Dealer Quid Pro Quo in the Municipal Bond Market*

Casey Dougal[1], Daniel A. Rettl[2], and Vasiliy Yakimenko[2]

[1]Florida State University
[2]University of Georgia

May 30, 2025

Abstract

Dealers intermediate trades in OTC markets through trading networks. In the municipal bond market, we document greater complexity than the typical core-periphery structure. Analyzing dealer reciprocity-the tendency to repay favors-we find reciprocity generally reduces markups. Dealers trade lower markups today for future liquidity. However, in small trading communities, reciprocity can foster collusion via quid-pro-quo agreements, inflating transaction chain markups. Among high-centrality dealers in large communities, high reciprocity lowers average markups by 80 basis points, while among low-centrality dealers in small communities, it raises markups by 72 basis points. Although only around 2% of transaction chains suggest collusive behavior, these significantly affect regression results, highlighting the importance of controlling for such outliers to accurately estimate centrality premiums or reciprocity discounts.

*Contact Information: Casey Dougal (cdougal@fsu.edu); Daniel Rettl (daniel.rettl@uga.edu); Vasiliy Yakimenko (vasiliy.yakimenko@uga.edu)

Dealers play a pivotal role in most OTC markets, connecting buyers and sellers, providing liquidity, and influencing price discovery through their trading networks. The decentralized nature of these markets complicates the acquisition of information on assets and counterparties, but dealer networks help address this challenge by pooling expertise and enhancing transparency, which leads to more accurate pricing and efficient trade matching (Brancaccio et al. (2017), Hendershott et al. (2020)). However, these networks also have potential downsides. Close ties within them can foster collusion, distorting prices and undermining competition.[1] Central to both cooperative and collusive behaviors is dealer reciprocity, which involves the mutual exchange of favors based on obligation or expectations of future repayment.[2] Reciprocity can lead to mutually beneficial transactions and stronger partnerships, promoting trust and operational efficiency. Conversely, it can also facilitate collusion, where dealers coordinate to limit competition and increase collective profits, ultimately harming market efficiency and raising prices for end buyers.

In this study, we investigate the role that dealer reciprocity plays in dealer networks and, in particular, the pricing of assets passing through these networks via inter-dealer transaction chains.[3] We hypothesize that reciprocity will manifest itself both positively as dealers trade lower markups today for future liquidity and lower markups tomorrow; and negatively as some dealers collude through quid pro quo arrangements to inflate transaction chain markups. We further hypothesize that the trade outcome will be a function of the dealer's

[1] We use "collusion" broadly to include informal, tacit coordination as well as explicit, potentially illegal price-fixing. No claim is made that the conduct we document necessarily violates antitrust or securities law.

[2] Or in the words of Sobel (2005),"Reciprocity refers to a tendency to respond to perceived kindness with kindness and perceived meanness with meanness and to expect this behavior from others....the hypothesis that reciprocity is an instrumental motivation for human behavior is overwhelming." Dealer reciprocity may be preference-based or arise as a means for dealers to "sustain a profitable long-term relationship or to obtain a (profitable) reputation for being a reliable associate" (Sobel, 2005). In either case, reciprocity gives way to both cooperation and collusion. Gouldner (1960) explores the role of reciprocity as a foundational norm that supports social stability, while Blau (1964) examines how reciprocal exchanges shape social structures and cohesion, implicitly addressing the formation and stability of networks through trust and favors.

[3] Transaction chains are a common feature of OTC markets wherein a dealer will buy an asset from a seller and then the asset passes through dealer-to-dealer transactions until it is placed with the ultimate buyer. Each successive transaction garners a price markup on the asset making the ultimate price paid by the buyer potentially much higher than if they were able to purchase the asset directly without the inter-dealer transaction chain.

local network size, with the latter outcome obtaining only in small communities where collusion is more easily sustained.[4] Similarly, we expect the outcome of reciprocity to depend on a dealer's centrality within the network. Dealers who occupy central positions and have more connections will be able to more efficiently leverage positive reciprocity to secure future liquidity or favorable pricing, whereas those in more peripheral positions with effectively less oversight will be more willing to resort to collusion to gain competitive advantages.

