution of the normalized prices $\sqrt{N_t(a)}\,(\hat{p}_t(a) - \mu(\theta_t, a))$ approximates that of a normal distribution with mean 0 and finite variance, which we denote by $\sigma^2(\theta_t, a)$. Importantly, the distribution of the normalized prices changes over time only through changes in θ_t.[64] Therefore, we can write the likelihood of observing prices $\{\hat{p}_t(a)\}_{t=0}^T$ as

$$\mathcal{L}\left(\{\hat{p}_t(a)\}_{t=0}^T | \mathbb{P}_\theta, \mathbb{P}_0, \mu, \sigma, \{N_t(a)\}_{t=0}^T \right) = \sum_{\theta_0, \cdots, \theta_T} \prod_{t=0}^T \phi\left(\sqrt{N_t(\text{s})} \frac{\hat{p}_t(\text{s}) - \mu(\theta_t, \text{s})}{\sigma(\theta_t, \text{s})} \right) \phi\left(\sqrt{N_t(\text{b})} \frac{\hat{p}_t(\text{b}) - \mu(\theta_t, \text{b})}{\sigma(\theta_t, \text{b})} \right) \mathbb{P}(\theta_0, \cdots, \theta_T | \mathbb{P}_\theta, \mathbb{P}_0), \tag{33}$$

where $\phi(\cdot)$ denotes the normal kernel.

We estimate the parameters $(\mathbb{P}_\theta, \mathbb{P}_0, \mu, \sigma)$ using an expectation-maximization algorithm to maximize the likelihood in Equation 33, while we use a Viterbi algorithm to recover the maximum likelihood estimates for the sequence of the unobserved demand state $\left\{\hat{\theta}_t\right\}_{t=0}^T$. Table 17 describes several statistics of the distribution of our estimates. Table 18 highlights the features of the demand process across different states.

As described in Section 2, illiquidity is the key driver of demand volatility in the market for municipal bonds. Consistent with this, we find that the demand process exhibits high persistence, as the common demand shock θ_t changes on average every six months. Using the same measure of persistence, we find that demand is more persistent in states where the municipal bonds market has more depth. Indeed, the right

[64] Also observe that, conditional on θ_t, the realizations of the average prices $\hat{p}_t(a)$ are independent over time and across types of investors.

panel of Table 17 shows that the expected time between changes in θ_t in a state correlates positively with the state's total outstanding municipal debt. To test directly the connection between demand persistence and market liquidity, we correlate the persistence of the demand shock across states with a standard measure for market illiquidity proposed by Amihud (2002). Following Acharya et al. (2013), we measure the monthly returns for an asset as the percentage change in the average price in a month compared to the average price in the previous month. Then, we measure the asset j's liquidity as

$$\text{Amihud}_j = \frac{1}{n_j} \sum_{t=1}^{n_j} \frac{|r_{jt}|}{Q_{jt}}, \tag{34}$$

where r_{jt} is the asset j's monthly return in month t, n_j is the number of months for which returns r_{jt} can be computed, and Q_j is the total trading volume in million \$. This measure proxies illiquidity by the ratio of absolute stock return to its dollar volume, and can be interpreted as "the daily price response associated with one dollar of trading volume, thus serving as a rough measure of price impact." A state's illiquidity consists of the average illiquidity of its assets. As shown in Table 17, we find that Amihud's liquidity indeed correlates negatively with the persistence of the demand process.

Table 17: Correlation between the persistence of demand and market characteristics

	Persistence of demand
Outstanding muni. debt	57.57%
Amihud Illiquidity	−66.06%
Share of Aaa rated assets	18.78%

Notes: The table correlates the persistence of the demand process across states with several indicators of the liquidity of the municipal bond market. Persistence is measured as the average time between changes to the demand state θ_t.

Table 18 highlights the features of the recovered parameters across different classes of assets. The three columns of the table show, for each group of assets, the average purchase and selling prices by state and realization of the aggregate demand shock θ_t. Finally the last column reports the average number of changes for θ_t within an year. Furthermore, the bid-ask spread is on average 5%, and it is larger for less liquid states. Finally, it is worth mentioning that the estimates for the mean prices $\mu(\theta, \mathrm{s})$ and $\mu(\theta, \mathrm{b})$ imply that dealers have substantial incentives to anticipate changes in demand and wisely time their trades. Indeed, a dealer's average margin from buying and selling an asset is around 1% if the asset is bought and sold in the same period, and increases to 4% if the asset is bought in a period of low demand

and sold in a period when demand is high. [65]

Table 18: Demand estimates

State	Average price for buyers			Average price for sellers			Changes per year
	θ_L	θ_M	θ_H	θ_L	θ_M	θ_H	
AR	0.996	1.013	1.023	1.005	1.021	1.030	4.966
MS	0.975	0.994	1.010	0.985	1.006	1.022	2.069
KS	0.997	1.010	1.034	1.013	1.028	1.053	2.897
NE	1.017	1.032	1.045	1.030	1.047	1.062	2.897
IN	0.959	0.984	0.999	0.965	0.990	1.006	2.069
WI	0.966	0.992	1.002	0.969	0.994	1.009	2.069
MI	0.995	1.023	1.036	1.002	1.029	1.044	1.241
NJ	1.001	1.018	1.039	1.009	1.026	1.049	2.897
NY	0.995	1.017	1.029	0.999	1.022	1.036	2.069

Notes: The table summarizes the estimates of the demand process. We cluster the assets based on that state of issuance and estimate the parameters independently across groups. The average price is expressed as a fraction of the average trade size.

