

Dealer Networks

Dan Li

Norman Schürhoff *

February 29, 2012

*Dan Li is with the Board of Governors of the Federal Reserve System and Norman Schürhoff 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 Rick Green and Andrew Ang for extensive comments and Darrell Duffie, Amy Edwards, Burton Hollifield, and Craig Lewis for illuminating discussions. Seminar participants at the SEC and at Carnegie Mellon University have provided valuable comments. The second author gratefully acknowledges research support from the Swiss Finance Institute and from NCCR FINRISK of the Swiss National Science Foundation.

Dealer Networks

We use the MSRB Transaction Reporting System audit trail to study the impact of dealer centrality on trading costs and liquidity provision in opaque over-the-counter markets. The dealership network in municipal bonds exhibits a hierarchical core-periphery structure with around 30 highly interconnected dealers at its core and several hundred peripheral dealer firms. Execution quality varies significantly across dealers depending on their centrality within the trading network. Central dealers charge larger trading costs to investors and face lower loss probabilities than peripheral dealers. Yet, more investor orders flow through central dealers. Informational efficiency increases with centrality, mitigating adverse selection risk. Dealers experience significant liquidity spillover from connected dealers, leading central dealers to hold larger and more volatile inventories on average and to keep bonds longer than peripheral dealers. Centrality thereby improves dealers' negotiating position with investors. This explains why competition is fierce at the periphery but not at the core of the decentralized market. These results demonstrate the trade-offs investors face when trading in over-the-counter markets, which may guide financial market design.

JEL Classification: G12, G14, G24

Key words: Municipal bond market, over-the-counter financial market, market quality, liquidity spillover, network analysis

Efficient allocation of resources requires well functioning markets for exchanging financial claims. Many financial securities, including municipal bonds, are traded through decentralized and opaque networks of financial intermediaries. In a decentralized dealership market, financial intermediaries form a network characterized by repeat interactions and long-term relations to facilitate the provision of liquidity to investors, the sharing of inventory risk, and the flow of information. But concentration of order flow reduces resilience to shocks and allows central dealers to exploit their advantage when interacting with investors. As a result, order execution quality and liquidity provision not only vary across trading venues and time but also depend systematically on dealers' centrality within the market. The following market commentary from New York Fed officials, who regularly talk to bond dealers, supports this notion:

“...Lastly, we have spoken about illiquidity in the market. It can be difficult to quantify liquidity in the market, and the color we hear from people in the market can really depend on the shop they are trading at. We’ve gotten the sense that trading may have been normal at larger dealers, but significantly light at smaller dealers.”

In this paper, we use the Municipal Securities Rulemaking Board’s (MSRB) proprietary Transaction Reporting System audit trail to study how market quality varies across the dealerships in municipal bonds and, in particular, how dealer centrality affects trading costs, liquidity provision, and price discovery. How well dealership markets perform in terms of market quality is largely an open question. Existing literature compares market quality across market structures or market-wide measures over time.¹ The impact on market quality of networked trading within a market is an important and understudied area.

The structure of the financial markets is pivotal for execution quality. When financial markets are underdeveloped, search and contractual frictions impose punitive transaction costs and hinder efficient price formation. By reducing transaction costs, financial intermediaries and trading facilities enable better allocation of securities among investors and improve risk sharing. The trading

¹Chowdhry and Nanda (1991) analyze cross-market trading by informed investors. Chordia, Roll, and Subrahmanyam (2000) and Hasbrouck and Seppi (2001) provide evidence for liquidity commonality across stocks. Hatch and Johnson (2002) and Comerton-Forde et al. (2010) study time variation in liquidity. Coughenour and Saad (2004) study liquidity spillovers across NYSE market-makers. Edwards, Harris, and Piwowar (2005) and Bessembinder, Maxwell, and Venkataraman (2006) study liquidity externalities and trading costs in corporate bonds. In recent work, Jovanovic and Menkveld (2010) study the liquidity impact of financial intermediaries' entry. Cespa and Foucault (2011) analyze liquidity spillovers across securities.

mechanism also enhances price discovery, or how efficiently new information is incorporated into prices. Efficient price discovery contributes to better real resource allocation, benefiting issuers and society at large by improving real investment decisions and increasing welfare.

The municipal bond market is the largest and most important capital market for state and municipal issuers. Its effective functioning is crucial for the provision of public services and the welfare of society. Trading of the bonds is organized like a typical decentralized broker-dealer market with limited pre- and post-trade transparency. More than 2,000 broker-dealer firms (all of which are obliged to register with the regulatory body MSRB) are actively trading in municipals.

We start by characterizing the bond market's structure in terms of the network of inter-dealer trading relations and its evolution and stability over time. We measure dealer centrality by their direct and indirect connections with other dealers through inter-dealer trading relations. The centrality measures that we use are borrowed from the literature on network analysis. We then document how the terms of trade for investors, the provision of liquidity by dealers, and the efficiency of prices depend on the financial intermediaries' position within the topology of the market. This provides insights into the incentives and market forces faced by broker-dealers in financial markets, the determinants of market quality, and the efficiency of price formation.

Our main empirical findings on the relation between trading costs, inventories, price efficiency, and dealer centrality are as follows:

1. The dealership network in municipal bonds exhibits a hierarchical core-periphery structure with around 30 highly connected dealers at its core and several hundred peripheral dealer firms. There is strong persistence in trading relations between dealers and in dealer ranks.
2. Dealers' average markups increase with the network centrality of the dealer(s) intermediating the trades. Central dealers charge significantly larger spreads in the secondary market than peripheral dealers, amounting up to 80% more for mid-size trades. Dealers' loss probability decreases with the network centrality of the dealer.
3. The informational efficiency of transaction prices rises with the centrality of the dealer(s) intermediating the trade.
4. Dealers experience significant liquidity spillovers from connected dealers. Central dealers provide more liquidity to customers than peripheral dealers; they trade more often and in larger aggregate volume. Central dealers also take more inventory risk as measured by daily inventory changes, inventory durations, and propensity for prearranging trades. Trading with central dealers reduces the number of inter-dealer transactions needed to complete a round-trip transaction.

5. Dealers’ bargaining position in trading with customers is stronger for central dealers than for peripheral dealers.

These findings suggest that competition is fierce at the periphery but not at the core of the decentralized market because of opacity and search frictions, giving rise to network effects. In networked dealership markets, investors face a trade-off between transactions cost and intermediation service quality (liquidity provision, price efficiency). Concentration of order flow with few central dealers leads to more efficient aggregation of new information about asset values and yields economies of scale in transaction processing and risk management—lowering transactions cost. On the other hand, concentration may reduce financial market stability and resilience to shocks, increasing risk and costs. In addition, dealers will be in a superior position to observe aggregate order flow and learn about the motives for trade. This will allow them to exploit their informational advantage when interacting with market participants, yielding them market power and raising trading costs. As a result, highly connected dealers at the center of the market offer accurate and efficient prices at a cost. By contrast, dealers on the periphery offer more competitive bid-ask spreads at the expense of lower liquidity and less informationally efficient prices. We provide support for these predictions in a proprietary sample of trades in municipal bonds. The results, more generally, shed light on the trade-offs investors face when trading in over-the-counter markets, which may guide financial market design.

The remainder of the paper is organized as follows. Section 1 describes the data sources. Section 2 documents the microstructure of the municipal bond market in terms of the trading relations between dealer firms. Section 3 documents how execution quality varies across dealers. Section 4 explores the relation between order flow, dealer inventory behavior, and dealer centrality. In Section 5, we provide evidence that networked trading allows dealers to charge discretionary markups. Section 6 concludes.

1 Institutional Background and Data

The municipal bond market is the largest and most important capital market for state and municipal issuers. It is a typical decentralized broker-dealer market with limited pre- and post-trade transparency. All trades are intermediated by broker-dealers who are registered with the Municipal Securities Rulemaking Board (MSRB). The MSRB is the self-regulatory body for the municipal bond market. Trade execution mostly occurs manually through sequential bilateral negotiations, by phone or electronic communication (see Biais (1993) and Yin (2005) for theoretical analysis). More than 2,000 broker-dealer firms are actively trading in municipal bonds. They provide liquidity by prearranging trades between customers or taking bonds into inventory.

Data sources: Our main data source is the proprietary MSRB Transaction Reporting System audit trail recorded by the MSRB. In an effort to improve market transparency, the MSRB requires all dealers in municipal debt to register with the MSRB and report all trades conducted in any municipal security. The data is thus comprehensive. Unlike the publicly available version of historical municipal bond transactions, our data provides identification of the dealer firms intermediating customer trades; for inter-dealer trades the data identify the dealers on each side of the trade.² The transactions data cover the 13 year period between February 1998 and July 2011. In addition to the complete transactions data, we obtain reference information on all outstanding bonds, including issuance date, maturity, coupon, taxable status, ratings, call features, issue size, and issuer characteristics from the Securities Data Company (SDC) Global Public Finance database.

We filter the transactions data to eliminate data errors and ensure data completeness. For a bond to be in our sample, we require availability of reference data in SDC and require the bond to have a fixed coupon. Green, Hollifield, and Schürhoff (2007) and Schultz (2012) document that trading and liquidity in newly issued bonds are markedly different from seasoned issues. For our transaction level analysis, we therefore remove all trades during the first 90 days after issuance.

Our final sample consists of approximately 60 million transactions in 1.4 million different bond issues. The trades are intermediated by a total of 2,078 dealer firms. Out of all transactions, 16 million are trades between dealers and the remainder are trades between investors and dealers.

²The data do not provide identifiers for the dealers' customers. See Hendershott and Madhavan (2011) for a recent study employing customer identifiers.

Network measures for dealer centrality: The bond dealers’ trading relations and relative positions in the trading network can be described by various network characteristics. There are six measures of centrality that are widely used in network analysis: degree centrality, k -core, eigenvector centrality, betweenness, closeness, and cliquishness.

