Learning by Trading:

The Case of the U.S. Market for Municipal Bonds*

Giulia Brancaccio[†], Dan Li[‡], Norman Schürhoff[§]

December 2020

Abstract

We study information acquisition as a motive for trade in the market for municipal bonds, a primary financing source for local governments. The market is decentralized and regional dealers intermediate trades among retail investors. Dealers' pricing behavior suggests that they learn by trading. We propose and estimate a dynamic model of OTC trading that accounts for dealers' learning and reveals that they are willing to pay 12% of the intermediation spread for the information acquired through trade. Dealers' learning incentives benefit investors and issuers, as they strengthen trading in an illiquid market, and interact with policies on post-trade disclosure.

Keywords: Experimentation, over-the-counter markets, search frictions, bargaining, market transparency.

*Brancaccio is deeply indebted to her 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, University of Lausanne, CH-1015 Lausanne, Switzerland. Email: norman.schuerhoff@unil.ch.

1 Introduction

In many markets, there exist incentives to privately acquire information about the state of market fundamentals. Information acquisition is particularly relevant in markets with decentralized trade, where public information about prices and trading volume is limited. In such opaque markets, negotiations with others reveal information about the counterparty's asset valuation. Aggregating the information gathered from trading is, in turn, is informative of the state of market fundmentals. In this paper, we empirically study how the incentives to learn about market fundamentals through trading shape price formation and trading behavior of financial institutions participating in a decentralized market with trades totaling over $3 trillion a year: the secondary market for U.S. municipal bonds.

Throughout the United States, states and local governments depend on the municipal bond market to raise funds for investments in schools, highways, and other public projects. There is no central marketplace to exchange municipal securities and, 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 dealers acquire information about the state of market fundamentals through their trading activity. Since this information is valuable for dealers, obtaining it is itself a motive for trade and an essential driver of liquidity in the market. We study how dealers' learning incentives interact with policies to improve access to public information about trade activity.[1] These policies weaken dealers' motives to acquire information and, ultimately, reduce their incentives to provide liquidity which hurts investor welfare. Our results highlight an unintended consequence of market transparency that may limit the benefits of these policies.

The opacity of the municipal bond market makes it a natural candidate to study the interaction between trading and information acquisition. The lack of a centralized marketplace for municipal securities implies that learning about trade prices requires participating in a trade directly. A large number and variety of bonds are outstanding at any given time, and each asset includes a variety of special provisions that complicate pricing and make information acquisition a first-order issue. A large number of goods and financial markets are prone to similar issues, including the markets for corporate bonds, agency-backed securities, asset-backed securities, and 144As.

We leverage a detailed dataset of transactions on the municipal bonds' secondary market to show that financial intermediaries acquire information about the state of market fundamentals through their trading

[1] Intended to improve liquidity in decentralized financial markets, these policies are the object of a lively debate and intense regulatory activity both in the U.S. and abroad. See as an example Vames (2003).

activity. The data reveal that the correlation in the pricing behavior of two dealers increases after they have traded with one another. A variety of placebo tests confirm that this result is not spurious which suggests that information acquisition is the likely driver of the dealers' behavior change. Trade, therefore, is a source of valuable information.

Driven by these observations, we construct a dynamic model of trade in decentralized markets that explicitly accounts for dealers' learning through trade. We build on an standard inventory model of dealer behavior in which forward-looking dealers accumulate costly inventory of municipal bonds by trading with other dealers and short-lived retail investors, subject to search frictions. Investor valuations for the asset change over time due to a persistent common demand shock. Dealers earn profits from buying the asset when demand is low and selling when demand is high. Importantly, the aggregate demand shock is unobserved and there exists no public information about it. Dealers acquire information about investors' demand for the assets through their trading activity since trade prices are informative about the unobserved demand. When trading with investors, each trade has the same information content and trading with more investors allows a dealer to sample more observations. When trading with other dealers, prices are informative about the counterparty's valuation of the asset, which reveals what the dealer knows about the state of the market. To capture the decentralized nature of trade, we allow dealers to only observe a summary statistic of their peers' past trading activity, which we call experience.

Learning about demand creates an *information acquisition motive* for trade. When deciding whether to engage in trade, dealers not only consider factors such as the expected purchase and resale price, but they also recognize that by trading, they will acquire valuable information. Information acquisition, therefore, strengthens dealers' incentives to trade and provide liquidity.

We estimate the model using data on trading in the market for municipal bonds between 2000 and 2003, a period when little information about market fundamentals was publicly available. The estimation proceeds in two steps. A dealer's trading decisions depend on his type, including his information about the demand shock and his experience. In the first step of our estimation, we recover the unobserved dealers' types. We exploit variation in trade prices in the inter-dealer market to identify dealers' experience. Our model suggests that dealers offer different prices to counterparties with different experience to account for factors such as the trade's information content and the probability of rejection. Our baseline specification compares how prices for trades for a given asset executed by a given seller, in a given month, change depending on the buyer's trading history. 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' beliefs about demand, we assume that dealers only acquire information through trade or public signals observed by everyone, the econometrician included. This assumption allows us to directly trace out the dealers' information sets at every point in time. We exploit a Hansen-Sargan test for overidentifying 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 finding confirms that learning activities in the market for municipal bonds are strongly connected to realized trade, justifying our approach. We estimate the remaining primitives, which include dealers' trading and inventory costs, as well as investors' valuations and their entry costs. Following the dynamic discrete choice literature (Rust (1987)), we recover dealers' costs from their optimal choice probabilities via the method of simulated moments. Moreover, we obtain investors' valuations directly from observed prices.

The estimated model allows us to quantify the impact of dealers' information acquisition motives for trade on the municipal bond market liquidity. For the average dealer, the information content of a trade with an investor is worth 12% 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 find that the value of information hinges on a dealer's incentives to time the sale of the asset to match changes in demand. For instance, we find that it is higher when search costs are small, when fluctuations in investors' valuations are large, and when the marginal cost of inventory is small. This analysis illustrates the main principles driving dealers' information-seeking behavior that generalize to markets other than that for municipal bonds.

We find that dealers' learning incentives benefit both investors and issuers, as they strengthen dealers' incentives to trade in a notoriously illiquid market. Investors' welfare would be 5% lower, corresponding to more than $1 billion, in the counterfactual equilibrium where dealers ignore the informational content of their trading activity. The cost of capital for issuers would increase by 7 basis points. These results are driven by a drop in trading volume of 9%.

Two channels are responsible for the effect. Around half of the decline is driven directly by the weaker information acquisition motives for trade. Additionally, if dealers ignored the informational content of their trades the precision of their forecast of asset demand would fall by up to 30%. This indirect effect, in turn, worsens dealers' ability to correctly time the sale and purchase of the assets, decreasing dealers' profits and weakening their incentives to make a market, further depressing market liquidity.

We use the estimated model to explore the impact on dealers' incentives to trade on policies increasing the availability of public information about trade activity. Such policies have increasingly been discussed and implemented across different countries and markets. We first look at the impact of these policies on market liquidity, since it has historically been the target of policy makers in decentralized financial markets. We find that improving market transparency can reduce the volume of trade up to 6%.

Importantly, the model traces out how the predicted change in trade activity affects welfare. Following the introduction of transparency investor welfare declines by up to 6%, and 2% on average. In a transparent market, dealers exercise more market power by curbing their trading activity. Therefore, some investors who would profitably exchange an asset are not served in a transparent market. By counteracting this effect, information acquisition motives for trade benefit investors and reduce asset misallocation.

In conclusion, we shed light on the role of experimentation in decentralized financial markets. In a variety of markets trade is a source of valuable information about market fundamentals, so that obtaining this information becomes a motive for trade. While information acquisition motives for trade are often overlooked, they alter the effectiveness of regulatory intervention in decentralized markets by creating a channel through which market transparency has an unintended impact on market liquidity. Naturally, a richer model would be necessary to derive clear policy implications. As an example, to focus on dealers' information acquisition motives for trade, we abstract from investors' search behavior and their dynamic incentives to trade, as well as the role of learning about default risk. Incorporating these elements in a full-fledged policy evaluation could be an interesting avenue for future research.

Related literature. This paper is at the intersection of three strands of literature. The basic trade-off between learning and sacrificing immediate payoff is focal in the literature on strategic experimentation. Experimentation has long been studied in economics, mostly from a theoretical standpoint (for a survey, see Hörner and Skrzypacz (2016)). Several papers in this literature 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 empirical. For this reason, we strip the incentives to experiment to their minimal components which makes the agents' problem tractable, allowing us to bring the model to the data.

Several papers, including Leach and Madhavan (1993) and Bloomfield and O'Hara (1999, 2000), discuss the implications of experimentation for the trading behavior of agents in financial markets. Wolinsky (1990), Golosov et al. (2014), and Blouin and Serrano (2001) explore the link between trading and information diffusion in decentralized markets with private information. Their objective is to study under what conditions all relevant information gets revealed. We contribute to this literature by employing a tractable analytical framework to empirically quantify the role of experimentation as a motive for trade.

We merge the literature on experimentation with studies on the trading behavior in decentralized markets. Our model draws from a rich tradition of papers using search models to study asset prices and allocations in decentralized markets. Early papers include Gehrig (1993), Spulber (1996), and Rust and Hall (2003). Recent papers build on the framework of Duffie et al. (2005, 2007), to study the implication of search frictions for the functioning of decentralized markets. Important works include Gavazza (2011b,a), Hugonnier et al. (2012), Hugonnier et al. (2014), 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.[2] For this reason, we borrow the structure of these models and enrich it with incomplete information and learning. The decision to trade then depends on experimentation in addition to search costs and inventory management.[3] Traditionally this literature models meeting among agents as random. Instead, dealers in our model direct their offer to a specific counterparty.

Our paper is also related to the growing empirical literature on decentralized financial markets. Bessembinder et al. (2006), Harris and Piwowar (2006), and Green et al. (2006) study transactions cost and price discovery in decentralized markets. Di Maggio et al. (2017), Hollifield et al. (2017), and Li and Schürhoff (2019) document a number of stylized facts both about the structure of networks defined by inter-dealer trades and how dealers' prices depend on their network centrality. While we do not target these stylized facts, in Appendix F.1, we show that our model gives rise in equilibrium to these facts in response to dealers' learning incentives.[4]

Our paper also draws on the literature on industry dynamics (e.g., Hopenhayn (1992), Ericson and Pakes (1995)) that characterizes Markov-perfect equilibria in entry, exit, and investment choices given uncertainty in the evolution of the states of firms and their competitors. We model agents' problem as a series of trading and pricing decisions. Since agents interact with one another repeatedly, the problem generates a high-dimensional state space. To simplify the computational burden we follow Weintraub et al. (2008) and assume that, given the large number of agents in the market, the distribution of their type is perfectly forecastable by agents, conditional on the aggregate demand shocks they are trying to learn. 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 the primitive of interest.

[2] Duffie and Manso 2007 have a similar focus, but they study information diffusion rather than information acquisition.

[3] Few papers in this literature allow for asymmetric information. Lester et al. (2003) and Glode and Opp (2016) are exceptions.

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

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 operations and finance capital projects. A total of $3.7 trillion municipal bonds are outstanding and $300 billion of municipal debt are issued every year. The market's importance cannot be overstated: municipal bonds are the main source of funding the public investment in infrastructure in the United States, representing 75% of the total as of 2017.

We focus on the secondary market for municipal bonds where they are traded after issuance. There is no central exchange for municipal securities. Instead, the market is decentralized, and financial institutions registered with the SEC act as dealers intermediating 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.[5] Every year there are more than 2,000 active broker-dealers.

The municipal bond market is particularly interesting to the extent that the large majority of the municipal debt outstanding is held by retail investors either directly or through mutual funds.[6] The reason is that, 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 tax advantage and the resulting ownership structure, trades in the secondary market are small: the median trade is worth $25,000, and 80% of trades have a value of less than $100,000.

Dealers interact with investors in a customer segment and with other dealers in an inter-dealer segment. Investors trade in the secondary market mostly to cover liquidity needs—factors such as retirement, college fees, unemployment, or the purchase of property. Apart from these private liquidity needs that correlate across investors, municipal bonds are considered buy-and-hold assets, as investors tend buy them at issuance and hold them until maturity. Municipal bonds are also considered to be a relatively safe, which limits speculative trade.[7]

The secondary market for municipal bonds is illiquid and volatile despite the low default risk. As an example, for the average dealer the purchase 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

[5] 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.

[6] 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.

[7] As an example, for Aa- and A- rated municipal bonds, the 10-year cumulative default rate is 0.03% compared to 0.8% for corporate bonds.

to fluctuations in the price of the asset corresponding to almost twice the average yearly return for a municipal bond. The main driver of illiquidity is the lack of depth. Demand for municipal bonds is segmented geographically since 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.[8] 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.

Naturally, dealers aim to exploit variation in investors' valuations of municipal bonds to "buy low and sell high." To this end, dealers need to anticipate the liquidity needs of a diverse group of investors.[9] We study how dealers' incentives to learn "what their clients want" affect their trading behavior.

Public information on market fundamentals. Historically, agents operating in the market for municipal bonds have not had access to rich public information about prices and trading activity in the market, due to the lack of a centralized trading mechanism and the lack of standardization. The available market indexes have been coarse and did not effectively reduce uncertainty about the pricing of individual bonds. Access to public information about trade activity has improved since the end of the 1990s, due to a push from the Securities and Exchange Commission (SEC) and the Municipal Securities Rulemaking Board (MSRB) to improve market transparency.