F Model Fit

Figure 7: Model fit. The left panel depicts the observed and predicted probabilities of trading with investors for high experience dealers. The left panel depicts the same objects for low experience dealers.

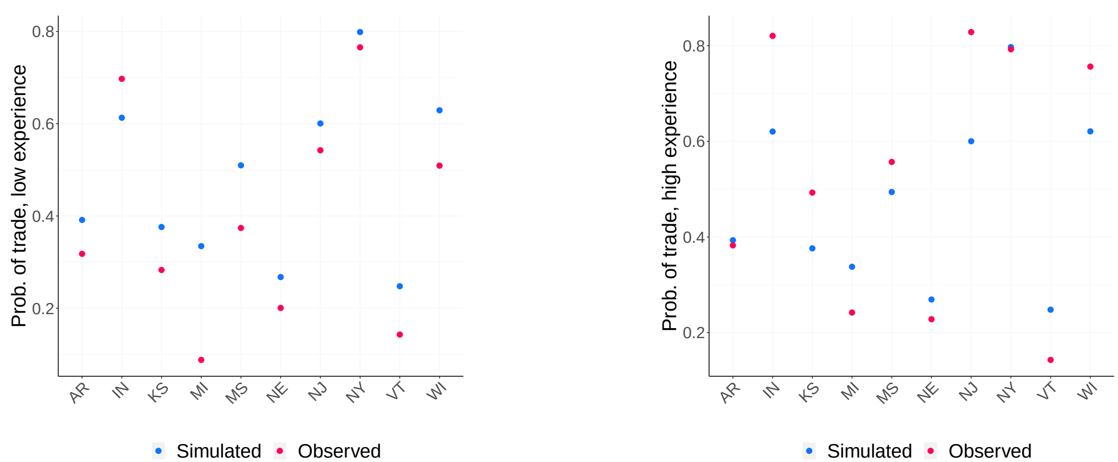

[65] These estimates are comparable with the markup recovered by a number of studies that analyze the market for municipal bonds. As an example, Li and Schürhoff (2019) find that average markups are 1.8%, Harris and Piwowar (2006) find that effective spreads range between 1% and 2%, as does Green et al. (2006).

Figure 8: Model fit. The left panel depicts the observed and predicted probabilities of selling an asset to investors for high experience dealers. The left panel depicts the same objects for low experience dealers.

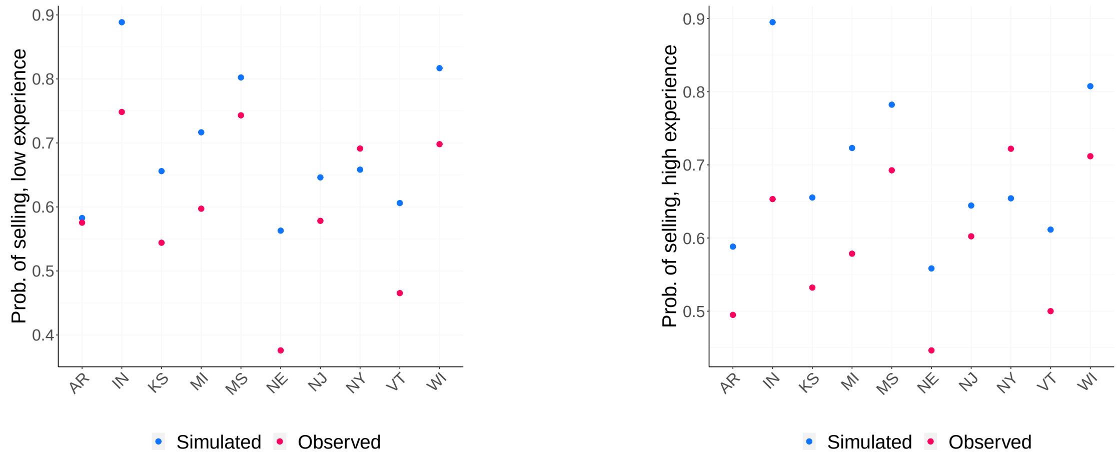

Figure 9: Model fit. The left panel depicts the observed and predicted probabilities of trading an asset in the inter-dealer market high experience dealers. The left panel depicts the same objects for low experience dealers.

