Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.

Dealer Networks

Dan Li and Norman Schurhoff

2014-95

Dealer Networks

Dan Li Norman Schürhoff *

November 10, 2014

*First draft April 2011. Dan Li is with the Board of Governors of the Federal Reserve System, Washington D.C. and Norman Schürhoff is with the University of Lausanne, Swiss Finance Institute, and CEPR (Phone: +41 (21) 692-3447. E-mail: norman.schuerhoff@unil.ch. Postal: Ecole des HEC, Université de Lausanne, Extranef 239, CH-1015 Lausanne, Switzerland). We thank Andrew Ang, Bruno Biais, Jean-Edouard Colliard, Fany Declerck, Darrell Duffie, Amy Edwards, Burton Hollifield, Laurence Lescourret, Craig Lewis, Semyon Malamud, Sophie Moinas, Michael Rockinger, and our discussants Jeff Harris, Larry Harris, Terry Hendershott, Christian Julliard, Bernd Mack, Tarun Ramadorai, Elvira Sojli for extensive comments and illuminating discussions. We are deeply indebted to Rick Green for advice. Seminar audiences at the SEC, Carnegie Mellon University, BI Oslo, University of Bern, London Business School, Aarhus University, Erasmus University, BIS, University of Mannheim, University of Frankfurt, National Bank of Serbia, ESSEC, Toulouse School of Economics, CU Boulder, Federal Reserve Board, CFTC and conference participants at the TI–SoFiE 2012 Conference, 2012 Municipal Finance Conference, ESSFM Gerzensee 2012, 2012 EFA, New York Fed 2012 Workshop on Inter-Bank and Money Markets, Finance Down Under 2013, Four Nations Cup 2013, SFI Research Day 2013, 3rd International Conference on "The Industrial Organization of Securities and Derivatives Markets," 3rd International Moscow Finance conference, Banque de France 2014 Workshop on "OTC Markets: Recent Advances in Research," 2014 LSE Economic Networks Conference, and 2014 WFA have provided valuable comments. The second author gratefully acknowledges research support from the Swiss Finance Institute and the Swiss National Science Foundation.

Abstract

Dealers in over-the-counter securities form networks to mitigate search frictions. The audit trail for municipal bonds shows the dealer network has a core-periphery structure. Central dealers are more efficient at matching buyers and sellers than peripheral dealers, which shortens intermediation chains and speeds up trading. Investors face a tradeoff between execution speed and cost. Central dealers provide immediacy by pre-arranging fewer trades and holding larger inventory. However, trading costs increase strongly with dealer centrality. Investors with strong liquidity need trade with central dealers and at times of market-wide illiquidity. Central dealers thus serve as liquidity providers of last resort.

JEL Classification: G12, G14, G24

Key words: Municipal bonds, over-the-counter financial market, trading cost, liquidity, immediacy, transparency, decentralization, market quality, network analysis

Over-the-counter (OTC) markets—including those for municipal, corporate, agency bonds, and derivatives—are often depicted as "wild west trading" (Alloway and Mackenzie, 2014) due to their decentralized structure. The opacity and fragmentation of OTC markets create search frictions that slow down trading and prohibit any single financial intermediary from matching all buyers and sellers. Recent regulatory initiatives, including Dodd-Frank, the Volcker Rule, and MiFID I/II, and technological innovations such as electronic trading platforms aim at improving the functioning of OTC markets. However, without an understanding of how financial intermediaries facilitate trading and compete for order flow it is impossible to assess the impact of any market intervention.

In this paper, we show broker-dealers form long-lived trading relations with one another to share risks and improve liquidity in OTC markets. We document the structure of the dealer network and its impact on market quality. While some recent theoretical literature touches on these issues (Duffie, Gârleanu, and Pedersen, 2005, 2007; Babus, 2012; Lester, Rocheteau, and Weill, 2014; Malamud and Rostek, 2014), surprisingly little is known empirically.

To conduct our empirical study, we use a proprietary dataset from the regulatory body for the municipal bond market, the Municipal Securities Rulemaking Board (MSRB), that identifies the dealers in each customer-to-dealer, interdealer, and dealer-to-customer trade in municipal bonds between 1998 and 2012. The dealer identifiers allow us to construct the trading network at each point in time, trace how bonds flow through the network, and link the terms of trade to the centrality of the intermediating dealers. We use network analysis to construct alternative measures of dealer centrality based on the past number of trading connections and the order flow between dealers.

The municipal bond market is a natural venue to study search frictions in OTC markets and their effects on price formation and liquidity provision. It is a typical OTC financial market in which adverse selection risk is considered low (75% of bonds are AAA rated). Search frictions are high for investors as the market is fragmented into more than 1.5 million different bonds, each with its idiosyncracies and infrequent trading. Virtually all matching of buyers and sellers occurs through dealer intermediation, with more than 700 broker-dealer firms actively trading in municipals in any given month.[1]

[1]All municipal broker-dealer firms are obliged to register with the Municipal Securities Rulemaking Board

We find the dealer network in municipal bonds has a core-periphery structure with 10 to 30 highly interconnected dealers at the center and several hundred sparsely connected peripheral broker-dealer firms. Dealers' trading relations and centrality rankings are highly persistent, yielding a stable network over time. Bonds are intermediated in chains of up to seven dealers—reflecting significant search and matching frictions. The penultimate dealer in these chains earns the largest markup, which confirms the notion in the industry that buyers for municipal bonds are more difficult to find than sellers (Schultz, 2012)

The high level of heterogeneity in dealers' connectedness reflects differences in search efficiency. Well-connected central dealers have a superior ability to locate counterparties compared to peripheral dealers. Consistent with the hypothesis that higher search efficiency facilitates trade (Babus, 2012), central dealers act as hubs by redistributing order flow from the periphery and selling to customers more often than peripheral dealers, which shortens intermediation chains.

The search heterogeneity also affects the liquidity services that dealers offer their clients and the way intermediaries compete for order flow. Central dealers use their search advantage to provide immediacy to investors, rather than to compete on cheapest execution. While peripheral dealers tend to pre-arrange trades, central dealers do fewer such riskless principal transactions and instead take bonds into inventory and hold them overnight. Central dealers afford to hold larger and more volatile inventories and provide investors with faster execution.

The consequence for execution costs is that, consistent with theories of asset pricing in search markets (Duffie, Gârleanu, and Pedersen, 2005, 2007; Babus, 2012; Condorelli and Galeotti, 2012a; Lester, Rocheteau, and Weill, 2014), markups increase with the matching rate of the dealers intermediating the trades.[2] Central dealers charge sizably larger round-trip spreads, up to twice what peripheral dealers charge. At the same time, central dealers are not obviously more exposed to adverse selection, as they are less likely to lose money on trades and as prices exhibit less reversal at the core than the periphery.

(MSRB). About 2,200 broker-dealer firms are registered with the MSRB.

[2]Under price competition for order flow among dealers, one would expect central dealers that face lower search frictions than their peripheral competitors to provide execution at lower cost. Neklyudov (2013) shows that when investors are homogeneous in their motives for trade and search randomly for best execution, price competition leads to trading costs that decline with dealer centrality. Hollifield, Neklyudov, and Spatt (2013) provide evidence for this negative relation in the institutional ABS/MBS markets.

Our findings show that central and peripheral dealers in municipal bonds compete over different terms of trade by differentiating their services (see Mussa and Rosen (1978) for a seminal model on consumer goods). In an opaque buyer's market like the one for municipal bonds, finding counterparties is a difficult task and, therefore, immediacy is valuable to investors. Demsetz (1968) was first to illustrate that while an investor willing to wait might trade at the price envisioned in the Walrasian setting, an impatient investor can pay a price for immediacy, or speed. This results in deviations from the law-of-one-price. The consequence is systematic price dispersion across dealers, with up to one third of all markup variation specific to the dealer.

The trade-off between execution costs and speed raises the question how costly is immediacy to investors in OTC markets. To back out the terms of trade across dealer tiers, we formalize dealer competition in a parsimonious model of investor-dealer matching. The model estimates confirm that liquidity-constrained investors with high willingness to pay for immediacy, such as municipal bond funds, trade with central dealers to save on execution delay, while price-sensitive investors pre-arrange round-trips with peripheral broker-dealers. The price of immediacy is substantial. Our estimates suggest it costs investors up to 0.6-0.7% of transaction value to speed up execution by a day.

Counterfactual analysis shows investors trade more with central dealers when liquidity is low among all dealers, as occurred during the financial crisis, when dealers exited the market, and when mutual funds faced large outflows. The introduction of post-trade transparency in 2005 coincided with a decline in the price of immediacy across all dealers, but it impacted central and peripheral dealers differently. Trading costs fell and liquidity rose across the market, but the price of immediacy dropped only at central and mid-tier dealers. Transparency thus benefitted investors by eroding the markups of central dealers.

The paper is related to several strands of literature, most notably to the growing field of studies on OTC markets. Existing theories, such as Duffie, Gârleanu, and Pedersen (2005, 2007), Lagos and Rocheteau (2007, 2009), Feldhütter (2012), Zhu (2012), Neklyudov (2013) analyze the impact of search frictions on asset prices. One of the main predictions from these models is that trading costs depend on dealers' bargaining power. Prior empirical studies have focused on heterogeneity across investors in their bargaining power, proxied by trade size (Green, Hollifield, and Schürhoff, 2007) or price rounding (Li, 2007). Little is known empirically, however, about how bargaining power differs among financial intermediaries and how it relates to search efficiency.[3] We document significant heterogeneity among dealers in their search efficiency and terms of trade. More generally, our findings highlight the systemic role of central dealers as liquidity providers of last resort. Regulations that aim at improving market stability by limiting the influence of key players in financial markets may come at the cost of increased bid-ask spreads and reduced immediacy, especially at times when market liquidity is scarce.

The paper also relates to the growing literature on networks in financial markets. Upper and Worms (2004), Cocco, Gomes, and Martins (2009), Bech and Atalay (2010), Afonso and Lagos (2012), Cohen-Cole, Kirilenko, and Patacchini (2012), and Craig and von Peter (2014) empirically study the network topology in interbank money markets. This paper also examines network topology, albeit in OTC bond markets that are dominated by both retail and institutional investors. Moreover, our focus is on the relation between dealers' network centrality and execution quality. Leitner (2005), Gale and Kariv (2007), Condorelli and Galeotti (2012b), and Babus (2012) develop models of OTC network formation. While our study does not directly address the origin of networks, we explore the benefits and costs associated with relations in OTC markets and thus motives for forming networks. The literature on price formation in network settings is mostly theoretical. Kondor and Babus (2013) and Malamud and Rostek (2014) study trading and price formation in OTC networks, and Gofman (2011), Colliard and Demange (2014), and Glode and Opp (2014) develop theories of intermediation chains. We are the first to empirically study individual dealer behavior in OTC bond markets.

The remainder of the paper is organized as follows. Section 1 describes the data. Section 2 documents the dealer network and Section 3 the search advantage of central dealers over their peripheral competitors. Sections 4 through 6 explore the consequences for transaction costs

[3] Bessembinder, Maxwell, and Venkataraman (2006), Edwards, Harris, and Piwowar (2007), Green, Hollifield, and Schürhoff (2007), and Green, Li, and Schürhoff (2010) study transactions cost and price discovery in OTC markets. Menkveld and Jovanovic (2010), Cespa and Foucault (2011), Lagos, Rocheteau, and Weill (2011), and Julliard, Denbee, Li, and Yuan (2013) study liquidity provision and liquidity spillovers across securities. Biais (1993) and Hendershott and Madhavan (2014) consider the trade-off between bilateral and electronic trading.

and liquidity. Section 7 explores robustness, and Section 8 concludes.

1 Data and Intermediation Chains

The municipal bond market is the main source of financing for state and local governments in the United States, with around \$4 trillion in debt outstanding and \$300-\$400 billion annual gross issuance in recent years. Municipal bonds trade in decentralized OTC networks of more than 2,000 broker-dealer firms that intermediate between buyers and sellers and actively trade among themselves. The market is fragmented and opaque, as there are about 55,000 issuers and 1.5 million different issues outstanding, each with its own idiosyncracies and limited pre- and post-trade transparency. Municipal bonds, typically held in buy-and-hold investment portfolios, are attractive investments for both institutional and retail investors, as their interests are tax exempt and default rates have historically been low. Municipal bond dealers thus serve the liquidity needs of investors with vastly different market access, trade motives, and liquidity preferences.

1.1 Data

Our main data source is the proprietary Transaction Reporting System audit trail from the MSRB. In an effort to improve market transparency, the MSRB has required dealers to report all transactions in municipal securities since 1998. Unlike the publicly available version of historical municipal bond transactions, our data provide identifiers for the dealer firms intermediating each trade.[4] For customer trades, the data identify the dealer buying and the dealer selling the bond. For interdealer trades, the data identify the dealers on each side of the trade.[5] The transactions data cover the 15-year period from February 1998 to December 2012.

In addition to the comprehensive transactions data, we obtained reference information

[4]The MSRB designation that we use aggregates subsidiaries of the same financial institution, with separate executing broker symbols (EBS), into a single MSRB identifier (see `http://www.msrb.org/msrb1/trsweb/dealers.asp`). For instance, Morgan Stanley Smith Barney LLC has one MSRB ID for the three EBS IDs DEAN, MSPW, and MSSB.

[5]The data do not provide identifiers for the dealers' customers. See Hendershott and Madhavan (2014) and Afonso, Kovner, and Schoar (2014) for recent studies employing customer identifiers.

on all municipal bonds, including issuance date, maturity, coupon, taxable status, ratings, call features, issue size, and issuer characteristics from Mergent.[6] Weekly municipal mutual fund flows are obtained from the Investment Company Institute.

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 (time-stamped after February 1998, four-letter alphabetic dealer identifier, no MSRB indicator for away-from-market price, par value $\geq$\$5K) and bond-specific filters (fixed or zero coupon, non-derivative, non-warrant, not puttable, maturity $\geq$1 year, \$5K denomination). Green, Hollifield, and Schürhoff (2007) and Schultz (2012) document that trading and liquidity in newly issued bonds are markedly different from seasoned issues. Since the focus of this paper is on the secondary market, we remove all trades during the first 90 days after issuance and less than one year away from maturity. Finally, we eliminate price outliers by truncating the distribution at 0.5% and 99.5%, separately for zero-coupon and other bonds. Appendix A provides a detailed description of the number of affected observations in each step.

Our final sample consists of 72.2 million transactions in 1.28 million bonds with different CUSIPs. The trades are intermediated by a total of 2,257 broker-dealer firms (as identified by their MSRB designation).[7] Order splitting is common. On average, each customer-to-dealer trade involves about one interdealer trade and two dealer-to-customer trades. A total of 17.9 million trades are customer(sell)-to-dealer(buy) trades, 20.2 million are between dealers, and 34.1 million are dealer(sell)-to-customer(buy) trades.

1.2 Round-trip chains and markups

Municipal bonds typically trade infrequently, and when they trade, they do so in round-trips. In a round-trip transaction, an investor sells bonds to a dealer (*CD* leg), and then the dealer sells the same bonds to another investor (*DC* leg) or other dealers (*DD* leg(s)). The dealer may sell all of the bond at once (we call this a *Nonsplit* trade) or split into smaller lots

[6] We have cross-checked with the Securities Data Company (SDC) Global Public Finance database and, while coverage is not the same, the results do not change materially if we switch to SDC reference data.

[7] Out of the 2,257 MSRB-registered broker-dealer firms with trading records in our sample, 63 do not engage in interdealer trading and 117 are interdealer brokers.

(a *Split* trade). The offsetting trades can occur through the same dealer, yielding a *CDC* round-trip, or through different dealers, after a sequence of interdealer transactions, yielding a *C(N)DC* round-trip.

Our round-trip matching algorithm, described in Internet Appendix Appendix IA.I, extends Green, Hollifield, and Schürhoff (2007). We look for matching legs of *CD*, *DD*, and consecutive *DC* trades (ordered by their time stamps) in the same CUSIP with matching dealer IDs and same par size.[8] Since trading in most issues occurs infrequently, the algorithm matches 82% of all customer-to-dealer trades to corresponding dealer-to-dealer and dealer-to-customer trades.[9] The remaining trades are part of incomplete *C(N)D* chains in which we cannot unambiguously assign each leg to a round-trip, which occurs in very liquid issues with frequent trading and in very illiquid issues when the bond remains in the dealers' inventory for more than a month. Our results are unaffected when we use alternative matching schemes.

We measure trading costs by the (gross and net) markups charged by the dealers on round-trips. The markups help compensate dealers for their search cost, inventory carrying cost, and to a lesser extent in this market, adverse selection cost. The dealers' total markup on round-trip i (with size $Par_i = \sum_j Par_{ij}^{DC}$) is defined as the par-weighted selling prices P_i^{DC} in the *DC* transactions, net of the purchase price P_i^{CD} from the customer in the initial *CD* transaction, all normalized by the original purchase price:[10]

$$\text{Markup } M_i = \frac{\frac{1}{Par_i}\sum_j Par_{ij}^{DC} P_{ij}^{DC} - P_i^{CD}}{P_i^{CD}} * 100. \tag{1}$$

When multiple dealers intermediate *C(N)DC* round-trips in chains, we compute interdealer markups using the dealer-to-dealer transaction prices P_i^{DD}. This tells us how the dealers split the total markup.

[8] For *Split* trades, we only allow order splitting in the last leg of the chain, the *DC* leg, as otherwise the matching becomes ambiguous.

[9] Prior studies do not identify dealers and thus focus only on the two legs of trade that are between dealers and customers (*CD* and *DC*). Our round-trip concept explicitly follows bonds through dealer network, which improves matching accuracy and allows us to observe interdealer markups and directly measure the complexity of intermediation.

[10] Net markup is gross markup minus the return on the Bond Buyer index. The net differs from the gross markup by less than 3 bps on average with a correlation of 0.9. They are close since inventory times are short and interest rates do not move much from one day to the next. We report results for gross markups.

1.3 Network construction and centrality measures

In constructing the trading network, we focus on trades between dealers in the recent past, usually within the past 30 days. The direction of the link follows the bond from the selling dealer to the buying dealer. The strength of a trading relation is the number of times they have traded or, alternatively, the number of bonds exchanged.

The dealers' connectedness to other dealers can be measured by network statistics that capture the local connectivity or global importance of a dealer. Five measures of centrality are widely used in network analysis: degree centrality, eigenvector centrality, betweenness, closeness, and cliquishness. Appendix A provides a detailed description of each centrality measure and discusses how we construct weighted and unweighted variants.

Degree centrality measures the local connectivity of a dealer in the network by computing the fraction of dealers in the network with which the dealer trades directly. Eigenvector centrality measures global importance by assigning scores to all dealers in the network based on the principle that connections to high-scoring dealers contribute more to the score of the dealer than equal connections to low-scoring dealers. Eigenvector centrality takes all direct and indirect trading partners into account. Betweenness is similar in that it counts the number of shortest existing trading chains linking any two dealers in the network that pass through the dealer. Closeness is the inverse of the average number of steps that a dealer needs to take within the network to reach or be reached by any other dealer, based on prior trading relations. Cliquishness is the likelihood that two associates of a dealer are themselves associated. The correlation between degree and cliquishness captures the hierarchical structure of the network, which we discuss in detail in Section 2.

We aggregate the network variables to a single centrality index, denoted Net_{jt}, for each dealer $j = 0, ..., J$ and time t, by taking the first principle component across all centrality measures, as described in Appendix A.[11] For robustness, we construct equal- and value-weighted centrality measures, where we weight each connection either by the number of trades or the value of the order flow between the dealers. Our results are robust to the exact

[11]We compute an aggregate index as our focus is on characterizing dealer heterogeneity instead of exploiting differences between centrality measures. Our results are robust to taking the individual statistics or the aggregate statistic Net, since they are highly correlated. Cliquishness cc is not included in the aggregation, as it requires at least two neighbors.

definition of dealer centrality and using the aggregate index or the disaggregated measures.