- Degree centrality measures connectivity of a dealer in the network (a local property) by computing the fraction of dealers in the network to which the dealer firm is directly connected through bilateral trades (direct neighbors). Derived from degree is the k -core, which measures centrality as the maximal sub-network in which each dealer has at least degree k .
- Eigenvector centrality measures importance of a dealer in the network (a global property) by assigning 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 than equal connections to low-scoring dealers.
- Betweenness measures absolute position by taking into account the connections beyond the first neighbors (indirect neighbors). Betweenness is computed by counting the number of shortest paths linking any two dealers in the network that pass through the dealer firm. Like eigenvector centrality, betweenness captures a dealer’s overall importance.
- Closeness measures influence with respect to centrality by computing 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.
- Cliquishness measures local connectivity by computing the likelihood that two associates of a dealer are associates themselves. The correlation between degree and cliquishness determines the hierarchical structure of the network.

Table 1 provides a more detailed description of the centrality measures. We aggregate all network variables to a single index, denoted Net , by taking the first principle component across the measures as described in Table 1. For robustness, we construct both equal-weighted and value-weighted centrality measures, where we weight each connection by the order flow between the dealers. Our results are robust to taking the individual statistics or the aggregate statistic Net .

[Table 1 about here]

The network properties for each dealer are calculated using all inter-dealer transactions between February 1998 and July 2011. For every day during our sample period, we compute a directed network based on transactions during the past 30 calendar days. The number of trades between each pair of dealers and the total par amount of trade between them are recorded for the different weighting schemes. Following Milbourn (2003) and others, we apply an empirical cumulative distribution function (ecdf) transformation for each network variable to reduce skewness in the variable and diminish the impact of outliers. The ecdf transformation also facilitates interpretation of the economic magnitude of the results. As distribution functions take values between 0 and 1, a change from 0 to 1 in the network variables corresponds to moving from the least central position to the most central position across dealers.

Table 2 provides summary statistics for the variables describing the dealers' network characteristics. We track all 2,078 dealer firms over 3,400 trading days, yielding 2,498,266 dealer-day observations. On average, 800 to 900 dealers are actively trading on any given day. Table 3 reports the correlation coefficients between the dealers' network-related characteristics and *Net*. Panel A reports correlation coefficients for the equal-weighted centrality measures, and Panel B for the value-weighted variants. Across columns, we vary between the raw and standardized network measures. While the various network characteristics measure different aspects of network relations, there is a sizeable common component. The correlations between the variables and *Net* are significantly positive but typically less than unity. Cliquishness *cc* is the exception, in that it correlates negatively with most other measures.

[Tables 2 and 3 about here]

2 OTC market structure

In the following, we describe the microstructure of the municipal bond market in terms of the trading relations between dealer firms. We also provide descriptive statistics on connectedness, hierarchical structure, and shock resilience in the municipal bond market.

2.1 Network of inter-dealer trading relations

We first describe the structure of the dealership market by measuring the order flow between dealers. Some dealers interact frequently, others rarely or never. The strength of a relation between a pair of dealers can be measured by the number of times or, alternatively, by the number of bonds they trade with each other. One can also assign a direction to a relation depending on who buys and who sells. In the following, a dealer firm is identified by its MRSB registration. We pool all transactions over the sample period. We later study the time-series dynamics of trading relations.

Figure 1 illustrates the network structure of dealers in the municipal bond market in terms of order flow between dealers. Each node represents a dealer. Each arrow represents directed order flow between two dealers (we only consider order flows that exceed a minimum of \$5,000 in par value). In Panel A, we impose the restriction that order flow between two dealers exceeds 10,000 transactions over the sample period. This allows us to focus on the most connected dealers, forming the core of the municipal bond market. In Panel B, we plot the dealer network using all transactions. The plots in Panel A and B are generated using multidimensional scaling. The figure suggests that the municipal bond market has a hierarchical core-periphery structure. Around 30 dealers are highly connected and trade heavily with other dealers. In contrast, the remaining several hundred dealer firms are peripheral in that they trade less frequently and with a more limited number of trading partners.

[Figure 1 about here]

2.2 OTC market connectedness, hierarchy, and resilience

We can use the dealers' network characteristics introduced in Section 1 to determine systematic patterns in inter-dealer relations and overall market structure. In a centralized market, each investor can trade with everybody else, so the market is perfectly connected. In decentralized markets, investors have preferred dealers and, in turn, dealers trade preferably with their associates. The question arises whether all dealers trade with each other in the municipal bond market or whether they form long-term relations, and how do these patterns vary across dealers?

Figure 2 documents the connectedness of the market. We plot the degree distribution across dealers in the network. The black dots correspond to the out-degrees, the red dots represent in-degrees. The dots trace out the inverse distribution function of degrees across dealers. For comparison, we add the degree distribution of a random trading network (blue dashed line) and a scale-free trading network (black and red dashed lines). The plot reveals that municipal bond dealers are much more connected with each other than suggested by random trading (a random trading network yields a Poisson distribution of degrees). There is a large number of weakly connected dealers but also a significant number of highly connected dealers—forming the core of the dealer network. The municipal bond market thus exhibits features of a scale-free network.³

Figure 3 explores a different aspect of the dealer network by documenting the hierarchical structure of the market. A natural question is whether the municipal bond market has a single market center (and a periphery of loosely connected dealers) or several local market centers? To answer this question, we plot the degree distribution across dealers in the network (horizontal axis) against the clustering coefficient, or cliquishness, of each dealer (vertical axis). The plot reveals that sparsely connected dealers are part of local markets (highly clustered sub-markets), with order flow between the different local markets being maintained by a few dealer hubs.⁴ Figures 2 and 3 combined reveal that the municipal bond market has a hierarchical core-periphery structure.

Concentration of order flow with few dealers leads to more efficient aggregation of new information about asset values and yields economies of scale in transaction processing and risk management—lowering transactions cost. On the other hand, concentration may reduce financial market stability and resilience to shocks, increasing risk and costs. We next explore the market’s resilience to shocks to the network structure by means of network comparative statics. The type of shocks we are interested in are defaults by dealers, or market exit.

Figure 4 documents the effect on the network structure of default by individual dealers. Default is defined here as a situation in which all order flow to and from the dealer disappears, and the

³Scale-free networks are characterized by a power-law degree distribution. The probability that a node has k links follows $Pr(k) \sim k^{-\gamma}$, where γ is the degree exponent. The probability that a node is highly connected is statistically more significant than in a random graph. As a result, the network’s properties are often determined by a relatively small number of highly connected nodes, which are known as hubs. See Erdős and Renyi (1960) and Albert, Reka, and Barabasi (1999).

⁴Hierarchical 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 et al. (2002).

network is otherwise held fixed.⁵ We borrow this type of comparative statics analysis from the networks literature (Albert, Jeong, and Barabási (2000)). We plot the relative size of the largest connected subgraph (the so-called giant component) as a function of the number of dealers that default. We consider two scenarios. The blue line corresponds to the network connectedness when dealers default at random. The red line corresponds to the network connectedness when the most connected dealers default first. In Panel A, the horizontal axis measures the number of defaulted dealers as a fraction of all dealers. In Panel B, the defaulted dealers are sorted on the horizontal axis according to their degree. The figure suggests that the municipal bond market is remarkably robust to random and targeted defaults of dealer firms. The reason is that the market has at its center several highly connected dealer hubs that can act as substitutes for each other, diversifying the risk of instability.

[Figures 2, 3 and 4 about here]

2.3 How stable are inter-dealer relations and dealer ranks?

The nature of financial intermediation—standing ready to provide liquidity—leads dealers to interact repeatedly with other dealers. Do such repeat interactions lead to long-term relations between dealers that benefit themselves and other market participants? Are relations formed and broken opportunistically when one dealer’s inventory matches the needs of another dealer and gains from trade exist, or are dealer relations formed strategically? To address these questions, we explore how stable are relations between dealers over time and what impact this has on the persistence in dealer ranks.

Table 4 shows the transition probabilities of the individual inter-dealer relationships. Conditional on a directional (that is, buy vs. sell) inter-dealer relation that existed in one month, the same directional relation exists with 62% probability in the next month. Ignoring trade directions, the probability that two dealers who traded in one month also trade in the next month is 65%. This

⁵That is, we take the network structure as given and abstract from the market’s endogenous response to dealer defaults. The descriptive statistics should therefore be interpreted as features of the network rather than economic responses of the dealers or the market as a whole. By contrast, we would need an exogenous source of variation in the network structure to measure the economic impact of dealer defaults.

suggests a high level of persistence in the trading relations between dealers and in the direction of the order flow.

Table 5 focuses on the persistence in dealer ranks, as measured by the ordering of the dealers' centrality *Net*. Dealer ranks are highly persistent from one month to the next. The top 10 dealers remain at the top of their league with 93% chance. Yet, there is sizeable downside potential. The chance of losing ranks is on average twice as high as the chance of winning ranks, as captured in the columns Pr(Up) and Pr(Down) . This is consistent with the notion that peripheral, lower ranked dealers compete aggressively to gain ranks, that is, to become more central.

[Tables 4 and 5 about here]

Having shown that dealers vary substantially in the amount of trading they do with each other, we now study the link between dealers' trading relations and (local) market quality.

3 Order execution quality and dealer centrality

Dealers intermediate bonds through round-trip trades. We consider three types of round-trip trades with varying dealer involvement. The first type of round-trip is the simplest way dealers intermediate trades. A dealer purchases a bond from a customer and then sells the same bond to another customer, where the original bond lot is not split into smaller order sizes. That is, there is only one dealer involved. We call such transactions *CDC-Nonsplits*, where *CDC* indicates the bond went from *C*ustomer to *D*ealer and then to another *C*ustomer. There are a total of 3,332,104 *CDC-Nonsplit* round-trips in our sample.