The MSRB has published since 1995 information about the volume of trade and average price for assets traded more than four times during the previous day, which covered just over 5% of the assets traded. On June 23, 2003, the MSRB started distributing daily summaries on the trading activity in the market during the previous day ("next-day reporting"). Starting on January 31, 2005, the MSRB mandated that details of all transactions in U.S. municipal bonds be reported on a timely basis and posted online within fifteen minutes. Investors have embraced the new source of information: 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 (Schultz, 2012).[10]

The quality of public information about trade activity interacts with dealers' incentives to acquire infor-

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

[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] In Appendix A, we analyze the impact of this policy intervention and we show that information acquisition motives for trade explain 51% in the variation of the effect of market transparency across states.

mation. The policy intervention in June 2003 offers a unique laboratory to observe information acquisition motives for trade, which we exploit in Section 3. The regulatory push toward market transparency is a widespread trend across different markets. Since 2002, similar provisions were imposed by the Financial Industry Regulatory Authority (FINRA) in the markets for U.S. corporate bonds, agency-backed securities, asset-backed securities, and 144A transactions. In July 2018, FINRA began requiring its member firms to report transactions in Treasury securities. European regulators have also been in the process of instituting a post-trade transparency regime for a broad range of instruments. This motivates us to study, in Section 8, how dealers' information acquisition motives for trade interact with market transparency.

2.2 Data

Our main data source is the proprietary Transaction Reporting System data 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 data, however, became public only after 2003. 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 trade, including trade 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), and an asset identifier. Importantly, 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 Mergent. Finally, we obtain the time series for market bond indexes and 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 Mergent 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. As a result of these filters, 65% of transactions remain from 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.

Table 1 provides summary statistics. 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

Table 1: Summary statistics

	Mean	St. Dev	Median	Min	Max
Monthly trade with investors ($\$\times 10^7$)	3.45	0.51	3.38	2.35	4.85
Monthly inter-dealer trade ($\$\times 10^7$)	1.50	0.25	1.48	1.02	2.25
Trade Price ($)	99.48	10.68	101.52	36.64	116
Intermediation spread (%)	2.1	1.55	1.19	-0.23	6.8
Trade size ($\$1,000$)	72.05	190.92	25	5	2,245
Dealer's market share	0.043	0.40	0.00026	2^{-7}	0.116
Share of inter-dealer trades with 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 U.S. 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.

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 in Section 2.1, the trade size is on average $70,000, median $25,000, and institutional size trades above $1 million represent just 1% of 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 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. The highest market share by state of issuance is on average 11%.

Interactions on the inter-dealer market are frequent: the number of inter-dealer trades is one-third that of trades with investors. As the last row of Table 1 reveals, on average, across dealers, one of every two transactions on the inter-dealer market involve a new counterparty.[11] This is consistent with dealers trading in the inter-dealer market to acquire information.

3 Learning by trading

In markets where public information about prices and trading activity is sparse, agents rely on their private interactions with other agents to aggregate information dispersed among market participants. Through their negotiations to trade, dealers can acquire information about their counterparties' valuation for the

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

asset. Information about private asset valuations, in turn, provides valuable information about the overall state of market fundamentals.

We use data on inter-dealer trades to provide evidence that dealers acquire information through their trading activity and that this information leads dealers to persistently change their trading behavior. We focus on the pricing behavior of pairs of dealers that have recently traded with one another in the inter-dealer market. Suppose that the bargaining process revealed dealers information about their counterparty's assessment of the market. In this case, after the trade, the dealers' pricing strategies should change to account for the new information. To support this notion, we test whether dealers price assets more similarly after having traded with one another.

Consider inter-dealer trade i 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 the inter-dealer trade. The price $\hat{p}_{a,d,t}$ is constructed as a par size-weighted average:

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

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

We pool all inter-dealer trades and estimate how the dummy variable $\text{d2d}_{i,t-k}$, indicating whether inter-dealer trade i took place in $t-k$, affects price differences between dealers, based on the regression

$$\left|\hat{p}_{a_i,d_i,t} - \hat{p}_{a_i,\tilde{d}_i,t}\right| = \phi_{a_i,d_i,\tilde{d}_i} + \sum_{k=-5}^{10} \psi_k \cdot \text{d2d}_{i,t-k} + \epsilon_i. \tag{2}$$

The coefficient ψ_k tells us how inter-dealer trade affects the difference in dealers' pricing strategies. We include dealer pair×asset fixed effects, $\phi_{a_i,d_i,\tilde{d}_i}$, to absorb any initial differences between dealers' pairs, and we cluster standard errors at the dealer pair and asset level.

The left panel of Figure 1 plots the coefficients ψ_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 first column of Table 2 reports the estimates of the cumulative effect of trading: 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

Notes: The left panel of the figure plots the regression coefficients and 95% confidence intervals from estimating Equation (1) using the sample of inter-dealer trades between January 2000 and June 2003. The right panel of the figure plots the regression coefficients and 95% confidence intervals from estimating Equation (1) using the sample of inter-dealer trades between June 2003 and January 2005, after market transparency was introduced. 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. The standard errors are clustered at the asset-pair level.

Falsification test. A natural way to falsify the result in the left panel of 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 increased market transparency and made information about trade prices more easily accessible for market participants. Under "next-day reporting," we conjecture that the informative content of a single inter-dealer trade becomes negligible and the price response documented in the left panel of Figure 1 should disappear. To test this, we reestimate Equation (3) using inter-dealer trades taking place both before and after the policy intervention. To control for the transparency regime, we introduce the dummy variable opaque_i that equals one for inter-dealer trades taking place before the policy intervention, and zero otherwise. The right panel of Figure 1 plots the coefficient ψ_k from the regression

$$\left|\hat{p}_{a_i,d_i,t} - \hat{p}_{a_i,\tilde{d}_i,t}\right| = \phi_{a_i,d_i,\tilde{d}_i} + \sum_{k=-5}^{10} \psi_k \cdot \text{d2d}_{i,t-k} + \sum_{k=-5}^{10} \gamma_k \cdot \text{d2d}_{i,t-k} \times \text{opaque}_i + \epsilon_i. \tag{3}$$

and confirms that the price adjustment is only present in the opaque market. Column II in Table 2 reports the cumulative price adjustment obtained from estimating

$$\left|\hat{p}_{a_i,d_i,t} - \hat{p}_{a_i,\tilde{d}_i,t}\right| = \phi_{a_i,d_i,d_i'} + \psi \cdot \text{post}_{i,t} + \gamma \cdot \text{post}_{i,t} \times \text{opaque}_i + \epsilon_{i,t}. \tag{4}$$

Table 2: Placebo tests

	I	II	III
		$\left\|\hat{p}_{\mathrm{a},t,d} - \hat{p}_{\mathrm{a},t,\tilde{d}}\right\|$	
$\text{post}_{\mathrm{a},d,d',t}$	-0.173^{***}	0.012	-0.042
	(0.040)	(0.054)	(0.027)
$\text{post}_{\mathrm{a},d,d',t} \times \text{opaque}_{\mathrm{a},d,d'}$		-0.186^{***}	
		(0.067)	
$\text{post}_{\mathrm{a},d,d',t} \times \text{obs}_{\mathrm{a},d,d'}$			-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. Standard errors are clustered at the asset-pair level.

The coefficient γ is large and significantly different from zero which confirms that inter-dealer trades have a substantially larger effect on dealers' pricing strategies in the opaque than the transparent regime. This finding supports the notion that information acquisition drives the effect displayed in the left panel.

Placebo trades. We construct a set of placebo trades to rule out that the results in Figure 1 are driven by the release of public information. 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 $\tilde{f}_i$. 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. 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 pool all observed and placebo trades and estimate Equation (5) 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:

$$\left|\hat{p}_{a_i,d_i,t} - \hat{p}_{a_i,\tilde{d}_i,t}\right| = \phi_{a_i,d_i,\tilde{d}_i} + \psi \cdot \text{post}_{i,t} + \gamma \cdot \text{post}_{i,t} \times \text{obs}_i + \epsilon_{i,t}. \tag{5}$$

Column III in Table 2 shows that once again the coefficient γ remains significantly different from zero even after controlling for market-level movements in prices. The parameter ψ is not significant, suggesting that the difference in trade price does not fall for dealers who have not traded with each other.

To sum up, dealers price assets more similarly after having traded with each other on the inter-dealer market, and the resulting price effect is persistent.

4 Model of learning by trading

In this section, we present a model of trade in decentralized markets that explicitly accounts for dealers' incentives to learn through trade. We start by presenting a simplified version of our main model of trading and learning to illustrate the main mechanisms at play and to aid the interpretation of the main model. Then, in Section 4.2, we enrich the simplified model to make it amenable to empirical analysis.

4.1 Simple model of trading and learning

Time $t \in \{0, 1, \ldots, T, \ldots\}$ is discrete and infinite. Two types of risk-neutral agents populate the market: short-run investors and long-run dealers. The agents trade a single asset that represents a municipal bond.

Investors live for one period and are characterized by their valuation of the asset v and liquidity need $u \in \{-u_{\mathrm{s}}, u_{\mathrm{b}}\}$.[12] Valuations v vary across investors to reflect heterogeneity in liquidity needs, portfolio holdings, and tax advantages. The distribution of investor valuations changes over time due to a persistent common demand shock $\theta \in \Theta$ that represents common factors in investors' willingness to pay for the asset, such as the profitability of alternative investments or budget tightness. In every period the valuations of investors with liquidity need u, conditional on demand shock θ, are independently drawn from distribution $F_v(v|\theta, u)$.[13] Demand shock θ 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, he does not observe the common demand state θ_t. This means that investors do not know how their valuation correlates with that of other investors. 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, dealers can search for investors interested in trading the asset. In particular, each dealer first decides whether to search for buyers or sellers, $a \in \{b, s\}$, subject to a logit shock. Then, the dealer

[12] Investors can either be interested in buying u_b units or selling u_s units of the asset.

[13] For simplicity, here we treat the distribution of investors' valuations as an exogenous object. However, we endogenize this object in Section 4.2.

draws search costs $c \sim F_a(c)$ sequentially. After every draw, the dealer can decide to pay the cost and contact an investor of the desired type. Whether the meeting with the investor takes place or not, with probability $1-\gamma$ the dealer has the opportunity to draw another cost and possibly meet another investor, while with probability γ he moves on to the next period. Assets bought and not sold accumulate over time and form the dealer's inventory x. Carrying inventory is costly, due to frictions that prevent dealers from increasing their balance-sheet size, such as limits to leverage or exposure to risk. In particular, every period a dealer with inventory x pays cost $\kappa(x) = \sum_{h=1}^{H} \kappa_h x^h$, where $\kappa_0, \ldots, \kappa_H$ are parameters.

When meeting an investor, dealer and investor play a one-sided incomplete information bargaining game:[14] The dealer immediately learns the investor's valuation and makes a price offer. If the investor accepts, dealer and investor trade the asset at the proposed price. If the investor rejects the offer, the negotiation can either break down exogenously or continue with a counteroffer by the investor. Dealer and investor alternate offers in this fashion until either offer is accepted or the bargaining breaks down.

Dealers' search and trading decisions. A dealer's trading decision depends on his inventory $x \in \{0, 1, \ldots, \bar{x}\}$ and his current beliefs $\pi \in \Delta(\Theta)$ about the unobserved demand state. The value function $V_0(\pi, x)$ at the beginning of the period satisfies

$$V_0(\pi, x) = -\kappa(x) + \max\{V_1(\pi, x; \mathrm{b}) + \epsilon_{\mathrm{b}}, V_1(\pi, x; \mathrm{s}) + \epsilon_{\mathrm{s}}\}, \tag{6}$$

where $V_1(\pi, x, a)$ denotes the value of a dealer who has decided to search investors of type $a \in \{\mathrm{b}, \mathrm{s}\}$. The dealer pays inventory cost $\kappa(x)$, observes action-specific shocks $\epsilon \in \mathbb{R}^2$ drawn from a type I extreme value (Gumbel) distribution with standard deviation σ_ϵ, and decides whether to search for buyers and obtain value $V_1(\pi, x; \mathrm{b})$, or to search for sellers and obtain value $V_1(\pi, x; \mathrm{s})$.

Next, consider the outcome of the bargaining between dealer and investor, which is derived in Appendix C.2 in detail. The trading price between an investor with liquidity need u and value v, and a dealer with beliefs π and inventory x can be written as

$$p(\pi, x, v, u) = (1-\rho)\rho_0(u) + \rho v, \tag{7}$$

[14] We assume that the dealer immediately learns the investor's valuation to sidestep the intricacies of solving a bargaining game with two-sided incomplete information and common values. To avoid distorting substantially the agents' incentives, we allow the dealer to leave the negotiations after observing the investor's valuation. However, this never happens in equilibrium. Importantly, this assumption captures the dealer's information advantage vis-à-vis the investor. We discuss the advantages of this bargaining protocol in detail in Section 4.3.

where $\rho_0(u)$ is an equilibrium function that depends only on the investor's type u, and $\rho \in (0,1)$ is a parameter that captures the dealer's bargaining power. Dealer and investor conclude the negotiations successfully with probability one and without delay. To simplify the notation, we drop the dependence of the trading price on the dealer's type and denote it by $p(v,u)$.

The value of a dealer who has decided to search for investors of type $a \in \{\mathrm{b},\mathrm{s}\}$ is

$$V_1(\pi,x;a) = \begin{cases} \mathbb{E}_c\left[\max\left\{-c + \mathbb{E}_v\left[-p(v,u_a) + V_2(\pi'(v,a,\pi), x-u_a;a)\right], V_2(\pi,x;a)\right\}\right] & \text{if } a=\mathrm{b} \\ \mathbb{E}_c\left[\max\left\{-c + \mathbb{E}_v\left[p(v,u_a) + V_2(\pi'(v,a,\pi), x+u_a;a)\right], V_2(\pi,x;a)\right\}\right] & \text{if } a=\mathrm{s} \end{cases}, \tag{8}$$

where $V_2(\pi,x;a) = \gamma V_1(\pi,x;a) + (1-\gamma)\beta V_0(\pi,x)$. The dealer draws search cost $c \sim F_a(c)$ and decides whether to contact an investor. If he does, he draws at random an investor with valuation $v \sim F_v(v|\theta,u_a)$ and pays or receives price $p(v,u_a)$. Not knowing the demand shock θ, the dealer forms expectations about the valuation of the investor that he will meet based on his belief π according to $\mathbb{E}_v[p(v,u_a)] = \sum_\theta \pi(\theta)\int p(v,u_a)\,dF(v|\theta,u_a)$. After trading with an investor, the dealer updates his beliefs to $\pi'(v,a,\pi)$ as he learns that the investor of type a has valuation v for the asset. His inventory evolves to account for the trade. Finally, after the dealer has concluded the negotiations with the investor, or if he has decided not to pay the search cost, he obtains value $V_2(\pi,x;a)$: with probability γ he can draw a new search cost and restart with valuation $V_1(\pi,x;a)$, while with probability $1-\gamma$ he moves on to next period, obtaining value $\beta V_0(\pi,x)$.