Our interdealer network is recalculated daily on a rolling basis. For every day during the sample period, we use all interdealer transactions during the past 30 calendar days. That is, Net_{jt} is computed over a 30-day rolling window and lagged one day. Following Milbourn (2003), we apply an empirical cumulative distribution function (ecdf) transformation to each network variable in order to normalize the centrality between 0 and 1 while keeping the original ordering.[12] The ecdf transformation diminishes the impact of skewness and outliers, and it facilitates interpretation of economic magnitudes. A unit change in centrality corresponds to moving from the most peripheral (Net=0) to the most central dealer (Net=1).

Table 1, Panel A provides summary statistics for the dealers' network characteristics. We track all 2,257 dealer firms over 3,757 trading days, yielding 2.87 million dealer-day observations (cc requires at least two neighbors, which reduces the number of observations). Panel B of Table 1 reports the correlation coefficients between each network characteristic and Net for the equal- and value-weighted variants. There is a sizable common component that is captured by Net. The correlations between the network characteristics and Net are less than one but positive (which is, as we will see, due to the core-peripheral structure).[13]

[Tables 1 about here]

2 Dealers' Trading Network

Since little is known about the characteristics of trading networks in OTC markets, we document topological features of the dealer network in municipal bonds in this section. We show that the number of dealer connections are heavily skewed with a fat right tail populated by several core dealers, that the overall market has a hierarchical core-peripheral structure, and that the trading network is resilient to dealer exits. We also document the persistence in dealers' trading relations.

[12] As an alternative, we use the share of interdealer trading volume as weight in the ecdf transformation, so that in this case a centrality equal to Net means that dealers with 1-Net market share are more central than the dealer. The results are not materially affected.

[13] Cliquishness cc is the exception in that it correlates negatively with most other measures. The negative relationship between cliquishness and degree is related to the hierarchical structure of the network, which is explored further in Section 2.

2.1 Network structure

2.1.1 Core-periphery structure

Figure 1 visualizes the network structure of interdealer trading in municipal bonds. In Panel A, two dealers are connected if they traded more than 10,000 times over the sample period. This threshold corresponds to about 2-3 trades per day and allows us to focus on the most central dealers that form the core of the municipal bond market.[14] In Panel B, we plot the dealer network using all transactions.

There is tremendous heterogeneity in dealers' connectedness. Some dealers have many trading partners and interact frequently, while many others have just a few counterparties. A small group of about 10-30 dealers are highly interconnected and trade heavily with many other dealers. In contrast to these core dealers, the remaining several hundred dealer firms are more peripheral and interact with a limited number of trading partners. The overall network is very sparse. Only 124,000 out of 5.1 million possible directed links, or 2.4%, exist over the 15-year period. The two plots combined indicate that the municipal trading network has a core-periphery structure.[15]

[Figure 1 about here]

To characterize the network topology more formally, we compare in Figure 2 the degree distribution across dealers to that of a random trading network. When all dealers trade anonymously and randomly through a centralized exchange, trading relations yield a Poisson distribution over degrees, which serves as a benchmark for comparison. We expect the degree distribution in decentralized markets to be very different from that of random trading, as dealers prefer to trade with their associates and some dealers are blacklisted by others for reputational reasons.[16]

Indeed, Panel A shows that there is a higher level of heterogeneity among dealers in terms of connectedness than suggested by random trading. Consistent with Figure 1, a large number of weakly connected dealers compete with a few core dealers. The dealership

[14] The plot is an artificial map generated by multidimensional scaling based on the criterion that the more trade links exist between two dealers, the closer is their location.

[15] In much of the theoretical work on network formation, star-mesh or hub-and-spoke arise as the Pareto efficient structure. See Jackson (2005) and Babus (2012).

[16] We thank our discussant Bernd Mack from Deutsche Börse for pointing this out.

network exhibits a heavy right tail, with 5% of dealers having more than 200 other dealers as trading partners. The out-degree is more fat-tailed than the in-degree.[17] This asymmetry suggests central dealers act more as redistributors than solely as intermediaries.

2.1.2 Hierarchical structure

The right plot in Panel A explores a more subtle aspect of the dealer network by documenting the local connectedness of dealers. The different tax treatment at the state level of in- and out-of-state bonds suggests that the municipal bond market may have not just a single nationwide center but multiple local sub-markets, mostly along state lines. One can infer the hierarchical structure of the market by plotting the degree distribution across dealers in the network (horizontal axis) against the clustering coefficient, or cliquishness, of each dealer (vertical axis). The negative relation between k and cc in the plot, again on a log-log scale, reveals that sparsely connected dealers are part of local markets and with order flow between the different local markets being maintained by a few central hubs.[18] We explore the impact of dealers' state-specific centrality further in Tables IA.3 and IA.4 in the Internet Appendix.

[Figures 2 about here]

2.1.3 Market resilience

We explore in Panel B of Figure 2 the resilience of the network structure to shocks that remove some dealers and their connections from the market, keeping the remainder fixed. Following Albert, Jeong, and Barabási (2000), we plot the size of the largest connected subgraph (so-called giant component), in percent of the remaining network size, as a function of the percent of dealers eliminated from the network.[19] We consider two scenarios, one in

[17]The municipal bond market thus exhibits features of an asymmetric scale-free network. Scale-free networks are characterized by a power-law degree distribution. See Erdős and Rényi (1960) and Barabási and Albert (1999).

[18]Hierarchical modularity yields scaling of the clustering coefficient, which follows $cc(k) \sim 1/k$ and, hence, traces out a straight line of slope -1 on a log-log plot. See Ravasz, Somera, Mongru, Oltvai, and Barabási (2002). Craig and von Peter (2014) document tiering in the German interbank market.

[19]Elimination in this context means all order flow to and from the dealer is canceled while the network is otherwise held fixed. The statistic should be interpreted as a descriptive feature of the network topology rather than an economic response of the dealers or the market as a whole. Chain reactions and other endogenous responses to dealer exits are ignored by this measure.

which dealers exit at random and one in which the most connected dealers are eliminated first.

The dealer network is highly resilient to random exits of dealers, as suggested by the slow decay in the size of the giant component as more and more dealers are randomly removed (circles). The intuition is that a random exit removes a "typical" dealer. In a core-peripheral network, a typical dealer is likely one of those peripheral dealers that are only sparsely connected, making the impact of dealer exit small.

The network also appears robust to targeted removal of the most connected dealers. For the network to disintegrate, one has to sequentially eliminate the top 20% of all dealers (triangles).[20] The average path length in this case peaks at eight dealers before the network disintegrates (right plot). This is because, at the core of the network, there are a number of dealers that are fully connected to the rest of the core, so that they can act as substitutes for one another. Such redundancy allows the network to sustain even the selected removal of the most connected dealers.

2.2 Stability in dealer relations and ranks

The nature of financial intermediation—standing ready to provide liquidity—forces dealers to interact repeatedly. These interactions may lead to long-term relations that benefit the dealers themselves and other market participants. Alternatively, relations may be formed and broken opportunistically when, for instance, a dealer's inventory matches the needs of another dealer and gains from trade exist.

Table 2, Panel A compares the trading relationship between any given pair of dealers from one month to the next. It shows that if two dealers did not trade last month, there is about an 85% chance that they again do not trade with each other this month. Conditional on a directional trade relation existing in the prior month, trades in the same direction occur with 62% probability in the next month. Ignoring trade directions, the probability that two dealers who traded last month also trade this month is 65%. To put these numbers

[20] In terms of robustness to targeted shocks, the dealer network compares favorably to some known large scale empirical networks such as the World-Wide Web, which has a threshold of 6.7% for disintegration (see Albert, Jeong, and Barabási (2000)).

in perspective, the probability would be 1.4% in a random network with the same number of dealers and interdealer trades. These results suggest high persistence in the trading relations between dealers and in the direction of order flow from month to month. The persistence in dealers' order flow is not surprising in an opaque market, in which suitable counterparties are difficult to identify.

Panel B of Table 2 focuses on the persistence in dealer ranks, as measured by the ordering of the dealers' centrality Net. Dealer ranks are also highly persistent, at least from one month to the next. The top 10 dealers remain at the top of their league 93% of the times. This is to be expected in an environment in which it is valuable, yet costly, to become and remain interconnected.

[Table 2 about here]

3 Dealers' Network Centrality and Search Efficiency

Given the evidence in Section 2 that dealers form long-term trading relations, centrality in the trading network should be related to the efficiency with which dealers search and match counterparties. This section documents how the propensity of dealers to sell bonds to investors, rather than to pass bonds on to other dealers, depends on the centrality and other characteristics of the dealers. We also show how dealer characteristics affect the complexity and length of intermediation chains (Colliard and Demange, 2014; Glode and Opp, 2014).

3.1 Intermediation chains and order flow routing

Bond dealers intermediate OTC trading by matching buyers and sellers in round-trip trades. Our sample comprises 11.4 million complete round-trip chains (and 0.4 million incomplete chains).[21] Dealer involvement varies across these chains. The simplest and most common way in which a dealer intermediates round-trips is by purchasing bonds from a customer and then selling the same bond to another customer, where the original bond lot is not split into smaller order sizes. There are 6.3 million such *CDC-Nonsplit* transactions, where *CDC*

[21] Throughout our analysis, we eliminate trades between customers and dealers in which a dealer acts in the capacity of agent, as opposed to principal (see Appendix A). The reason is that dealers acting as agents are compensated through commission, not markup. Agency trades account for 6% of the sample.

indicates the bond went from *C*ustomer to *D*ealer and then to another *C*ustomer in two sequential trades. Alternatively, the dealer purchasing the bond can split the lot and sell each piece to a different customer. We call the 2.5 million such transactions *CDC-Splits*, as there is still only one dealer involved. In total, there are 8.8 million *CDC* round-trips.

Alternatively, 23% of the time, that is, in 2.6 million round-trips, dealers involve other dealers in the intermediation. We find a maximum of 7 dealers in the sequence of trades (there are very few cases involving 8 or more dealers). Such longer chains start with a dealer purchasing a bond from a customer, followed by one or several interdealer trades that move the bond from the head dealer to the tail dealer, and end with the tail dealer selling the bond to one or more customers. We call this type of round-trip *C(N)DC*, where *(N)* indicates that multiple dealers may be involved.[22]

Table 3 reports the average network centrality of each dealer in the intermediation chain for *C(N)DC* trades. $Net = 0$ (1) corresponds to the least (most) central dealer. As one may expect, interdealer trading occurs systematically through the assistance of central dealers. Bonds flow from periphery to center and partially back. Dealers in the middle of the chains are more central than either the dealer purchasing the bond from customer or the dealer ultimately selling the bond to customer. Dealer centrality peaks with the second and the penultimate dealers for all types of *C(N)DC* trades. Central dealers thus act as hubs by redistributing the order flow (see also Section 2.1).

[Table 3 about here]

3.2 Complexity of intermediation chains

To more directly assess whether search and matching efficiency give rise to the importance of central dealers as hubs, we check how the complexity of intermediation, measured by the number of dealers in a chain, depends on the network centrality of the dealers in the trade sequence. Centrality should be a significant determinant for intermediation chain length so

[22] We explore these three classifications separately in our empirical analysis. The baseline sample consists of all *CDC-Nonsplit* round-trips. For this sample, we are the most certain that the bid- and ask-side trades pair up exactly. The second sample allows order splitting, thus includes all *CDC* round-trips. This sample is more representative of a typical trade but may add some noise when split orders are wrongly assigned to be part of the same round-trip. The last sample are all the *C(N)DC* trades, which include all matched transactions chains involving more than one dealer.

long as central dealers are better able to locate investors. We explore this issue in two ways.

We first examine the tendency of dealers to sell directly to investors instead of passing bonds on to other dealers, and how this propensity depends on dealer centrality. We estimate a panel probit specification for dealer-to-customer (DC Trade$_i = 1$) versus interdealer (DC Trade$_i = 0$) trades, allowing for state and month fixed effects and adjusting standard errors for heteroskedasticity and clustering by issuer and time:

$$Pr(DC \text{ Trade}_i \mid \text{Trade}_i) = \Phi(\alpha + \delta Net_i + X_i'\beta + \varepsilon_i), \tag{2}$$

where Φ is the cumulative distribution of the standard normal distribution, Net_i is the centrality of the dealer in each bilateral trade i (here i denotes each trade in a round-trip chain), and X_i are the explanatory variables compiled in Appendix B.

The explanatory variables include information specific to the trade, the bond, the issuer and the dealer. The trade characteristics are the par value of the trade interacted with trade size group dummies. Bond and issuer characteristics include credit quality, remaining maturity in years, age of the bond issue, issue size, call feature, general obligation bonds, sinking fund, bank qualification, tax status, and whether the bond is subject to AMT.[23] Dealer characteristics include a dummy for being the underwriter of the bond, a dummy for being a primary dealer, a dummy for having headquarter in New York City, and the total asset base of the dealer.

Table 4, columns (1) and (2) report the determinants of how the dealers route the order flow. Central dealers are significantly more likely than peripheral dealers to find an investor to buy the bond. The baseline probability of finding a buyer for an average bond is 41%. The coefficient 0.58 on Net corresponds to a marginal effect of 51%, raising this probability to about 60% for the most central dealers. The bond's underwriter, primary dealer, smaller dealers, non-NYC-headquartered dealers, and dealers with large abnormal inventory are also more likely to sell to investors.

[Table 4 about here]

We now directly explore how the length of intermediation chains varies with dealer cen-

[23] In alternate specifications without calendar time controls, we have included the VIX index, LIBOR rate, and aggregate new issuance volume during the prior week. The results are very similar and are omitted.

trality by estimating the following equation for the number N_i of dealers in round-trip i:

$$\text{Chain length } N_i = \alpha + \delta Net_i + X_i'\beta + \varepsilon_i, \tag{3}$$

where Net_i is the centrality of the head dealer and X_i are the explanatory variables compiled in Appendix B. We again allow for state and month fixed effects and adjust standard errors for heteroskedasticity by clustering at both the issuer-level and in time. Columns (3) and (4) in Table 4 use panel OLS regressions, and columns (5) and (6) report estimates from panel Poisson regressions.

The complexity of intermediation is negatively related to the centrality of the head dealer in all specifications. Intermediation chains are 0.9-1.2 dealers (columns 3 and 4), or 36-46% (columns 5 and 6) shorter at the center than the periphery, while the average chain length is about 1.5 with a minimum of 1 and a maximum of 7. These estimates are consistent with central dealers being more efficient at searching and matching investors than peripheral dealers. Central dealers can reach ultimate buyers more directly, which results in shorter and less complex intermediation chains. Similar effects hold for underwriters, primary dealers, and non-NYC-headquartered dealers.

4 Transaction Costs and Dealer Centrality

The evidence in the previous sections that dealers form networks to facilitate the matching of bond buyers and sellers, with central dealers acting as hubs, suggests that transaction costs should depend on the complexity of the trade and the identity of the dealer intermediating the trade. This section relates order execution costs to trade and dealer characteristics.

We measure trading costs by the total markup on a round-trip transaction charged by all the dealers involved. We also check how the dealers in the chain split the total markup among each other. The dealers' total markup M, defined in expression (1), is computed as the difference between the par-weighted average price at which they sold the bonds to customers and the price at which they purchased the bonds, as described in Section 1.[24]

[24] As interest rates do not move much from one day to the next, compared to municipal trading costs, we ignore interest rate movements in the computation of dealer markups. Results are unchanged when we

Markups are expressed in percent of the head dealer's original purchase price.

Table 5 reports descriptive statistics on trading costs.[25] Trading costs on municipal bonds are large on average, consistent with Harris and Piwowar (2006) and Green, Hollifield, and Schürhoff (2007). Average round-trip trading costs on non-splits are 1.8%, while dealers earn 2.1% on split round-trips. Average markups decline monotonically with transaction size, with \$25K trades costing 1.8% on average, whereas \$1M trades cost 24 basis points on average. Consistent with the evidence in Schultz (2012) on newly issued bonds, total costs are larger the more dealers are involved in the round-trip. The last split in Panel A shows that average markups increase monotonically but at declining rate with the number of dealers intermediating the chain, from 1.9% on average when one dealer is involved to 3.7% with seven dealers involved. The markups dealers charge other dealers are thus substantially smaller than the markups they charge investors.

[Table 5 about here]

4.1 Dealer markups and centrality

Dealer markups vary substantially across round-trips, even if the number of dealers involved and the size of the transaction are fixed. Are markups systematically related to the centrality of the dealer intermediating the trade? We first confirm this relationship using double-sorted univariate comparisons, and then show that it continues to hold in a multivariate regression setting, where a host of explanatory variables are controlled for.

Table 5, Panel B employs double-sorting across dealer tiers and trade-size groups for different round-trip samples. We construct dealer tiers by sorting the dealers each day on the centrality measure *Net*. We assign the dealers with highest centrality and cumulative market share of 25% (top quartile) to the central dealer tier (or, about 0.7% of all broker-dealer firms), the dealers in the next quartile to the mid-tier, and the remainder to peripheral dealers.[26] We map trade size into ranges that roughly proxy for investor type, with par sizes

compute markups net of interest rate changes.

[25] In reporting these numbers, we apply no data filters. For the regression analysis performed later, we eliminate extreme values by truncating the distribution at 0.5% and 99.5%.

[26] The variable *Net* for sorting the dealers' centrality is computed over the prior 30 days and lagged by an extra day to not overlap with the time of trade. The tier to which a dealer is assigned may thus change.

less than \$100K typically originated from retail investors, between \$100K and \$1M from municipal mutual funds, and \$1M plus blocks likely from taxable institutions.

We find that central dealers charge larger trading costs than peripheral dealers across all trade sizes, despite their higher matching efficiency. Average markups differ by up to 100%. For medium-sized trades of \$250K, markups at central dealers are 1.0%, double the 0.5% at peripheral dealers. For small and large trades, the difference is smaller but still positive and both economically and statistically significant.

Not only do average markups differ across dealers, the composition of trades also varies. Figure 3 plots the markup distribution for peripheral, mid-tier, and central dealers. While peripheral dealers intermediate many round-trips at positive but close to zero markups, central dealers rarely do such trades. They mostly trade at markups away from zero.[27]

[Figure 3 about here]

To fully account for other characteristics of the round-trips that may vary systematically across trade-size groups and dealer tiers, we run multivariate regressions that explore how trading cost M_i in round-trip i depends on dealer centrality and other explanatory variables:

$$\text{Markup } M_i = \alpha + \delta Net_i + X_i'\beta + \varepsilon_i, \tag{4}$$

where Net_i is the centrality of the first dealer in the chain and X_i are the control variables described in Appendix B that include information specific to the trade, the bond, the issuer, and the dealer. Markups that dealers charge may include compensation for additional services that broker-dealers provide, which could include pricing, underwriting, and research services. In addition, theory predicts that inventory management should affect bid-ask spreads. To capture these effects, we include several dealer characteristics: a dummy variable for whether the dealer is the lead underwriter of the bond, indicators for primary dealers and NYC headquarters, respectively, dealer size, and the dealer's abnormal aggregate inventory. We also allow for state and month fixed effects and adjust standard errors for heteroskedasticity and clustering by issuer and time.

The key variable *Net* is lagged to not overlap with trades in the round-trip. The coefficient

[27] The spikes at round markups in Panel A reflect dealers' propensity to trade at round prices (Li, 2012).

δ measures the difference in trading costs between the most peripheral (Net=0) and the most central dealer (Net=1). Across columns, we vary the centrality measure and, respectively, the trade categories. In columns (1), (3), and (5) the dealer centrality measure *Net* is defined as the first principle component of the equal-weighted centrality proxies. In the remaining columns, we use the order flow-weighted centrality proxies. The regression samples are *CDC-Nonsplit*, *CDC*, and *C(N)DC*.

The estimates in Table 6 reveal that, consistent with Figure 3, dealer characteristics—their centrality, in particular—are important sources of variation for transaction costs. Holding trade size and bond characteristics constant, trading costs are positively related to dealer centrality. Highly interconnected dealers are able to charge customers 0.4-0.7% larger markups, up to twice as large, on average, than peripheral dealers.

We also find that being lead underwriter on the bond issue has little impact on trading costs for investors. As one would expect, primary and larger dealers charge lower transaction costs. However their impact is small. There is also evidence for inventory effects. Dealers charge less the larger their abnormal inventory. This is consistent with Huang and Stoll (1997), who show that dealers adjust their bid and ask to revert inventory toward a target. Inventory effects alone, however, do not explain the positive sign on centrality, as central dealers hold larger inventories than peripheral dealers. The last two columns in Table 6 show, consistent with Table 5, that transaction costs are larger the more dealers are involved in the chain, but that this effect is smaller the more central is the head dealer.