Alternatively, the dealer purchasing the bond can split the bond lot into smaller sizes and sell each piece to a different customer. We call such transactions *CDC-Splits*, as there is still only one dealer involved. There are a total of 1,236,766 *CDC-Split* round-trips in our sample, for a total number of 4,568,870 *CDC* round-trips.

Dealers can, alternatively, involve other dealers in the intermediation by using the inter-dealer market. Such round-trips start with a dealer purchasing a bond from a customer, followed by one or several inter-dealer trades that move the bond from the head dealer to the tail dealer, and end with

sale from the tail dealer to a customer. In order to be able to trace the flow of a bond across the dealers with reasonable certainty, we consider only unsplit round-trips. We allow for a maximum of 6 dealers in the sequence of trades (there are very few cases involving 7 or more dealers). We call this type of round-trip *C(N)DC-Nonsplit*, where *(N)* indicates that multiple dealers may be involved. The *C(N)DC-Nonsplit* sample comprises 3,635,309 round-trips. Among these, 8.3% (or, a total of 303,205) involve more than one dealer.

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

According to this classification, we consider three samples for our empirical analysis. The baseline sample consists of all *CDC-Nonsplit* round-trips. For this sample, we are the most certain that the same dealer handles both the bid- and ask-side trades. The second sample includes all split orders, comprising 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 the *C(N)DC-Nonsplit* trades which include all round-lot transactions flowing through the dealership network.

3.1 Trading costs and dealer centrality

We can now relate order execution costs to trade types and dealer centrality. We measure trading costs by the markup on round-trip transactions charged by dealers. The dealers' markups are 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, scaled by the purchase price.

Table 6 and Figure 5 report descriptive statistics for dealer markups on round-trip transactions across trade categories. Markups are measured in percent of the dealer's original purchase price from customer. In reporting these numbers, we apply no data filters. For the regression analysis performed later, we winsorize the round-trip costs at 0.5% and 99.5%. Average round-trip trading costs on non-splits are 1.77% (1.66% at median), while dealers earn an average of 2.00% (2.01% at median) on split round-trips. Round-trip costs vary widely across transactions within category,

from about 0.73% at the lower quartile to 2.75% at the upper quartile of the distribution for *CDC* round-trips. Average markups decline monotonically with transaction size, as illustrated by Panel A of Figure 5 (see Green, Hollifield, and Schürhoff (2007) and Harris and Piwowar (2006)). Panel B of Table 6 shows that average markups increase monotonically with the number of dealers intermediating the trade. The total markup for round-trips involving six dealers peaks at 4.19%. This trading cost is roughly equal to the annual coupon payment on an average bond.

[Table 6 and Figure 5 about here]

Total round-trip costs when more than one dealer is involved reveal another feature. Panel B of Table 6 reports descriptive statistics for dealer markups on round-trip transactions with varying number of dealers involved. *CDC* are customer-dealer-customer transactions without inter-dealer trading. *CDDC*, *CDDDC*, *CDDDDC*, *CDDDDDC* and *CDDDDDDC* are round-trip pairs intermediated by two, three, four, five and, respectively, six dealers. We restrict the sample to non-splits, yielding 3,635,309 observations. Markups are again measured in percentage of the dealer's purchase price from customer. Trading costs rise with the degree of dealer involvement. Average round-trip costs are 1.77% when one dealer handles the bond lot. This number rises to 4.19% when six dealers intermediate before the bond reaches a customer. In the extreme, transactions involving more than four dealers exceed 7% costs in more than 5% of cases.

Having documented that total trading costs rise with the number of dealers involved, the natural question is how much does each dealer in the chain of intermediaries earn? This allows examining how the total surplus from financial intermediation is split among dealers. In particular, does the first, middle, or last dealer earn more than the remaining dealers?

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 round-trip transaction. We restrict the sample to non-splits. 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.

[Table 7 about here]

Dealer markups on round-trip transactions vary substantially across trades, as documented in Table 6. Are they systematically related to the centrality of the dealer intermediating the trade?

Figure 5, Panel B documents average trading costs across dealers for all *CDC-Nonsplit* transactions. Central dealers charge larger average trading costs than peripheral dealers across all trade sizes. Average bid-ask spreads differ by up to 80%. For medium-sized trades, bid-ask spreads at central dealers are 40% higher than the average, while they are 40% lower than the average at peripheral dealers. For small and large trades, the difference is smaller but still positive.

Table 8 documents using multivariate regressions how total round-trip trading costs depend on the network position of the intermediating dealer. Across columns, we vary the centrality measure and, respectively, the trade categories. In columns (1)-(4), the dealer centrality measure *Net* is defined as the first principle component of the equal-weighted centrality proxies. In columns (5)-(8), the dealer centrality measure *Net* is defined as the first principle component of the value-weighted centrality proxies. The regression samples are *CDC-Nonsplit*, All *CDC*, and *C(N)DC-Nonsplit*, respectively. The estimates are obtained from panel regressions with issuer fixed effects. Standard errors are adjusted for heteroskedasticity and clustering.

The estimates reveal that trading costs are positively related to dealer centrality and other measures of connectedness of dealers. Highly connected dealers, due to their central network position, are able to charge larger markups to investors than peripheral dealers. Being connected to a central dealer may offer advantages to investors in need of selling a bond. Yet, central dealers are in a better position to charge large spreads than peripheral dealers. Last, the *C(N)DC-Nonsplit* sample reveals that dealers' total profits are more sensitive to the network centrality of the tail dealer than to the head dealer.

[Table 8 about here]

3.2 Intermediation risk and dealer centrality

Central dealers charge larger markups than peripheral dealers, as documented in Figure 5 and Table 8. Are central dealers, in turn, taking on more risk? We next ask how variable are central

dealers' profits relative to those of peripheral dealers? The natural way to examine this is by looking at the probability of taking a trading loss on a round-trip transaction.

Table 9 documents the determinants of trading losses. In each of our three round-trip samples, dealers lose money in less than 2% of the round-trips. Still, the loss probability depends strongly on the dealers' relative position in the network, on bond characteristics, and other explanatory variables. We use a panel probit model with issuer fixed effects to estimate the determinants of dealers' losses. In each of the three samples, central dealers are less likely to lose on round-trips than peripheral dealers. Thus, the profits of connected dealers are larger on average and less risky.

[Table 9 about here]

3.3 Price efficiency across the dealer network

While Table 9 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 or processing skill. On the other hand, highly connected dealers, due to their central network position, observe a larger fraction of a bond's aggregate order flow and order flow in more bonds than peripheral dealers. Central dealers can therefore better filter liquidity motives for trade and aggregate fundamental information than dealers with little exposure to aggregate order flow. As a result, one would expect that bond prices are more efficient at central dealers.

To determine informational price efficiency, we adopt the Hasbrouck (1993) and Hotchkiss and Ronen (2002) settings and adjust their market quality measures to our setting with decentralized, infrequent trading. Hasbrouck's local trend model in which price movements are decomposed into permanent and transitory components provides a parametric estimate for market quality, MQ . Market quality is highest when prices are martingales so that price changes are uncorrelated. Autocorrelation in returns can therefore be viewed as a proxy for market quality. In Hasbrouck's parametric model, market quality is captured by the first-order lag in autocorrelation. In decentralized markets with infrequent trading, however, market imperfections can affect price dynamics beyond the first lag. We therefore construct a non-parametric market quality measure that is robust to return autocorrelation with an unspecified structure.

Time-series variation in returns can be decomposed into permanent and transitory components by quantifying predictable and, respectively, unpredictable variation. The standard estimator for return variance ignores autocorrelation in returns and measures total price variation:

$$\sigma_0^2 = \frac{1}{N} \sum_{n=1}^N (\Delta p_n)^2, \quad (1)$$

where N is the number of trading days in a given bond. In contrast to the standard estimator, the Newey and West (1987) HAC estimator for variance is robust to autocorrelated disturbances with an unspecified structure. A robust estimator for return variance captures the predictable variation through return autocorrelation up to lag $L \geq 1$:

$$\sigma_L^2 = \sum_{l=-L}^L w_l s_l^2, \quad (2)$$

where $w_l \in (0, 1]$ is a weight with $\sum_{l=-L}^L w_l = 1$ ($w_l = 1 - l/(L + 1)$ in Newey and West (1987)) and the autocovariance of lag l is defined as $s_l^2 = \frac{1}{N-|l|} \sum_{n=1+|l|}^N (\Delta p_n \Delta p_{n-|l|})$.

The difference between the standard and the robust estimator of return variance provides a non-parametric estimate of market quality by quantifying the autocorrelation in returns for any $L \geq 1$. Information efficiency can therefore be measured by

$$MQ_L = \sigma_L^2 / \sigma_0^2. \quad (3)$$

$MQ_L = 1$ corresponds to a situation in which all return movements are permanent, so that $\sigma_L^2 = \sigma_0^2$. The market is then considered perfectly informationally efficient. $MQ_L = 0$ corresponds to a situation in which all return movements are transitory and eventually reversed, so that $\sigma_L^2 = 0$. The market quality measure MQ_1 (that is, $L = 1$) corresponds to the model in Hasbrouck (1993) and Hotchkiss and Ronen (2002), where autocorrelation in returns is measured up to the first lag.

In our empirical specification, we compute four alternative MQ_L measures. We set $L = 1$ to be in line with the prior literature and use either transaction-by-transaction data or daily midpoint data. Alternatively, we set $L = 10$ (around $T^{\frac{1}{4}}$ as stipulated by Newey and West (1987)). The specifications with $L = 10$ capture general return autocorrelation patterns.