Motives for trade. To gain insight into dealers' motives to trade, it is useful to analyze the search decision described in Equation (8). The dealer chooses to pay the search cost if his draw is lower than threshold $\bar{c}(\pi,x,a)$ that satisfies

$$\bar{c}(\pi,x,a) = \begin{cases} \mathbb{E}_v\left[-p(v,u_s) + V_2(\pi'(v,s,\pi), x+u_s;s) - V_2(\pi,x;s)\right] & a=s \\ \mathbb{E}_v\left[p(v,u_b) + V_2(\pi'(v,b,\pi), x-u_b;b) - V_2(\pi,x;b)\right] & a=b \end{cases}. \tag{9}$$

The thresholds $\bar{c}(\pi,x,a)$ can be rewritten as

$$\begin{aligned} \bar{c}(\pi,x,\mathrm{s}) &= \underbrace{V_2(\pi,x+u_\mathrm{s};\mathrm{s}) - V_2(\pi,x;\mathrm{s}) - \mathbb{E}_v[p(v,u_s)]}_{\text{expected return } R(\pi,x;\mathrm{s})} + \mathbb{E}_v\underbrace{\left[V_2(\pi'(v,\mathrm{s}), x+u_s;\mathrm{s}) - V_2(\pi,x+u_\mathrm{s};\mathrm{s})\right]}_{\text{value of information } I(\pi,x;\mathrm{s})}, \\ \bar{c}(\pi,x,\mathrm{b}) &= \underbrace{V_2(\pi,x-u_\mathrm{b};\mathrm{b}) - V_2(\pi,x;\mathrm{b}) + \mathbb{E}_v[p(v,u_b)]}_{\text{expected return } R(\pi,x;\mathrm{b})} + \mathbb{E}_v\underbrace{\left[V_2(\pi'(v,\mathrm{b}), x-u_\mathrm{b};\mathrm{b}) - V_2(\pi,x-u_\mathrm{b};\mathrm{b})\right]}_{\text{value of information } I(\pi,x;\mathrm{b})}. \end{aligned} \tag{10}$$

The terms in (10) capture the two fundamental motives for trade in our model: *inventory management* and *information acquisition*. The first terms $R(\pi, x; a)$, $a \in \{b, s\}$ capture dealers' inventory management motives for trade in the tradition of Ho and Stoll (1981). The second terms $I(\pi, x; a)$ capture the dealers' information acquisition motives for trade.

The terms $R(\pi, x; \mathrm{s})$ and $R(\pi, x; \mathrm{b})$ represent the expected return of trading with investors under the dealer's optimal future behavior, net of the cost of trade. Dealers want to stock up inventory when demand is low to be able to sell the asset when demand is high.[15] Consistent with this motive, the first two plots in Figure 2 show that the term $R(\pi, x; \mathrm{b})$ prompts dealers to search for and trade with buyers when demand is high. The expected return from meeting a buyer is higher if the dealer expects demand and, hence, prices to be higher. Vice versa, $R(\pi, x; \mathrm{s})$ prompts dealers to search for sellers when demand and prices are low. The term $R(\pi, x; \mathrm{s})$ is decreasing in the dealer's inventory since dealers have stronger incentives to accumulate inventory when inventory is low. In this case, they run the risk of finding themselves in a period of high demand without assets to sell. Symmetrically, $R(\pi, x; \mathrm{b})$ is increasing in inventory.

The terms $I(\pi, x; \mathrm{s})$ and $I(\pi, x; \mathrm{b})$ capture dealers' information acquisition motives for trade. By meeting with investors and observing their value for the asset, dealers acquire information about the current realization of the demand state, since investors' valuations are noisy signals of θ_t. This information is valuable, since it allows dealers to forecast changes in demand. The terms $I(\pi, x; \mathrm{s})$ and $I(\pi, x; \mathrm{b})$ capture the improvement in the dealer's future payoff from his ability to forecast demand more accurately. Therefore, the more valuable is the information a dealer can acquire through trade, the stronger are his incentives to search and trade with investors. The third plot in Figure 2 shows how $I(\pi, x; \mathrm{s})$ depends on the dealer's inventory x and prior beliefs π.[16] Importantly, information is more valuable if it has the potential to improve a dealer's future trading decisions. For this reason, dealers value new information more when they are more uncertain about the current realization of the demand state. By the same token, the value of information is higher when inventory is lower. With a large inventory position the dealer wants to offload his inventory regardless of what he knows about the demand state. Therefore, in this case, the value of information is small.

While the simple model illustrates the main mechanisms at play, it is too stylized and incomplete to identify and measure the interplay between information acquisition and dealers' trading behavior in the data, which is the goal of this paper. For this reason, in the next section we enrich the model in two

[15] Note that short selling in this market is infeasible. Therefore, when demand is high, the dealer can only sell the assets he already owns.

[16] For simplicity we focus on $I(\pi, x; \mathrm{s})$. Plotting $I(\pi, x; \mathrm{b})$ would deliver an analogous intuition.

Figure 2: Dealers' trading decisions

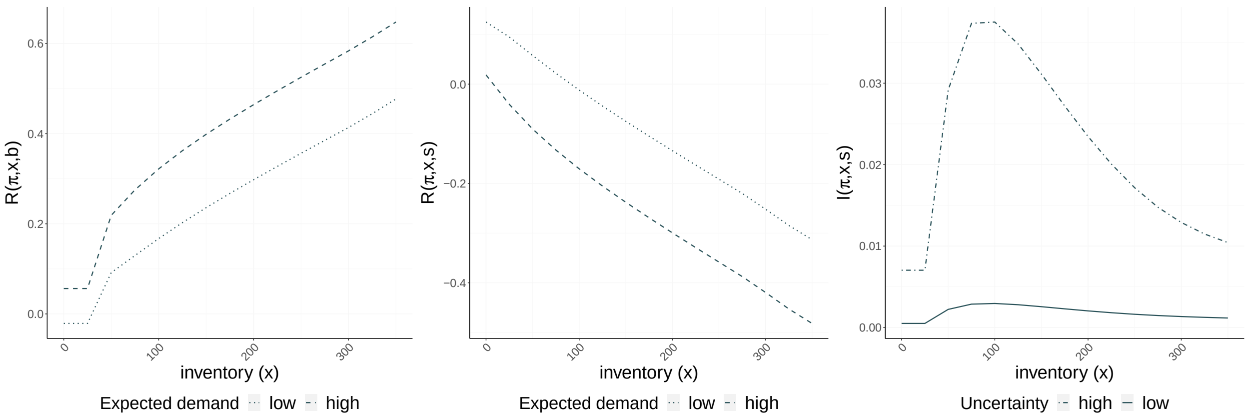

Notes: The Figure illustrate the properties of dealers' search (and trading) decisions. The first two panels show how the expected return of searching for a buyer (and seller) depend on the dealer's inventory and prior belief. Analogously, the third panel show how the value of information depends on inventory and prior beliefs.

important dimensions. First, we model explicitly the entry behavior of investors. Second, we enrich the model with inter-dealer trade. Inter-dealer trade is an important channel through which information about demand percolates among dealers in the market. Therefore it is a potentially important feature both for identifying and quantifying the role of information acquisition in determining trading behavior.

4.2 Empirical model of trading and learning in decentralized markets

4.2.1 Investors' entry decision.

We begin extending the model described in Section 4.1 by endogenizing the entry of investors. Endogenizing investor entry allows us to estimate the cost that investors face when they negotiate and trade with dealers. This, in turn, allows us to measure investor welfare.

At the beginning of each period, masses $\mathcal{I}_b$ and $\mathcal{I}_s$ of potential buyers and sellers observe their private valuation of the asset v and decide whether to enter the market. We denote by $\tilde{F}_v(v|u,\theta)$ the distribution of valuations among potential investors of type u, given the demand shock θ.[17] If an investor is contacted by a dealer, he incurs cost $\phi_u \geq 0$ to conduct the negotiations. If the negotiations are successful, the investor either pays or receives the agreed-upon price $p(v,u)$. Therefore, a potential entrant solves

$$\begin{cases} \max\left\{\lambda(u)\left(v - p(v,u) - \phi_u\right), 0\right\} & \text{if } u = u_b \\ v + \max\left\{\lambda(u)\left(p(v,u) - \phi_u - v\right), 0\right\} & \text{if } u = u_s \end{cases}, \tag{11}$$

[17] While each investor knows his own valuation, they do not observe the common shock θ_t.

where $\lambda(u)$ is the equilibrium probability that an investor with liquidity need u will be contacted by a dealer each period. Substituting for prices in Equation (7), investors' entry decisions satisfy

$$\begin{cases} \mathbb{I}\{(1-\rho)(v-\rho_0(u_b))-\phi_{u_b}\geq 0\} & u=u_b \\ \mathbb{I}\{(1-\rho)(\rho_0(u_s)-v)-\phi_{u_s}\geq 0\} & u=u_s \end{cases}. \tag{12}$$

Given Equation (12), buyers enter if their valuation is large enough. Vice versa, sellers enter if their valuation is large enough. We denote by $F_v(v|u,\theta)$ the equilibrium distributions of valuations among the buyers and sellers active in the market, conditional on the demand shock θ. Crucially, Equation (12) implies that all meetings between a dealer and an investor result in a trade. While we allow for the possibility that dealers meet investors only to learn their valuation, without the intention of trading the asset, Equation (12) ensures that this never happens in equilibrium.

4.2.2 Inter-dealer trade

Next, we extend the model in Section 4.1 to allow dealers to trade among themselves, after having traded with investors. Similar to trade in the investors' market, trade in the dealers' market allows dealers both to exchange an asset and to acquire information about their counterparty's valuation for the asset. Trade in the dealers' market differs from trade in the investors' market in two important dimensions. First, contacts among dealers are not random. Instead, we allow dealers to direct their search towards a specific counterparty based on the quality of the information that they expect to obtain from them. Second, we introduce a different bargaining protocol for inter-dealer trade. The bargaining protocol between dealers and investors accounts for the informational advantage of dealers. While this assumption is natural in the investors' market, we do not want to build in such an asymmetry in the dealers' market.

Inter-dealer trade proceeds as follows. A constant share γ_{d2d} of the dealers are randomly selected to be potential sellers, while the remaining are labeled as potential buyers.[18] Each potential seller can make a take-it-or-leave-it offer $q\in\mathbb{R}_+$ to sell u_{d2d} units of the asset to a potential buyer, where u_{d2d} is an exogenous constant.[19] Potential buyers decide whether to accept the offer they receive or not to trade.[20]

[18]The role of γ_{d2d} is similar to that of the number of potential entrants in standard entry games, since γ_{d2d} sets an upper bound on the total volume of trade in the market. Note, however, that the decision of a dealer to buy or sell an asset in the inter-dealer market remains endogenous, since dealers can decide not to engage in trade and buyers can reject the offer received.

[19]For tractability we don't model the decision on the quantity exchanged in the inter-dealer marker. An intuitive consequence of information acquisition as a motive for trade in the inter-dealer market is that dealers should trade small amounts among themselves. This prediction is confirmed in the data, where the median size of inter-dealer trades is $50,000, a negligible share of the dealers' inventory of the asset.

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

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, the agreed-upon price, the negotiation process itself, and the casual interaction between the parties. To keep the model tractable, we model this information exchange by assuming that, if the offer is accepted, after trading dealers communicate their posterior beliefs to each other.[21] This assumption allows us to leave aside dealers' strategic decision about what information to reveal. As we discuss in Section 4.3, it seems a reasonable assumption given the structure of the municipal bond market.

Contacts between buyer and seller are not random. Instead, dealers direct their offer to a specific counterparty based on their *experience* $e \in \{\bar{e}_1, \ldots, \bar{e}_E\}$. Experience is a variable that captures the quality of a dealer's information about the demand for municipal bonds.[22] If a potential seller decides to make an offer, he chooses the experience level $\tilde{e} \in \{\bar{e}_1, \ldots, \bar{e}_E\}$ of the counterparty, subject to logit shocks. Then, he proceeds to make the offer to a random potential buyer with the chosen experience level, who can either accept or reject the offer. In case of trade, the seller pays a trading cost $c_{\tilde{e}}^{\text{d2d}}$ that captures the time and effort put into finalizing the trade as well as external motives for inter-dealer trade.

A dealer's experience is a proxy for the precision of a dealer's information about the demand for municipal bonds. Dealers accumulate information and experience through their trading activity. In particular, after a dealer with experience e trades with n investors, his experience evolves to $e' = e + n$, and after he trades on the dealers' market with a counterparty with experience $\tilde{e}$, his experience evolves to $e' = e + \alpha\tilde{e}$, where $\alpha \in (0, 1)$ captures how information depreciates. This means that experience grows as dealers trade and it grows faster if they trade with more experienced counterparties. Finally, a dealer's stock of experience depreciates over time at rate $\delta \in (0, 1)$. Indeed, a dealer's stock of experience is constantly being eroded as information becomes less relevant over time. For tractability we assume that the parameters δ and α are exogenous and fixed.