The takeaway from this analysis is that the identity of the broker-dealer intermediating bond trades is an important determinant for transactions costs. Section 7.2 provides checks for how robust the positive relation between markups and dealer centrality is to alternative specifications. Why would investors choose to trade with the more expensive central dealers? What is the source of bargaining power that central dealers have against investors? We explore the differences in the types of liquidity services that central and peripheral dealers provide in Section 5. Before doing so, we explore the source of the trade surplus.

[Table 6 about here]

4.2 How do dealers split markups?

A natural follow-up question to the systematic price dispersion across dealers is how much each dealer earns in the chain of intermediaries. This reveals how dealers split the total surplus from financial intermediation. Presumably, the dealer earning the largest share adds the most value (Lester, Rocheteau, and Weill, 2014).

Table 7 reports average markups per dealer on round-trip transactions with varying degree of dealer involvement. Total dealer markups are broken down by the number of dealers (across rows) and by each dealer (across columns) in the sequence of dealers intermediating the chain, from dealer one through seven. We find that the dealer closest to the ultimate buyer earns the largest share of the overall profits, irrespective of how many dealers are involved. This pattern reflects the common perception in the industry that bonds are not bought, they are sold. The municipal bond market is known to be a buyer's market (Schultz, 2012), and identifying investors willing to buy is considered to be dealers' most crucial and costly task. If it is the case that central dealers are better able to identify investors willing to buy bonds and, potentially, other dealers who do know those investors (see Section 3.2), then the markup split explains some of the markup differences across dealers documented in Section 4.1. By being better at matching buyers and sellers, central dealers can internalize the extra costs of intermediating long chains.

[Table 7 about here]

5 Liquidity and Dealer Centrality

Central dealers are more efficient at locating counterparties and tend to act as hubs in round-trip bond trades, as documented in the previous sections. Central dealers, however, do not charge investors smaller bid-ask spreads. How are they able to charge larger markups than peripheral dealers? In this section, we explore the possibility that central dealers use their search advantage to provide more liquidity, in terms of immediacy or in other ways, to investors.

Bond dealers provide liquidity in various ways, by pre-arranging trades between customers (similar to a limit order) or taking bonds into inventory (as in a market order). Highly interconnected dealers, due to their central network position, are able to spread inventory risk across their affiliates better than peripheral dealers. Central dealers can therefore provide immediacy to investors and afford greater inventory risk. We thus expect central dealers to pre-arrange fewer trades than peripheral ones and to hold larger and more volatile inventory.

5.1 Pre-arranged trades

Municipal bond dealers typically intermediate round-trip trades in one of two ways. Upon receiving a sell order, they face the choice between taking the bond into inventory or asking the seller to wait until a matching buyer is found. The former choice provides immediacy—that is, faster execution to the seller—but it entails inventory risks to the dealer. The latter choice, typically called pre-arranged or "riskless principal" transaction, allows the dealer to be purely a match-maker who commits no capital. The seller, however, bears the cost of execution delay.

Table 8 and Table IA.1 in the Internet Appendix relate the propensity of pre-arranging trades to the centrality of the intermediating dealer. While the data does not identify pre-arranged trades directly, we can infer them using the time stamps associated with the legs of a round-trip. We define pre-arranged round-trips as those whose buy and sell legs have the same time stamps, dealer identifiers, and par size in the same CUSIP (immediate matches). As an alternative, we require the same par size buy and sell trades through the same dealer(s) to occur on the same calendar day (same-day matches). Dealers face a capital charge on bonds they take into overnight inventory. Same-day matches, while riskier for the dealer than immediate matches, are cheaper to intermediate than overnight round-trips. We estimate the model, assuming the errors are normally distributed (probit), allowing for state and month fixed effects and adjusting standard errors for heteroskedasticity and clustering by issuer and time:

$$Pr(\text{Riskless principal Trade}_i \mid \text{Trade}_i) = \Phi(\alpha + \delta Net_i + X_i'\beta + \varepsilon_i). \quad (5)$$

We find in Table 8 (immediate matches) and Table IA.1 (same-day matches) that central dealers are significantly less likely to pre-arrange trades than peripheral dealers, suggesting that central dealers provide more immediacy to investors than their peripheral competitors. Trading with a central (peripheral) dealer is, in this regard, like submitting a market (limit) order. The behavior is similar for dealers with large inventory. By contrast, all else equal, dealers with larger balance sheets or dealers who are headquartered in NYC are more likely to pre-arrange trades than other dealers. The results are mixed for primary dealers. The effect of centrality on the choice of trade type is thus distinct from a pure size effect.

[Table 8 about here]

5.2 Inventory risk taking

Figure 4 documents the distribution in inventory durations across trade size and dealer centrality. We measure a bond's inventory duration as the number of days, hours, and minutes it takes a dealer to find a buyer. When bond lots are split, a bond's inventory duration is the average time, weighted by the size of the split orders, that it takes to resell the entire bond lot. We plot histograms for retail (<\$100K), mutual fund (\$100K-\$1M), and institutional (≥\$1M) sizes across the three panels, and for three dealer tiers across columns. We again assign the dealers with highest centrality and cumulative market share of 25% (top quartile) to the central dealer tier, the dealers in the next quartile to the mid-tier, and the remainder to peripheral dealers. Consistent with Tables 8 and IA.1 and across all trade sizes, central dealers have longer overall inventory durations than peripheral dealers since they pre-arrange fewer trades, which shows up in Figure 4 as a spike at zero inventory time.

[Figure 4 about here]

Table 9 explores the link between inventory durations and dealer centrality more systematically. We estimate the following equation in a panel Tobit model where the lower limit is set at zero, since inventory time cannot be negative.

$$\text{Inventory time } \Delta T_i = \alpha + \delta Net_i + X_i'\beta + \varepsilon_i. \tag{6}$$

Coefficients δ and β capture the sensitivity of inventory duration to dealer centrality and characteristics of the bond and trade. We find that dealers, on average, hold bonds in their inventory for 3.3 days, and the median inventory holding time is 0.92 days. Central dealers hold bonds in their inventory longer. The difference in average inventory time between the most central and the most peripheral dealers is more than a day, ranging from 1.1 to 2.0 days across specifications.

[Table 9 about here]

Not only are central dealers more willing to take bonds into their inventory, and hence provide more immediacy per round-trip, they also supply the lion's share of the liquidity in the market. Figure 5, Panel A illustrates how aggregate order flow varies across dealers. We measure dealers' order flow by the total number of trades that a dealer conducts (left chart) and, respectively, aggregate dollar volume (right chart). Not surprisingly, central dealers trade more often with investors and in larger volumes than peripheral dealers. The top 0.7%, 1.6%, and 5.4% of dealers conduct, 25%, 50%, and 75% of all transactions, respectively.

Dealers also differ markedly in their inventory risk taking behavior at the aggregate level. Figure 5, Panel B provides statistics for the variability in dealer inventory. We use two measures for inventory changes, the absolute value of daily dollar changes in inventory (left chart) and proportional changes (right chart). We define proportional changes as the daily change normalized by the average inventory over the past 30 days, $|\Delta inv_t/\frac{1}{30}\sum_{i=1}^{30} inv_{t-i}|$. We approximate dealer inventories by the cumulative sum of all purchases net of all sales that occur at least 90 days after the original sale date from the underwriter. As one may expect, central dealers have more variable inventory than peripheral dealers, both in absolute and relative terms. Table IA.2 in the Internet Appendix confirms this positive relation between dealer centrality and inventory risk taking in multivariate regressions.

[Figure 5 about here]

6 Liquidity Preferences and Dealer Competition

The prior sections have documented systematic differences in the terms of trade across dealers. The evidence suggests that investors choose bond dealers based on a tradeoff between transaction costs and execution speed. One alternative explanation for why central dealers charge larger bid-ask spreads than peripheral dealers, yet handle most of the order flow, is that investors are captive or unaware of the tradeoffs. Another possibility is that investors search for cheapest execution based solely on price, but central dealers handle bonds that are more costly to intermediate. To disentangle these explanations, we build an investor-dealer matching model that endogenizes investors' choices of broker-dealers. The model allows us to estimate the price of immediacy (Demsetz, 1968) and the component in the bid-ask spread that is due to dealers' liquidity provision.

6.1 Model for investor-dealer matching

Suppose dealers $j = 0, ..., J$ compete for order flow by differentiating the terms of trade, liquidity L_j and trading cost M_j, that they offer so that, for simplicity, L_j and M_j increase with j. Like in vertical differentiation models, all investors prefer more liquidity (i.e., services of higher quality) but differ in their willingness to pay.[28] Assume the gains from trade for an investor are $v_i = u(M, L; \theta_i) = \theta_i L - M$, where the parameter θ_i captures the investor's liquidity preference. The special case $\theta_i = 0$ corresponds to price competition for order flow.

The threshold at which investors are indifferent between dealer tiers $j-1$ and j (sorted according to their terms of trade) is the willingness to pay Θ_j for which $u(M_{j-1}, L_{j-1}; \Theta_j) = u(M_j, L_j; \Theta_j) = \Pi_j$. Once θ_i or, equivalently, v_i crosses this threshold, the investor directs the order to the more central dealer offering faster execution. Directed search models naturally lead to such assortative matching equilibria (Lester, Rocheteau, and Weill, 2014). Let D_i be the index of the transacting dealer. Investors and dealers then match as follows:

$$\text{Investor } i \text{ trades with dealer } D_i = j \iff \Pi_j \leq v_i < \Pi_{j+1}. \tag{7}$$

To operationalize the model, we assume that the trading cost $M_{ij} = X_i\beta_j + \varepsilon_{ij}$ charged in transaction i by dealer j depends on the characteristics of the bond and dealer and on market conditions, summarized by the set of explanatory variables X described in Appendix B. The liquidity $L_{ij} = X_i\gamma_j + \eta_{ij}$ offered by dealer j depends similarly on the X.[29] As in Sections 4

[28] Hotelling (1929) and Mussa and Rosen (1978) provide early differentiated goods models. The following derivations are general in that L may be a vector of liquidity features. In the empirical implementation, we focus on immediacy, which we capture by relative execution speed measured as the difference in inventory times across dealers.

[29] The data we observe are the terms of trade of the dealers chosen by the investors, $M_i = \sum_{j=1}^{J} M_{ij}(D_i = j)$ and $L_i = \sum_{j=1}^{J} L_{ij}(D_i = j)$. We do not observe, by contrast, the counterfactual terms of trade offered by

and 5, the residuals ε and η capture unobservable characteristics that affect the costs of the trade and the liquidity provided by the dealers, such as inventory effects.

Under standard distributional assumptions, the empirical specification we obtain is a multivariate extension of endogenous switching models (Maddala, 1983). The unobserved determinants e_i of investors' utility, $v_i = X_i\delta + e_i$, are naturally correlated with the unobservable characteristics of dealers' terms of trade (ε, η). Assuming joint normality,

$$\begin{aligned} \lambda_{ij} &= E[e_i | X_i, \text{ investor } i \text{ transacts with dealer } D_i = j] \\ &= -\frac{\phi(\Pi_{j+1} - X_i\delta) - \phi(\Pi_j - X_i\delta)}{\Phi(\Pi_{j+1} - X_i\delta) - \Phi(\Pi_j - X_i\delta)}, \end{aligned} \tag{8}$$

where Φ and ϕ denote the normal distribution and density function, respectively. Taking investors' dealer choices into account, the markup expected at dealer j is given by

$$E[M_{ij} | D_i = j] = X_i\beta_j + \rho_j^\varepsilon \sigma_j^\varepsilon \lambda_{ij}, \tag{9}$$

where ρ_j^ε is the correlation between e_i and ε_{ij} and σ_j^ε is the standard deviation of ε_{ij}.[30] The first term in (9) is the quoted markup at the dealer, and the second term captures the endogenous selection of investors to trade with the dealer depending on their willingness to pay and the markup charged by the dealer.[31] The selection term is absent if investors are captive or unaware of the markup differences across dealers.

The liquidity at the transacting dealer j can similarly be decomposed as

$$E[L_{ij} | D_i = j] = \underbrace{X_i\gamma_j}_{\text{Liquidity supply}} + \underbrace{\rho_j^\eta \sigma_j^\eta \lambda_{ij}}_{\text{Liquidity preferences}}, \tag{10}$$

where ρ_j^η is the correlation between e_i and η_{ij} and σ_j^η is the standard deviation of η_{ij}. The first term in (10) are the determinants of dealers' liquidity offering, and the second term (the difference between expected and observed liquidity) tells us in which bonds and under what market conditions investors trade with different types of dealers. The selection term is

all other dealers that were inferior from investors' perspective.

[30] In the bivariate case with two types of dealers, $\lambda_{i1} = \phi(-X_i\delta)/(1 - \Phi(-X_i\delta))$ is the inverse Mills ratio. As a by-product we obtain the market share of dealer j as $\Pr(D_i = j) = \Phi(\Pi_{j+1} - X_i\delta) - \Phi(\Pi_j - X_i\delta)$.

[31] The markup expected in the data can be written $E[M_i] = \sum_{j=0}^{J} E[M_{ij} | D_i = j] \Pr(D_i = j)$.

absent if investors select dealers purely based on lowest execution cost.

6.2 Who trades with central dealers?

The empirical implementation requires making several choices. First, we need to specify dealer tiers. It seems unlikely that investors have or would even want to have access to all 2,200 broker-dealers at the same time. Investors can, however, establish accounts with several dealers to be able to pick the best terms of trade. Dealers can be sorted by their centrality into $J+1$ tiers $j = 0, 1, ..., J$. For simplicity, we pick $J + 1 = 3$ (our results are robust to larger J values). Tier 2 includes the most central and well connected core dealers in the market, while Tier 0 includes the peripheral dealers with low search efficiency. The remainder are mid-tier dealers. As before, we assign the top quartile of the most connected dealers (with cumulative market share of 25%) to the central tier, the next quartile to the middle tier, and the remainder to peripheral dealers.[32] We compute centrality rolling over 30 days and lagged by an extra day to not overlap with the time period of trade. The tier of a dealer can thus vary over time.

Second, the terms of trade are captured by a bundle of trading cost and liquidity, (M, L). Trading costs are the round-trip markups.[33] Our liquidity proxy is the time the bond is at risk in the dealer's inventory, $L_i = \Delta T_i$ (Section 5.2). This proxy corresponds to the time savings for an investor selling a bond to a dealer instead of waiting for the ultimate buyer.[34]

To achieve better identification, we construct instrumental variables that affect investors' choice of dealer independently from M and L. One would expect investors to exhibit a need for immediacy when dealers previously exited the market, when mutual funds are faced with large outflows, and when market accessibility is low. We therefore construct alternate instruments, defined as either the total market share (in terms of number of trades or dollar volume) of dealers that disappeared over the prior six months, or the aggregate number by each state.[35] In addition, we exploit municipal mutual fund in- and outflows during the prior

[32]Our results do not depend on these exact thresholds.

[33]In the *C(N)DC* sample, we compute execution costs and inventory times for the head dealer.

[34]Alternatively, we have used $L_i = -exp(-\Delta T_i)$, measuring investors' time value of money, and obtained similar results. For an investor receiving \$1 at time T, the utility gain from immediate execution is $1 - e^{-r\Delta T} = 1 + L^r$, where r is the rate of time preference.

[35]There are 177 such control transactions during the sample period. Major transactions, largely sparked

week (obtained from Investment Company Institute) and calendar time effects as proposed by Hendershott and Madhavan (2014). Most central dealers are located on the East Coast. Liquidity supply is therefore restricted in after-hours trading (which can be captured by hour-of-day effects). Day-of-week and month-end effects provide additional variation.

The first column in Table 10 reports the determinants for investors' dealer choices from an ordered probit regression for (7). As indicated by the positive, flat, and negative coefficients on small, medium, and large trade sizes, respectively, central dealers more often intermediate medium-sized trades that predominantly originate from municipal mutual funds that, at times, face strong liquidity needs. Investors trade with central dealers when liquidity is low, as it occurs when dealers exit the market (in the bond's state of issuance and other states), and when mutual funds face large outflows. Investors also tend to choose central dealers for more seasoned, higher rated, investment grade, general obligation, callable, non-bank qualified bonds. It is thus not the case that central dealers systematically intermediate bonds that are a priori more costly to intermediate, such as junk bonds. These same characteristics determine trading costs and liquidity offerings. We now investigate this relation by estimating conditions (9) and (10) using OLS, with the regressors X_i augmented by λ_{ij}.

[Table 10 about here]

6.3 Liquidity supply and market conditions

The remaining columns in Table 10 report the estimates for the determinants of markups and liquidity, controlling for the endogenous trade choices. The coefficients are generally consistent with Tables 6 and 9. Trade characteristics, including order size, bond and issuer characteristics affect markups and liquidity differently across dealer tiers. Investors can therefore expect the terms offered by different dealers to vary across bonds and market conditions. The coefficients on the selection terms λ, defined in (8), are economically and statistically significant. This suggests investors are aware, at least broadly, about differences in the terms of trade, and they trade off trading costs with immediacy when selecting their broker-dealer (instead of searching for best execution based solely on price).

by the financial crisis, include Bear Stearns (bought by JP Morgan), Lehman (by Barclays), Merrill Lynch (by BoA), Washington Mutual (by JPM-Chase), Wachovia (by Wells Fargo), and National City (by PNC).

Table 11 reports the expected markup and liquidity at different dealer tiers, computed as the average predicted values of $X\widehat{\beta}_D$ and $X\widehat{\gamma}_D$, $D \in \{P, M, C\}$, and compares the markup and liquidity at transacting dealers to the counterfactual (quoted) terms of trade offered at non-transacting dealers. Expected markup and liquidity are computed based on the estimates for β and γ in Table 10 and reported in Panel A. The differences with investors' actual choices, $E[M_D|D]$–$X\widehat{\beta}_D$ and $E[L_D|D]$–$X\widehat{\gamma}_D$, reflect the endogenous selection to trade based on price and liquidity, respectively. We can predict the unobserved (counterfactual) markups and liquidity for any of the non-chosen dealers $j \neq D$, assuming all investors were to switch to dealer j and taking into account the dealer that was actually chosen. Expected markup and liquidity at dealer j when dealer D_i was chosen in transaction i are given by

$$E[M_{ij}|D_i] = X_i\beta_j + \rho_j^{\varepsilon}\sigma_j^{\varepsilon}\lambda_{iD_i}, \quad (11)$$

$$E[L_{ij}|D_i] = X_i\gamma_j + \rho_j^{\eta}\sigma_j^{\eta}\lambda_{iD_i}, \quad (12)$$

with λ_{iD_i} calculated as in (8), using the actual D_i.

Panel A in Table 11 shows that expected markups, analogous to quoted bid-ask spreads, and expected execution speeds vary systematically with dealer centrality. Central dealers quote larger markups than peripheral dealers (more than 3% at top tier dealers compared with 1.2% at the periphery, on average). They also provide consistently more speed (4.6 days at top-tier compared to 1.6 at bottom-tier dealers), corresponding to about 3 day faster execution at core dealers. Investors do not automatically trade with the cheapest dealer. Investors choose to trade with central dealers (who are the most expensive to trade with) when liquidity is below average. This can be seen in the comparison between expected and observed liquidity.

Counterfactual experiment 1, reported in Panel B of Table 11, corroborates this notion. It shows that investors direct trades to central dealers when liquidity is scarce. Central dealers provide little liquidity, but other dealers provide even less liquidity in these situations. The liquidity improvement offered by central dealers over peripheral ones is 2.2 (0.7) days compared to peripheral (mid-tier) dealers, which is significant at less than 1% and corresponds to tripling the execution speed. By contrast, investors choose peripheral counterparties (Counterfactual experiment 3) in situations when liquidity is abundant. This occurs when peripheral dealers offer a minimum level of liquidity, at the lowest price. On average, across all trades, the observed target speed of execution is around 2-3 days. Investors trade with peripheral dealers when they are capable of offering investors' required liquidity. Similarly, mid-tier dealers are investors' preferred choice when central dealers offer more liquidity but peripheral dealers offer too little liquidity.