Table 10 documents the link between the informational efficiency of bond prices and dealer centrality. Information inefficiency is measured by MQ_L with $L = 1$ (columns 1-4) and, respectively, $L = 10$ (columns 5-8). MQ_L is calculated separately for each bond and then aggregated at the dealer level. Across columns, we vary the construction of the market quality measure MQ_L and the network centrality measure Net . MQ_L is computed either based on daily midpoint or trade-by-trade price changes. In the EW columns, the dealer centrality measure Net is defined as the first principle component of the equal-weighted centrality proxies. In the VW columns, the dealer centrality measure Net is defined as the first principle component of the value-weighted centrality proxies. The estimates are obtained from OLS regressions with year fixed effects. Standard errors are adjusted for heteroskedasticity and clustering at dealer level. We find consistently across specifications that prices are more efficient at central dealers. Central dealers thus seem to be better informed than peripheral dealers.

[Table 10 about here]

4 Liquidity provision across the dealer network

Bond dealers provide liquidity in two ways, by prearranging trades between customers (similar to a limit order) or taking bonds into inventory (similar to a market order). Highly connected dealers, due to their central network position, are better able to spread inventory risk across their neighbors than peripheral dealers. They can provide liquidity more efficiently to investors and afford greater inventory risk. We would therefore expect that central dealers hold larger and more volatile inventory, have longer inventory durations, and exhibit a lower propensity to prearrange trades than peripheral dealers. We now explore the relation between order flow volume, dealer inventory behavior, and dealer centrality.

Table 11 reports how aggregate order flow varies across the dealer network. We measure the dealers' order flow by the daily number of trades that a dealer conducts (columns 1 and 3) and, alternatively, by the dollar volume traded (columns 2 and 4). Across columns, we again vary the definition for the network measures between equal- and value-weighted cross-sectional averages. Central dealers trade more often and overall larger volume than peripheral dealers.

[Table 11 about here]

4.1 Inventory risk taking

We now turn to dealers' inventory risk taking behavior. Dealers' inventories are computed separately for each bond using all trades in bonds that occurs at least 90 days after the original sale date from the underwriter. We use two measures for inventory changes, absolute daily inventory changes in \$K and, alternatively, percentage absolute daily inventory changes in %. We define percentage absolute daily inventory changes as $|\Delta inv_t / \frac{1}{30} \sum_{i=1}^{30} inv_{t-i}|$.

Table 12, Panel A provides summary statistics for the dealers' inventory variables that we employ. Panel B documents the determinants of the variability in dealers' inventory. The sample consists of all dealer inventories in each bond issue on each day during the sample period. The estimates are obtained from panel regressions with issuer fixed effects. Standard errors are adjusted for heteroskedasticity and clustering. As one may expect, central dealers have more variability in their inventories than peripheral dealers, both in absolute and relative terms.

Table 13 explores the link between inventory durations and dealer centrality. Columns (1) and (2) in the table document the determinants for the duration of time that a bond (which is part of a round-trip transaction) spends in a dealer's inventory. We measure a bond's inventory duration by the number of days it takes for a dealer to find one or several customers to take the inventory. When bond lots are split, we compute a bond's inventory duration as the average number of days, weighted by the size of the split orders, that it takes to resell the entire bond lot. Inventory times exceeding 30 days are truncated. The sample consists of all *CDC* transactions. The estimates are obtained from panel regressions with issuer fixed effects. Standard errors are adjusted for heteroskedasticity and clustering. We find that central dealers hold bonds longer in inventory than peripheral dealers.

Central dealers can have longer inventory durations partly because they are less reluctant to take on inventory than peripheral dealers. To check whether this is the case we focus on prearranged trades, which are similar to limit orders in equity markets. In a prearranged trade, an investor indicates trading interest to the dealer. The dealer then searches for a counterparty. Only once two trading parties are found or likely to be found, the dealer takes the bond off the seller's hand. In this case, the dealer provides intermediation services without committing capital.

Columns (3) through (6) in Table 13 relate the propensity of prearranged trades to the centrality of the intermediating dealer. We identify prearranged trades in the data by the time stamps associated with the two legs of a round-trip. The columns entitled $\text{Pr}(\text{Immediate Match})$ identify prearranged trades as round-trip trades with same time stamp on the buy and sell trade. The columns entitled $\text{Pr}(\text{Same Day Match})$ only require that the buy and sell trade occur on the same calendar day. We estimate the model assuming the errors are normally distributed (probit). We find that central dealers are significantly less likely to intermediate prearranged trades than peripheral dealers, suggesting central dealers supply more liquidity than peripheral dealers.

[Tables 12 and 13 about here]

4.2 Order flow routing

Table 14 reports the average network centrality of each dealer in the intermediation chain for $C(N)DC\text{-}Nonsplit$ trades. We again measure dealer centrality by the first principle component of all equal-weighted network variables in Table 1, normalized by the ecdf transformation. As one would expect, inter-dealer trading occurs systematically through the assistance of central dealers. Consistently across rows, the dealers in the middle have more central network positions than either the dealer purchasing the bond from a customer or the dealer ultimately selling the bond to a customer. Dealer centrality peaks with the second dealer for all types of $C(N)DC\text{-}Nonsplit$ trades. Dealers at the end of the chain, the “tail” dealers, have higher levels of centrality than dealers at the beginning of the chain, the “head” dealers.

[Table 14 about here]

How does the complexity of dealer intermediation, as measured by the number of dealers in the sequence of trades, depend on the network centrality of the first dealer and last dealer (“head” and “tail”) in the trade sequence?

We explore this question in two ways. First, if we fix the identity of the head and tail dealers, we can calculate the frequency with which one through five inter-dealer trades happen between these dealers. Table 15, Panel A shows the summary statistics of the frequencies. In the overall sample, two dealers are connected through one inter-dealer trade with probability 44.8%, two inter-dealer

trades with probability 45.4%, conditional on two dealers being connected in the sample. For dealer pairs that interact more often (at least 10 trades during the sample period), the probability that they are connected by only one inter-dealer trade is 52.6%.

Second, we can look at all sequences of dealer trades and regress the number of dealers in the round-trip on the centrality of the head and tail dealers. Columns (1) and (2) in Panel B of Table 15 include sequences of trades that have at least one dealer, and columns (3) and (4) require the round-trip to have at least two dealers (i.e., head and tail). The complexity of intermediation is negatively related to the centrality of both the head dealer and the tail dealer.

[Table 15 about here]

4.3 Liquidity spillover

The question remaining is why central dealers take on more inventory risk? One explanation is that inventory risk sharing with connected dealers reduces inventory costs. If this is the case, we would expect to observe positive inventory spillovers across dealers.

Table 16 documents spillover effects in dealer inventories from connected dealers. The model we consider for the inventory decision y_i of dealer i is:

$$y_i = \alpha + \lambda \sum_{j \neq i} w_{ij} y_j + \beta' X_i + \varepsilon_i, \quad (4)$$

where w_{ij} equals the connection strength between dealers i and j , and X_i is a set of explanatory variables. The coefficient λ measures inventory spillovers across dealers. The dependent variables are constructed as average values over the sample period. The model is estimated using maximum likelihood (estimates from GMM/IV are similar and omitted). The estimate for λ is significantly positive in all specifications, suggesting strong positive inventory spillover effects. Large dealer inventories (levels, changes, and volatilities) cause connected dealers to increase their own inventories.

[Table 16 about here]

5 Strategic dealer pricing

The natural question deriving from the previous sections is why can central dealers charge larger and less variable spreads than peripheral dealers? One hypothesis is that the dealers' network position allows them to exploit centrality in bargaining with customers. To explore this hypothesis, we adapt the framework for intermediation in dealership markets developed in Green, Hollifield, and Schürhoff (GHS, 2007) to our network setting. When customers and dealers negotiate over the surplus from trade, the equilibrium markup on transaction i is determined as follows:

$$Markup_i = \alpha + \beta Net_i + \gamma' X_i + \epsilon_i + u_i, \quad (5)$$

where Net_i denotes the network centrality of the dealer intermediating trade i , X_i is a set of explanatory variables, ϵ_i is a normally distributed variable, and $u_i \geq 0$ is a one-sided error drawn from an exponential distribution with parameter $\lambda_i = \exp(\alpha_\lambda + \beta_\lambda Net_i + \gamma'_\lambda Z_i)$, where the Z_i denote a set of conditioning variables. We also allow the standard deviation of ϵ_i to be log-linear in Net_i and the conditioning variables Z_i : $\sigma_i = \exp(\alpha_\sigma + \beta_\sigma Net_i + \gamma'_\sigma Z_i)$.

In this model, markups can rise with dealer centrality for two reasons. Physical intermediation costs affect markups deterministically by alternating the intercept term in a stochastic frontier regression (captured by β). A dealer's bargaining position, by contrast, has a stochastic effect on markups. Better bargaining raises the mean of the one-sided random component in markups (captured by β_λ). This stochastic frontier model for dealers' round-trip markups can be estimated using maximum likelihood. The estimated parameters for the one-sided component can be interpreted as a measure for dealers' bargaining power.

Table 17 shows the estimation results for the stochastic frontier model. Across columns, we again vary the definitions of the network variables and the estimation sample. The estimates for the intermediation cost function in Panel A suggest that the dealers' intermediation cost rise with centrality, presumably because they keep larger inventory positions and have longer inventory durations. The estimates for the determinants of dealer bargaining power (one-sided error component variance) reveal that dealers' bargaining power is positively related to centrality. This is true for all centrality measures and data samples, with the exception of the $C(N)DC\text{-}Nonsplit$ sample which

includes round-trips involving more than one dealer. This evidence suggests the profits of highly connected dealers are large on average and little variable, because their central market position improves their bargaining position with customers.