For computational reasons we consider a coarse discretization of the dealer's experience, and assume that a dealer's experience can only take finite values $e \in \{\bar{e}_1, \ldots, \bar{e}_E\}$. Let e^* be the experience level that would be achieved using the process described in the previous paragraph, $\bar{e}_d$ the largest experience state less than e^*, and $\bar{e}_u$ the smallest experience state greater than e^*. Then the dealer's experience at the end of the period is $e' = \arg\min\{|e^* - \bar{e}_u|, |e^* - \bar{e}_l|\}$.

Finally, we make two assumptions to simplify the inference problem faced by dealers. First, borrowing

at random.

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

[22] Beyond experience, inter-dealer trade is *anonymous*: dealers do not keep track of the identity of their trading counterparty. Therefore, dealers only know their peers' experience, while they do not know their inventory level or beliefs.

from the literature on social learning, 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 counterparty's experience.[23] 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. These assumptions, driven by empirical concerns, are discussed in Section 4.3.

Inter-dealer trading behavior. Compared to the simple model described in Section 4.1, a dealer's trading decisions depend not only on his inventory and beliefs but also on his experience e. Indeed, while a dealer's experience is not directly payoff-relevant it affects its "standing" on the inter-dealer market.[24] It is straightforward to generalize value functions (6)-(9) to account for inter-dealer trade.[25] In the following, we focus on the value functions that describe dealers' decisions on the dealers' market.

Consider the situation when a potential buyer with type (π, x, e) 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 rejecting the offer:

$$W_{\mathrm{b}}\left(\pi, x, e; \tilde{q}, \tilde{e}\right) = \max\left\{-\tilde{q} + \mathbb{E}\left[\beta V_0\left(\pi'\left(\tilde{\pi}, \tilde{q}, \tilde{e}\right), x + u_{\mathrm{d2d}}, e'\right)|\tilde{e}, \tilde{q}\right], \beta V_0\left(\pi'\left(\tilde{q}, \tilde{e}\right), x, e\right)\right\}. \tag{13}$$

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 $\pi'\left(\tilde{\pi}, \tilde{q}, \tilde{e}\right)$ in order to account for the new information he observed, namely (i) the offer $\tilde{q}$ he received, (ii) his counterparty's experience $\tilde{e}$, and (iii) the posterior belief $\tilde{\pi}$ that his counterparty communicates. When deciding whether to accept the offer, the dealer computes the expectation over his counterparty's posterior belief $\tilde{\pi}$, conditional on the offer $\tilde{q}$ received as well as on the counterparty's experience $\tilde{e}$. Importantly, the offer $\tilde{q}$ depends both on the counterparty's inventory and his own 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 $\tilde{e}$ 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 with

[23] 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 than standard Bayesian updating.

[24] This is consistent with empirical evidence in Li and Schürhoff (2019) that dealers terms of trade in the inter-dealer market depend on their network centrality in the dealers' network.

[25] See Appendix C.1 for the analogous of value functions (6)-(9) under this more general model.

beginning-of-period 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:

$$W_{\mathrm{s}}(\pi, x, e; \tilde{e}) = \max_{q \geq 0} \mathbb{E}\Big[\mathbb{I}\{\text{acc. } q\}\left(q + \beta V_0\left(\pi'(\tilde{\pi}, \tilde{e}, \text{acc. } q), x - u_{\mathrm{d2d}}, e'\right) + c_{\tilde{e}}^{\mathrm{d2d}}\right) \\ + \mathbb{I}\{\text{rej. } q\} \beta V_0\left(\pi'(\tilde{e}, \text{rej. } q), x, e\right)\Big]. \tag{14}$$

If the offer is accepted, the seller pays cost $c_{\tilde{e}}^{d2d}$ and receives price q, and his inventory and experience change to account for the trade. Moreover, he observes his counterparty's posterior $\tilde{\pi}$ and his beliefs evolve to $\pi'(\tilde{\pi}, \tilde{e}, \text{acc. } q)$ to account for the fact that (i) offer q was accepted by a dealer with experience $\tilde{e}$, and (ii) his counterparty 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 with beginning-of-period value βV_0.[26]

It is worth emphasizing that Equation (14) captures several factors that affect the offer that a dealer makes to buyers with different experience. Most notably, more experienced dealers have, on average, better information about the common demand shock θ. First, this implies that trading with more experienced buyers allows dealers to acquire more precise information about the demand shock θ. Second, being more informed, experienced buyers value the information conveyed by other dealers less. These factors imply that the probability that an offer is rejected is higher if the buyer is more experienced. Third, since experienced buyers have more precise information, trading with them may involve adverse selection. Indeed, when an experienced buyer accepts an offer it may signal that the asset's demand is higher than the seller thought. Vice-versa, a rejection may be a signal that demand is low.[27]

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

$$W_s(\pi, x, e) = \mathbb{E}\left(\max\left\{\max_{\tilde{e} \in \{\bar{e}_1, \ldots, \bar{e}_E\}}\left\{W_s(\pi, x, e; \tilde{e}) + \xi^{\tilde{e}}\right\}, \beta V_0(\pi, x, e) + \xi^0\right\}\right). \tag{15}$$

The expression shows that he can either wait for the next period, when he restarts with value βV_0, or make an offer to a dealer with experience $\tilde{e}$ in which case he pays cost $c_{\mathrm{d2d}}^{\tilde{e}} + \xi^{\tilde{e}}$ and decides what to offer, as captured by $W_s(\pi, x, e; \tilde{e})$. We assume that the shocks $\xi \in \mathbb{R}^{E+1}$ are draws from a double-exponential

[26] Under the assumption that dealers only update based on successful trades, the last terms in Equations (13) and (14) are $\beta V_0(\pi, x, e)$.

[27] Importantly, the buyers' decision to accept the offer depends both on his inventory and on his beliefs. For this reason, a no-trade theorem does not apply and the extent of the adverse selection is limited.

distribution with standard deviation σ_ξ.

4.2.3 Timing and equilibrium.

The timing each period is as follows: the unobserved common demand shock is realized and dealers pay a cost that depends on their accumulated inventory. Potential investors observe their private valuation of the asset and their liquidity needs, and they make their entry decision. Dealers start searching for investors to either buy or sell the asset. The prices for these trades depend on the current value of the common demand state. Finally, dealers trade among themselves.

We consider a long-run equilibrium which consists of (i) distribution $F^*(\omega|\theta)$ of the dealers' types $\omega = (\pi, x, e)$ conditional on the demand shock θ; and (ii) distribution $F_v(v|\theta, u)$ of active investors' valuations conditional on the demand shock θ and their liquidity need, such that:

E1 The distribution $F_v(v|\theta, u)$ of investors' valuations is consistent with investors' entry decision (12);

E2 Dealers' search decisions in the investors' market solve (6) and (8);

E3 Offers and replies in the inter-dealer market maximize (13)-(15);

E4 Conjectures in (6)-(15) are correct, given $F^*(\omega|\theta)$, $F_v(v|\theta, u)$, and the agents' optimal behavior;

E5 $F^*(\omega|\theta)$ is stationary. In particular, let $TF^*(\omega|\theta)$ be the end-of-period distribution of dealers' types, conditional on demand shock θ, implied by the initial distribution $F^*(\omega|\theta)$; the agents' optimal behavior; and the experience process. Moreover, let $\mathbb{P}^*_\theta(\theta', \theta) = P(\theta_{t-1} = \theta'|\theta_t = \theta)$ be the forward probability.[28] Then F^* satisfies

$$F^*(\omega|\theta) = \sum_{\theta'} \mathbb{P}^*{}_\theta(\theta', \theta)\, TF^*(\omega|\theta'). \tag{16}$$

It is worth discussing the stationarity condition **E5** in more detail. $F^*(\omega|\theta)$ describes the distribution of dealers' types at the beginning of the period, conditional on the demand shock θ. As the agents trade during the period according to their equilibrium strategies, their types evolve. 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, Equation (16) requires that the distribution $F^*(\omega|\theta)$ at the beginning of the period, conditional on the period's demand shock θ, coincides

[28] Denoting by $\mathbb{P}^*_\theta(\theta)$ the stationary distribution associated to transition matrix $\mathbb{P}_\theta$, we have $\mathbb{P}^*_\theta(\theta', \theta) = P(\theta_{t-1} = \theta', \theta_t = \theta)\, \mathbb{P}^*_\theta(\theta)$.

with 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.

4.3 Discussion of key assumptions

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

Trade in the investors' market. We begin with the bargaining protocol between dealers and investors. We model this interaction as a one-sided incomplete information game, 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. Moreover, the assumption is consistent with the perception of dealers' informational advantage vis-à-vis investors, a perception that has shaped government intervention in this market (see Section 2.1 for details). Note that, importantly, the dealer does not commit to trading the asset when calling an investor. Therefore, we allow for the possibility that the dealer decides not to trade the asset after observing the investor's valuation.

Also related to the trading between dealers and investors, for tractability we assume that the quantity traded is fixed. Note, however, that we allow the dealer to decide with how many investors to trade. As a result, dealers' trading decisions entail the same trade-off present in a model where the quantity traded is endogenously determined, and larger trades offer more precise information about the state of demand.

Finally, 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 for municipal bonds where we observe a "comparable" amount of sales and purchases for only 7% of dealer-month-issuer pairs.[29,30]

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. To keep the model tractable, we capture these different ways of learning by assuming that, after trading, dealers communicate their posterior beliefs to each other in a non-strategic way. We believe this is a reasonable approximation of what happens in reality. In particular, dealers' communication happens

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

[30] This is not surprising; indeed 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%.

after trade and, therefore, it only affects the strategic interactions of two dealers in the future. The empirical bite of this assumption is limited, since 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 (Green et al., 2010).

In inter-dealer trade we assume that dealers only update their beliefs based on realized trade. This implies that, 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' 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. In Appendix B we use a Hansen-Sargan test for over-identifying restrictions 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 content of this assumption is limited (we use the specification test in Dickstein and Morales (2015)).

Dealer experience. Dealers are ex-ante homogeneous in the model. 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 counterparty's inventory. Fully modeling dealers' expectations about their peers' trading histories is cumbersome, since they are high-dimensional objects. The theoretical literature on trade in decentralized markets has largely sidestepped this issue, assuming that dealers' types are observable and meetings between dealers are random. The conventional approach is too restrictive for our setup since 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 noting that, given this setup, the concept of dealer experience is related to that of dealer centrality in the inter-dealer network. For this reason, our model speaks to the growing empirical literature describing the features of the inter-dealer network, and how dealers' trade prices and allocations depend on their position in the network (Li and Schürhoff, 2019). While we do not target these stylized facts explicitly, Appendix F.1 shows that they arise in equilibrium in response to dealers' incentives to acquire information.

5 Dealer types: estimation and results

The empirical strategy consists of two steps. In the first step, described in this section, we estimate the equilibrium distribution of the dealers' types, which consists of their experience and their beliefs about the unobserved demand shock. In the second step, described in Section 6, the estimates of dealers' experience and beliefs are used together with dealers' behavior to estimate the model primitives. Each step of the estimation described below is followed by the results.

The estimation is based on a comprehensive data set on trading activity in the market for municipal bonds between January 2000 and June 2003. [31]Importantly, the data includes dealer identifiers, that allow us to build a detailed trading history for each dealer that we leverage for identification. See Section 2 for a detailed description of the data.

We group assets based on the state of issuance and estimate the model separately for ten states. To select these, we divide the states of issuance 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 York, and Michigan. This approach allows us to analyze the data more parsimoniously and exclude states without individual income tax, as due to the different tax treatment the market for bonds issued by these states is integrated at the national level.

Appendix D discusses how we recover the sequence $(\theta_t)_{t=0}^T$ of realizations of the demand shock and their transition matrix $\mathbb{P}_\theta$. Here, we treat these objects as known and assume that the demand shock θ_t can take three values.

5.1 Experience

Based on the process described in Section 4.2.2 for a dealer with experience e_t who trades with n_t investors and a counterparty with experience $\tilde{e}_t$,[32] experience evolves according to

$$e_{t+1} = \delta e_t + n_t + \alpha \tilde{e}_t. \tag{17}$$

[31]After June 2003 the policy described in Section 2.1 muted dealers' learning incentives. In Section A, we leverage this policy to validate our model.

[32]For tractability, 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.

As discussed in Section 4.2, sellers offer different prices to buyers with different experience to account for factors like the different information content of the trade as well as the different rejection probabilities. Following this observation, we exploit the variation in trading prices in the inter-dealer market to identify dealers' experiences. In particular, we recover parameters α and δ by estimating the baseline specification

$$\log(p_i) = \phi_{s_i \times \mathrm{m}_i \times a_i} + \phi_{b_i} + \phi_0 \log\left(e_{b_i,\mathrm{m}_i}(\delta, \alpha)\right) + \phi_1 x_{b_i,\mathrm{m}_i} + \upsilon_i, \tag{18}$$

where s_i, b_i, a_i and m_i denote, respectively, seller, buyer, asset traded, and month in which trade happens.[33] The parameter ϕ_0 measures the impact of a dealer's experience on the price he pays on the inter-dealer market.[34] We estimate the parameters in Equation (18) using non-linear least squares.

Identification of the parameters in Equation (18) relies on comparisons of inter-dealer prices in trades for specific asset a_i, seller s_i, and month m_i. Especially, Equation (18) attributes systematic differences in prices 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 $\phi_{s_i \times \mathrm{m}_i \times a_i}$ absorbs market-wide shocks to prices, as well as the seller's persistent heterogeneity that might affect prices, while the fixed effect ϕ_{b_i} absorbs the buyers' persistent heterogeneity. Finally, we control for the buyers' inventory in x_{b_i,m_i}.

Note that the experience process described in Equation (17) is real valued. To bridge the gap with the model, we consider a coarse discretization of experience, and we run a k-means algorithm with three centers to identify the clusters.

Results. The first panel of Table 3 summarizes the distribution of the estimates for parameters α and δ across the U.S. states, with detailed estimates described in Table A.17. The parameter δ captures the depreciation of a dealers' experience as the information gathered becomes less relevant over time, with an average value across states of $\delta = 53\%$. This implies that only 2% of the experience that dealers accumulate lives through six months. Therefore, our estimates suggest that information in the market is short-lived.