[Table 11 about here]

These outcomes can be related to hedonic pricing. With directed search, the market functions as if there is a price, in terms of trading costs, for every level of liquidity. We can get an estimate from the model for the average price of liquidity across the dealers in the market. The price of immediacy, or liquidity, corresponds to how much investors have to pay for each day of faster execution:

$$p_j = \frac{E[X_i\beta_j]}{E[X_i\gamma_j]}. \tag{13}$$

Table 11, Panel A reports the price per unit of liquidity for the three dealer tiers. The average cost is between 0.6% and 0.7% per day. These execution costs are punitive to investors in need of immediacy, given that municipal bonds carry smaller monthly returns.

6.4 Liquidity provision in the financial crisis

We repeat the analysis in the previous section for different time periods. This allows us to check whether the markup-speed tradeoff is robust over time, and it illustrates the effect of transparency and the financial crisis on liquidity provision in the municipal bond market. We split the sample into three periods that we title dark, transparent, and crisis. The time periods are chosen to reflect the introduction of real-time reporting in February 2005 and the beginning of the financial crisis in September 2008.

Table 12 documents expected markups and liquidity across the three time periods. The liquidity supply and price of liquidity vary strongly by time period. Markups dropped across all dealers with the introduction of post-trade transparency (from 1.3-3.8% to 1.2-2.7%). At the same time, liquidity improved, predominantly at central and mid-tier dealers. As a result, the price of liquidity dropped by more than half at core dealers. Transparency thus seems to have benefitted investors by eroding the inventory holding costs and markups of central dealers.

The situation reversed during the financial crisis, when markups rose and liquidity dropped across all dealers. The resulting rise in the price of liquidity is most pronounced at peripheral dealers, who lost significant market shares during the financial crisis. Overall, the results documented in the previous sections are robust to different cutoffs for the three time periods and varying market conditions.

[Table 12 about here]

7 Robustness

In this section, we perform additional tests to check alternative explanations for the larger markups observed at central dealers.

7.1 Price risks in intermediation

Are the larger markups at central dealers compensation for taking on more price risks? This can happen if central dealers intermediate riskier bonds with less certain demand and higher price volatility. We have partly addressed this concern in our investor-dealer matching model, where we have shown that investors do not choose central dealers for bonds that are particularly hard to trade a priori. In this subsection, we explicitly test this hypothesis by looking at the probability of dealers taking a trading loss on a round-trip.

Table 5, Panel B provides univariate results for the most common trade sizes, and Table 13 documents the determinants of trading losses using a panel probit model with state and month fixed effects and adjusting standard errors for heteroskedasticity and clustering by issuer and time. In each of the three round-trip samples, dealers lose money in less than 2% of all trades. Still, the loss probability depends strongly on the dealers' relative position in the network, but with the opposite sign predicted by the risk-return tradeoff. Central dealers are significantly less likely than peripheral dealers to lose on round-trips. Thus, the profits of more connected dealers are larger, on average, and less risky. Therefore the larger markups that central dealer charge are not obviously to compensate for price risks.

[Table 13 about here]

7.2 Unobserved dealer characteristics

Despite our efforts to control for a list of observable dealer characteristics, there may still be unobservable characteristics that potentially drive the cross-sectional results on dealer markups and inventory duration. To sharpen the identification, we modify our baseline specification along two dimensions. First, we allow for dealer-fixed effects that absorb all dealer characteristics, so long as they are time invariant. Second, to capture the hierarchical structure documented in Section 2.1.2, we construct the trading network for each state separately and use state-specific centrality measures in the regression model. If dealer search efficiency matters, then state-specific centrality should matter beyond aggregate centrality.

Tables IA.3 and IA.4 in the Internet Appendix summarize the results from various specifications that aim at identification and robustness. Column (1) provides the baseline specification for the *C(N)DC* sample. In column (2), we use the individual centrality measure *ev*, which is the dealers' eigenvector centrality. Column (3) replaces the centrality measure by the dealers' state-specific centrality State-Net. In column (4), we add dealer fixed effects to the specification, so that the coefficient on Net captures the impact of time-series variation in dealer centrality. In column (5), we include Net, State-Net, and dealer fixed effects. In column (6), we perform an instrumental variables regression with Net instrumented by the same liquidity condition variables used in Section 6 and listed in Appendix B.

Regardless of whether aggregate dealer centrality or state-specific centrality is used, dealer centrality is positively and significantly related to markups (in Table IA.3), even if dealer-fixed effects are included. When the two centrality measures are included, both are significant, suggesting that dealers' search efficiency both at the aggregate market level and the state level matter. These results suggest that the dealers' network position itself, instead of unobservable dealer characteristics, drives the dispersion in transaction costs. The results for inventory duration in Table IA.4 suggest that state-specific dealer centrality is even more important than aggregate centrality in explaining cross-sectional differences in inventory duration (Column (5)).

7.3 Price efficiency across the dealer network

While Table 13 reveals that central dealers incur smaller trading losses *ex post*, it could be that they widen spreads to mitigate adverse selection risk from investors with superior information. If so, the price impact of a trade should be stronger at the center than the periphery of the OTC market. Alternatively, highly interconnected dealers, due to their central network position, may be better at filtering liquidity-based motives for trade and aggregate fundamental information than dealers with little exposure to aggregate order flow. In this case, one would expect central dealers' trades to have less price impact, and bonds with more central dealers trading on them to be more price efficient (Kondor and Babus, 2013).

To determine price impact, we adopt the Hasbrouck (1993) and Hotchkiss and Ronen (2002) setting. In Hasbrouck's (1993) local trend model, price movements are decomposed into permanent and transitory components and market quality MQ is captured by the first-order lag in autocorrelation. Market quality is highest when prices are martingales so that price changes are uncorrelated. $MQ = 1$ corresponds to a situation in which all price movements are permanent. The market is then considered perfectly informationally efficient. $MQ = 0$ corresponds to a situation in which all price movements are transitory and eventually reversed.

The last two columns of Table IA.2 in the Internet Appendix document the link between the informational efficiency of bond prices and dealer centrality. Information inefficiency is measured by MQ based on trade-by-trade price changes. MQ is calculated separately for each bond and then aggregated at the dealer level. The estimates are obtained from OLS regressions with year fixed effects. Standard errors are adjusted for heteroskedasticity and double clustering by dealer and year. We find consistently across specifications that prices are more efficient at central dealers. Central dealers thus seem to be better informed than peripheral dealers.

8 Conclusion

Demsetz (1968) showed that the structure of the financial markets is an important determinant of trading costs, liquidity, and price discovery. Many financial securities, including municipal bonds, are traded through decentralized and opaque networks of financial intermediaries. We show in a comprehensive sample of municipal bond trades that the dealership network diminishes search frictions and facilitates trading, but that it also alters dealer competition for order flow in several important ways.

The dealership network in municipal bonds exhibits a stable core-periphery structure with 10 to 30 highly interconnected core dealers and several hundred peripheral broker-dealer firms. Central dealers have a superior ability to locate counterparties and are therefore more willing to supply liquidity. They provide immediacy to investors by taking inventory risk when their peripheral competitors prefer to pre-arrange trades. As a result, dealer centrality is negatively related to the complexity of intermediation chains and the likelihood of pre-arranging trades, but positively related with dealer inventories and the markups they charge investors. Consistent with asset pricing theories in search markets (Duffie, Gârleanu, and Pedersen, 2005, 2007; Lester, Rocheteau, and Weill, 2014, and others), municipal trading costs reflect the dealers' search efficiency. The introduction of post-trade transparency has benefitted investors by eroding the markups at central dealers.

Our findings highlight the systemic role of central dealers as liquidity providers of last resort. Investors with a need for fast execution, such as municipal bond funds, tend to trade more with central dealers and at times of market-wide illiquidity. Regulations that aim at improving market stability by limiting the influence of key players in financial markets may come at the cost of increased bid-ask spreads and reduced immediacy, especially at times when market liquidity is scarce. Our results, more generally, shed light on the trade-offs investors face when trading in over-the-counter markets, which may guide financial market design.

References

Afonso, Gara, Anna Kovner, and Antoinette Schoar, 2014, Trading Partners in the Interbank Lending Market, Discussion paper, Federal Reserve Bank of New York.

Afonso, Gara, and Ricardo Lagos, 2012, Trade dynamics in the market for federal funds, Discussion Paper, 549.

Albert, R., H. Jeong, and A.-L. Barabási, 2000, Error and attack tolerance of complex networks, *Nature* 406, 378–82.

Alloway, Tracy, and Michael Mackenzie, 2014, US corporate bond traders go electronic, *Financial Times,* *http://www.ft.com/intl/cms/s/0/4c39439c-425e-11e4-9818-00144feabdc0.html.*

Babus, Ana, 2012, Endogenous intermediation in over-the-counter markets, Discussion paper, Federal Reserve Bank of Chicago.

Barabási, Albert-László, and Réka Albert, 1999, Emergence of scaling in random networks, *science* 286, 509–12.

Bech, Morten L, and Enghin Atalay, 2010, The topology of the federal funds market, *Physica A: Statistical Mechanics and its Applications* 389, 5223–46.

Bessembinder, Hendrik, William Maxwell, and Kumar Venkataraman, 2006, Market transparency, liquidity externalities, and institutional trading costs in corporate bonds, *Journal of Financial Economics* 82, 251–88.

Biais, Bruno, 1993, Price formation and equilibrium liquidity in fragmented and centralized markets, *Journal of Finance* 48, 157–85.

Cespa, Giovanni, and Thierry Foucault, 2011, Learning from Prices, Liquidity Spillovers, and Market Segmentation, Discussion paper, Centre for Economic Policy Research.

Cocco, Joao F, Francisco J Gomes, and Nuno C Martins, 2009, Lending relationships in the interbank market, *Journal of Financial Intermediation* 18, 24–48.

Cohen-Cole, Ethan, Andrei A. Kirilenko, and Eleonora Patacchini, 2012, How Your Counterparty Matters: Using Transaction Networks to Explain Returns in CCP Marketplaces, Discussion paper, University of Maryland.

Colliard, Jean-Edouard, and Cash Providers Demange, Gabrielle, 2014, Asset Dissemination Over Intermediation Chains, Discussion paper, Paris School of Economics.

Condorelli, Daniele, and Andrea Galeotti, 2012a, Bilateral Trading in Networks, Discussion paper, University of Essex.

——— , 2012b, Endogenous Trading Networks, Discussion paper, University of Essex.

Craig, Ben, and Goetz von Peter, 2014, Interbank tiering and money center banks, *Journal of Financial Intermediation* pp. –.

Demsetz, Harold, 1968, The cost of transacting, *Quarterly Journal of Economics* 82, 33–53.

Duffie, Darrell, Nicolae Gârleanu, and Lasse Heje Pedersen, 2005, Over-the-counter markets, *Econometrica* 73, 1815–47.

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

Edwards, Amy K., Lawrence E. Harris, and Michael S. Piwowar, 2007, Corporate bond market transaction costs and transparency, *Journal of Finance* 62, 1421–51.

Erdős, Paul, and Alfréd Rényi, 1960, On the evolution of random graphs, *Magyar Tud. Akad. Mat. Kutató Int. Közl* 5, 17–61.

Feldhütter, Peter, 2012, The Same Bond at Different Prices: Identifying Search Frictions and Selling Pressures, *Review of Financial Studies* 25, 1155–1206.

Gale, Douglas M, and Shachar Kariv, 2007, Financial networks, *American Economic Review* 97, 99–103.

Glode, Vincent, and Christian Opp, 2014, Adverse selection and intermediation chains, Discussion paper, University of Pennsylvania - The Wharton School.

Gofman, Michael, 2011, A network-based analysis of over-the-counter markets, Discussion paper, University of Wisconsin - Madison.

Green, Richard C., Burton Hollifield, and Norman Schürhoff, 2007, Financial intermediation and the costs of trading in an opaque market, *Review of Financial Studies* 20, 275–314.

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

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

Hasbrouck, Joel, 1993, Assessing the quality of a security market: A new approach to transaction- cost measurement, *Review of Financial Studies* 6, 191–212.

Hendershott, Terrence, and Ananth Madhavan, 2014, Click or call? auction versus search in the over-the-counter market, *Journal of Finance* p. forthcoming.

Hollifield, Burton, Artem Neklyudov, and Chester Spatt, 2013, Bid-ask spreads and the pricing of securitizations: 144a vs. registered securitizations, Discussion paper, Carnegie Mellon University.

Hotchkiss, Edith S., and Tavy Ronen, 2002, The informational efficiency of the corporate bond market: An intraday analysis, *Review of Financial Studies* 15, 1325–54.

Hotelling, Harold, 1929, Stability in competition, *Economic Journal* 39, 41–57.

Huang, Roger D, and Hans R Stoll, 1997, The components of the bid-ask spread: A general approach, *Review of Financial Studies* 10, 995–1034.

Jackson, Matthew O., 2005, A Survey of Models of Network Formation: Stability and Efficiency, Discussion Paper, 0303011 Princeton University Press.

Julliard, Christian, Edward Denbee, Ye Li, and Kathy Yuan, 2013, Network risk and key players: A structural analysis of interbank liquidity, Discussion paper, London School of Economics.

Kondor, Peter, and Ana Babus, 2013, Trading and Information Diffusion in Over-the-Counter Markets, Discussion Paper, 792 Society for Economic Dynamics 2013 Meeting.

Lagos, Ricardo, and Guillaume Rocheteau, 2007, Search in asset markets: Market structure, liquidity, and welfare, *American Economic Review* 97, 198–202.

———, 2009, Liquidity in asset markets with search frictions, *Econometrica* 77, 403–26.

———, and Pierre-Olivier Weill, 2011, Crises and liquidity in over-the-counter markets, *Journal of Economic Theory* 146, 2169–205.

Leitner, Yaron, 2005, Financial Networks: Contagion, Commitment, and Private Sector Bailouts, *Journal of Finance* 60, 2925–53.

Lester, Benjamin R., Guillaume Rocheteau, and Pierre-Olivier Weill, 2014, Competing for order flow in otc markets, Discussion paper, Federal Reserve Bank of Philadelphia.

Li, Dan, 2007, Price clustering and dealer's market power in otc markets, Discussion paper, Carnegie Mellon University.

Maddala, Gangadharrao S, 1983, *Limited-Dependent and Qualitative Variables in Economics* (Cambridge University Press: New York).

Malamud, Semyon, and Marzena Rostek, 2014, Decentralized exchange, Discussion paper, Swiss Finance Institute.

Menkveld, Albert J., and Boyan Jovanovic, 2010, Middlemen in Limit Order Markets, Discussion Paper, 955 Society for Economic Dynamics 2010 Meeting.

Milbourn, Todd T., 2003, CEO reputation and stock-based compensation, *Journal of Financial Economics* 68, 233–62.

Mussa, Michael, and Sherwin Rosen, 1978, Monopoly and product quality, *Journal of Economic Theory* 18, 301–17.

Neklyudov, Artem, 2013, Bid-ask spreads and the over-the-counter interdealer markets: Core and peripheral dealers, Discussion paper, Carnegie Mellon University.

Ravasz, Erzsébet, Anna Lisa Somera, Dale A Mongru, Zoltán N Oltvai, and A-L Barabási, 2002, Hierarchical organization of modularity in metabolic networks, *Science* 297, 1551–55.

Schultz, Paul, 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.

Upper, Christian, and Andreas Worms, 2004, Estimating bilateral exposures in the German interbank market: Is there a danger of contagion?, *European Economic Review* 48, 827–49.

Zhu, Haoxiang, 2012, Finding a good price in opaque over-the-counter markets, *Review of Financial Studies* 25, 1255–85.

A Measures of Dealer Centrality

Several measures of centrality are widely used in network analysis. They describe either the local connectivity of a dealer or its global importance. The following provides a brief description of each measure and its variants, and it explains how we aggregate them to a single index.

- Degree *dg*: A measure of the local connectivity of a dealer. The degree of a dealer is computed as the sum of all direct relations that a dealer has with other dealers in the network, divided by the total number of dealers in the network. Using information on the direction of order flow, one can calculate in-degree *dgin* and, respectively, out-degree *dgout*. Using information on the volume of order flow, one can calculate several weighted variants:
  - *dgoutwntrade* = Out-degree, weighted by number of trades.
  - *dginwntrade* = In-degree, weighted by number of trades.
  - *dgoutwpar* = Out-degree, weighted by total par amount.
  - *dginwpar* = In-degree, weighted by total par amount.
- Related to the degree is the K-core *kcore*. The K-core is defined as the maximal sub-network in which each dealer has at least degree k. For directed graphs, one differentiates between *kcoreout* and *kcorein*, the largest k-cores the dealer belongs to, counting only out-links or in-links.
- Eigenvector centrality *ev*: A measure of the overall importance of a dealer firm in the network. It assigns relative scores to all dealers in the network based on the principle that connections to high-scoring dealers contribute more to the score of the dealer firm than equal connections to low-scoring dealers. For weighted graphs, one can calculate several weighted variants:
  - *evwntrade* = Eigenvector centrality, weighted by number of trades.
  - *evwpar* = Eigenvector centrality, weighted by total par amount.
- Betweenness *bt*: A measure of the absolute position of a dealer in the network. The betweenness of a dealer is computed as the number of shortest paths linking two dealers in the network that pass through the dealer firm. Betweenness measures the connections beyond the first neighbors, and it takes into account the connections of second and higher-order neighbors. We use the directed version of betweenness.
- Closeness *cl*: A measure of influence with respect to centrality. The closeness of a dealer is computed as the inverse of the average number of steps that a dealer needs to take within the network to reach or be reached by any other dealer firm. It captures the connection to highly influential dealers. For directed graphs, *clout* (*clin*) is based on out-links (in-links) only.

- Cliquishness cc: A measure of the likelihood that two associates of a dealer are associates themselves. A higher value indicates a greater cliquishness. cc is also called clustering coefficient or transitivity.

We define an aggregate centrality measure Net for each dealer and time period by aggregating the above network measures. We compute Net as the first principle component of the above individual centrality measures, in equal- and value-weighted variants:

- Net (EW): First principle component of the equal-weighted network measures $dgout$, $dgin$, $kcoreout$, $kcorein$, bt, $clout$, $clin$, ev (cc is dropped since it requires at least two neighbors).

- Net (VW): First principle component of the value-weighted network measures $dgoutwntrade$, $dginwntrade$, $dgoutwpar$, $dginwpar$, $evwntrade$, $evwpar$.

B Explanatory Variables

Dealer characteristics:

- Underwriter: Dummy variable that equals one if the dealer is the underwriter of the bond, and zero otherwise.

- Primary dealer: Dummy variable for primary dealers

- NYC dealer: Dummy variable for NYC-headquartered dealers

- Dealer size: Size of the dealer firm in terms of assets (in billions)

- Dealer inventory: Aggregate dealer inventory on the day prior to the trade, standardized by subtracting the average dealer inventory and dividing by its standard deviation.

Trade characteristics:

- Natural log of par value interacted with a trade size dummy. Trade size dummies are for Par <\$100K, \$100K$\leq$ Par <\$1M and, respectively, Par $\geq$\$1M.

Bond characteristics:

- Maturity: natural logarithm of the time until the bond matures, expressed in years

- Age: natural logarithm of the time since the bond was issued

- Issue size: natural logarithm of the bond's issue size

- Credit quality: indicator variable for rating category (1=AAA, 2=AA+, etc.); dummy for high-yield rated bonds; dummy for unrated bonds

- GO bonds: Dummy variable for general obligation bond

- Callable: Dummy variable for call feature

- Sinkable: Dummy variable for sinking fund feature
- Bank qualified: Dummy variable for bank qualified bonds that commercial banks can purchase with tax benefits
- Taxable: Dummy variable for taxable bonds
- AMT: Dummy variable for alternative minimum tax

Liquidity conditions:

- Dealer exits: Dealer firms exiting from sample during the prior month, measured as fraction of total state trading volume
- Mutual fund flows: Aggregate municipal bond mutual fund out- and inflows during the prior week
- Calendar time controls: Beginning and end of week dummy; end-of-month dummy that equals one for the last three trading days of the month, and zero otherwise; hour-of-day dummies

Table 1: Summary statistics for dealer centrality

Panel A reports descriptive statistics for the dealer centrality measure Net and the individual centrality measures in the pooled dealer-day sample. Net (EW) is the equal-weighted, and Net (VW) the order flow-weighted centrality measure. Panel B reports correlations between the centrality measure Net and the individual centrality measures in the pooled dealer-day sample. The number of observations for the clustering coefficient cc is 1,639,422. For all other variables, the number of observations is 2,498,266.