[Table 17 about here]

6 Conclusion

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. The dealership network facilitates the sharing of inventory risk and the flow of information, but the concentration of order flow reduces resilience to shocks and allows central dealers to exploit their advantage when interacting with investors. We provide evidence for these tradeoffs in a comprehensive sample of trades in municipal bonds. We find that the dealership network exhibits a hierarchical core-periphery structure with around 30 highly interconnected dealers at its core and several hundred peripheral dealer firms. There is strong persistence in dealers' trading relations and in dealer ranks. Dealers' average markups increase with their network centrality. Central dealers charge up to 80% larger spreads than peripheral dealers, while facing a smaller probability of a trading loss. The informational efficiency of transaction prices rises with the centrality of the intermediating dealer. Dealers are exposed to significant liquidity spillovers from connected dealers. Central dealers provide more liquidity to customers than peripheral dealers; they trade more often and in larger aggregate volume. Central dealers also take more inventory risk. Dealers' bargaining position in trading with customers is stronger for central dealers than for peripheral dealers. These findings suggest that competition is fierce at the periphery but not at the core of the decentralized market because of opacity, search frictions, and network effects. 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

- Albert, Réka, Hawoong Jeong, and Albert-László Barabási, 2000, "Error and attack tolerance in complex networks," *Nature* 406, 378-382.
- Albert, Réka and Albert-László Barabási, 1999, "Emergence of Scaling in Random Networks," *Science* 286, 509-512.
- Bessembinder, Hendrik, William Maxwell, and Kumar Venkataraman, 2006, "Market transparency, liquidity externalities, and institutional trading costs in corporate bonds," *Journal of Financial Economics* 82, 251-288.
- Biais, Bruno, 1993, "Price formation and equilibrium liquidity in fragmented and centralized markets," *Journal of Finance* 48, 157-185.
- Cespa, Giovanni and Thierry Foucault, 2011, "Learning from Prices, Liquidity Spillovers, and Market Segmentation," CEPR Discussion Paper No. 8350.
- Chordia, Tarun, Richard Roll and Avanidhar Subrahmanyam, 2000, "Commonality in liquidity," *Journal of Financial Economics* 56, 3-28.
- Chowdhry, Bhagwan and Vikram Nanda, 1991, "Multimarket trading and market liquidity," *Review of Financial Studies* 4, 483-511.
- Comerton-Forde, Carole, Terrence Hendershott, Charles M. Jones, Pamela C. Moulton, and Mark S. Seasholes, 2010, "Time variation in liquidity: The role of market maker inventories and revenues," *Journal of Finance* 65, 295-331.
- Coughenour, Jay F. and Mohsen M. Saad, 2004, "Common market-makers and commonalities in liquidity," *Journal of Financial Economics* 73, 37-69.
- Edwards, Amy, Lawrence Harris, and Michael Piwowar, 2005, "Corporate bond markets transparency and transaction costs," *Journal of Finance* 62, 1421-1451.
- Erdős, Paul and Alfréd Rényi, "On the evolution of random graphs,"

- 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(2), 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?,” *The Journal of Finance* 65(5), 1669-1702
- Harris, Lawrence E. and Michael S. Piwowar, 2006, “Secondary Trading Costs in the Municipal Bonds 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.
- Hasbrouck, Joel and Duane J. Seppi, 2001, “Common factors in prices, order flows, and liquidity,” *Journal of Financial Economics* 59, 383-411.
- Hatch, Brian C. and Shane A. Johnson, 2002, “The impact of specialist firm acquisitions on market quality,” *Journal of Financial Economics* 66, 139-167.
- Hendershott, Terrence and Ananth Madhavan, 2011, “Click or Call? Auction versus Search in the Over-the-Counter Market,” UC Berkeley Working Paper.
- Hotchkiss, Edith S. and Tavy Ronen, 2002, “The Informational Efficiency of the Corporate Bond Market: An Intraday Analysis,” *Review of Financial Studies* 15(5), 1325-1354.
- Jovanovic, Boyan and Albert J. Menkveld, 2010, “Middlemen in limit order markets,” Working Paper, New York University.
- Milbourn, Todd T., 2003, “CEO reputation and stock-based compensation,” *Journal of Financial Economics* 68, 233-262.
- Newey, Whitney K. and Kenneth D. West, 1987, “A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica* 55(3), 703-708.
- Ravasz, E., A. L. Somera, D. A. Mongru, Z. N. Oltvai, and A.-L. Barabási, 2002, “Hierarchical organization of modularity in metabolic networks,” *Science* 297, 1551-1555.

Schultz, Paul, 2012, “The Market for New Issues of Municipal Bonds: The Roles of Transparency and Limited Access to Retail Investors,” University of Notre Dame Working Paper.

Yin, Xiangkang, 2005, “A Comparison of Centralized and Fragmented Markets with Costly Search,” *Journal of Finance* 60, 1567-1590.

Table 1: Description of dealer centrality measures

Variable	Description
Degree dg	<p>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. Degree captures the order flow and information to which a dealer is exposed, because it measures the fraction of dealers to which the dealer firm is connected. For directed graphs, one can calculate in-degree $dgin$ and, respectively, out-degree $dgout$. For weighted graphs, one can calculate weighted variants:</p> <p>$dgoutwntrade$ = 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.</p>
Eigenvector centrality ev	<p>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 weighted variants:</p> <p>$evwntrade$ = Eigenvector centrality, weighted by number of trades. $evwpar$ = Eigenvector centrality, weighted by total par amount.</p>
Betweenness bt	<p>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 the neighbors and the neighbors' neighbors. A dealer with a high degree of betweenness is in a critical position where a large flux of order flow and information pass through; they are called "hubs." We use the directed graph version of betweenness.</p>
Closeness cl	<p>A measure of influence with respect to centrality, rather than information or order flow. 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, $clout$ measures out-links only; $clin$ measures in-links only.</p>
K-core $kcore$	<p>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.</p>
Cliquishness cc	<p>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.</p>
Net	<p>Aggregate centrality measure of a dealer, computed as the first principle component of the above individual centrality measures and either equal- or value-weighted:</p> <p>Net (EW) = First principle component of equal-weighted network measures $dgout$, $dgin$, $kcoreout$, $kcorein$, bt, $clout$, $clin$, ev (cc is dropped since it requires at least two neighbors, reducing the number of observations).</p> <p>Net (VW) = First principle component of value-weighted network measures $dgoutwntrade$, $dginwntrade$, $dgoutwpar$, $dginwpar$, $evwntrade$, $evwpar$.</p>

Table 2: Summary statistics of dealer centrality measures

The table reports the descriptive statistics for the dealer centrality measures in the pooled dealer-day sample. In Panel A, we summarize the equal-weighted centrality measures. In Panel B, we summarize the order flow value-weighted measures. 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.

	Mean	S.D.	Min.	Max.
Panel A: Equal-weighted centrality measures				
<i>Net</i> (EW)	0.000	2.248	-1.721	18.868
<i>dgout</i>	0.014	0.033	0.000	0.338
<i>dgin</i>	0.014	0.027	0.000	0.272
<i>ev</i>	0.099	0.173	0.000	1.000
<i>bt</i>	0.002	0.007	0.000	0.169
<i>clout</i>	0.008	0.005	0.001	0.022
<i>clin</i>	0.005	0.002	0.001	0.014
<i>kcoreout</i>	4.812	7.178	0.000	31.000
<i>kcorein</i>	6.049	7.573	0.000	33.000
<i>cc</i>	0.505	0.289	0.000	1.000
Panel B: Value-weighted centrality measures				
<i>Net</i> (VW)	0.000	2.107	-0.445	60.536
<i>dgoutwntrade</i>	0.110	0.659	0.000	34.073
<i>dginwntrade</i>	0.110	0.621	0.000	31.515
<i>dgoutwpar</i>	17.62	93.786	0.000	4164.433
<i>dginwpar</i>	17.62	87.021	0.000	3061.083
<i>evwntrade</i>	0.011	0.066	0.000	1.000
<i>evwpar</i>	0.014	0.074	0.000	1.000

Table 3: Correlation matrix of dealer centrality measures

The table reports the correlations between the aggregate centrality measure Net and the individual centrality measures in the pooled dealer-day sample. The first set of columns report correlations between the raw measures. The second set of columns report correlations between the measures after standardization by the ecdf transformation. Across columns, Net (EW) is the equal-weighted aggregate centrality measure, and Net (VW) is the order flow value-weighted aggregate centrality measure. In Panel A, we summarize the equal-weighted centrality measures. In Panel B, we summarize the order flow value-weighted measures. 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.

	Raw centrality measures		Standardized centrality measures	
	Net (EW)	Net (VW)	Net (EW)	Net (VW)
Panel A: Equal-weighted centrality measures				
Net (EW)	1.00	0.62	1.00	0.89
$dgout$	0.95	0.69	0.90	0.80
$dgin$	0.96	0.65	0.85	0.78
ev	0.98	0.64	0.94	0.89
bt	0.74	0.54	0.86	0.78
$clout$	0.29	0.10	0.70	0.69
$clin$	0.11	0.03	0.14	0.24
$kcoreout$	0.88	0.43	0.90	0.80
$kcorein$	0.84	0.38	0.85	0.78
cc	-0.26	-0.17	-0.25	-0.22
Panel B: Value-weighted centrality measures				
Net (VW)	0.62	1.00	0.89	1.00
$dgoutwntrade$	0.57	0.89	0.87	0.86
$dginwntrade$	0.57	0.84	0.84	0.85
$dgoutwpar$	0.58	0.90	0.83	0.87
$dginwpar$	0.62	0.93	0.81	0.90
$evwntrade$	0.53	0.73	0.85	0.90
$evwpar$	0.57	0.85	0.82	0.92

Table 4: Stability in inter-dealer relations

The table 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.

Order flow this month	Order flow next month		Order flow in same direction next month	
	= 0	> 0	= 0	> 0
= 0	85.11%	14.89%	85.90%	14.10%
> 0	34.72%	65.28%	37.58%	62.42%

Table 5: Persistence in dealer ranks

The table 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 inter-dealer trades during the past 30 trading days.