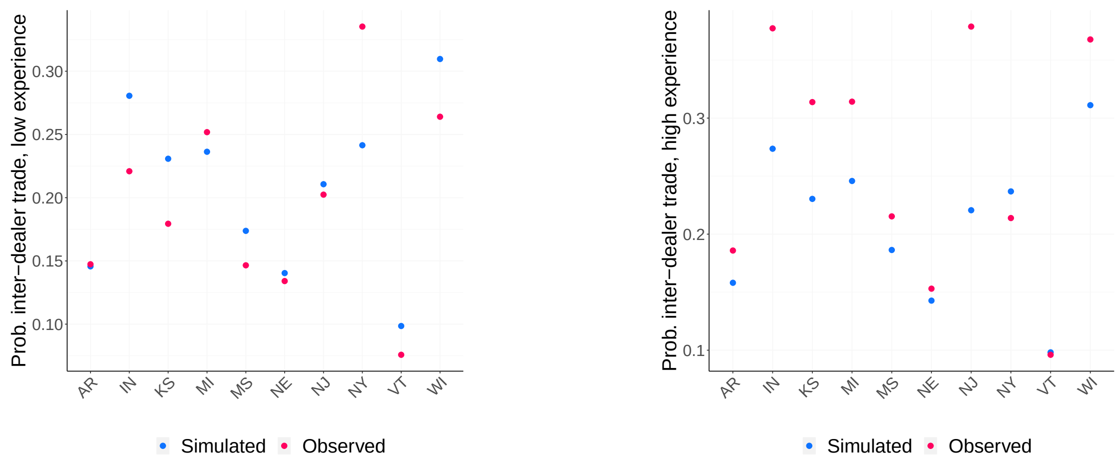

F.1 Stylized facts on networks

Recent empirical studies have documented a number of stylized facts about the intermediation process in OTC markets. For one, these studies highlight that OTC markets exhibit a core-periphery inter-dealer network: few central dealers trade frequently and with many dealers, while the majority of peripheral dealers trade sparsely and with few dealers. While our estimation approach doesn't explicitly target this stylized facts, the core-periphery network arises in equilibrium. To showcase this feature of our setup, we use the estimated model to simulate dealers' behavior in the market for one hundred periods. In accordance with the literature, we use the simulated data to define the inter-dealer network where two dealers are connected if they traded the asset with one another.

Figure 10: Core-periphery structure

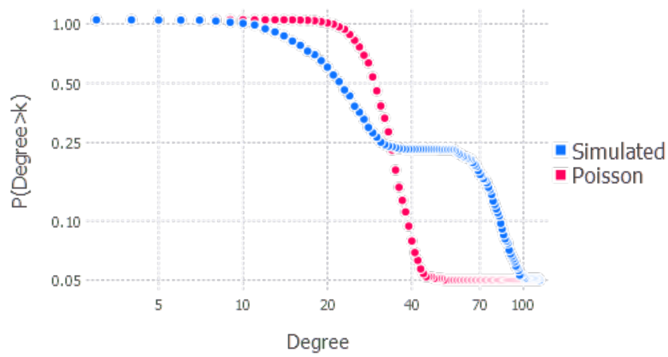

Notes:

Figure 10 shows the degree distribution across dealers in the simulated market. as a benchmark for comparison Figure 10 also shows the degree distribution that would arise if dealers traded anonymously and randomly through a centralized exchange. Consistent with the empirical literature on trading in OTC market, we find that the simulated network displays a higher level of heterogeneity among dealers in terms of connectedness than suggested by random trading. Moreover, the simulated network displays a core-periphery structure: a large number of weakly connected dealers compete with a few core dealers. The dealership network exhibits a heavy right tail, with 5% of dealers having more than 100 other dealers as trading partners. The overall network is very sparse, as only 2% of the possible links are formed.

Another robust finding in the empirical literature of OTC markets is that dealers trading prices vary systematically with their centrality in the network.[66] Figure 19 shows that our model replicates this facts, especially for inter-dealer trades. We divide the dealers into deciles based on their degree centrality. Each point in in the left hand panel of Figure 19 corresponds to a different decile of the dealers' degree centrality, and compares the average centrality of those dealers to the average price they obtain in the inter-dealer market. The plot shows that there is a clear positive correlation between dealers' centrality and their selling price in the inter-dealer market. We confirm this relationship on the right hand panel of Figure 19, where we regress inter-dealer prices on the seller's centrality.

[66] Empirical evidence on the market for municipal bonds shows that there is a centrality premium: more active dealers charge up to 80% higher bid-ask spread for medium-size customer trades (Li and Schürhoff (2019)), while on the market for asset-backed securities and non-agency collateralized mortgage obligations there is a centrality discount: more active dealers charge smaller bid-ask spreads to customers (Hollifield et al. (2014)).

Table 19: Centrality Premium

	Inter-dealer trading price (log)	
	II	II
Seller's centrality (log)	0.064**	0.063**
	(0.009)	(0.009)
Seller's inventory (log)		-0.057
		(0.041)
N	15,510	15,510
Level	Trade	

$^{**}p \leq 0.05$, $^{*}p \leq 0.1$