Panel A: Summary statistics

Statistic	Mean	S.D.	Min.	Max.
Equal-Weighted Centrality Measures				
dgout	0.01	0.03	0.00	0.34
dgin	0.01	0.03	0.00	0.27
ev	0.10	0.17	0.00	1.00
bt	0.00	0.01	0.00	0.17
clout	0.00	0.01	0.00	0.02
clin	0.01	0.00	0.00	0.01
kcoreout	4.81	7.18	0.00	31.00
kcorein	6.05	7.57	0.00	33.00
Trade-Weighted Centrality Measures				
dgoutwntrade	0.11	0.66	0.00	34.07
dginwntrade	0.11	0.62	0.00	31.52
dgoutwpar	17.62	93.79	0.00	4,164.43
dginwpar	17.62	87.02	0.00	3,061.08
evwntrade	0.01	0.07	0.00	1.00
evwpar	0.01	0.07	0.00	1.00
Cliquishness *cc*	0.51	0.29	0.00	1.00

Panel B: Correlation matrix

Statistic	Net (EW)	Net (VW)
Net (EW)	1.00	0.62
dgout	0.95	0.69
dgin	0.96	0.65
ev	0.98	0.64
bt	0.74	0.54
clout	0.29	0.10
clin	0.11	0.03
kcoreout	0.88	0.43
kcorein	0.84	0.38
Net (VW)	0.62	1.00
dgoutwntrade	0.57	0.89
dginwntrade	0.57	0.84
dgoutwpar	0.58	0.90
dginwpar	0.62	0.93
evwntrade	0.53	0.73
evwpar	0.57	0.85
Cliquishness *cc*	-0.26	-0.17

Table 2: Stability in dealer relations and persistence in dealer ranks

Panel A reports the transition probability matrix for dealer relations from one month to the next. The transition matrix is calculated separately for unconditional relations between dealers and relations conditional on the direction of the order flow. The row headings indicate if a pair of dealers traded with each other in a given month or did not. The column headings indicate if the same trade relation persists in the next month. Panel B documents the persistence on dealer ranks across time. We report the transition matrix of dealer rank categories from one month to the next. Dealer ranks are measured by the ordering of their centrality measure *Net*. To compute the dealer rank in a given month, we use all interdealer trades during the past 30 trading days.

Panel A: Stability in dealer relations

	Order flow next month		Order flow in same direction next month	
Order flow this month	$= 0$	> 0	$= 0$	> 0
$= 0$	0.85	0.15	0.86	0.14
> 0	0.35	0.65	0.38	0.62

Panel B: Persistence in dealer ranks

		Rank month $t+1$					
		Top 10	11-20	21-50	51-100	101-200	>200
Rank month t	Top 10	0.93	0.07	0.00	0.00	0.00	0.00
	11-20	0.07	0.78	0.14	0.00	0.00	0.00
	21-50	0.00	0.05	0.81	0.14	0.00	0.00
	51-100	0.00	0.00	0.08	0.79	0.13	0.00
	101-200	0.00	0.00	0.00	0.06	0.79	0.15
	>200	0.00	0.00	0.00	0.00	0.03	0.97

Table 3: Order flow routing

The table reports the average network centrality for each dealer in round-trip chains of different length. We restrict the sample to round-trips that involve no more than seven dealers. Agency trades, in which dealers act in the capacity of agent for an investor, are eliminated. We measure the dealer centrality Net by the first principle component of the centrality proxies described in Appendix A, standardized by the empirical cdf. $Net = 0$ (1) corresponds to the most peripheral (central) dealer. Columns correspond to the position of the dealer in the round-trip chain. Rows correspond to round-trip chains of given length. In parentheses we report t-tests for the difference in centrality between two consecutive dealers. Standard errors are adjusted for heteroskedasticity and double clustering by issuer and time. Significance levels are indicated by * (10%), ** (5%), *** (1%).

		Dealer centrality in round-trip chain						
Trade type	N	#1	#2	#3	#4	#5	#6	#7
CDC	8,808,119	0.960	.	.	.	.	.	.
			.	.	.	.	.	.
CDDC	1,511,196	0.918	0.924	.	.	.	.	.
			(0.01***)	.	.	.	.	.
CDDDC	866,450	0.899	0.965	0.896	.	.	.	.
			(0.07***)	(-0.07***)	.	.	.	.
CDDDDC	173,579	0.850	0.940	0.940	0.845	.	.	.
			(0.09***)	(-0.00)	(-0.10***)	.	.	.
CDDDDDC	37,229	0.858	0.954	0.893	0.945	0.832	.	.
			(0.10***)	(-0.06***)	(0.05***)	(-0.11***)	.	.
CDDDDDDC	6,866	0.848	0.943	0.911	0.854	0.949	0.785	.
			(0.09***)	(-0.03***)	(-0.06***)	(0.10***)	(-0.17***)	.
CDDDDDDDC	882	0.835	0.953	0.870	0.898	0.878	0.945	0.794
			(0.12***)	(-0.08***)	(0.03**)	(-0.02*)	(0.07***)	(-0.15***)

Table 4: Search efficiency and dealer centrality

The table documents the dealers' propensity to trade with customers and the length of intermediation chains and relates them to the dealers' network centrality. In the first two columns, we report how the frequency of trading with a customer as opposed to another dealer depends on the centrality of the dealer. The estimates are obtained from probit regressions with state and month fixed effects. In the last four columns, we report how the path length of round-trip chains is related to the network centrality of head dealer. The estimates are obtained from panel regressions and, alternatively, from poisson regressions with state and month fixed effects. The dealer centrality measure Net is the first principal component of the network variables in Appendix A. The EW (VW) columns employ the equal-weighted (value-weighted) dealer centrality measures. The sample consists of all *C(N)DC* round-trip transactions. Standard errors are adjusted for heteroskedasticity and double clustering by issuer and time. Significance levels are indicated by * (10%), ** (5%), *** (1%).

	(1)	(2)	(3)	(4)	(5)	(6)
			No. of dealers in round-trip chain			
	Pr(Customer sell | Trade)		OLS		Poisson	
Determinant	EW	VW	EW	VW	EW	VW
Dealer centrality	0.58***	0.50***	-1.23***	-0.95***	-0.46***	-0.36***
Underwriter	0.09***	0.09***	-0.16***	-0.16***	-0.07***	-0.07***
Primary dealer	0.04***	0.05***	-0.07***	-0.09***	-0.03***	-0.04***
NYC dealer	-0.03***	-0.03***	0.08***	0.07***	0.03***	0.03***
Dealer size	-0.26***	-0.25***	0.79***	0.78***	0.26***	0.26***
Dealer inventory	0.02***	0.02***	-0.05***	-0.05***	-0.02***	-0.02***
$log(Par)$*Small	-0.01***	-0.01***	0.02***	0.02***	0.01***	0.01***
$log(Par)$*Medium	-0.01***	-0.01***	0.00**	0.00**	0.00**	0.00***
$log(Par)$*Large	0.01***	0.01***	-0.02***	-0.02***	-0.01***	-0.01***
Maturity	-0.00***	-0.00***	0.00	0.00	0.00	0.00
Seasoning	-0.01***	-0.01***	0.02***	0.02***	0.01***	0.01***
Issue size	0.00***	0.00***	-0.00***	-0.00***	-0.00***	-0.00***
Rating	-0.00***	-0.00***	0.01***	0.01***	0.00***	0.00***
Junk rated	-0.03***	-0.03***	0.03	0.03*	0.01	0.01*
Unrated	-0.03***	-0.03***	0.03***	0.03***	0.01***	0.01***
Insured	-0.00	-0.00	0.02***	0.02***	0.01***	0.01***
General obligation	0.00	0.00	0.01***	0.01***	0.00***	0.00***
Callable	0.01***	0.01***	-0.02***	-0.02***	-0.01***	-0.01***
Sinking fund	0.01***	0.01***	-0.03***	-0.03***	-0.01***	-0.01***
Bank qualified	-0.04***	-0.04***	0.03***	0.04***	0.01***	0.02***
Taxable bond	0.05***	0.04***	-0.13***	-0.10***	-0.05***	-0.04***
Subject to AMT	0.02***	0.02***	-0.04***	-0.03***	-0.02***	-0.01***
Constant	-0.41***	-0.36***	2.83***	2.65***	1.01***	0.95***
R^2	.	.	0.116	0.096	.	.
N	27,749,874	27,749,874	11,404,321	11,404,321	11,404,321	11,404,321

Table 5: Dealer markups on round-trip transactions

The table reports descriptive statistics for dealer markups on round-trip transactions for different types of trades. Panel A documents average trading cost and dispersion in trading costs for different samples, trade sizes, and intermediation chain lengths. The base sample, *CDC-Nonsplit*, comprises all round-trips without interdealer trading and order splitting. The *CDC* sample pools all round-trips without interdealer trading. The largest sample, *C(N)DC*, combines all round-trips that involve no more than seven dealers. Agency trades in which dealers act as customers' agent are eliminated. All markups are measured in percent of the head dealer's purchase price from customer, as defined in expression (1). Panel B documents trading cost and loss probabilities by the centrality of the dealer intermediating the trade. We report average markups on round-trip transactions and the fraction of round-trip transactions with markdowns for different dealer tiers. We construct the dealer tiers by sorting the dealers on the equal-weighted centrality measure Net. We assign the top quartile of dealers to the central tier, the next quartile to the mid-tier, and the remainder to peripheral dealers. Significance levels are indicated by * (10%), ** (5%), *** (1%).

Panel A: Trading cost distribution by trade size and chain length

Trade type	N	Mean	S.D.	5%	25%	50%	75%	95%
CDC-Nonsplit	6,294,447	1.85	1.39	0.06	0.75	1.79	2.78	4.14
CDC-Split	2,513,672	2.08	1.39	0.11	1.08	2.08	2.95	4.24
CDC	8,808,119	1.92	1.39	0.08	0.84	1.90	2.84	4.17
C(N)DC	11,404,321	2.02	1.54	0.10	0.90	1.93	2.92	4.50
CDC-Nonsplit, split by trade size:								
25	909,451	1.81	1.26	0.20	0.87	1.69	2.60	3.98
50	537,488	1.45	1.24	0.10	0.48	1.19	2.20	3.63
100	309,612	1.09	1.17	0.04	0.22	0.72	1.70	3.27
250	46,545	0.60	0.82	0.00	0.10	0.29	0.86	2.24
500	36,283	0.46	0.71	0.00	0.07	0.20	0.62	1.80
1,000	43,615	0.24	0.54	-0.03	0.04	0.10	0.28	1.12
C(N)DC, split by chain length:								
CDC	8,808,119	1.92	1.39	0.08	0.84	1.90	2.84	4.17
CDDC	1,511,196	2.23	1.71	0.22	0.98	1.94	3.12	5.26
CDDDC	866,450	2.45	1.96	0.31	1.13	2.08	3.31	5.81
CDDDDC	173,579	3.12	2.52	0.41	1.46	2.65	4.13	7.39
CDDDDDC	37,229	3.44	2.83	0.45	1.64	2.85	4.54	8.29
CDDDDDDC	6,866	3.62	3.05	0.54	1.77	2.97	4.66	8.86
CDDDDDDDC	882	3.68	3.08	0.53	1.82	2.96	4.62	9.20
C(N)D	573,549	0.79	1.16	0.00	0.13	0.50	1.06	2.62

Panel B: Trading costs and loss probabilities by dealer centrality

	Average markups (in percent)				Loss probability (in percent)			
Trade type	Peripheral dealers	Mid-tier dealers	Central dealers	Central – Peripheral	Peripheral dealers	Mid-tier dealers	Central dealers	Central – Peripheral
CDC-Nosplit	1.73	1.81	2.15	0.42***	1.73	1.46	0.82	-0.90***
CDC-Split	1.96	2.01	2.36	0.39***	3.24	2.10	1.49	-1.75***
CDC	1.79	1.87	2.21	0.42***	2.13	1.65	1.03	-1.10***
C(N)DC	1.94	1.92	2.22	0.29***	1.89	1.66	1.08	-0.81***
CDC-Nonsplit, split by trade size:								
25	1.68	1.76	2.08	0.40***	1.40	1.30	0.86	-0.54***
50	1.18	1.52	1.83	0.65***	1.87	1.28	0.95	-0.92***
100	0.79	1.27	1.57	0.78***	2.78	1.62	1.19	-1.59***
250	0.50	0.75	0.97	0.47***	2.56	2.31	1.59	-0.97***
500	0.42	0.54	0.63	0.21***	3.14	2.53	2.42	-0.72**
1,000	0.24	0.29	0.31	0.07***	6.23	6.68	6.59	0.36

Table 6: Trading costs and dealer centrality

The table reports the determinants of round-trip trading costs. The estimates are obtained from panel regressions with state and month fixed effects. The determinants include the dealer characteristics, trade characteristics, issue and issuer characteristics described in Appendix B. The dealer centrality measure is the first principal component of the network variables in Appendix A. The EW (VW) columns employ the equal-weighted (value-weighted) dealer centrality measures. For the *C(N)DC* sample, the dealer centrality measure is defined as the centrality of the head dealer. We vary the regression sample across columns, considering three types of trades with varying dealer involvement. *CDC-Nonsplits* are round-trips intermediated by a single dealer where the original bond lot is not split. The *CDC* sample includes all round-trips intermediated by a single dealer. *C(N)DC* are round-trips intermediated by one or several dealers. Standard errors are adjusted for heteroskedasticity and double clustering by issuer and time. Significance levels are indicated by * (10%), ** (5%), *** (1%).

	(1)	(2)	(3)	(4)	(5)	(6)
	CDC-Nonsplit		*CDC*		*C(N)DC*	
Determinant	EW	VW	EW	VW	EW	VW
Dealer centrality	0.55***	0.42***	0.71***	0.53***	0.68***	0.57***
Chain length					1.06***	1.15***
Chain length*Centrality					-0.78***	-0.87***
Primary dealer	-0.02	-0.01	-0.09***	-0.08***	-0.09***	-0.09***
NYC dealer	0.01	0.01	0.09***	0.09***	0.04***	0.04***
Underwriter	0.02	0.01	0.02	0.02	0.01	0.01
Dealer size	-0.52***	-0.53***	-0.58***	-0.59***	-0.43***	-0.46***
Dealer inventory	-0.01	-0.01	-0.02**	-0.02**	0.00	0.00
$log(Par)$*Small	-0.34***	-0.34***	-0.24***	-0.24***	-0.27***	-0.27***
$log(Par)$*Medium	-0.37***	-0.37***	-0.26***	-0.26***	-0.29***	-0.29***
$log(Par)$*Large	-0.35***	-0.36***	-0.30***	-0.30***	-0.31***	-0.31***
Maturity	0.66***	0.66***	0.75***	0.75***	0.79***	0.79***
Seasoning	-0.06***	-0.06***	-0.06***	-0.06***	-0.05***	-0.05***
Issue size	-0.00**	-0.00**	-0.00***	-0.00***	-0.01***	-0.01***
Rating	0.07***	0.07***	0.07***	0.07***	0.08***	0.08***
Junk rated	-0.05	-0.05	-0.05	-0.05	0.04	0.04
Unrated	0.26***	0.26***	0.28***	0.28***	0.37***	0.37***
Insured	0.08***	0.08***	0.08***	0.08***	0.05***	0.05***
General obligation	-0.01	-0.01	-0.01	-0.01	-0.04***	-0.04***
Callable	-0.24***	-0.24***	-0.24***	-0.24***	-0.26***	-0.26***
Sinking fund	-0.09***	-0.09***	-0.08***	-0.09***	-0.08***	-0.08***
Bank qualified	0.08***	0.08***	0.03***	0.03***	0.03***	0.03***
Taxable bond	0.04	0.03	0.07*	0.06	0.11***	0.10***
Subject to AMT	0.11***	0.10***	0.11***	0.10***	0.13***	0.13***
Constant	1.17***	1.27***	0.63***	0.76***	0.65***	0.71***
R^2	0.369	0.368	0.358	0.357	0.354	0.354
N	6,294,447	6,294,447	8,808,119	8,808,119	11,404,321	11,404,321

Table 7: How do dealers split markups?

The table reports average markups per dealer on round-trip transactions with varying degree of dealer involvement. Total dealer markups are broken down by the number of dealers (across rows) and by each dealer (across columns) in the sequence of dealers intermediating the round-trip transaction. We restrict the sample to non-splits. Markups are measured in percentage of the first dealer's purchase price from customer. No additional data filters are applied.

		Markup of each dealer in round-trip chain						
Trade type	Total markup	#1	#2	#3	#4	#5	#6	#7
CDC	1.85	1.85	.	.	.	.	.	.
		(100%)	.	.	.	.	.	.
CDDC	1.94	0.84	1.10	.	.	.	.	.
		(43%)	(57%)	.	.	.	.	.
CDDDC	2.26	0.66	0.52	1.08	.	.	.	.
		(29%)	(23%)	(48%)	.	.	.	.
CDDDDC	2.92	0.64	0.60	0.55	1.13	.	.	.
		(22%)	(21%)	(19%)	(39%)	.	.	.
CDDDDDC	3.26	0.63	0.30	0.82	0.40	1.11	.	.
		(19%)	(9%)	(25%)	(12%)	(34%)	.	.
CDDDDDDC	3.57	0.60	0.27	0.45	0.85	0.27	1.14	.
		(17%)	(8%)	(13%)	(24%)	(8%)	(32%)	.
CDDDDDDDC	3.71	0.62	0.23	0.43	0.53	0.47	0.29	1.15
		(17%)	(6%)	(12%)	(14%)	(13%)	(8%)	(31%)

Table 8: Riskless principal trades and dealer centrality

The table reports the determinants of immediate matches in round-trip transactions (1) versus trades entering dealer inventory (0). The estimates are obtained from panel regressions with state and month fixed effects. The determinants include the dealer characteristics, trade characteristics, issue and issuer characteristics described in Appendix B. The dealer centrality measure is the first principal component of the network variables in Appendix A. The EW (VW) columns employ the equal-weighted (value-weighted) dealer centrality measures. For the *C(N)DC* sample, the dealer centrality measure is defined as the centrality of the head dealer. We vary the regression sample across columns, considering three types of trades with varying dealer involvement. *CDC-Nonsplits* are round-trips intermediated by a single dealer where the original bond lot is not split. The *CDC* sample includes all round-trips intermediated by a single dealer. *C(N)DC* are round-trips intermediated by one or several dealers. Standard errors are adjusted for heteroskedasticity and double clustering by issuer and time. Significance levels are indicated by * (10%), ** (5%), *** (1%).