		Rank month $t + 1$						Pr(Up)	Pr(Down)
		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	0.00	0.07
	11-20	0.07	0.78	0.14	0.00	0.00	0.00	0.07	0.15
	21-50	0.00	0.05	0.81	0.14	0.00	0.00	0.05	0.14
	51-100	0.00	0.00	0.08	0.79	0.13	0.00	0.08	0.13
	101-200	0.00	0.00	0.00	0.06	0.79	0.15	0.06	0.15
	>200	0.00	0.00	0.00	0.00	0.03	0.97	0.03	0.00

Table 6: Dealer markups on round-trip transactions

The table reports descriptive statistics for dealer markups on round-trip transactions of different types of trades. Agency trades in which dealers act as customers' agent instead of principle are eliminated. We first restrict the sample to round-trips with no inter-dealer trading. We call this data set the *CDC* sample. We then look at round-trips that involve no more than 6 dealers with no order splitting. We call this data set the *C(N)DC-Nonsplit* sample. All markups are measured in percentage of the (first) dealer's purchase price from customer.

	Obs	Mean	S.D.	Skew	Kurt.	5%	25%	50%	75%	95%
Panel A: Round-trips without inter-dealer trading										
All CDC	4,568,870	1.833	1.398	5.153	762.355	0.075	0.728	1.767	2.749	4.086
CDC-Nonsplit	3,332,104	1.769	1.387	4.943	686.578	0.065	0.647	1.659	2.668	4.061
CDC-Split	1,236,766	2.003	1.415	5.819	978.614	0.104	0.979	2.005	2.897	4.155
Panel B: Round-trips without order splitting										
All C(N)DC-Nonsplit	3,635,309	1.809	1.463	5.423	596.156	0.073	0.677	1.682	2.705	4.156
CDC-Nonsplit	3,332,104	1.769	1.387	4.943	686.578	0.065	0.647	1.659	2.668	4.061
CDDC-Nonsplit	165,506	1.896	1.666	5.602	285.597	0.152	0.755	1.601	2.714	4.603
CDDDC-Nonsplit	123,808	2.601	2.279	5.430	234.937	0.297	1.175	2.214	3.565	6.190
CDDDDC-Nonsplit	12,011	3.242	2.997	5.567	83.708	0.487	1.561	2.727	4.153	7.636
CDDDDDC-Nonsplit	1,721	3.408	5.061	8.776	164.380	0.099	1.244	2.522	4.403	9.524
CDDDDDDC-Nonsplit	159	4.194	5.335	4.669	30.986	0.410	1.797	2.824	4.700	11.547

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.

	Total markup	Dealer #1	Dealer #2	Dealer #3	Dealer #4	Dealer #5	Dealer #6
CDC	1.769	1.769 (100%)
CDDC	1.896	0.752 (40%)	1.144 (60%)
CDDDC	2.601	0.654 (25%)	0.652 (25%)	1.295 (50%)	.	.	.
CDDDDC	3.242	0.606 (19%)	0.532 (16%)	0.857 (26%)	1.247 (38%)	.	.
CDDDDDC	3.408	0.603 (18%)	0.362 (11%)	0.861 (25%)	0.425 (12%)	1.158 (34%)	.
CDDDDDDC	4.194	0.545 (13%)	0.463 (11%)	0.820 (20%)	0.706 (17%)	0.511 (12%)	0.511 (12%)

Table 8: Round-trip trading costs and dealer centrality

The table reports the determinants of round-trip trading costs. 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 *All CDC* sample includes all round-trips intermediated by a single dealer. *C(N)DC-Nonsplit* are round-trips intermediated by one or several dealers where the original bond lot is not split. The dealer centrality measure *Net* is the first principal component of the network variables in Table 1. The EW (VW) columns employ the equal-weighted (value-weighted) dealer centrality measures. For the *C(N)DC-Nonsplit* sample, *Net* is defined as the head or, alternatively, the tail dealer's centrality (indicated in the column header). The estimates are obtained from panel regressions with issuer fixed effects. Standard errors are adjusted for heteroskedasticity and clustering.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	CDC-Nonsplit		All CDC		All C(N)DC-Nonsplit			
					Head Dealer		Tail Dealer	
	EW	VW	EW	VW	EW	VW	EW	VW
<i>Net</i>	0.62*** (30.63)	0.54*** (25.82)	0.95*** (37.88)	0.87*** (33.87)	0.14*** (9.66)	0.15*** (8.85)	0.33*** (18.99)	0.22*** (12.32)
logpar_retail	-0.42*** (-177.78)	-0.42*** (-177.52)	-0.29*** (-116.02)	-0.29*** (-115.80)	-0.42*** (-176.25)	-0.42*** (-176.27)	-0.42*** (-176.60)	-0.42*** (-176.47)
logpar_medsize	-0.44*** (-244.30)	-0.44*** (-244.39)	-0.30*** (-148.85)	-0.30*** (-148.76)	-0.44*** (-244.08)	-0.44*** (-243.80)	-0.44*** (-243.70)	-0.44*** (-243.88)
logpar_lgsize	-0.37*** (-219.68)	-0.37*** (-222.28)	-0.29*** (-199.34)	-0.29*** (-202.15)	-0.37*** (-226.10)	-0.37*** (-227.02)	-0.37*** (-225.03)	-0.37*** (-226.82)
isgo	-0.03* (-1.90)	-0.03* (-1.88)	-0.04* (-1.90)	-0.04* (-1.89)	-0.04** (-1.99)	-0.04** (-1.99)	-0.04** (-2.00)	-0.04** (-1.99)
taxable	0.01 (0.29)	0.01 (0.35)	0.01 (0.54)	0.01 (0.62)	0.01 (0.39)	0.01 (0.42)	0.01 (0.46)	0.01 (0.45)
amt	0.21*** (10.99)	0.21*** (10.96)	0.24*** (10.84)	0.23*** (10.79)	0.21*** (11.18)	0.21*** (11.18)	0.21*** (11.21)	0.21*** (11.19)
Rating	0.00** (2.15)	0.00** (2.18)	0.00 (1.46)	0.00 (1.57)	0.00** (2.56)	0.00*** (2.63)	0.00*** (2.84)	0.00*** (2.74)
logamt	0.14*** (45.85)	0.14*** (45.69)	0.17*** (48.45)	0.17*** (48.23)	0.14*** (45.80)	0.14*** (45.76)	0.14*** (45.85)	0.14*** (45.77)
callable	0.32*** (31.36)	0.32*** (31.35)	0.40*** (33.14)	0.40*** (33.13)	0.33*** (32.09)	0.33*** (32.09)	0.33*** (32.09)	0.33*** (32.10)
cons	2.00*** (80.62)	2.08*** (82.41)	1.30*** (43.04)	1.37*** (45.07)	2.49*** (122.23)	2.47*** (110.14)	2.30*** (102.96)	2.41*** (106.68)
N	2,933,867	2,933,867	4,023,515	4,023,515	3,184,913	3,184,913	3,186,180	3,186,180

Table 9: Loss probability and dealer centrality

The table reports the determinants for the probability that dealers take a loss on a round-trip transaction. 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 *All CDC* sample includes all round-trips intermediated by a single dealer. *C(N)DC-Nonsplit* are round-trips intermediated by one or several dealers where the original bond lot is not split. The dealer centrality measure *Net* is the first principal component of the network variables in Table 1. The EW (VW) columns employ the equal-weighted (value-weighted) dealer centrality measures. For the *C(N)DC-Nonsplit* sample, *Net* is defined as the head or, alternatively, the tail dealer's centrality (indicated in the column header). The estimates are obtained from panel probit regressions with issuer fixed effects.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	CDC-Nonsplit		All CDC		All C(N)DC-Nonsplit			
					Head Dealer		Tail Dealer	
	EW	VW	EW	VW	EW	VW	EW	VW
<i>Net</i>	-0.98*** (-52.62)	-1.00*** (-54.04)	-1.02*** (-61.52)	-1.04*** (-63.24)	-0.72*** (-41.75)	-0.83*** (-47.65)	-0.84*** (-47.12)	-0.87*** (-48.88)
logpar_retail	0.11*** (31.37)	0.11*** (31.49)	0.09*** (28.35)	0.09*** (28.73)	0.10*** (31.24)	0.10*** (31.39)	0.10*** (31.03)	0.10*** (31.12)
logpar_medsize	0.13*** (56.59)	0.13*** (57.04)	0.11*** (55.53)	0.12*** (56.14)	0.13*** (57.63)	0.13*** (57.84)	0.13*** (56.86)	0.13*** (57.22)
logpar_lgsize	0.15*** (96.85)	0.16*** (98.80)	0.14*** (99.50)	0.14*** (101.51)	0.16*** (100.12)	0.16*** (101.45)	0.15*** (99.29)	0.16*** (101.04)
isgo	0.05*** (9.21)	0.06*** (9.33)	0.06*** (10.91)	0.06*** (11.11)	0.06*** (9.67)	0.06*** (9.85)	0.06*** (9.75)	0.06*** (9.90)
taxable	-0.09*** (-5.43)	-0.10*** (-5.59)	-0.09*** (-6.05)	-0.09*** (-6.22)	-0.09*** (-5.52)	-0.10*** (-5.76)	-0.09*** (-5.70)	-0.10*** (-5.84)
amt	-0.12*** (-8.51)	-0.12*** (-8.57)	-0.13*** (-10.56)	-0.13*** (-10.58)	-0.12*** (-8.81)	-0.12*** (-8.91)	-0.12*** (-8.88)	-0.12*** (-8.96)
Rating	0.00 (0.86)	0.00 (0.13)	0.00 (1.43)	0.00 (0.51)	0.00* (1.88)	0.00 (0.57)	0.00 (0.90)	0.00 (0.14)
logamt	-0.02*** (-15.71)	-0.02*** (-14.96)	-0.03*** (-23.70)	-0.03*** (-22.84)	-0.03*** (-17.01)	-0.02*** (-16.21)	-0.03*** (-16.66)	-0.02*** (-15.97)
callable	-0.13*** (-18.08)	-0.13*** (-18.10)	-0.16*** (-26.95)	-0.17*** (-27.05)	-0.13*** (-18.64)	-0.13*** (-18.67)	-0.13*** (-18.70)	-0.13*** (-18.70)
cons	-1.72*** (-78.09)	-1.70*** (-78.03)	-1.62*** (-82.63)	-1.60*** (-82.49)	-1.97*** (-95.71)	-1.86*** (-89.70)	-1.85*** (-87.17)	-1.82*** (-86.14)
N	2,958,219	2,958,219	4,060,191	4,060,191	3,219,080	3,219,080	3,220,393	3,220,393