The estimates across states show that there is substantial heterogeneity in the persistence of experience. Indeed, differences in the estimates for δ reflect institutional differences in the market for municipal bonds across states. Consistent with the interpretation of δ as depreciation in dealers' information, the first row in the second panel of Table 3 shows that δ is higher and hence information depreciates more slowly in

[33] Note that, given parameters α and δ, Equation (17) pins down dealers' experience given an initial condition. We impose that dealers' experience at the beginning of the sample is zero, and we drop the first year of data to estimate the parameters.

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

states where the demand process is more persistent, based on our estimates.

Illiquidity is a key driver of demand volatility in the market for municipal bonds. For this reason, we expect the heterogeneity across states in the estimates for δ to reflect differences in liquidity of the municipal bond market. In Table 3 we test the connection between the persistence of experience and liquidity by correlating our estimate for δ with (i) market depth, as measured by the states' total outstanding municipal debt, and (ii) a standard measure for market illiquidity proposed by Amihud (2002).[35] Table 3 shows that information depreciates more slowly in states where the market has more depth and appears more liquid.

Table 3: Estimates of the experience process

	Average	Standard Deviation
δ	0.531	0.231
α	0.352	0.269

	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 (17). 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.

5.2 Dealers' beliefs

The next step of our empirical strategy is to recover the equilibrium distribution $F^*(\pi|e,\theta)$ of dealers' beliefs π conditional on the demand shock θ and dealers' own experience e. We leverage the equilibrium stationarity condition **E5** and identify $F^*(\pi|e,\theta)$ as the solution to the fixed point

$$F^*(\pi|e,\theta) = \sum_{\theta'} P_\theta^*(\theta',\theta)\, TF^*(\pi|e,\theta'), \tag{19}$$

where T maps the equilibrium distribution $F^*(\pi,e|\theta)$ to the distribution of dealers' beliefs at the end of the period, after they have updated their beliefs based on the information they gathered about the demand state.[36] Crucially, our approach allows to recover the equilibrium distribution $F^*(\pi|e,\theta)$ without simulating agents' equilibrium behavior.

[35] In particular, we measure asset j's illiquidity as $\text{Amihud}_j = \frac{1}{n_j}\sum_{t=1}^{n_j}\frac{|r_{jt}|}{Q_{jt}}$, where r_{jt} is 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 millions \$. A state's illiquidity consists of the average illiquidity of its assets.

[36] In the estimation, we assume that dealers observe a public signal correlated with last period's trading activity. To construct this signal, we regress the average market price of municipal bonds in each state on (i) interest rate for 1, 5, and 10 year Treasuries, (ii) public performance indexes for the municipal bond market, and (iii) total new municipal debt issued.

Updating. To operationalize Equation (19), it is useful to analyze the dealers' updating rule after trading with investors and other dealers. First, consider a dealer with prior beliefs π who trades with an investor of type (v, u). Note that the dealer's posterior beliefs can be written as

$$\pi_{\text{inv}}(v, u, \pi; F_{\text{p}}) = \frac{dF_{\text{p}}(p(v,u)|\theta, u)\,\pi(\theta)}{\sum_{\theta'} dF_{\text{p}}(p(v,u)|\theta, u)\,\pi(\theta')}, \tag{20}$$

where $p(v, u)$ is the trade price and $F_{\text{p}}(p|\theta, u)$ is the distribution of prices conditional on the demand state and type of trade. Equation (20) relies on the simple observation that in the investors' market the trade price is a strictly increasing and deterministic function of the investor's valuation (see Appendix C.3 for details). For this reason, the dealer's updating rule can be written as a function of the trade price p and its distribution, rather than a function of the investor's valuation and its distribution. This substantially simplifies the estimation. Indeed, while trade prices and their distribution are observed, investors' valuations and their distribution are not.

Next, we consider a dealer with prior beliefs π and experience e who trades with a dealer with experience $\tilde{e}$ and prior belief $\tilde{\pi}$. The dealers' posterior beliefs do not depend on the trade price $\tilde{q}$ and satisfy

$$\pi_{\text{d2d}}(\tilde{\pi}, \tilde{e}, \pi; F^*_{\text{inv}}) = \frac{dF^*_{\text{inv}}(\tilde{\pi}|\tilde{e}, \theta)\,\pi(\theta)}{\sum_{\theta'} F^*_{\text{inv}}(\tilde{\pi}|\tilde{e}, \theta)\,\pi(\theta')}, \tag{21}$$

where $F^*_{\text{inv}}(\tilde{\pi}|\tilde{e}, \theta)$ is the equilibrium distribution of dealers' beliefs after trading with investors, conditional on demand shock θ and the dealer's own experience. To gain intuition for Equation (21), consider the case of a dealer purchasing an asset in the inter-dealer market. The dealer updates his beliefs based on (i) the purchasing price that he is offered and (ii) the posterior belief that his counterparty communicates after the trade. Equation (21) relies on the observation that the dealer's update does not depend on the offer received. While the offer conveys information about what his counterparty knows about the demand shock θ, the beliefs communicated after the trade are a sufficient statistic for this information.

A fixed-point algorithm. Equations (20) and (21) allow us to exploit the following fixed-point algorithm to derive the solution to Equation (19). Each iteration of the algorithm starts with a guess for beliefs $\pi^{(m)} = \left(\pi^{(m)}_{dt}\right)_{d\in D, t\in\{0,\ldots,T\}}$ for each dealer d and period t in our data set, together with an estimate of their distribution $F^{(m)}(\pi|e, \theta)$ across dealers.[37] Note that the distribution $F^{(m)}(\pi|e, \theta)$ corresponds to

[37] We initialize $\pi^{(0)}$ to the dealers' beliefs 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 (20) to compute initial guess $\pi^{(0)}$. Remember that at this stage, each dealer's experience is known.

a guess for the equilibrium distribution $F^*(\pi|e,\theta)$. Next, we update this guess in three steps.

The first step consists of updating the dealers' beliefs based on trades with investors. For every observed trade between a dealer with imputed beliefs π and an investor of type u at price p, we update each dealer's beliefs according to $\pi_{\text{inv}}(p,u,\pi;\hat{F}_p)$, where $\hat{F}_p$ is an estimate of the distribution of trading prices F_p. This allows us to recover the distribution $F^{(m)}_{\text{inv}}(\pi|e,\theta)$ of dealers' beliefs after trade with investors, implied by the dealers' observed behavior and initial guess $\pi^{(m)}$.

The second step consists of updating the dealers' beliefs based on trades with other dealers. For every observed trade between two dealers with imputed beliefs π and $\tilde{\pi}$, we update each dealer's beliefs according to $\pi_{\text{d2d}}(\pi,\tilde{\pi},\tilde{e};F^{(m)}_{\text{inv}})$. Finally, the third step consists of updating dealers' beliefs to account for the evolution of θ. In particular, if a dealer has belief π after the updates in the previous two steps, his beliefs at the end of the period are $\pi'(\theta)=\sum_{\theta'}\mathbb{P}_\theta(\theta,\theta')\,\pi(\theta)$.

Note that the distribution of dealers' beliefs after the three updating steps described above corresponds to the end-of-period distribution $TF^{(m)}(\pi|e,\theta')$ implied by the initial guess $F^{(m)}(\pi|e,\theta)$ and dealers' observed behavior. In light of this, the algorithm ends if $F^{(m)}$ satisfies the criterion

$$\left\| F^{(m)}(\pi|e,\theta) - \sum_{\theta'} \hat{P}_\theta\left(\theta',\theta\right) TF^{(m)}\left(\pi|e,\theta'\right) \right\| < \epsilon. \tag{22}$$

Otherwise, we set a new guess for each dealer's beliefs equal to

$$\pi^{(m+1)}_{d0}(\theta) = \mathbb{P}^*_\theta(\theta), \quad \pi^{(m+1)}_{dt}(\theta) = \pi'^{(m)}_{d,t-1}(\theta), \tag{23}$$

where $\mathbb{P}^*_\theta$ is the stationary distribution associated with transition matrix $\mathbb{P}_\theta$. Once the new guess for the dealers' beliefs is computed, the algorithm moves to the following iteration.

Results. The estimates of dealers' beliefs allow us to measure the contribution of experimentation to dealers' information about the state of demand in the market for municipal bonds. To illustrate the results, we compute the precision of dealers' forecasts of average sale prices:

$$\mathit{precision}(\pi) = \sqrt{\mathbb{E}\left(\left[\mathbb{E}_\pi\left(p\left(v_{it},u_{\text{s}}\right)\right) - \mathbb{E}\left(p\left(v_{it},u_{\text{s}}\right)|\theta_t\right)\right]^2\right)}, \tag{24}$$

where $\mathbb{E}_\pi(p(v_{it},\text{s}))$ is the expected market selling price for the asset for a dealer with belief π, and $\mathbb{E}(p(v_{it},\text{s})|\theta_t)$ is the average selling price for the asset, given the realization of demand shock θ_t. Table 4 reports the average forecast precision for dealers with different experience levels.

Table 4: Precision of information

	Uninformed Dealers	**Market Average**	**Inexperienced Dealers**	**Experienced Dealers**
Precision of dealers' price forecasts	1.588	1.344	1.422	1.112
Percentage	100%	84.5%	89.4%	69.9%

Notes: The table reports the average precision of dealers' price forecast for different classes of dealers. For each state we compute compute the average precision based on the equilibrium distribution of dealers' types, the table reports the average of this measure across states. The first column reports the measure for dealers who ignore the information content of trading activity, and only update based on public information. The second column reports the average precision across all dealers, and the last two columns distinguish among experienced and inexperienced dealers.

We find that experimentation allows dealers to substantially improve the precision of their estimate. Comparing the first two columns of Table 4 reveals that the average dealer has 16% higher precision compared to an "uninformed" dealer who ignores the information content of his own trades. For experienced dealers the improvement is even more striking. As shown in the third column of Table 4, information acquired through trade improves the precision of experienced dealers' predictions by 31%. Finally, we find that more experienced dealers have better information, as expected. Experimentation improves inexperienced dealers' forecasts only by 11%.

6 Model primitives: estimation and results

This section describes how we estimate the model primitives and presents the results. On the dealer's side there are four sets of primitive parameters in the model: (1) inventory cost κ, (2) distribution of search costs faced by dealers $F_c(\cdot|a)$, (3) search costs for inter-dealer trade $c^{\mathrm{d2d}} = \{c_{\tilde{e}}\}_{\tilde{e}=1}^{E}$, and (4) standard deviations $(\sigma_\epsilon, \sigma_\xi)$ associated with the preference shocks. Additionally, we recover the distribution of investors' valuations and their entry costs. In both cases, we allow the parameters to vary across municipal bonds issued in different U.S. states.

We calibrate the discount factor to $\beta = 0.999$, the probability that a dealer can search for multiple investors to $\gamma = 0.95$, dealers' bargaining power to $\rho = 0.75$, and the fraction of potential sellers to $\gamma_{\mathrm{d2d}} = 0.5$.[38] Finally, we set the size of different types of trades, u_{s}, u_{b}, and u_{d2d} equal to their sample medians.[39] As in Section 5, the estimation leverages data on dealer behavior in the market for municipal

[38] The results are similar for $\rho = 0.5$ and $\rho = 0.9$.

[39] The values of u_{s}, u_{b}, and u_{d2d} vary across states. On average, $u_{\mathrm{s}} = \$23,888$, $u_{\mathrm{b}} = \$21,667$, and $u_{\mathrm{d2d}} = \$41,389$.

bonds between January 2000 and June 2003.

Estimation of dealer costs. We assume a flexible parameterization of the inventory costs $\kappa(x)$, imposing that inventory costs are a third-order polynomial of a dealer's inventory x with parameters $(\kappa_j)_{j=1}^{3}$.[40] 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.

Our 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 predictions for dealers' behavior, described in Section 4.2.2, 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 concern dealers' behavior vis-à-vis investors. We match the probability that a dealer with type (π, x, e) buys an asset, sells an asset, or does not trade. Moreover, we match the probability that a dealer with type (π, x, e) trades one additional unit, having already traded n units with investors of type $a \in \{\text{b}, \text{s}\}$. The second set of moments concern dealers' behavior on the inter-dealer market. In particular, we match the probability of observing a dealer with type (π, x, e) selling an asset to a dealer with experience level $\tilde{e}$. It is worth noting that, to compute each one of these moments, it is essential that we observe dealer identifiers.

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

$$\hat{\tau} = \arg\min \left(\hat{\mathbb{P}} - \Psi(\tau)\right)' \Sigma \left(\hat{\mathbb{P}} - \Psi(\tau)\right), \tag{25}$$

where $\hat{\mathbb{P}}$ is the stacked vector of observed choice probabilities, $\Psi(\tau)$ is the vector of respective predicted choice probabilities, and Σ is a weighting matrix.[41]

Identification of dealers' value of information. While the data are used jointly to identify the parameters, we can consider what variation specifically allows us to identify dealers' value of the information that they acquire through trading. When deciding with whom to trade on the inter-dealer market, dealers

[40] As a robustness we run the estimation with a fourth- and fifth-order polynomial and find very similar results.
[41] As a weighting matrix we use the covariance matrix of the estimated choice probabilities $\hat{\mathbb{P}}$.

face the following trade-off: selling an asset to a more experienced counterparty is less profitable, but it conveys more information about the common demand shock. How a dealer solves this trade-off as a function of his prior belief π pins down his value of information. Indeed, if a dealer is more uncertain about the demand shock, then information is weakly more valuable (this is a general property of the value of information). 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 a dealer's value of information.