The empirical setting for our study is the US municipal bond market. This market consists of debt issued by states, counties, or municipalities to finance public projects such as schools, roads, bridges, hospitals, or other infrastructure developments. It is the US's largest and most important source of capital for infrastructure investments, with outstanding debt as of 2024 of about $4 trillion. Municipal bonds trade over the counter. Daily trading volume in the US averages $12.5 billion, with about a sixth of this volume attributed to inter-dealer trades.[5] The Municipal Securities Rulemaking Board (MSRB) has expressed concerns about transaction chain markups, noting that while each dealer in the chain does not generate excessive profits individually, the cumulative markup can be substantial.[6] Schultz (2012) confirms these concerns and finds longer transaction chains result in higher prices for retail buyers, while Piwowar (2007) argues that inter-dealer trading can create opportunities for quid pro quo arrangements, where dealers might favor certain counterparties in exchange for future benefits.

Municipal bond dealers typically operate on a principal basis, purchasing bonds from sellers and holding them until a buyer is found, with sales often precipitated for liquidity reasons such as funding a major purchase (Green et al., 2007)[7][8]. Once a bond is bought, the

[4]The notion that coordination becomes more challenging as the number of participants in collusion schemes increases has been addressed in the theoretical IO literature (e.g. Stigler, 1964; Tirole, 1988). For example, Stigler (1964) shows the incentive for deviation from collusion rises with the number of competitors in oligopolistic markets, suggesting that larger groups may face greater difficulties in maintaining collusive agreements.

[5]`https://www.sifma.org/resources/research/us-municipal-bonds-statistics/`

[6]`https://www.sec.gov/files/rules/sro/msrb/2012/34-66625.pdf`

[7]"A well-known adage in the municipal business is that bonds are not bought, they are sold. That is, investors seldom approach dealers to purchase a specific bond." (Schultz 2012)

[8]In the municipal bond market, principal and agency trades represent two distinct approaches to trans-

dealer then has three options: retain the bond in inventory, sell it to a customer, or sell it to another dealer, with retention involving risk and sales requiring a search. Prior research has found that the municipal bond dealer network has a core-periphery structure, where central "hub" dealers are more efficient at finding buyers, leading to faster trades but higher markups compared to peripheral dealers, thus presenting a trade-off between execution speed and cost (Li and Schürhoff, 2019).

In this context, consider a reciprocal dealer that maintains a high degree of mutuality with its network of customers and other dealers. This dealer will also have lower search costs since others in their network know that if they offer liquidity now, then the tacit expectation is that the next time they need a buyer this dealer will be ready and willing to provide them liquidity at a good price. We measure a dealer's reciprocity using their node reciprocity, which is a network-based measure calculated as the ratio of bidirectional connections to all possible connections within a dealer's network over the past 30 days. Higher values of node reciprocity indicate a greater tendency for nodes to reciprocate interactions within the network, reflecting stronger mutual relationships. Consistent with our hypotheses, we find that after controlling for a dealer's centrality, higher dealer reciprocity generally leads to lower bond markups. For example, the average round-trip transaction chain markup for chains where the first dealer is a central dealer[9] is 1.99% for central dealers with low reciprocity and 1.19% for central dealers with high reciprocity, a statistically significant difference of 80 basis points. We find the negative relationship between reciprocity and markups holds controlling for bond and dealer characteristics, including dealer centrality (Li and Schürhoff, 2019), past dealer relationships (Di Maggio et al., 2017), and measures of dealer market share (Griffin

action execution. In principal trades, the dealer acts as the counterparty, buying and holding the bonds in their own inventory before selling them to clients. This allows the dealer to take on the price risk of the bond and often leads to the inclusion of a markup in the final sale price, which may not be explicitly disclosed to the client. In contrast, agency trades involve the dealer facilitating the transaction between the buyer and seller without taking ownership of the bond. The dealer acts as an intermediary and charges a commission for the service, with fees clearly disclosed to the client. The main difference lies in the dealer's risk exposure and the transparency of costs for the investor, with principal trades potentially obscuring hidden costs, while agency trades provide more clarity on fees.

[9]A dealer in the upper tercile of eigenvector centrality.

et al., 2023).