	(1)	(2)	(3)	(4)	(5)	(6)
	CDC-Nonsplit		*CDC*		*C(N)DC*	
Determinant	EW	VW	EW	VW	EW	VW
Dealer centrality	-1.52***	-1.49***	-1.61***	-1.57***	-2.11***	-1.87***
Chain length					0.47***	0.28***
Chain length*Centrality					0.18***	0.41***
Primary dealer	0.03	0.02	0.04**	0.03*	-0.03	-0.05
NYC dealer	0.00	0.02	-0.01	0.01	0.07**	0.03
Underwriter	-0.08*	-0.06	-0.07	-0.06	-0.08**	-0.08*
Dealer size	0.65***	0.62***	0.65***	0.62***	0.35***	0.27***
Dealer inventory	-0.07***	-0.07***	-0.07***	-0.07***	-0.13***	-0.12***
$log(Par)$*Small	0.16***	0.16***	0.05***	0.05***	-0.02***	-0.02***
$log(Par)$*Medium	0.15***	0.15***	0.04***	0.04***	-0.03***	-0.03***
$log(Par)$*Large	0.11***	0.12***	0.05***	0.05***	-0.01	0.00
Maturity	-0.02***	-0.02***	-0.07***	-0.07***	-0.08***	-0.08***
Seasoning	-0.04***	-0.04***	-0.02***	-0.01***	0.01**	0.01***
Issue size	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***
Rating	-0.00	-0.00	-0.01***	-0.01**	-0.01***	-0.01***
Junk rated	0.15**	0.15**	0.16**	0.16**	0.01	0.01
Unrated	-0.01	-0.01	-0.04*	-0.04*	-0.07***	-0.07***
Insured	-0.15***	-0.15***	-0.17***	-0.16***	-0.08***	-0.07***
General obligation	-0.08***	-0.08***	-0.08***	-0.08***	-0.03***	-0.03***
Callable	-0.00	-0.00	0.01	0.01	0.04***	0.04***
Sinking fund	-0.03**	-0.02**	-0.03**	-0.02**	-0.03***	-0.03***
Bank qualified	-0.10***	-0.10***	-0.03**	-0.03***	-0.15***	-0.13***
Taxable bond	0.34***	0.37***	0.30***	0.34***	0.16***	0.21***
Subject to AMT	0.07***	0.08***	0.04**	0.05***	0.02*	0.03**
Constant	-0.53***	-0.58***	-0.14	-0.19*	0.09	0.05
N	6,294,444	6,294,444	8,808,115	8,808,115	11,404,316	11,404,316

Table 9: Inventory duration and dealer centrality

The table reports the determinants of inventory durations. The estimates are obtained from panel Tobit regressions (lower limit=0) with state and month fixed effects. The determinants include the dealer characteristics, trade characteristics, issue and issuer characteristics described in Appendix B. The dealer centrality measure is the first principal component of the network variables in Appendix A. The EW (VW) columns employ the equal-weighted (value-weighted) dealer centrality measures. For the *C(N)DC* sample, the dealer centrality measure is defined as the centrality of the head dealer. We vary the regression sample across columns, considering three types of trades with varying dealer involvement. *CDC-Nonsplits* are round-trips intermediated by a single dealer where the original bond lot is not split. The *CDC* sample includes all round-trips intermediated by a single dealer. *C(N)DC* are round-trips intermediated by one or several dealers. Standard errors are adjusted for heteroskedasticity and double clustering by issuer and time. Significance levels are indicated by * (10%), ** (5%), *** (1%).

	(1)	(2)	(3)	(4)	(5)	(6)
	CDC-Nonsplit		*CDC*		*C(N)DC*	
Determinant	EW	VW	EW	VW	EW	VW
Dealer centrality	1.52***	1.09***	2.04***	1.40***	1.73***	1.31***
Chain length					1.34***	1.42***
Chain length*Centrality					-0.60***	-0.71***
Primary dealer	0.31***	0.32***	0.16**	0.18***	0.26***	0.27***
NYC dealer	0.36***	0.36***	0.43***	0.44***	0.33***	0.34***
Underwriter	0.07*	0.07	0.08**	0.08*	0.10***	0.10***
Dealer size	-2.16***	-2.21***	-2.57***	-2.65***	-1.94***	-1.95***
Dealer inventory	0.14***	0.14***	0.11***	0.11***	0.07***	0.07***
$log(Par)$*Small	-0.55***	-0.55***	-0.06***	-0.06***	-0.15***	-0.15***
$log(Par)$*Medium	-0.62***	-0.62***	-0.06***	-0.06***	-0.18***	-0.18***
$log(Par)$*Large	-0.55***	-0.56***	-0.19***	-0.19***	-0.23***	-0.23***
Maturity	0.17***	0.17***	0.38***	0.38***	0.31***	0.31***
Seasoning	-0.14***	-0.14***	-0.24***	-0.24***	-0.23***	-0.23***
Issue size	-0.01***	-0.01***	-0.01***	-0.01***	-0.02***	-0.02***
Rating	0.02***	0.02***	0.05***	0.05***	0.03***	0.03***
Junk rated	-0.29**	-0.30***	-0.14	-0.15	-0.16	-0.16
Unrated	0.04	0.04	0.16***	0.16***	0.05	0.05
Insured	-0.06***	-0.06***	-0.09***	-0.08***	-0.13***	-0.13***
General obligation	0.03	0.03	-0.03	-0.03	-0.04	-0.03
Callable	-0.46***	-0.46***	-0.61***	-0.62***	-0.66***	-0.66***
Sinking fund	-0.46***	-0.46***	-0.51***	-0.51***	-0.54***	-0.54***
Bank qualified	0.40***	0.39***	0.18***	0.16***	0.24***	0.22***
Taxable bond	-0.12	-0.14	0.06	0.02	-0.00	-0.03
Subject to AMT	-0.49***	-0.50***	-0.47***	-0.47***	-0.55***	-0.55***
Constant	3.64***	3.95***	2.48***	2.94***	3.46***	3.71***
N	6,294,447	6,294,447	8,808,119	8,808,119	11,404,321	11,404,321

Table 10: Dealer choice and terms of trade

The table documents the determinants of investors' dealer choice and, respectively, dealer markups and liquidity supply. The determinants include the dealer, trade, issue and issuer characteristics described in Appendix B. The first column reports the determinants of investors' dealer choice $D_i \in$ {Peripheral, Mid-tier, Central} from expression (7) and the thresholds Π_j. We construct the dealer tiers by sorting the dealers on the equal-weighted centrality measure Net. We assign the top quartile of dealers to the central tier, the next quartile to the mid-tier, and the remainder to peripheral dealers. The estimates are from panel ordered probit regressions with state fixed effects. The remaining columns report the coefficients β on the markup equation in expression (9) and, respectively, the coefficients γ for liquidity supply in expression (10) for each dealer tier. The estimates are from panel regressions with state and month fixed effects. The sample comprises all *C(N)DC* round-trips. Standard errors are adjusted for heteroskedasticity and robust to clustering by issuer and time. Significance levels are indicated by * (10%), ** (5%), *** (1%).

	Dealer	Peripheral		Mid-tier		Central	
Determinant	Choice	Markup	Liquidity	Markup	Liquidity	Markup	Liquidity
Underwriter		0.24***	0.54***	0.09***	-0.06**	0.06***	0.11***
Primary dealer		-0.19***	-0.07	-0.09***	1.12***	0.04	-0.33***
NYC dealer		0.15***	0.63***	-0.12***	-0.05	-0.17***	-0.00
Dealer size		-0.80***	-1.63***	0.00	0.00	0.00	0.00
Dealer inventory		0.03***	0.14***	0.05***	0.21***	0.02	0.12***
$log(Par)$*Small	0.05***	-0.29***	-0.11***	-0.21***	-0.23***	-0.23***	-0.10***
$log(Par)$*Medium	-0.01*	-0.29***	-0.10***	-0.19***	-0.12***	-0.18***	-0.02
$log(Par)$*Large	-0.03***	-0.26***	-0.11***	-0.21***	-0.15***	-0.23***	-0.06***
Maturity	0.05***	0.54***	0.14***	0.60***	0.27***	0.68***	0.27***
Seasoning	0.04***	-0.05***	-0.25***	-0.12***	-0.32***	-0.11***	-0.32***
Issue size	-0.00***	-0.00***	-0.01***	0.00	-0.00*	0.00***	-0.00
Rating	-0.00**	0.05***	0.01**	0.05***	0.04***	0.06***	0.04***
Junk rated	-0.09**	-0.02	0.08	-0.15*	-0.37***	-0.13**	-0.23*
Unrated	0.07***	0.16***	0.02	0.06**	-0.14***	0.06**	-0.08
Insured	0.07***	0.00	-0.20***	0.07***	-0.17***	0.05***	-0.14***
General obligation	0.02***	-0.03***	-0.01	-0.01	-0.11***	-0.01	-0.08***
Callable	0.03***	-0.14***	-0.42***	-0.26***	-0.78***	-0.31***	-0.70***
Sinking fund	-0.07***	0.07***	-0.27***	-0.09***	-0.44***	-0.09***	-0.49***
Bank qualified	-0.25***	0.07***	0.10***	0.13***	0.30***	0.21***	0.42***
Taxable bond	-0.28***	0.15***	0.37***	0.34***	0.54***	0.43***	0.35***
Subject to AMT	-0.21***	0.22***	-0.16***	0.23***	-0.17**	0.33***	-0.11**
Dealer exits (state)	0.30***						
Dealer exits (total)	1.30**						
Mutual fund flows	-27.21**						
Monday	0.00						
Friday	-0.02***						
End-of-month	-0.00						
λ		-0.37***	-0.87***	-0.63***	-0.62***	-1.26***	-1.29***
Constant		1.34***	3.80***	-0.07	3.08***	2.25***	6.14***
Preference thresholds:							
Π_1	1.56***						
Π_2	2.26***						
R^2	.	0.326	0.028	0.312	0.039	0.257	0.033
N	11,404,321	5,702,598	5,702,598	2,852,809	2,852,809	2,848,914	2,848,914

Table 11: Counterfactuals analysis

The table reports the expected markup and liquidity at different dealers, split by their centrality into three tiers, and it compares the markup and liquidity at transacting dealers to the counterfactual quotes at non-transacting dealers. Expected markup and liquidity are computed based on the estimates for β and γ in Table 10 and reported in Panel A. The markups and liquidity observed based on investors' actual choices $D \in \{P, M, C\}$, $E[M_D|D]$ and $E[L_D|D]$, are from expression (9). The selection on markup $E[M_j|D{=}j]{-}E[X\beta_j]$ and, respectively, liquidity $E[L_j|D{=}j]{-}E[X\gamma_j]$ are reported in parenthesis. The counterfactual markups and liquidity by dealers that the investor did not trade with, $E[M_j|D]$ and $E[L_j|D]$ for $j \neq D$, are computed using expression (12) and reported in Panel B. Standard errors are adjusted for heteroskedasticity and robust to clustering by issuer and time. Significance levels are indicated by * (10%), ** (5%), *** (1%).

Statistic	Expression	Dealer tier: Peripheral	Middle	Central
Panel A: Markups and liquidity across dealer tiers				
Markup expected	$E[X\beta_j]$	1.17	1.87	3.52
Markup observed	$E[M_j\|D{=}j]$	1.45	1.69	1.95
(Δ Markup)	$E[M_j\|D{=}j]{-}E[X\beta_j]$	(0.28***)	(-0.18***)	(-1.57***)
Liquidity expected	$E[X\gamma_j]$	1.64	3.23	4.64
Liquidity observed	$E[L_j\|D{=}j]$	2.30	3.05	3.04
(Δ Liquidity)	$E[L_j\|D{=}j]{-}E[X\gamma_j]$	(0.66***)	(-0.18***)	(-1.60***)
Liquidity price p_j	$E[X\beta_j]/E[X\gamma_j]$	0.72	0.58	0.76
Panel B: Counterfactuals analysis				
Counterfactual 1: Quoted markups and liquidity when investors trade with central dealers				
Markup quoted	$E[M_j\|D{=}C]$	0.84	1.13	1.95
(Δ Markup)	$E[M_j{-}M_C\|D{=}C]$	(-1.11***)	(-0.82***)	–
Liquidity offered	$E[L_j\|D{=}C]$	0.82	2.31	3.04
(Δ Liquidity)	$E[L_j{-}L_C\|D{=}C]$	(-2.22***)	(-0.74***)	–
Counterfactual 2: Quoted markups and liquidity when investors trade with mid-tier dealers				
Markup quoted	$E[M_j\|D{=}M]$	1.09	1.69	3.09
(Δ Markup)	$E[M_j{-}M_M\|D{=}M]$	(-0.60***)	–	(1.40***)
Liquidity offered	$E[L_j\|D{=}M]$	1.67	3.05	4.54
(Δ Liquidity)	$E[L_j{-}L_M\|D{=}M]$	(-1.38***)	–	(1.49***)
Counterfactual 3: Quoted markups and liquidity when investors trade with peripheral dealers				
Markup quoted	$E[M_j\|D{=}P]$	1.45	2.38	4.46
(Δ Markup)	$E[M_j{-}M_P\|D{=}P]$	–	(0.92***)	(3.01***)
Liquidity offered	$E[L_j\|D{=}P]$	2.30	3.51	5.91
(Δ Liquidity)	$E[L_j{-}L_P\|D{=}P]$	–	(1.21***)	(3.61***)

Table 12: Sub-period analysis

The table reports the expected markup and liquidity for different time periods and at different dealers, split by centrality into three tiers. The time periods are chosen to reflect the introduction of real-time reporting in February 2005 and the beginning of the financial crisis in September 2008. The markups and liquidity observed based on investors' actual choices $D \in \{P, M, C\}$, $E[M_D|D]$ and $E[L_D|D]$, are from expression (9). The selection on markup $E[M_j|D{=}j]{-}E[X\beta_j]$ and, respectively, liquidity $E[L_j|D{=}j]{-}E[X\gamma_j]$ are reported in parenthesis. Standard errors are adjusted for heteroskedasticity and robust to clustering by issuer and time. Significance levels are indicated by * (10%), ** (5%), *** (1%).

Statistic	Peripheral	Mid-tier	Central
Panel A: Dark period (Feb98-Jan05)			
Markup expected	1.31	2.18	3.80
Markup observed	1.77	2.07	2.11
(Δ Markup)	(0.46***)	(-0.12***)	(-1.69***)
Liquidity expected	2.50	2.79	3.53
Liquidity observed	2.84	2.83	2.95
(Δ Liquidity)	(0.34***)	(0.04***)	(-0.58***)
Liquidity price p_j	0.52	0.78	1.08
Panel B: Transparent period (Feb05-Aug08)			
Markup expected	1.16	1.65	2.68
Markup observed	1.20	1.47	1.72
(Δ Markup)	(0.04***)	(-0.19***)	(-0.96***)
Liquidity expected	1.47	3.85	6.90
Liquidity observed	2.14	3.53	3.72
(Δ Liquidity)	(0.67***)	(-0.32***)	(-3.18***)
Liquidity price p_j	0.78	0.43	0.39
Panel C: Crisis period (Sep08-Dec12)			
Markup expected	1.47	1.46	2.47
Markup observed	1.34	1.42	1.83
(Δ Markup)	(-0.12***)	(-0.04***)	(-0.64***)
Liquidity expected	1.43	3.20	4.30
Liquidity observed	1.94	2.86	2.91
(Δ Liquidity)	(0.51***)	(-0.34***)	(-1.39***)
Liquidity price p_j	1.03	0.46	0.57

Table 13: Loss probability and dealer centrality

The table reports the determinants for the probability that dealers take a loss on a round-trip transaction. The estimates are obtained from panel regressions with state and month fixed effects. The determinants include the dealer characteristics, trade characteristics, issue and issuer characteristics described in Appendix B. The dealer centrality measure is the first principal component of the network variables in Appendix A. The EW (VW) columns employ the equal-weighted (value-weighted) dealer centrality measures. For the *C(N)DC* sample, the dealer centrality measure is defined as the centrality of the head dealer. We vary the regression sample across columns, considering three types of trades with varying dealer involvement. *CDC-Nonsplits* are round-trips intermediated by a single dealer where the original bond lot is not split. The *CDC* sample includes all round-trips intermediated by a single dealer. *C(N)DC* are round-trips intermediated by one or several dealers. Standard errors are adjusted for heteroskedasticity and double clustering by issuer and time. Significance levels are indicated by * (10%), ** (5%), *** (1%).

	(1)	(2)	(3)	(4)	(5)	(6)
	CDC-Nonsplit		*CDC*		*C(N)DC*	
Determinant	EW	VW	EW	VW	EW	VW
Dealer centrality	-0.32***	-0.42***	-0.35***	-0.41***	-0.32***	-0.38***
Chain length					-0.45***	-0.48***
Chain length*Centrality					0.46***	0.49***
Primary dealer	0.02	0.02	0.06**	0.05**	0.08***	0.08***
NYC dealer	0.17***	0.18***	0.07***	0.08***	0.05***	0.06***
Underwriter	-0.12***	-0.11***	-0.10***	-0.09***	-0.08***	-0.08***
Dealer size	-0.48***	-0.49***	-0.25*	-0.26**	-0.28***	-0.27***
Dealer inventory	-0.10***	-0.10***	-0.07**	-0.07**	-0.05**	-0.05**
$log(Par)$*Small	0.02***	0.02***	0.03***	0.03***	0.01***	0.01***
$log(Par)$*Medium	0.06***	0.06***	0.08***	0.08***	0.07***	0.07***
$log(Par)$*Large	0.10***	0.10***	0.12***	0.12***	0.11***	0.11***
Maturity	-0.00	-0.00	-0.02**	-0.02**	-0.03**	-0.02**
Seasoning	-0.06***	-0.06***	-0.05***	-0.05***	-0.05***	-0.05***
Issue size	0.00	0.00	0.00***	0.00***	0.00***	0.00***
Rating	-0.01***	-0.01***	-0.01***	-0.01***	-0.01***	-0.01***
Junk rated	0.03	0.03	0.06*	0.06*	0.06**	0.06**
Unrated	-0.03*	-0.03*	-0.06***	-0.06***	-0.07***	-0.07***
Insured	-0.04***	-0.04***	-0.06***	-0.06***	-0.05***	-0.05***
General obligation	-0.00	-0.00	-0.00	-0.00	0.00	0.00
Callable	-0.17***	-0.17***	-0.20***	-0.20***	-0.20***	-0.20***
Sinking fund	-0.03***	-0.03***	-0.04***	-0.04***	-0.05***	-0.05***
Bank qualified	-0.03**	-0.04**	-0.07***	-0.07***	-0.06***	-0.07***
Taxable bond	-0.08***	-0.08***	-0.07***	-0.07***	-0.08***	-0.08***
Subject to AMT	-0.03*	-0.03*	-0.07***	-0.07***	-0.08***	-0.08***
Constant	-0.47***	-0.43***	-0.62***	-0.59***	-0.64***	-0.60***
N	6,294,444	6,294,444	8,808,115	8,808,115	11,404,316	11,404,316

Panel A: Order flow among most active dealers

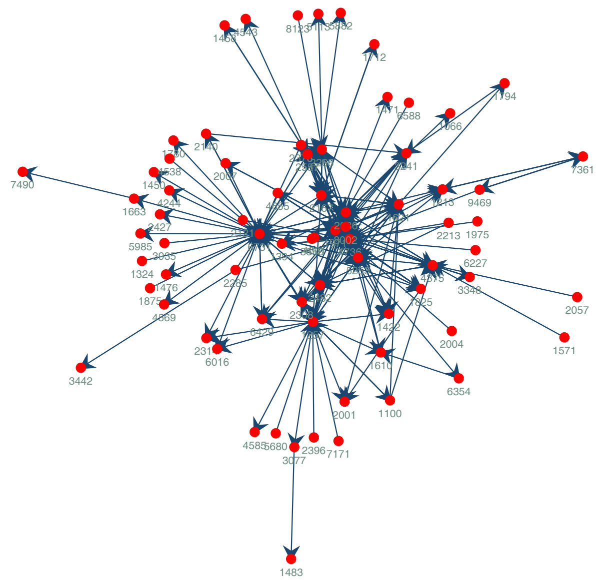

Panel B: Order flow in entire network

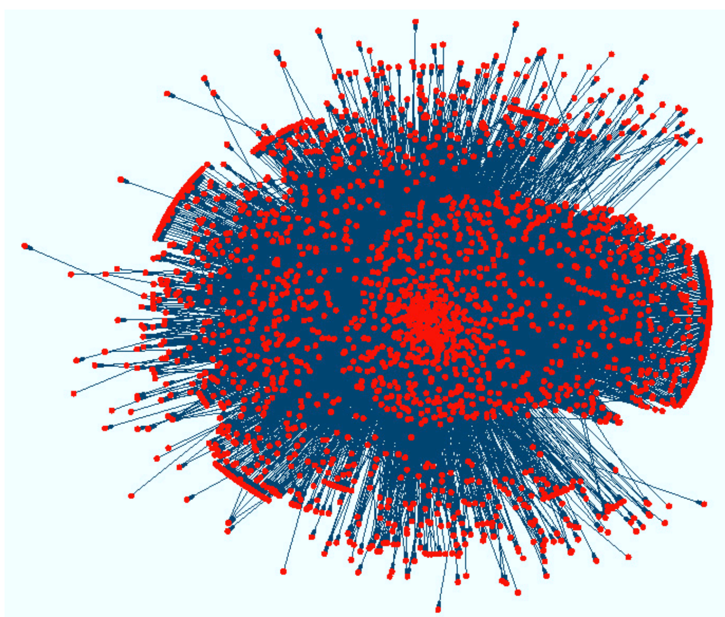

Figure 1: Topology of the dealer network

The figure illustrates the network structure of dealers in the municipal bond market in terms of the order flow between the dealers. Each node represents a dealer firm. Each arrow represents directed order flow between a pair of dealers. In Panel A, we impose the restriction that order flow between two dealers exceeds 10,000 transactions over the sample period. In Panel B, we plot the dealer network using all transactions. The plots are generated using multidimensional scaling based on the criterion that the more trade links exist between two dealers, the closer is their location on the map.