Table 5 confirms that the expected pattern holds in the data. Columns (I) and (II) in Table 5 show that more experienced buyers pay lower prices in the inter-dealer market, after controlling for a rich set of fixed effects.[42] Columns (III) and (IV) of Table 5 show that when a dealer has a more uncertain prior, as measured by Shannon's entropy, $entropy(\pi) = \sum_{\theta} \pi(\theta) \log(\pi(\theta))$, he tends to substitute into the more informative yet expensive trade with experienced dealers.

Table 5: Identification of the value of information

	Inter-dealer price (log)		counterparty's experience (log)	
	(I)	(II)	(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×issuer 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' counterparty 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

Results. Table 6 reports the baseline estimates for dealer's cost parameters averaged across U.S. states, with the full set of estimates reported in Table A.19 in Appendix E. Consistent with industry narratives and prior studies, search costs are large and trading municipal bonds requires lengthy negotiations. Our estimates imply that a dealer faces average search costs that range from 12% to 16% of trade size, while the

[42]Importantly, identification does not rely on the fact that experienced buyers pay lower prices on the inter-dealer market. Rather, the identification relies on how, given the prices, the sellers substitute across trade with buyers with different experience level as a function of their prior beliefs.

average search cost payed by dealers ranges from 1.3% to 2.2% of trade size (the difference arises because the dealer only trades for low realizations of the cost shock). This corresponds to an average of \$238 to find an investor interested in selling and \$619 to find an investor interested in buying an average trade size of \$25,000. The asymmetry is consistent with the notion that finding buyers is more difficult than finding sellers; indeed, as industry practitioners say, "municipal bonds are sold, not bought" (Feldstein et al., 2008). 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. These estimates can be interpreted as a measure of model fit and suggest that preference shocks are not the key driver of dealers' trading decisions.

Based on our estimates, there is substantial heterogeneity in the shape and curvature of the inventory cost function across assets issued in different states. The estimates suggest that the heterogeneity reflects differences in credit rating. Rating explains 35% of the variation in inventory cost across states. Doubling the share of assets with Aaa rating decreases the cost of inventory by 40% across states, for the average dealer. This is consistent with dealers facing risk constraints when managing inventory.

Table 6: 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.

Investors' valuations. We use data on trade prices between investors and dealers to recover investors' valuations point-wise. In particular, for every investor with type (v, u) trading with a dealer at price p, we can rewrite Equation (7) as

$$v = \frac{p - (1-\rho)\,\rho_0\,(u)}{\rho}, \tag{26}$$

where $\rho \in (0,1)$ is the parameter capturing dealers' bargaining power and $\rho_0\,(u)$ is an equilibrium object that only depends on the estimates of dealers' valuations. Note that we can leverage Equation (26) to directly recover investors' valuations, after we calibrate ρ. Indeed, the term $\rho_0\,(u)$ can be readily recovered based on the estimates of the dealers' cost parameters.

We exploit investors' entry conditions (12) to estimate investors' bargaining costs, $\phi_{\rm b}$ and $\phi_{\rm s}$. To recover the distribution of investors' valuations before the entry decision, we assume that among potential investors of type a, the valuations are log-normally distributed with mean $\mu_v(\theta, a)$ and standard deviation $\sigma_v(\theta, a)$. These parameters are recovered through maximum likelihood. The full set of estimates is reported in Table A.20 in Appendix E.

7 Impact of information acquisition motives for trade

This section quantifies dealers' information acquisition motives for trade and their impact on market liquidity and welfare. Our results reveal that dealers' incentives to acquire information benefit both investors and issuers as they strengthen dealers' incentives to trade in a notoriously illiquid market.

Furthermore, we use our model to trace out how the equilibrium forces and model primitives determine the strength of dealers' information acquisition motives for trade across U.S. states. This analysis illustrates general principles that can apply to markets other than that for municipal bonds.

7.1 Value of information

How strong are dealers' incentives to experiment? How do the model primitives determine their prevalence? To answer these questions, in this section we quantify the value of information for dealers in the market for municipal bonds. Following Section 4.1, we define a dealers' value of information as

$$I(\pi, x, e; a) = \mathbb{E}_v \left[V_2 \left(\pi'(v, a, \pi), x, e; a \right) - V_2 (\pi, x, e; a) \right]. \tag{27}$$

The term I is the highest price that a dealer of type (π, x, e) is willing to pay to observe an investor's valuation.[43] When searching for investors, the dealer draws at random an investor with valuation $v \sim F_v(v|\theta, u_a)$ and updates his beliefs to $\pi'(v, a, \pi)$ as he learns that the investor of type a has valuation v for the asset. The term I captures the impact of this piece of information on the dealer's value function, holding his inventory and experience fixed. In other words, this term captures the dealer's future expected

[43] Similarly to the analysis for the simple model in Section 4.1, in the general model a dealer chooses to pay the search cost to trade with investors if his draw is lower than a threshold $\bar{c}(\pi, x, e, a)$. For $a = s$ (and analogously for $a = b$) the threshold satisfies

$$\begin{aligned} \bar{c}(x, \pi, e; \mathrm{s}) =& \mathbb{E}_v \left[-p(v, u_s) + V_2(\pi, x + u_s, e; s) - V_2(\pi, x, e; s) \right] \\ &+ \mathbb{E}_v \left[V_2 \left(\pi'(v, s, \pi), x + u_s, e'; s \right) - V_2 \left(\pi, x + u_s, e'; s \right) \right] \\ &+ V_2 \left(\pi, x + u_s, e'; s \right) - V_2 (\pi, x + u_s, e; s). \end{aligned}$$

payoff derived from his ability to forecast demand more accurately. As discussed in Section 4.1, the term I influences dealers' incentives to search and trade with investors and offers a proxy for the strength of dealers' information acquisition motives for trade.

Table 7: Value of information

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

	Value of Information (log) (I)	Value of Information (log) (II)
Marginal cost of inventory (log)	−14.660	
Average search cost (log)	−1.877	
$\mathbb{E}_b(v\|\theta_H)$-$\mathbb{E}_b(v\|\theta_L)$ (log)	0.793	
Demand persistence		7.199
Demand persistence squared		−0.018
R-squared	0.21	0.46

Notes: The left panel of the figure reports the average value of $I(\pi, x, e; a)$, which measures the average value of information across states, and its standard deviation. The first column of the right panel of the figure reports the results of a regression of $I(\pi, x, e; a)$ on (i) the marginal cost of inventory, (ii) trading costs (iii) the change in buyers' valuation between low and high demand state. The second column of the right panel of the figure reports the results of a regression of $I(\pi, x, e; a)$ on the persistence of the demand process and its square. The persistence of the demand process is measured as the average time between changes in the demand state.

The left panel of Table 7 reports dealers' value of information averaged across states, as well as its standard deviation.[44] The estimated value of information is positive and economically large. For the average dealer, the information conveyed by trading with investors is worth 12% of the intermediation spread, corresponding to 7 basis points. The left panel of Table 7 reveals large heterogeneity across states in dealers' value of information. We explore this heterogeneity in the right panel of Table 7 that shows how the model primitives affect the dealers' incentives to acquire information.

The value of information hinges on a dealer's incentives to time the sale of the asset to match changes in demand. The first column in the right-hand panel of Table 7 regresses the value of information across states on several factors that affect these incentives. The table shows that the value of information is large when search costs are small and when fluctuations in investors' valuations are large. Indeed, in states where buyers' valuations change more in response to changes in demand, dealers have stronger incentives to sell the asset at the "right moment". This, in turn, increases the value of information. Moreover, the value of information is highest when inventory costs are small. When facing a large marginal cost of inventory, a dealer will want to offload inventory regardless of what he knows about the demand state. Therefore, the value of information is small in this case.

The second column in the right-hand panel of Table 7 explores how the persistence of the demand

[44]More precisely, in each state we compute the average value of $I(\pi, x, e; a)$ based on the equilibrium distribution of dealers' types. The table reports the average and standard deviation of this quantity across states.

process affects the value of information. It reveals that the value of information is highest for intermediate levels of the persistence of the demand state. There are two effects at play. On the one hand, the value of information for a dealer is highest when demand is persistent. Indeed, the dealer can leverage the information he has gathered only if he expects demand not to change in the following period. On the other hand, when demand is more persistent dealers tend to have more precise information in equilibrium, since information about the demand depreciates more slowly.[45] This equilibrium effect lowers the marginal value of information.[46] The interplay of these two forced implies that the value of information is hump-shaped with respect to demand persistence.

7.2 Volume of trade and welfare

In this section, we study how dealers' learning through trade impacts the market equilibrium.

Dealers' information acquisition through trade has both a direct and an indirect effect on the market equilibrium that are best showcased looking at dealers' decision to search for investors. When a dealer decides whether to trade with an investor, he contemplates two factors. First, the dealer considers the expected return on an asset, capturing factors such as the expected resale price, the price that he expects to pay for it, as well as the holding costs he will bear until selling. Second, the dealer anticipates that he can acquire valuable information by interacting with investors. Learning through trade shapes the dealers' trading behavior by interacting with both of these factors. First, if dealers can acquire valuable information by interacting with investors, they will have stronger incentives to trade and provide liquidity. Second, as shown in Table 4, dealers can improve the precision of their forecast about demand by up to 30% thanks to the information they acquire through trading. The availability of better information about demand indirectly affects dealers' motives for trade since it improves the return that a dealer expects to realize by trading the asset: Having more information, the dealers anticipate that they will be able to sell and buy the asset at more favorable times.

To isolate the direct and indirect impact of dealer's information acquisition on the market equilibrium we consider two counterfactual scenarios. First, we compute and simulate the long-run equilibrium for the market under the assumption that dealers ignore the informational content of their trades. Second, we shut down dealers' learning but hold the precision of dealers' information fixed.[47] By holding dealers'

[45] As an example, if the realizations of demand were i.i.d over time, dealers would have no information about the its current realization, despite learning.

[46] As explained in detail in Section 4.1, an additional piece of information is less valuable when a dealer is more confident about the current realization of demand.

[47] To operationalize the second counterfactual scenario, we assume that dealers do not acquire information when they trade. However, at the end of each period dealers observe a signal with the same information content that they would acquire, on

Table 8: Economic impact of dealers' learning

	Impact	
	Total	Direct
Liquidity	-11.92%	-4.39%
Investor welfare (%)	-4.49%	-3.05%
Investor welfare ($ millions)	-1,276	-720
Dealers' purchasing price	-7bp	-4bp
Dealers' selling price	11bp	9bp

Notes: The first column presents the total effect of the dealers' information acquisition on the market equilibrium, while the second column presents the direct effect that holds the average precision of dealers' information constant to the observed level.

information fixed, we can turn off the indirect effect of dealers' learning.

Table 8 succinctly describes the equilibrium by reporting several measures that showcase the impact of information acquisition motives for trade for both issuers and investors. The results reveal that dealers' learning though trade is a significant driver of market liquidity, which we measure as the number of investors who enter the market to sell an asset and are able to do so within a period. Indeed, as shown in the first row of Table 8, halting dealers' learning lowers liquidity by 9%. A comparison between the first and second column in Table 8 reveals that the weakened information-acquisition motive for trade is responsible for more than half of the decline in trading volume.

The decline in trading volume hurts both investors and issuers in the market. The second column of Table 8 shows that investor welfare falls by almost 5% in the scenario where dealers ignore the information content of their own trading activity. The total decline in investor welfare amounts to more than $1 billion every year.[48] The effect is large because the valuation, net of entry cost, of an average investor buying municipal bonds is 7% higher than the one selling bonds. As trading volume declines, asset misallocation increases, ultimately hurting investors. It is worth emphasizing how dealers' learning impacts the cost of raising capital for municipalities. As shown in Table 8, shutting down learning through trade lowers the price at which the dealer is willing to buy by 7 basis points, thus raising the cost of capital for

average, learning from their trading activity in the baseline equilibrium.

[48] We compute investors' welfare as

$$W(a) = \begin{cases} \sum_{\theta} n_t(\theta, u_{\mathrm{b}}) \int (v - p(v, u_b) - \phi_b)\, dF(v|u_b, \theta) & a = \mathrm{b} \\ \mathcal{I}_s \mathbb{E}(v|u_s) + \sum_{\theta} n_t(\theta, u_{\mathrm{s}}) \int (p(v, u_s) - v - \phi_s)\, dF(v|u_s, \theta) & a = \mathrm{s} \end{cases},$$

where $n(\theta, u)$ denotes the total number of investors of type $u$who trade the asset in a period, given demand shock θ, and $\mathcal{I}_s$ is the number of potential seller. We calibrate $\mathcal{I}_s$ to be equal to two times the largest number of trades we observe in any given period. Note, moreover, that the estimates of the $ change in investors' welfare do not depend on $\mathcal{I}_s$.

municipalities. In sum, dealers' learning benefits both investors and issuers as it increases trading volume.

8 Market transparency

Assets in decentralized markets are usually traded in an opaque environment with limited or no public information about market activity. Financial authorities in the U.S. and abroad have been steadily improving access to information about trade activity with the objective to increase market liquidity. One argument for the effectiveness of these policies is that dealers have an informational advantage vis-à-vis investors. By eroding this informational advantage, access to public information should encourage investor participation in the market and increase liquidity and welfare. These arguments ignore, however, dealers' incentives to trade which are a key driving force behind market liquidity.

The model allows us to quantify the effect of market transparency on dealers' incentives to trade. To capture a market with post-trade transparency, 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. Figure 3 plots the impact of the policy on market liquidity, which we measure as the number of investors are able to sell an asset within one period. For six out of ten states, the policy weakens dealers' incentives to trade. As a result in these states both market liquidity and investor welfare decline, reducing the positive effect of transparency on investor participation and liquidity. There is substantial heterogeneity across states in the effect of the policy. Liquidity is predicted to fall by as much as 6% in New Jersey while it increases by as much as 3% in Michigan.[49]

Two effects are at play that explain the heterogeneous market impact. On the one hand, transparency weakens the incentives to experiment. When information on trading activity is made public, the information conveyed by trade becomes irrelevant. This makes each trade less valuable for the dealers and implies that dealers trade less frequently. On the other hand, transparency lowers uncertainty, which increases trading volume and partially offsets the first effect. Figure 3 shows that the second effect dominates when demand is less volatile. In fact, the decline in trading activity is strongest in states where demand is most volatile, such as Arizona, and Mississippi. States where demand is more persistent, such as Michigan and Kansas, are predicted to experience an increase in overall trading activity.