Table 10: Informational price efficiency and dealer centrality

The table documents the determinants of price efficiency. Information efficiency is measured by the MQ_L measure described in Section 3. MQ_L is calculated separately for each bond using daily mid-price data; MQ_L is then collapsed at the dealer level. Across columns, we vary the construction of the market quality measure MQ_L and the network centrality measure Net . MQ_L is computed either based on daily midpoint or trade-by-trade price changes. In the EW columns, the dealer centrality measure Net is defined as the first principle component of the equal-weighted centrality proxies. In the VW columns, the dealer centrality measure Net is defined as the first principle component of the value-weighted centrality proxies. The estimates are obtained from OLS regression with year fixed effects. Standard errors are adjusted for heteroskedasticity and clustering at dealer level.

	$L = 1$				$L = 10$			
	Midpoint		Trade-by-trade		Midpoint		Trade-by-trade	
	EW	VW	EW	VW	EW	VW	EW	VW
Net	0.08*** (2.72)	0.12*** (3.84)	0.05*** (5.82)	0.07*** (7.31)	0.05 (1.33)	0.06 * (1.40)	0.08*** (7.42)	0.10*** (10.04)
cons	0.59*** (33.72)	0.60*** (34.04)	0.59*** (118.04)	0.59*** (122.88)	0.49*** (21.90)	0.49*** (21.59)	0.71*** (128.52)	0.72*** (136.46)
N	14,518	14,518	14,518	14,518	14,518	14,518	14,518	14,518

Table 11: Order flow and dealer centrality

The table documents the determinants of the daily order flow for each dealer. Number of trades and trade volumes are aggregated at dealer-day level. Trade volume is log-transformed. The dealer centrality measure Net is the first principal component of the network variables in Table 1. The EW (VW) columns employ the equal-weighted (value-weighted) dealer centrality measures. The estimates account for day fixed effect.

	Daily No. Trades		Daily Trade Volume	
	EW	VW	EW	VW
Net	78.29*** (399.75)	79.92*** (410.84)	8.21*** (1591.25)	8.36*** (1668.23)
cons	-24.17*** (-213.33)	-24.87*** (-221.20)	-1.55*** (-519.25)	-1.61*** (-556.55)
N	2,498,266	2,498,266	2,498,266	2,498,266

Table 12: Dealer inventory and centrality

The table documents the determinants of the variability in dealers' inventory. 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, $|\text{inv}_t / \frac{1}{30} \sum_{i=1}^{30} \text{inv}_{t-i}|$, truncated at 1,000 percent. Across columns, we vary the variable *Net* describing the dealers' network characteristics. The sample consists of all dealer inventories on each day. The estimates are obtained from panel regressions with issuer fixed effects. Standard errors are adjusted for heteroskedasticity and clustering.

Panel A: Descriptive statistics for dealer inventories

	Mean	S.D.	Min	Max	Obs.
Absolute daily inventory change [\$M]	0.503	5.814	0	5.577	2,381,530
Relative daily inventory change [%]	2.404	22.028	0	1,000	2,348,951

Panel B: Inventory variability

	Absolute daily inventory change [\$M]		Relative daily inventory change [%]	
	EW	VW	EW	VW
<i>Net</i>	2.39*** (5.63)	2.60*** (5.69)	4.32*** (10.46)	4.36*** (10.73)
cons	-0.74*** (-5.23)	-0.84*** (-5.35)	0.16 (1.10)	0.14 (0.97)
N	2,381,530	2,381,530	2,348,951	2,348,951

Table 13: Intermediation services and dealer centrality

The table documents the link between liquidity provision and dealer centrality. Columns (1)-(2) document the determinants of the duration of time (in terms of number of days) that a bond remains in a dealer's inventory. The sample consists of all customer-dealer-customer (CDC) transactions. The estimates are obtained from panel regressions with issuer fixed effects. Standard errors are adjusted for heteroskedasticity and clustering. Columns (3)-(6) document the determinants of the propensity of prearranged trades. We consider two types of prearranged trades, immediate and same day matches. Immediate matches have the same time stamp for dealer purchase and sale (columns (3) and (4)). Same day matches are round-trip transactions where the dealer purchase and sale occur on the same calendar day (columns (5) and (6)). The estimates account for year fixed effect. The dealer centrality measure *Net* is the first principal component of the network variables in Table 1. The EW (VW) columns employ the equal-weighted (value-weighted) dealer centrality measures.

	(1)	(2)	(3)	(4)	(5)	(6)
	Inventory Duration		Pr(Immediate Match)		Pr(Same Day Match)	
	EW	VW	EW	VW	EW	VW
<i>Net</i>	0.56*** (10.38)	0.33*** (6.08)	-1.58*** (-140.89)	-1.64*** (-148.68)	-1.58*** (-167.42)	-1.47*** (-159.81)
logpar_retail	-0.29*** (-56.66)	-0.29*** (-56.56)	0.16*** (89.98)	0.16*** (90.63)	0.04*** (36.35)	0.04*** (35.99)
logpar_medsized	-0.35*** (-85.80)	-0.35*** (-86.05)	0.15*** (119.66)	0.15*** (121.18)	0.10*** (126.06)	0.10*** (127.50)
logpar_lgsized	-0.26*** (-77.03)	-0.27*** (-77.92)	0.10*** (103.28)	0.11*** (108.27)	0.12*** (180.17)	0.13*** (185.87)
isgo	-0.03 (-1.08)	-0.03 (-1.06)	-0.00* (-1.84)	-0.00 (-1.52)	-0.00 (-1.00)	-0.00 (-0.78)
taxable	-0.56*** (-14.05)	-0.56*** (-14.12)	0.26*** (32.13)	0.26*** (31.61)	0.29*** (50.74)	0.29*** (50.28)
amt	-0.34*** (-9.04)	-0.34*** (-9.05)	0.05*** (10.20)	0.05*** (9.99)	0.26*** (79.57)	0.26*** (79.32)
Rating	0.00*** (6.30)	0.00*** (6.17)	0.00*** (36.58)	0.00*** (33.53)	0.00*** (33.35)	0.00*** (31.37)
logamt	-0.37*** (-76.17)	-0.37*** (-76.21)	0.06*** (80.28)	0.06*** (82.75)	0.10*** (220.71)	0.11*** (223.19)
callable	-0.09*** (-4.68)	-0.09*** (-4.66)	-0.08*** (-19.01)	-0.08*** (-19.03)	0.01*** (3.61)	0.01*** (3.61)
cons	3.37*** (54.58)	3.59*** (57.65)	-0.59*** (-46.02)	-0.55*** (-43.68)	0.60*** (59.72)	0.50*** (50.41)
N	2,929,570	2,929,570	2,917,162	2,917,162	2,917,162	2,917,162

Table 14: Inter-dealer trading and dealer centrality

The table reports the average network centrality for each dealer in the intermediation chain. Agency trades in which dealers act as customers' agent instead of principle are eliminated. We restrict the sample to round-trips that involve no more than 6 dealers and no order splitting, the $C(N)DC\text{-}Nonsplit$ sample. We measure dealer centrality by the first principle component $pca1$ of the centrality proxies described in Table 1, standardized by the empirical cdf.

	Dealer #1	Dealer #2	Dealer #3	Dealer #4	Dealer #5	Dealer #6
CDC	0.949
CDDC	0.880	0.919
CDDDC	0.880	0.977	0.901	.	.	.
CDDDDC	0.813	0.972	0.949	0.884	.	.
CDDDDDC	0.845	0.975	0.934	0.959	0.875	.
CDDDDDDC	0.862	0.972	0.946	0.929	0.948	0.883

Table 15: Order flow routing and dealer centrality

The table documents path lengths between pairs of dealers and relates them to the dealers' network centrality. For each round-trip intermediation chain with at least two dealers involved, the path length between them, measured by the number of inter-dealer transactions, is calculated (ranging from one to five). In Panel A, we report the frequency distribution of the path length for each possible pairing of the head and tail dealer. In Panel B, we report how the frequency of each path length (dependent variable) is related to the network centrality of head and tail dealer. The dependant variable is the number of dealer involved in the round-trip, with values from 1 to 6. The regressor *CNet* (Head Dealer) denotes the cumulative ranking of the first principle component of dealer network variables of the first dealer in the chain. The regressor *CNet* (Tail Dealer) denotes the cumulative ranking of the first principle component of the dealer network variables of the last dealer in the chain. Specifications (1) and (2) include all round-trips. Specifications (3) and (4) restrict the sample to round-trips with at least 2 dealers, so we can determine a head and a tail dealer. Specifications (1) and (3) are estimated using panel regressions with issuer fixed effects. Standard errors are adjusted for heteroskedasticity and clustering at the issuer level. Specifications (2) and (4) are estimated using Poisson regressions with robust standard errors clustered at the issuer level.