Panel A: Market connectedness (left: degree distribution, right: local clustering)

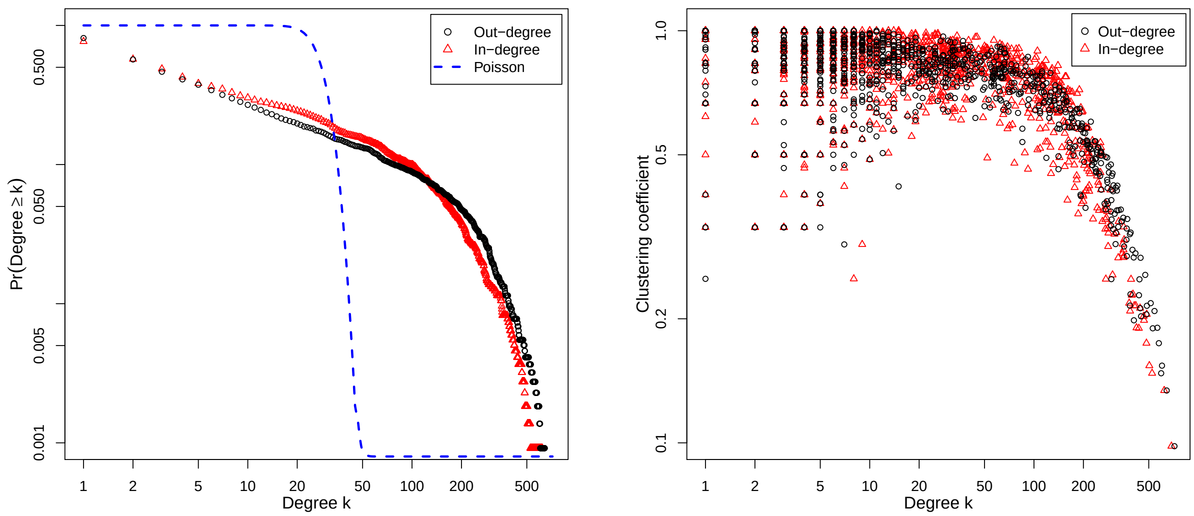

Panel B: Market resilience (left: giant component, right: average path length)

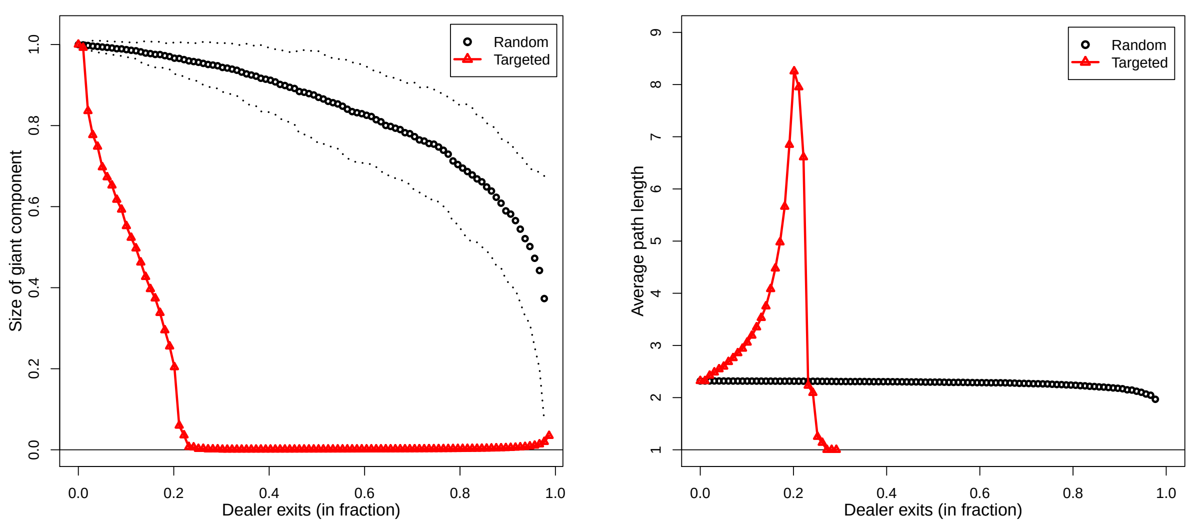

Figure 2: Market structure and resilience

The figure documents the structure and resilience of the municipal bonds market. Panel A explores the market connectedness and the (non-)randomness of trading relations between dealers. In the left plot, we show the inverse distribution function for the degree across dealers in the network on a log-log scale. The dots correspond to the out-degree. The triangles represent in-degrees. We add for comparison the degree distribution of a random trading network with the same average degree (dashed line). In the right plot, we explore the local clustering and hierarchical structure of the market. We plot the degree distribution across dealers in the network (horizontal axis) against the clustering coefficient of each dealer (vertical axis) on a log-log scale. Panel B documents the effect on the network structure of eliminating order flow by dealers. We plot the (residual) network connectedness as a function of the number of dealers that are being eliminated from the network. We consider two scenarios. The circles corresponds to the network connectedness when dealers exit at random. The triangles correspond to the network connectedness when the most connected dealers exit first. In the left plot, the vertical axis measures the size of the largest connected subgraph (so-called giant component) in proportion to the remaining dealers. In the right plot, the vertical axis measures the average path length between any two dealers.

Panel A: Par < \$100K

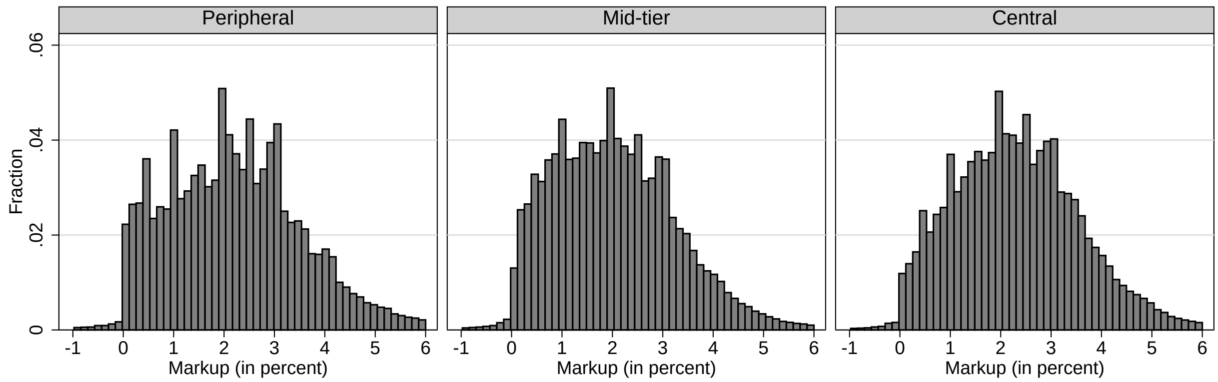

Panel B: \$100K $\leq Par <$ \$1M

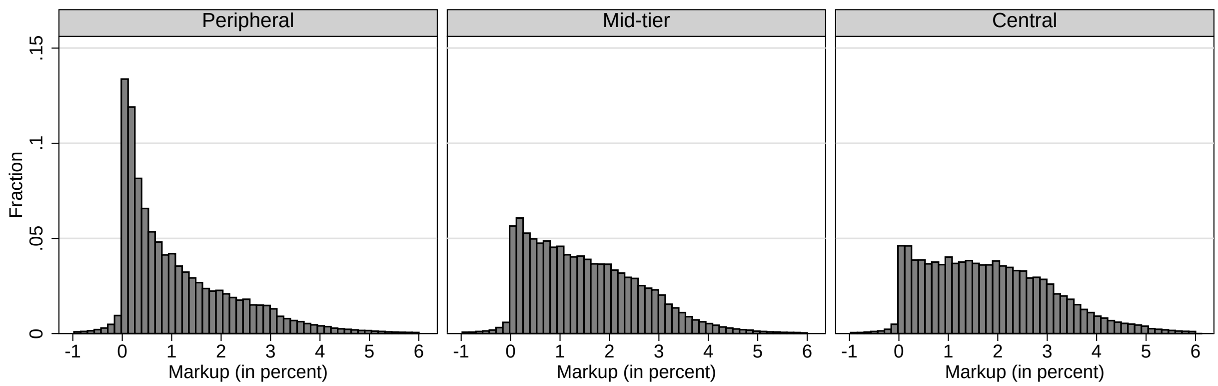

Panel C: $Par \geq$ \$1M

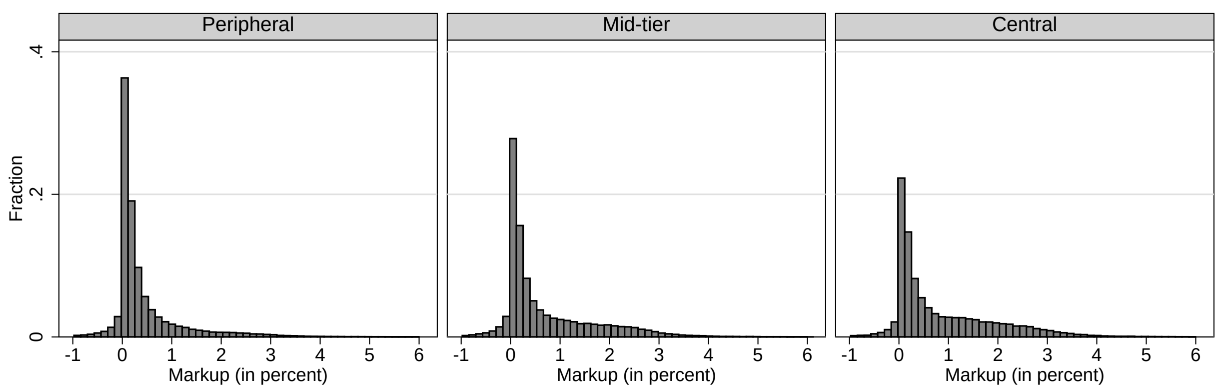

Figure 3: Trading costs and dealer centrality

The figure documents trading cost by trade size and the centrality of the dealer intermediating the trade. We plot the distribution of round-trip markups for different dealer tiers. We construct the dealer tiers by sorting the dealers on the centrality measure Net. We assign the top quartile of dealers to the central tier, the next quartile to the mid-tier, and the remainder to peripheral dealers. We plot the distribution over the range from -1% to 6%. The sample consists of all *C(N)DC* transactions.

Panel A: $Par < \$100K$

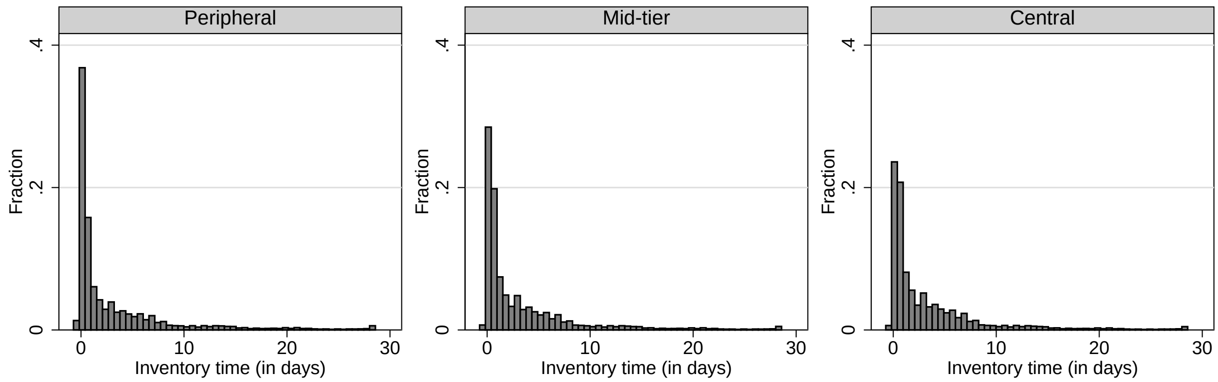

Panel B: $\$100K \leq Par < \$1M$

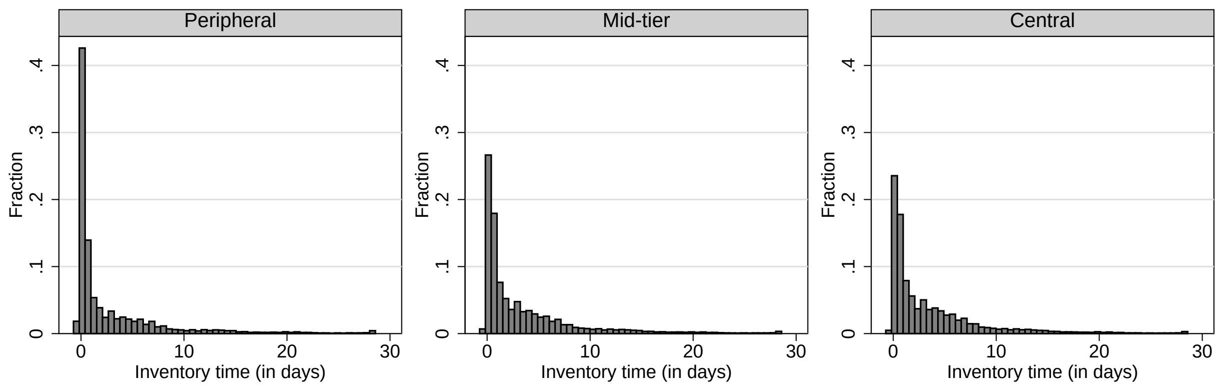

Panel C: $Par \geq \$1M$

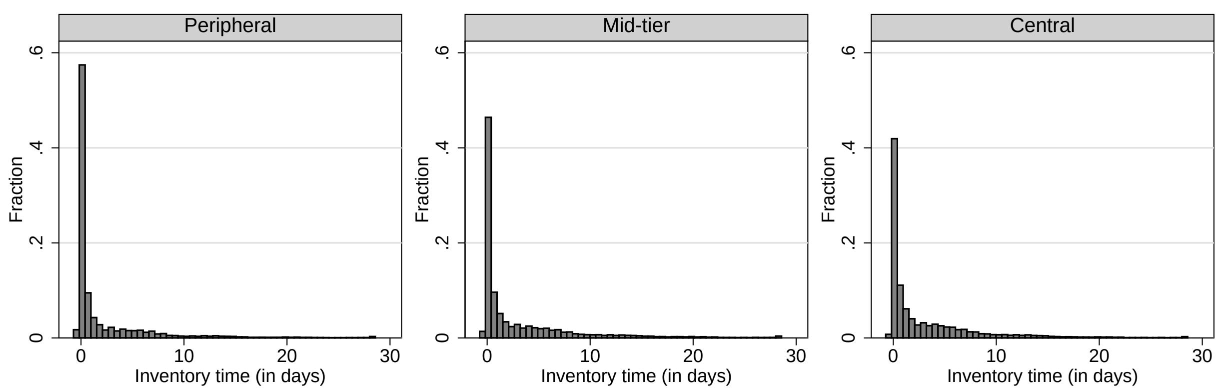

Figure 4: Inventory durations and dealer centrality

The figure documents inventory durations by trade size and the centrality of the dealer intermediating the trade. We plot the distribution of round-trip durations for different dealer tiers. We construct the dealer tiers by sorting the dealers on the centrality measure *Net*. We assign the top quartile of dealers to the central tier, the next quartile to the mid-tier, and the remainder to peripheral dealers. The sample consists of all *C(N)DC* transactions.

Panel A: Investor order flow (left: number of trades, right: par volume)

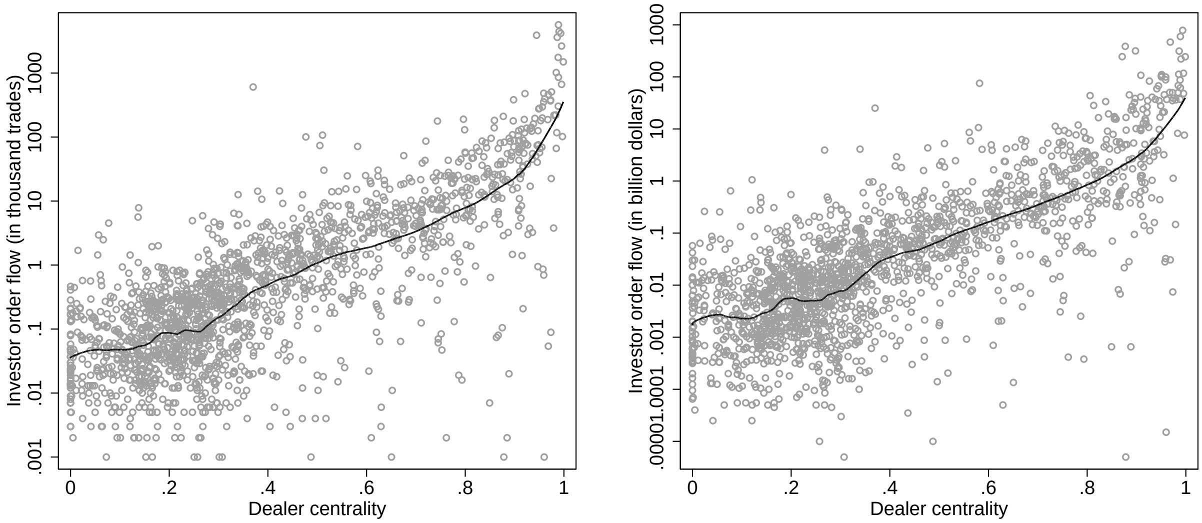

Panel B: Inventory volatility (left: absolute daily dollar changes, right: absolute daily changes relative to 30-day average)

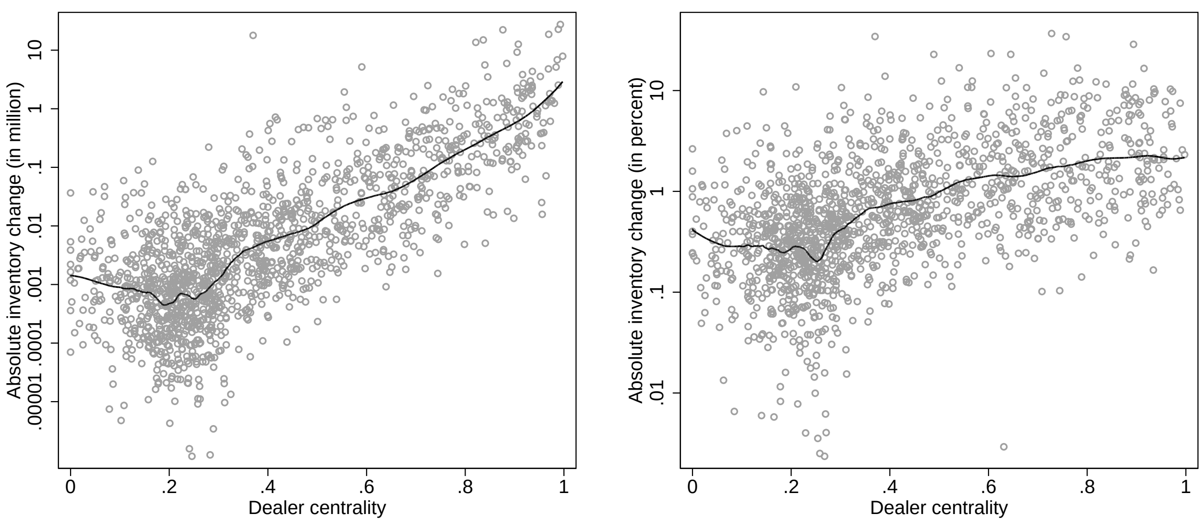

Figure 5: Order flow, inventory volatility and dealer centrality

The figure documents the relation between investor order flow, dealer inventories and dealer centrality. Panel A plots aggregate investor order flow over the sample period against centrality of the dealers (averaged over the sample period). Each circle corresponds to a dealer. In the left plot, we measure investor order flow by the total number of round-trip trades. In the right plot, we show the aggregate par volume of trade. Panel B documents the relation between dealer inventories and centrality. In the left plot, the vertical axis measures the daily dollar change in dealers' bond inventories. In the right plot, we normalize the daily changes by the average inventory over the past 30 days. The solid lines are the predictions from a local polynomial regression. The plots are on a log-linear scale.