Demand uncertainty explains only a small portion of the effect of market transparency predicted by

[49] As described in Section 2.1, 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. In order to validate our model, in Appendix A we show that our model predict 51% of the differences in the impact of the policy for assets issued by different states.

Figure 3: Impact of market transparency on market liquidity and investor welfare

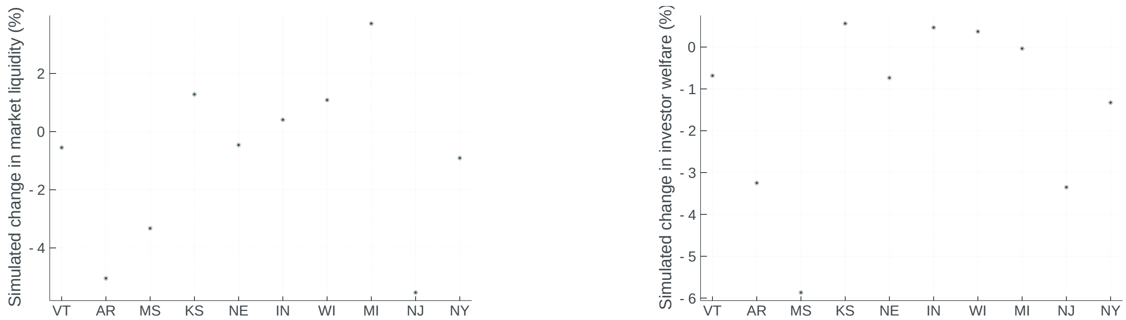

Notes: The left panel of the figure displays the change in market liquidity associated to an increase in market transparency across different classes of assets. We measure market liquidity as which the number of investors who enter the market to sell an asset and are able to do so within a period. The right panel of the figure displays the change in investor welfare associated to the introduction of market transparency.

the model, as it can be seen by looking at the impact of the policy in New Jersey. Indeed, demand for municipal bonds is relatively stable here, yet this state sees the largest decline in trading activity 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. Thus, as discussed in Section 7.2, information is particularly valuable here, as dealers face lower impediments to trade.

It is important to emphasize the distributional impact of the policy. To this end, we exploit our model to trace out how the predicted changes in trade activity translate into changes in investor welfare. The weakening of dealers' incentives to provide liquidity is harmful to investors. Indeed, under market transparency, the dealers exercise their market power more by curbing their trading activity. Market power, therefore, results in substantial misallocation. Counteracting this effect, information acquisition motives for trade benefit investors and reduce assets' misallocation in an opaque market. In the states most affected by market transparency, New Jersey, and Mississippi, investor welfare declines by 3.3% and 5.8%. Investor welfare, instead, increases in states like Michigan and Kansas where liquidity improves. Overall investor welfare decreases by 2%, or around $ 568 million per year.

We conclude our analysis with some caveat about the interpretation of results. Our model focuses on dealers' incentives to trade. Consistent with this, our counterfactual analysis capture how market transparency affects dealers' trading behavior, holding investors' behavior fixed. We do not model how investors' search behavior and information change in reaction to market transparency. Naturally, a richer model would be necessary to derive policy recommendations. Nevertheless, our approach allows us to highlight an important channel through which market transparency can have unintended impact on market liquidity and welfare. Policy makers should to take these repercussions into consideration to assess more comprehensibly the impact of market transparency in opaque decentralized markets.

9 Conclusion

This paper studies the role of learning by trading, or strategic experimentation in dealer-intermediated markets that are common in wholesale trade and financial markets. To measure the incentives to experiment, we use a regulatory dataset of transactions on the secondary market of municipal bonds. We show that information acquisition about investor demand is an important determinant of trading activity and price formation. Transparency regulations weaken dealers' learning incentives and, as a result, may hurt market liquidity and investor welfare. Our results highlight unintended consequences of market interventions in decentralized markets that may limit the benefits of these policies.

References

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

Aghion, P., M. P. Espinosa, and B. Jullien (1993): "Dynamic duopoly with learning through market experimentation," *Economic Theory*, 3, 517–539.

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

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

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

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

Bajari, P., C. L. Benkard, and J. Levin (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75, 1331–1370.

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

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

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

Binmore, K., A. Rubinstein, and A. Wolinsky (1986): "The Nash bargaining solution in economic modelling," *The RAND Journal of Economics*, 176–188.

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

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

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

Chandrasekhar, A. G., H. Larreguy, and J. P. Xandri (2012): "Testing models of social learning on networks: Evidence from a framed field experiment," *Work. Pap., Mass. Inst. Technol., Cambridge, MA.*

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

Dickstein, M. J. and E. Morales (2015): "What do exporters know?" *National Bureau of Economic Research.*

Duffie, D., N. Gârleanu, and L. H. Pedersen (2005): "Over-the-Counter Markets," *Econometrica*, 73, 1815–1847.

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

Duffie, D. and G. Manso (2007): "Information percolation in large markets," *The American economic review*, 97, 203–209.

Ellison, G. and D. Fudenberg (1993): "Rules of thumb for social learning," *Journal of Political Economy*, 101, 612–643.

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

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

Feldstein, S. G., F. J. Fabozzi, et al. (2008): *The handbook of municipal bonds*, vol. 155, John Wiley & Sons.

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

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

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

Glode, V. and C. Opp (2016): "Asymmetric information and intermediation chains," *American Economic Review*, 106, 2699–2721.

Golosov, M., G. Lorenzoni, and A. Tsyvinski (2014): "Decentralized trading with private information," *Econometrica*, 82, 1055–1091.

Green, R. C., B. Hollifield, and N. Schürhoff (2006): "Financial intermediation and the costs of trading in an opaque market," *The Review of Financial Studies*, 20, 275–314.

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

Grossman, S. J., R. E. Kihlstrom, and L. J. Mirman (1977): "A Bayesian approach to the production of information and learning by doing," *The Review of Economic Studies*, 533–547.

Grossman, S. J. and M. Perry (1986): "Perfect sequential equilibrium," *Journal of economic theory*, 39, 97–119.

Gul, F. and H. Sonnenschein (1988): "On delay in bargaining with one-sided uncertainty," *Econometrica: Journal of the Econometric Society*, 601–611.

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

Ho, T. and H. R. Stoll (1981): "Optimal dealer pricing under transactions and return uncertainty," *Journal of Financial economics*, 9, 47–73.

Hollifield, B., A. Neklyudov, and C. Spatt (2017): "Bid-ask spreads, trading networks, and the pricing of securitizations," *The Review of Financial Studies*, 30, 3048–3085.

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

Hörner, J. and A. Skrzypacz (2016): "Learning, Experimentation and Information Design," .

Hugonnier, J., B. Lester, and P.-O. Weill (2014): "Heterogeneity in decentralized asset markets," *National Bureau of Economic Research.*

——— (2018): "Frictional intermediation in over-the-counter markets," Tech. rep., National Bureau of Economic Research.

Hugonnier, J., S. Malamud, and E. Trubowitz (2012): "Endogenous completeness of diffusion driven equilibrium markets," *Econometrica*, 80, 1249–1270.

Kihlstrom, R., L. Mirman, and A. Postlewaite (1984): "Experimental Consumption and the 'Rothschild Effect'," *Bayesian Models of Economic Theory.*

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

Lester, B., A. Shourideh, V. Venkateswaran, and A. Zetlin-Jones (2003): "Market-making with search and information frictions," Tech. rep., National Bureau of Economic Research.

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

Menzio, G. (2005): "A cheap-talk theory of random and directed search," *Manscript, Univ. Pennsylvania.*

Mirman, L. J., L. Samuelson, and A. Urbano (1993): "Monopoly experimentation," *International Economic Review*, 549–563.

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

Pakes, A., M. Ostrovsky, and S. T. Berry (2007): "Simple Estimators for the Parameters of Discrete Dynamics Games," *The RAND Journal of Economics*, 38, 373–399.

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

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

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

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

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

Vames, S. (2003): "In Major Shift, NASD Expands Corporate Bond-Price Reports," *Wall Street Journal.*

Weintraub, G. Y., C. L. Benkard, and B. Van Roy (2008): "Markov perfect industry dynamics with many firms," *Econometrica*, 76, 1375–1411.

Wolinsky, A. (1990): "Information revelation in a market with pairwise meetings," *Econometrica*, 1–23.

Online Appendix

A 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.[50]

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.[51,52] 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 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 A.9, confirm that prices for uninsured assets are both more volatile over time and dispersed across trades.

We study the effect of market transparency on trading activity, which we measure as the number of

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

[51] An asset is considered 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.

[52] 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.

Table A.9: 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.

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

$$y_{it} = \phi_i + \kappa_t + \gamma post_t + \lambda post_t \times unins_i + \epsilon_{it}, \tag{28}$$

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 (28), 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 A.10 reports estimates of the parameters in Equation (28) for three different estimation windows covering 2, 4, and 6 months surrounding the policy intervention. Figure A.4 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.

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 in Table A.11. The results 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

Table A.10: Difference-in-difference estimates

	Number of trades with investors		
	2 Months	4 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.

the average level of trade before dissemination.

Table A.11: Difference-in-difference estimates of transparency on inter-dealer trade

	Number of inter-dealer 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.

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 trade price in trade j for asset a in period t, we estimate the following

Figure A.4: Event study for the number of trades with investors

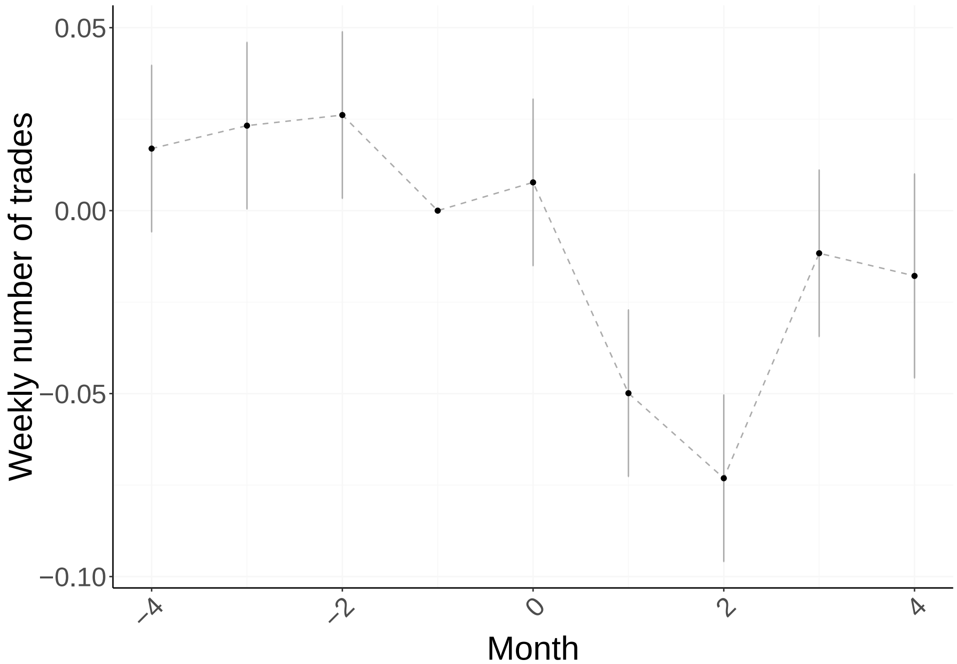

Notes: The figure plots the monthly regression coefficients and 95% confidence intervals from estimating Equation (3) for a four-month window around the introduction of "next day reporting." The outcome variable is the number of trades between dealers and investors. The coefficients are plotted relative 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.

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 A.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 A.12: Effect of market transparency on intermediation spread

	Trade 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.

A.1 Model's predictions and observed outcome

To validate our model, in this section we compare the observed outcome of the market transparency 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 the previous Section. In particular, we estimate Equation (28) separately for assets issued by different states, and we compare the estimated effect of market transparency with the our model predictions.

Table A.13 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. Furthermore, 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.

Table A.13: 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 municipal bonds. The policy experiment is included in our sample period, but not in the sample we use to estimate the model. The table 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.

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 A.14 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 trade 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 demand shock $\theta_{t,s}$, 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.

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

Table A.14: p-values for Hansen-Sargan test

C Derivation of the Equilibrium Objects

C.1 Value functions

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.

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 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 trade 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

$$V_1(\omega, a) = \begin{cases} \mathbb{E}\left[\max\left\{-c + \mathbb{E}\left[-p(\omega, v, a) + V_2(\omega'(v, a), a)\right], V_2(\omega, a)\right\}\right] & \text{if } a = \mathrm{b} \\ \mathbb{E}\left[\max\left\{-c + \mathbb{E}\left[p(\omega, v, a) + V_2(\omega'(v, a), a)\right], V_2(\omega, a)\right\}\right] & \text{if } a = \mathrm{s} \end{cases}, \tag{29}$$

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 $W(\omega)$.

C.2 Trade prices

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 $1 - \exp(-(1 - \rho)\Delta)$ after the dealer rejects an offer and $1 - \exp(-\rho\Delta)$ 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.