Panel A: Average frequencies of different path lengths

	N	1 Step	2 Steps	3 Steps	4 Steps	5 Steps
All	39,183	0.448	0.454	0.085	0.012	0.001
Frequency of trade between pair of dealers:						
< 4	28,995	0.422	0.464	0.099	0.014	0.001
4 – 9	5,951	0.523	0.423	0.046	0.007	0.001
≥ 10	4,237	0.526	0.426	0.041	0.007	0.001

Panel B: Number of dealers involved and network properties of head and tail dealers

	(1) FE regression	(2) Poisson regression	(3) FE regression	(4) Poisson regression
<i>CNet</i> (Head Dealer)	-0.94*** (-95.63)	-0.69*** (-97.10)	-0.21*** (-20.50)	-0.08*** (-20.90)
<i>CNet</i> (Tail Dealer)			-0.45*** (-33.02)	-0.17*** (-35.15)
logpar_retail	0.01*** (17.40)	0.01*** (19.38)	0.00** (2.37)	0.00** (1.97)
logpar_medsize	0.01*** (28.58)	0.01*** (34.32)	-0.00** (-2.25)	-0.00*** (-5.04)
logpar_lgsize	-0.00*** (-6.63)	0.00 (0.92)	0.01*** (5.07)	0.00** (2.09)
isgo	0.00 (0.99)	0.00*** (3.24)	-0.00 (-0.27)	-0.02*** (-10.86)
taxable	-0.04*** (-10.41)	-0.02*** (-8.39)	-0.03* (-1.90)	-0.01** (-2.32)
amt	-0.02*** (-5.37)	-0.01*** (-2.99)	0.00 (0.07)	0.01*** (3.17)
Rating	-0.00*** (-4.54)	-0.00** (-1.97)	0.00*** (8.46)	0.00*** (16.31)
logamt	-0.01*** (-18.57)	-0.01*** (-21.12)	-0.00*** (-3.57)	-0.00** (-2.09)
callable	-0.00 (-0.58)	0.00 (0.61)	0.07*** (10.92)	0.02*** (8.76)
cons	2.01*** (206.85)	0.76*** (108.88)	3.03*** (165.02)	1.12*** (179.60)
N	3,226,416	3,226,416	260,366	260,366

Table 16: Inventory spillovers across dealers

The table documents the extent of spillover effects in dealer inventories from connected dealers. The model we consider for the inventory decision y_i of dealer i is:

$$y_i = \alpha + \lambda \sum_{j \neq i} w_{ij} y_j + \beta' X_i + \varepsilon_i,$$

where w_{ij} equals the connection strength between dealers i and j , and X_i is a set of explanatory variables. The dependent variables are constructed as average values over the sample period. The model is estimated using maximum likelihood (estimates from GMM/IV are similar and omitted).

	Inventory	No. Trades	Inventory	SD(Inventory)
λ	0.10 (6.26)	0.04 (3.22)	0.09 (6.45)	0.13 (8.71)
Constant	6.61 (0.83)	2.50 (1.05)	6.04 (0.97)	-2.12 (-1.49)
σ	242.57 (25.21)	51.29 (25.21)	185.84 (25.21)	37.95 (25.21)
N	1,271	1,271	1,271	1,271

Table 17: Dealer bargaining power and centrality


The table documents the determinants of dealers' bargaining power. 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 *All CDC* sample includes all round-trips intermediated by a single dealer. *C(N)DC-Nonsplit* are round-trips intermediated by one or several dealers where the original bond lot is not split. The dealer centrality measure *Net* is the first principal component of the network variables in Table 1. The EW (VW) columns employ the equal-weighted (value-weighted) dealer centrality measures. For the *C(N)DC-Nonsplit* sample, *Net* is defined as the head or, alternatively, the tail dealer's centrality (indicated in the column header). The estimates are obtained from stochastic frontier regressions.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	CDC-Nonsplit		All CDC		All C(N)DC-Nonsplit			
					Head Dealer		Tail Dealer	
	EW	VW	EW	VW	EW	VW	EW	VW
Panel A: Intermediation cost function								
<i>Net</i>	0.39*** (47.44)	0.56*** (62.94)	0.11*** (17.98)	0.22*** (30.92)	0.28*** (39.62)	0.49*** (61.48)	0.34*** (45.37)	0.47*** (60.83)
logpar_retail	-0.50*** (-379.54)	-0.50*** (-379.20)	-0.03*** (-15.33)	-0.03*** (-14.67)	-0.49*** (-368.73)	-0.48*** (-370.63)	-0.49*** (-370.57)	-0.48*** (-370.25)
logpar_medsize	-0.46*** (-490.93)	-0.46*** (-488.70)	-0.27*** (-243.11)	-0.27*** (-241.86)	-0.46*** (-500.34)	-0.45*** (-499.85)	-0.46*** (-501.02)	-0.45*** (-499.87)
logpar_lgsize	-0.32*** (-456.13)	-0.32*** (-455.91)	-0.20*** (-249.02)	-0.20*** (-248.33)	-0.32*** (-469.37)	-0.32*** (-470.12)	-0.32*** (-469.53)	-0.32*** (-469.92)
isgo	-0.07*** (-80.91)	-0.07*** (-80.19)	-0.08*** (-97.43)	-0.08*** (-97.08)	-0.07*** (-84.20)	-0.07*** (-83.38)	-0.07*** (-84.08)	-0.07*** (-83.42)
taxable	0.07*** (23.16)	0.07*** (23.95)	0.09*** (36.73)	0.09*** (37.18)	0.07*** (25.11)	0.08*** (25.91)	0.07*** (25.15)	0.08*** (25.80)
amt	0.23*** (108.75)	0.23*** (108.56)	0.21*** (118.96)	0.21*** (118.75)	0.23*** (112.16)	0.23*** (112.10)	0.23*** (112.25)	0.23*** (112.11)
Rating	0.00*** (45.71)	0.00*** (47.10)	0.00*** (46.33)	0.00*** (47.68)	0.00*** (47.88)	0.00*** (50.10)	0.00*** (49.53)	0.00*** (50.11)
logamt	0.03*** (101.60)	0.03*** (97.85)	0.04*** (168.92)	0.04*** (165.45)	0.03*** (101.03)	0.03*** (97.56)	0.03*** (100.55)	0.03*** (97.61)
callable	0.16*** (111.28)	0.16*** (111.73)	0.20*** (137.91)	0.20*** (137.92)	0.16*** (116.75)	0.16*** (117.23)	0.16*** (116.91)	0.16*** (117.25)
Panel B: Dealer bargaining power (One-sided error component)								
<i>Net</i>	0.50*** (20.53)	0.15*** (6.33)	2.64*** (99.36)	2.21*** (77.89)	-0.14*** (-7.64)	-0.39*** (-20.11)	-0.01 (-0.54)	-0.33*** (-16.91)
logpar_retail	0.17*** (55.14)	0.16*** (53.08)	-0.59*** (-89.07)	-0.60*** (-88.39)	0.15*** (50.33)	0.15*** (48.92)	0.15*** (50.57)	0.15*** (49.05)
logpar_medsize	0.01*** (5.42)	0.00* (1.74)	-0.12*** (-38.06)	-0.13*** (-39.19)	0.01*** (4.98)	0.00** (2.01)	0.01*** (5.22)	0.01*** (2.49)
logpar_lgsize	-0.16*** (-81.38)	-0.17*** (-85.48)	-0.23*** (-97.40)	-0.24*** (-99.48)	-0.17*** (-88.03)	-0.17*** (-90.81)	-0.16*** (-87.88)	-0.17*** (-90.57)
Panel C: Resale price risk (Symmetric error component)								
<i>Net</i>	-0.24*** (-8.41)	0.20*** (6.54)	-1.17*** (-83.93)	-1.07*** (-78.38)	-0.27*** (-11.30)	0.11*** (3.92)	-0.26*** (-10.19)	0.10*** (3.72)
logpar_retail	-0.80*** (-201.79)	-0.80*** (-204.08)	0.15*** (54.49)	0.14*** (53.41)	-0.76*** (-197.17)	-0.76*** (-199.87)	-0.76*** (-198.82)	-0.76*** (-199.59)
logpar_medsize	-1.05*** (-339.22)	-1.04*** (-339.75)	-0.64*** (-324.07)	-0.64*** (-324.55)	-1.01*** (-351.34)	-1.01*** (-352.35)	-1.01*** (-351.70)	-1.01*** (-351.81)
logpar_lgsize	-0.65*** (-252.18)	-0.65*** (-252.97)	-0.45*** (-275.39)	-0.46*** (-275.68)	-0.64*** (-261.44)	-0.64*** (-262.07)	-0.64*** (-261.65)	-0.64*** (-261.87)
N	2,933,867	2,933,867	4,023,515	4,023,515	3,184,913	3,184,913	3,186,180	3,186,180

Figure 1: Dealer network topology

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.

Panel A: Order flow among most active dealers



Panel B: Order flow in entire network




Figure 2: Market connectedness

The figure documents the connectedness of the market. We plot the inverse distribution function for the degree across dealers in the network. The black dots correspond to the out-degree, the red dots represent in-degrees. For comparison, we add the degree distribution of a random trading network (blue dashed line) and a scale-free trading network (black and red dashed lines).




Figure 3: Market hierarchy

The figure documents the 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).





Figure 4: Market resilience

The figure documents the effect on the network structure of default by dealers. We plot the relative size of the largest connected subgraph (so-called giant component) as function of the number of dealers that default. We consider two scenarios. The blue line corresponds to the network connectedness when dealers default at random. The red line corresponds to the network connectedness when the most connected dealers default first. In Panel A, the horizontal axis measures the number of defaulted dealers as a fraction of all dealers. In Panel B, the defaulted dealers are sorted on the horizontal axis according to their degree.

Panel A: Market connectedness against % dealer defaults



Panel B: Market connectedness against degree of defaulted dealers





Figure 5: Trading costs and dealer centrality

The figure documents trading cost and volume by dealer and size. The sample consists of all *CDC-Nonsplit* transactions.

Panel A: Trading cost and volume by size



Panel B: Trading cost by dealer and size