Internet Appendix

Appendix IA.I Data Filters and Trade Matching

This appendix describes our data filters and the algorithm used to match customer-to-dealer and consecutive dealer-to-dealer and dealer-to-customer transactions into round-trip chains.

Data filters: We apply several data filters to clean the raw MSRB data from erroneous entries and missing information. The filters check the trades sequentially based on information about the trade, bond, price, and time. In the following table we provide a short description for each filter and report the number of trades and bonds remaining in the sample after applying each filter.

Description	No. of individual trades (in millions)
1. Trade-specific filters	
Keep trades after February, 1998 (only interdealer trades are reported before 1998)	
Keep trades for which the dealer identifier is a four letter alphabetic symbol	
Drop trades with MSRB indicator for away-from-market prices	
Keep trades with par value of at least \$5K	123.20
2. Bond-specific filters	
Keep bonds that have fixed or zero coupon (based on Mergent data)	
Keep bonds that are non-derivative and non-warrant, and not puttable	
Keep bonds that have at least one year to maturity at the time of issuance	
Keep bonds that have a denomination face value of \$5K (large majority of bonds)	100.65
3. Price- and time-based filters	
Keep only trades at least 90 days after issuance (seasoned bonds)	75.49
Drop bonds with less than a year to maturity (maturing bonds)	72.93
Eliminate price outliers by truncating the distribution at 0.5% and 99.5%, separately for zero-coupon and other bonds	72.20

We next search the filtered MSRB sample for trade sequences that correspond to round-trip chains using the algorithm documented below. The algorithm matches 82% (14.6 million) of all customer-to-dealer trades to corresponding dealer-to-dealer and dealer-to-customer trades. Of these, 13.3 million trades are associated with a complete round-trip chain, which ends in a dealer selling the bonds to a customer, constituting our *C(N)DC* sample. The remainder are incomplete intermediation chains, with a dealer keeping the bond in inventory or due to incomplete matching. We apply additional filters to the complete round-trip sample as described in the next table.

Description	No. of round-trip chains
Complete round-trip chains from filtered MSRB sample	13,253,694
1. Drop round-trip chains where head or tail dealer act as agent for the investor (agent trades)	11,819,422
2. Drop round-trip chains where first trade happened in March, 1998 (incomplete data)	11,791,684
3. Drop round-trip chains where first or last dealer do not have an MSRBID assigned, or it is unknown	11,454,377

Matching algorithm for round-trip chains: Our round-trip matching algorithm is a extension of the algorithm first used in Green, Hollifield, and Schürhoff (2007). The differences are due to the information we have on the identities of the dealers involved in each trade. This data allows us to trace the order flow of bonds through the dealer network with higher accuracy.

A typical trade in seasoned municipal bonds starts with an existing holder, being subject to a liquidity shock, wanting to liquidate the position. The investor will contact a broker-dealer in order to sell the bonds to a buyer with intermediation by the broker-dealer. We search for sequences of trades in each bond issue that resemble such bond orders flowing through the dealer network. The reverse chain of trades is more unlikely since short-selling of municipal bonds is very difficult and costly. Our round-trip chains therefore start with a time-stamped trade from a customer to a dealer, followed by either a trade from a dealer to a customer or to another dealer. The round-trip chain ends with a sale to a customer.

For each CUSIP and each customer-to-dealer trade, we sequentially look for one of two types of matches, non-split and split chains. For non-split round-trips, we restrict to matching trades of the same par sizes. We look for trades by the initial dealer within a calendar-day window from -10 to +30 days around the initial customer-to-dealer trade. If there is a trade of the same dealer selling the same amount of bonds to customers, this identifies a *CDC* round-trip. The round-trip chain is then removed from the trading data for future lookups. If no such direct match exists, we look for interdealer trades in which the initial dealer is a seller. If such an interdealer trade exists, we have a chain of *CDD* trades.

```
foreach CUSIP do
    while not the last trade do
        foreach Dealer buys from customer (Dealer = A, Par = X) do
            Run procedure to find matching trades (Dealer = A, Par = X)
            if Procedure to find matching trades returns finished chain, Finished=1 then
                Record round-trip, type='C(N)DC', Remove from data
            else
                if Find splitting dealer sells to customer (Dealer =Last dealer in the chain,
                Par < X) then
                    Record round-trip, type='C(N)DC-Split', remove from data
                else
                    Record unfinished chain, type='C(N)D', remove from data
```

Algorithm 1: Procedure to find round-trip chains

We then recursively look for matching dealer-to-customer trades or other interdealer trades, where each leg is no more than -10 to +30 calendar days apart. We allow for up to 7 dealers in a round-trip (very few matches occur with larger thresholds), and we require that each leg is at most 10 trades away in the original sequence of execution-time sorted MSRB trades. These restrictions are intended to assure accuracy in trade matching. The last restriction allows us to implement the lookup loop more efficiently and, in addition, diminishes erroneously matching unrelated trades. We look for order-splitting at the end of the chain if no matching dealer-to-customer trade can be found. For these split matches, we require that the total par amount of the dealer-to-customer trades are not greater than the original par size. We take the weighted average of the sale prices to compute a price for the last leg, and the weighted-average trade date is the trade date for the last leg.

Our matching algorithm matches a total of 13.3 million (74%) of all customer-to-dealer transactions to complete round-trip chains. The remaining customer-to-dealer transactions

```
return linked list of round-trip chain while not the last trade do
    if Dealer A sells to customer, Par = X then
        Add to the round-trip chain, Finished=1, Exit
    else
        if Dealer A sells to dealer B, Par = Y ≤ X then
            if Existing chain has more than 7 dealers then
                Finished=0, Exit
            else
                Add to round-trip chain, Run procedure to find matching trades (Dealer = B,
                Par = Y)
        else
            Next trade within [-10,30] window
```

Algorithm 2: Procedure to find matching trades ($Dealer = A$, $Par = X$)

are not matched to complete round-trip chains for the following potential reasons (we ignore these types of trades, since there is considerable noise in matching them). First, dealers may aggregate purchases from multiple customers and to a single customer sale. In this case, the dealer sale is in a par size larger than the original dealer purchase. Second, dealers may split orders at the interdealer stage. Third, any of the consecutive legs in a round-trip may be more than 30 days apart. Fourth, dealers may report their IDs erroneously.

Table IA.1: Robustness test: Overnight trades and dealer centrality

The table reports the determinants of same day matches in round-trip transactions (1) versus overnight trades (0). The estimates are obtained from panel regressions with state and month fixed effects. The determinants include the dealer characteristics, trade characteristics, issue and issuer characteristics described in Appendix B. The dealer centrality measure is the first principal component of the network variables in Appendix A. The EW (VW) columns employ the equal-weighted (value-weighted) dealer centrality measures. For the *C(N)DC* sample, the dealer centrality measure is defined as the centrality of the head dealer. We vary the regression sample across columns, considering three types of trades with varying dealer involvement. *CDC-Nonsplits* are round-trips intermediated by a single dealer where the original bond lot is not split. The *CDC* sample includes all round-trips intermediated by a single dealer. *C(N)DC* are round-trips intermediated by one or several dealers. Standard errors are adjusted for heteroskedasticity and double clustering by issuer and time. Significance levels are indicated by * (10%), ** (5%), *** (1%).

	(1)	(2)	(3)	(4)	(5)	(6)
	CDC-Nonsplit		*CDC*		*C(N)DC*	
Determinant	EW	VW	EW	VW	EW	VW
Dealer centrality	-0.83***	-0.79***	-1.06***	-0.97***	-1.18***	-1.05***
Chain length					2.32***	1.67***
Chain length*Centrality					-1.85***	-1.13***
Primary dealer	-0.41***	-0.42***	-0.31***	-0.32***	-0.32***	-0.33***
NYC dealer	0.05***	0.06***	0.01	0.02	0.02**	0.01
Underwriter	0.00	0.00	-0.01	-0.00	-0.03***	-0.02***
Dealer size	0.41***	0.39***	0.44***	0.43***	0.75***	0.71***
Dealer inventory	-0.05***	-0.05***	-0.04***	-0.05***	-0.06***	-0.06***
$log(Par)$*Small	0.10***	0.10***	-0.05***	-0.05***	-0.03***	-0.03***
$log(Par)$*Medium	0.16***	0.16***	-0.01***	-0.01***	0.00	0.01*
$log(Par)$*Large	0.18***	0.19***	0.06***	0.06***	0.06***	0.06***
Maturity	-0.06***	-0.06***	-0.12***	-0.12***	-0.08***	-0.08***
Seasoning	0.03***	0.03***	0.06***	0.06***	0.07***	0.07***
Issue size	0.00***	0.00***	0.00***	0.00***	0.00**	0.00***
Rating	-0.01***	-0.01***	-0.02***	-0.02***	-0.01***	-0.01***
Junk rated	0.02	0.02	0.04	0.04	0.07*	0.06*
Unrated	-0.02***	-0.03***	-0.06***	-0.07***	0.00	-0.00
Insured	-0.04***	-0.03***	-0.05***	-0.05***	-0.06***	-0.05***
General obligation	-0.04***	-0.04***	-0.03***	-0.04***	-0.05***	-0.05***
Callable	0.12***	0.12***	0.14***	0.14***	0.15***	0.15***
Sinking fund	0.11***	0.11***	0.11***	0.11***	0.11***	0.11***
Bank qualified	-0.10***	-0.10***	-0.01**	-0.01**	0.01*	0.02**
Taxable bond	0.16***	0.17***	0.14***	0.15***	0.13***	0.15***
Subject to AMT	0.24***	0.24***	0.22***	0.22***	0.21***	0.21***
Constant	0.16**	0.12	0.71***	0.63***	0.44***	0.36***
N	6,294,447	6,294,447	8,808,119	8,808,119	11,404,321	11,404,321

Table IA.2: Robustness test: Dealer inventory, market quality, and centrality

The table documents the determinants of the daily order flow, inventory variability, and market quality for each dealer. All dependent variables are calculated at dealer-year level. The absolute daily inventory change is defined as the absolute value of the daily change in each dealer's inventory level. The relative daily inventory change is defined as the absolute value of the daily percentage change in each dealer's inventory as a fraction of the daily change over the 30-day moving average, $|\Delta inv_t / \frac{1}{30}\sum_{i=1}^{30} inv_{t-i}|$, truncated at 1,000 percent. The market quality measure MQ is based on Hasbrouck (1993) and Hotchkiss and Ronen (2002). The dealer centrality measure is the first principal component of the network variables in Appendix A. The EW (VW) columns employ the equal-weighted (value-weighted) dealer centrality measures. Standard errors are adjusted for heteroskedasticity and double clustering at dealer and year. Yearly fixed-effects are included. Significance levels are indicated by * (10%), ** (5%), *** (1%).

	No. of trades (thousand)		Volume ($Billion)		Inventory variability ($M)		Inventory variability (%)		Market quality MQ	
Determinant	EW	VW	EW	VW	EW	VW	EW	VW	EW	VW
Dealer centrality	20.25***	20.10***	2.44***	2.50***	1.53***	1.59***	3.65***	3.51***	0.02***	0.02***
Primary dealer	30.66	30.04	12.58***	12.48***	8.04***	7.97***	0.76	0.69	0.01	0.01
NYC dealer	3.56	3.41	0.61*	0.59*	0.33	0.32	0.04	0.02	-0.01**	-0.01**
Dealer size	-0.00**	-0.00**	-0.00**	-0.00**	-0.00**	-0.00**	-0.00	-0.00	-0.00	-0.00
Constant	-4.95***	-4.88***	-0.64***	-0.65***	-0.40***	-0.41***	-0.06	-0.03	0.62***	0.62***
R^2	0.076	0.076	0.244	0.246	0.246	0.249	0.069	0.067	0.434	0.435
N	13,291	13,291	13,291	13,291	12,788	12,788	12,652	12,652	13,001	13,001

Table IA.3: Robustness test: Trading costs and dealer centrality

The table reports the determinants of round-trip trading costs. The estimates are obtained from panel regressions with state and month fixed effects, or indicated otherwise. The determinants include the dealer characteristics, trade characteristics, issue and issuer characteristics described in Appendix B. In column (1), the dealer centrality measure is the first principal component of the equal-weighted network variables in Appendix A. In column (2), the centrality measure is the dealers' eigenvector centrality *ev*. In column (3), the centrality measure is the dealers' state-specific centrality State-Net in the state of issuance of the bond being traded. In column (4), we add dealer fixed effects to the specification so that the coefficient on Net captures the impact of time-series variation in dealer centrality. In column (5), we include Net, State-Net, and dealer fixed effects. In column (6), we perform an instrumental variables regression which instruments Net by the liquidity condition variables listed in Appendix B. The regression sample are the *C(N)DC* round-trip chains. Standard errors are adjusted for heteroskedasticity and double clustering by issuer and time. Significance levels are indicated by * (10%), ** (5%), *** (1%).

Determinant	(1) Baseline	(2) Eigenvector centrality *ev*	(3) State-specific centrality	(4) Dealer FE	(5) State-Net + Dealer FE	(6) IV
Dealer centrality	0.68***	0.48***		0.56***	0.52***	2.19***
State-specific centrality			0.29***		0.06***	
Chain length	1.06***	0.66***	0.53***	0.99***	0.99***	0.43***
Chain length*Centrality	-0.78***	-0.54***	-0.26***	-0.73***	-0.73***	
Primary dealer	-0.09***	-0.12***	-0.09***	0.00	0.00	-0.13***
NYC dealer	0.04***	0.02***	0.04***	-0.07***	-0.07***	-0.05***
Underwriter	0.01	0.01	0.00	0.06***	0.06***	-0.01***
Dealer size	-0.43***	-0.62***	-0.38***	0.00	0.00	-0.35***
Dealer inventory	0.00	-0.00	0.01	0.02***	0.02***	0.02***
$log(Par)$*Small	-0.27***	-0.27***	-0.27***	-0.25***	-0.25***	-0.27***
$log(Par)$*Medium	-0.29***	-0.29***	-0.30***	-0.26***	-0.26***	-0.29***
$log(Par)$*Large	-0.31***	-0.31***	-0.31***	-0.28***	-0.28***	-0.31***
Maturity	0.79***	0.79***	0.79***	0.74***	0.74***	0.79***
Seasoning	-0.05***	-0.05***	-0.05***	-0.06***	-0.06***	-0.05***
Issue size	-0.01***	-0.01***	-0.01***	-0.00***	-0.00***	-0.01***
Rating	0.08***	0.08***	0.08***	0.07***	0.07***	0.08***
Junk rated	0.04	0.05	0.03	0.03	0.03	0.07***
Unrated	0.37***	0.37***	0.36***	0.31***	0.31***	0.37***
Insured	0.05***	0.04***	0.05***	0.05***	0.05***	0.04***
General obligation	-0.04***	-0.04***	-0.04***	-0.04***	-0.04***	-0.04***
Callable	-0.26***	-0.26***	-0.26***	-0.27***	-0.27***	-0.25***
Sinking fund	-0.08***	-0.08***	-0.08***	-0.10***	-0.10***	-0.08***
Bank qualified	0.03***	0.04***	0.03***	0.03***	0.03***	0.10***
Taxable bond	0.11***	0.11***	0.10***	0.14***	0.14***	0.19***
Subject to AMT	0.13***	0.13***	0.13***	0.13***	0.13***	0.13***
Constant	0.65***	1.12***	1.24***	0.74***	0.78***	-0.06*
R^2	0.354	0.358	0.352	0.399	0.399	0.316
N	11,404,321	11,404,321	11,404,321	11,404,321	11,404,321	11,404,321

Table IA.4: Robustness test: Inventory duration and dealer centrality

The table reports the determinants of inventory durations in round-trips. The estimates are obtained from panel regressions with state and month fixed effects, or indicated otherwise. The determinants include the dealer characteristics, trade characteristics, issue and issuer characteristics described in Appendix B. In column (1), the dealer centrality measure is the first principal component of the equal-weighted network variables in Appendix A. In column (2), the centrality measure is the dealers' eigenvector centrality *ev*. In column (3), the centrality measure is the dealers' state-specific centrality State-Net in the state of issuance of the bond being traded. In column (4), we add dealer fixed effects to the specification so that the coefficient on Net captures the impact of time-series variation in dealer centrality. In column (5), we include Net, State-Net, and dealer fixed effects. In column (6), we perform an instrumental variables regression which instruments Net by the liquidity condition variables listed in Appendix B. The regression sample are the *C(N)DC* round-trip chains. Standard errors are adjusted for heteroskedasticity and double clustering by issuer and time. Significance levels are indicated by * (10%), ** (5%), *** (1%).

	(1)	(2)	(3)	(4)	(5)	(6)
		Eigenvector	State-specific		State-Net +	
Determinant	Baseline	centrality *ev*	centrality	Dealer FE	Dealer FE	IV
Dealer centrality	1.73***	1.39***		0.25	0.21	0.46*
State-specific centrality			1.15***		0.06**	
Chain length	1.34***	1.02***	0.95***	1.20***	1.20***	0.71***
Chain length*Centrality	-0.60***	-0.36***	-0.23***	-0.54***	-0.54***	
Primary dealer	0.26***	0.17***	0.26***	0.00	0.00	0.22***
NYC dealer	0.33***	0.19***	0.32***	-0.08	-0.08	0.33***
Underwriter	0.10***	0.12***	0.09**	0.09***	0.09***	0.07***
Dealer size	-1.94***	-2.36***	-1.94***	0.00	0.00	-1.29***
Dealer inventory	0.07***	0.07***	0.07***	0.07***	0.07***	0.06***
$log(Par)$*Small	-0.15***	-0.16***	-0.16***	-0.14***	-0.14***	-0.12***
$log(Par)$*Medium	-0.18***	-0.17***	-0.17***	-0.13***	-0.13***	-0.13***
$log(Par)$*Large	-0.23***	-0.22***	-0.23***	-0.15***	-0.15***	-0.18***
Maturity	0.31***	0.31***	0.31***	0.30***	0.30***	0.28***
Seasoning	-0.23***	-0.24***	-0.23***	-0.25***	-0.25***	-0.23***
Issue size	-0.02***	-0.02***	-0.02***	-0.01***	-0.01***	-0.01***
Rating	0.03***	0.03***	0.03***	0.03***	0.03***	0.03***
Junk rated	-0.16	-0.17	-0.16	-0.09	-0.09	-0.04*
Unrated	0.05	0.06	0.04	0.03	0.03	0.07***
Insured	-0.13***	-0.15***	-0.12***	-0.17***	-0.17***	-0.19***
General obligation	-0.04	-0.04*	-0.04	-0.08***	-0.08***	-0.07***
Callable	-0.66***	-0.67***	-0.67***	-0.64***	-0.64***	-0.62***
Sinking fund	-0.54***	-0.52***	-0.54***	-0.52***	-0.52***	-0.53***
Bank qualified	0.24***	0.28***	0.21***	0.17***	0.17***	0.22***
Taxable bond	-0.00	-0.03	-0.02	0.06	0.06	0.14***
Subject to AMT	-0.55***	-0.52***	-0.54***	-0.33***	-0.33***	-0.48***
Constant	3.46***	4.61***	5.09***	5.21***	5.26***	4.79***
R^2	.	.	.	0.051	0.051	0.030
N	11,404,321	11,339,421	11,404,321	11,404,321	11,404,321	11,404,321