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). Menzio (2005) shows that, under the standard restrictions of stationarity of equilibrium strategies and a 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

$$p(\omega, v, u) = \begin{cases} \rho v + (1-\rho)\max_{\omega'} \Delta V_2(\omega', \mathrm{b}) & \text{if } u = u_{\mathrm{b}} \\ (1-\rho)\min_{\omega'} \Delta V_2(\omega', \mathrm{s}) + \rho v & \text{if } u = u_{\mathrm{s}} \end{cases}, \tag{30}$$

where the term $\Delta V_2(\omega', \mathrm{b})$ denotes the dealers' the value of the asset

$$\Delta V_2(\omega, a) = \begin{cases} V_2((\pi, x, e), \mathrm{b}) - \mathbb{E}V_2((\pi, x - u_{\mathrm{b}}, e'), \mathrm{b}) & \text{if } a = \mathrm{b} \\ \mathbb{E}V_2((\pi, x + u_{\mathrm{s}}, e'), \mathrm{s}) - V_2((\pi, x, e), \mathrm{s}) & \text{if } a = \mathrm{s} \end{cases}. \tag{31}$$

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

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 trade price is independent of the dealer's type.

C.3 Dealer's updating

Consider first a dealer with prior π who meets an investor of type u with valuation v. Denote by $\pi_{\text{inv}}(v, u, \pi)$ the dealer's posterior belief after this interaction and $p(v, u)$ the corresponding trade price. As discussed above, the trade price is a strictly increasing function of the investor's valuation. For this reason, we can write the dealer's posterior beliefs $\pi_{\text{inv}}(v, u, \pi)$ as a function of the trade price p and its distribution, which we denote by $F_{\mathrm{p}}\left(p|\theta,a\right)$:

$$\pi_{\mathrm{inv}}\left(v,u,\pi\right)=\frac{dF_{\mathrm{p}}\left(p\left(v,u\right)|\theta,u\right)\pi\left(\theta\right)}{\sum_{\theta'}dF_{\mathrm{p}}\left(p\left(v,u\right)|\theta',u\right)\pi\left(\theta'\right)}. \tag{32}$$

Next, 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 C.1, 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. To write the dealer's posterior belief after an inter-dealer interaction we leverage the assumptions that (i) trade is anonymous beyond experience and (ii) dealers treat the information as independent on what they already know. Indeed, this allows us to write the dealer's posterior beliefs as

$$\pi_{\mathrm{d2d}}\left(\tilde{\pi},\tilde{e},\tilde{q},\pi\right)=\frac{\mathbb{P}\left(\{\text{offer }\tilde{q}\text{ from }\tilde{e}\}\cap\{\text{observes }\tilde{\pi}\text{ from }\tilde{e}\}|\theta\right)\pi\left(\theta\right)}{\sum_{\theta'}\mathbb{P}\left(\{\text{offer }\tilde{q}\text{ from }\tilde{e}\}\cap\{\text{observes }\tilde{\pi}\text{ from }\tilde{e}\}|\theta'\right)\pi\left(\theta'\right)}. \tag{33}$$

To simplify Expression 33, note that the dealer's update will not depend on the action played by his counterparty ((i) above). Indeed, while the offer received by the dealer conveys information about what his counterparty knows about θ_t, the beliefs $\tilde{\pi}$ communicated by the counterparty are a sufficient statistic for this information. In light of this, we can write the dealer's posterior beliefs as

$$\begin{aligned}\pi_{\mathrm{d2d}}\left(\tilde{\pi},\tilde{e},\pi\right)&=\frac{\mathbb{P}\left(\{\text{observes }\tilde{\pi}\text{ from }\tilde{e}\}|\theta\right)\pi\left(\theta\right)}{\sum_{\theta'}\mathbb{P}\left(\{\text{observes }\tilde{\pi}\text{ from }\tilde{e}\}|\theta'\right)\pi\left(\theta'\right)}\\&=\frac{dF^{*}_{\mathrm{inv}}\left(\tilde{\pi}|\tilde{e},\theta\right)\pi\left(\theta\right)}{\sum_{\theta'}dF^{*}_{\mathrm{inv}}\left(\tilde{\pi}|\tilde{e},\theta\right)\pi\left(\theta'\right)},\end{aligned}$$

where F^{*}_{inv} is the equilibrium distribution of dealers' beliefs after trade with the investors. The analysis proceeds similarly for the case in which a dealer's offer is accepted by the buyer.

C.4 Choice probabilities

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

$$\mathbb{P}\left(a|\omega\right)=\frac{\exp\left(\sigma_{\epsilon}^{-1}V_{1}\left(\omega,a\right)\right)}{\exp\left(\sigma_{\epsilon}^{-1}V_{1}\left(\omega,\mathrm{b}\right)\right)+\exp\left(\sigma_{\epsilon}^{-1}V_{1}\left(\omega,\mathrm{s}\right)\right)}. \tag{34}$$

Next, 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 he draws a cost shock c that satisfies

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

Then, we have

$$\begin{aligned}
\mathbb{P}(\text{no trade}|\omega, \text{s}, n) &= \sum_{k \geq 0} \mathbb{P}(\text{can contact } k \text{ investors})\left(1 - F_s(K(\omega, \text{s}))\right)^k \\
&= \sum_{k \geq 0} (1-\gamma)\gamma^k \left(1 - F_s(K(\omega, \text{s}))\right)^k \qquad (35) \\
&= \frac{(1-\gamma)}{1 - \gamma\Big(1 - F_s(K(\omega, \text{s}))\Big)}. \qquad (36)
\end{aligned}$$

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

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

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

D Demand transition: estimation and results

We specify the evolution of the demand shock 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 $\hat{p}_t(u)$ denote the average price at which dealers trade with investors of type $u \in \{u_\text{b}, u_\text{s}\}$ in period t. Given Equation (30), we have:

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

and

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

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

Note that, for large numbers of trades $N_t(u)$, the average price $\hat{p}_t(u)$ converges in probability to its expected value, which we denote by $\mu(\theta_t, u)$. Similarly, for large numbers of trades $N_t(u)$ the distribution of the normalized prices $\sqrt{N_t(u)}(\hat{p}_t(u) - \mu(\theta_t, u))$ approximates that of a normal distribution with mean 0 and finite variance, which we denote by $\sigma^2(\theta_t, u)$. Importantly, the distribution of the normalized prices changes over time only through changes in θ_t.[53] Therefore, we can write the likelihood of observing prices $\{\hat{p}_t(u)\}_{t=0}^{T}$ as

$$\mathcal{L}\left(\{\hat{p}_t(u)\}_{t=0}^{T} | \mathbb{P}_\theta, \mathbb{P}_0, \mu, \sigma, \{N_t(u)\}_{t=0}^{T}\right)$$
$$= \sum_{\theta_0, \cdots, \theta_T} \prod_{t=0}^{T} \phi\left(\sqrt{N_t(u_s)}\frac{\hat{p}_t(u_s) - \mu(u_{\theta t}, u_s)}{\sigma(\theta_t, u_s)}\right) \phi\left(\sqrt{N_t(u_b)}\frac{\hat{p}_t(u_b) - \mu(\theta_t, u_b)}{\sigma(\theta_t, u_b)}\right) \mathbb{P}(\theta_0, \cdots, \theta_T | \mathbb{P}_\theta, \mathbb{P}_0), \quad (37)$$

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 (37), while we use a Viterbi algorithm to recover the maximum likelihood estimates for the sequence of the unobserved demand shocks $\left\{\hat{\theta}_t\right\}_{t=0}^{T}$.

Results. Table A.15 describes several statistics of the distribution of our estimates. Table A.16 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 panel of Table A.15 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

$$\text{Amihud}_j = \frac{1}{n_j}\sum_{t=1}^{n_j}\frac{|r_{jt}|}{Q_{jt}}, \quad (38)$$

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}

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

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 A.15, we find that Amihud's liquidity indeed correlates negatively with the persistence of the demand process.

Table A.15: 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 shock θ_t.

Table A.16 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. [54]

[54] 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).

Table A.16: 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.

E Estimates

Table A.17: Experience estimates

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

Notes: The table shows the estimates of the experience process defined in Equation (17). 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 (18). 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. The standard errors are computed from 100 bootstrap samples with the resampling done at the dealers' level.

Table A.18: 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. 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.

Table A.19: 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. 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.

Table A.20: Investors' cost estimates

	ϕ_{b}	ϕ_{s}	$\mathbb{E}_{\mathrm{b}}(v\|\theta_H)$	$\mathbb{E}_{\mathrm{b}}(v\|\theta_M)$	$\mathbb{E}_{\mathrm{b}}(v\|\theta_L)$	$\mathbb{E}_{\mathrm{s}}(v\|\theta_H)$	$\mathbb{E}_{\mathrm{s}}(v\|\theta_M)$	$\mathbb{E}_{\mathrm{s}}(v\|\theta_L)$
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. 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.

F Model Fit

Figure A.5: 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.

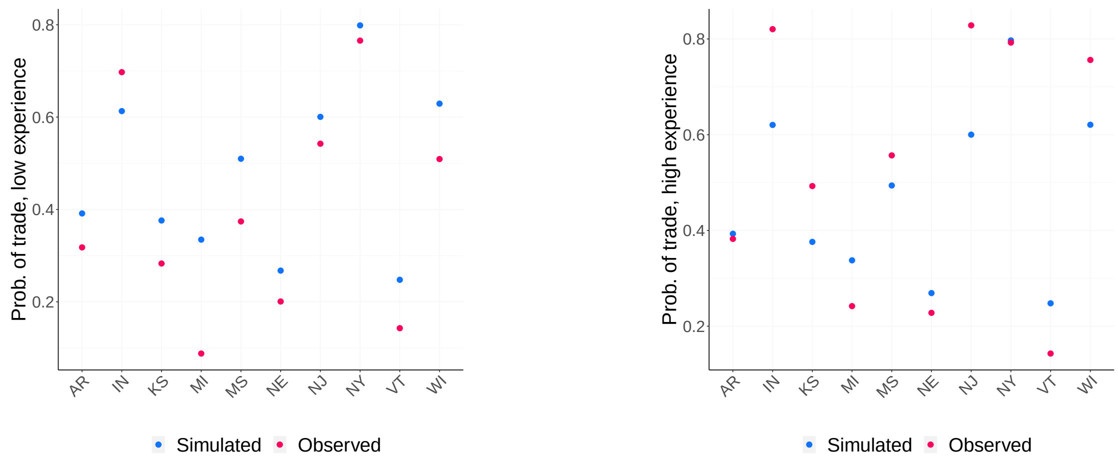

Figure A.6: Model fit. The left panel depicts the observed and predicted probabilities of selling an asset to investors for high experience dealers. The left panel depicts the same objects for low experience dealers.

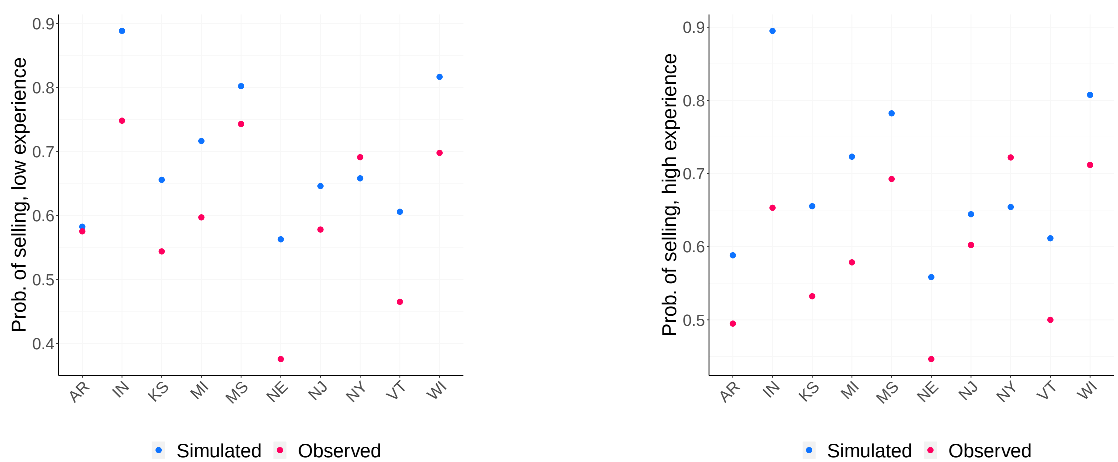

F.1 Stylized facts on networks

Recent empirical studies have documented a number of stylized facts about the intermediation process in decentralized financial markets. For one, these studies highlight that decentralized financial 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

Figure A.7: 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.

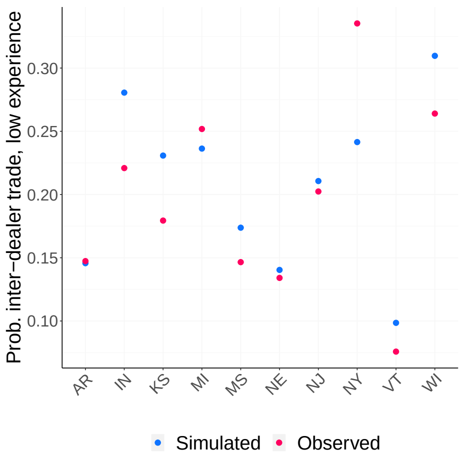

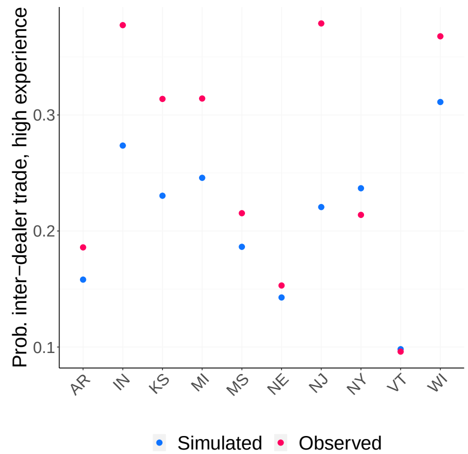

Figure A.8: Core-periphery structure

Notes:

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 A.8 shows the degree distribution across dealers in the simulated market. as a benchmark for comparison Figure A.8 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' trade prices vary systematically with their centrality in the network.[55] Figure A.21 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 A.21 corresponds to a different decile of the dealers' degree

[55]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. (2017)).

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 A.21, where we regress inter-dealer prices on the seller's centrality.

Table A.21: Centrality Premium

	Inter-dealer trade 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	

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