Dealer-level reciprocity is typically associated with lower cumulative markups, yet we identify a small cohort of highly reciprocal, peripheral dealers that systematically charges notably higher markups, on average 72 basis points more than their less-reciprocal counterparts. The evidence suggests that tacit collusion within these tightly knit networks inflates trading costs: these dealers operate in networks roughly one-sixth the size of other peripheral dealers, but their transaction chains are nearly twice as long, indicating repeated trading within the same closed circle. Their transactions also feature slightly smaller trade sizes and bonds with more complex contractual features, consistent with a strategic focus on retail investors, who are more susceptible to cognitive biases or limited financial literacy.[10]

We also observe what appears to be a higher incidence of anomalous trading and strategic pricing by these dealers. We find their deals have an unusually high incidence of "round-trip" chains in which the initiating and concluding dealers are identical, raising concerns about pump-and-dump strategies that shift costs onto retail investors.[11] We also find that trades routed through low-centrality, high-reciprocity networks are far more likely to employ coarse price increments (e.g., eighths or whole numbers) and less likely to use fine yield quotations, consistent with higher markups and reduced transparency for retail investors (Griffin et al., 2023). Taken together, these findings suggest that peripheral, high-reciprocity dealers use strategic behavior to obscure true transaction costs from less sophisticated retail clients.

An alternative explanation for our findings could be that the higher markups and longer transaction chains observed among highly reciprocal, peripheral dealers stem from the nature of the bonds they trade, specifically, illiquid or complex securities that demand specialized

[10]For example, see Harris and Piwowar (2006); Green et al. (2007); Edwards et al. (2007); and Brancaccio and Kang (2022) which treat small trade size as a proxy for retail participation and bond complexity as a source of cognitive strain.

[11]For example, suppose a bond is trading at 100 and dealers charge a 1% markup. Consider two scenarios. Scenario 1: Dealer 1 buys the bond for 100 and resells it for 101 to a customer. Dealer 1 makes \$1. Scenario 2: Dealer 1 buys the bond for 100 and resells it to Dealer 2, who then resells it to the customer. Here, Dealer 1 makes \$1 and Dealer 2 makes \$1.1, and the customer pays \$1.1 more than in Scenario 1. Importantly, even though Dealer 1 makes the same in both scenarios, Dealer 1 is better off in Scenario 2 since he is now owed a favor from Dealer 2. Thus, the customer is worse off in Scenario 2, while Dealer 1 and Dealer 2 are better off. For other examples, see Piwowar (2007).

expertise and incur higher transaction costs. These dealers may operate in niche markets where pricing and trading are more challenging, leading to longer chains and justified higher markups due to increased risk, effort, and the value of their specialized services. Moreover, the strong relationship-based trading within smaller networks could account for the repeated interactions, potentially reflecting trust and reliability rather than opportunistic behavior. This seems less likely given the large number of bond characteristics we control for in our regression analysis. However, we find that our results persist even after controlling for bond fixed effects, indicating that bond characteristics alone do not fully explain the observed patterns.

Consistent with our hypothesis that reciprocal dealer behavior is driven by the expectation of quid pro quo, our analysis finds that past dealer reciprocity predicts future quid pro quo behavior. Specifically, we show that reciprocal dealers are more likely to engage in reverse or reciprocating trades within one day. For instance, if dealer 1 sells to dealer 2 today, the probability that dealer 2 sells back to dealer 1 in the near future, what we define as a reverse or reciprocal trade, significantly increases. A standard deviation increase in dealer reciprocity raises the likelihood of a reverse trade by 7% within the day following the initial trade. This effect is more pronounced among low-centrality, high-reciprocity dealers, where the probability of a reciprocal deal occurring within a day is 34% higher than for those in the middle-centrality, middle-reciprocity group. Additionally, consistent with the idea of reciprocal dealers returning favors, we observe that transaction chains involving reciprocal dealers are longer and less likely to result in direct customer placements. We attribute this to the inclusion of extra dealers in the chain as part of favor exchanges. A standard deviation increase in dealer reciprocity corresponds to a 6% decrease in the probability of placing the bond with an ultimate customer and a 9% increase in chain length, with these effects being even stronger among low-centrality, high-reciprocity dealers.

We further explore the role of reciprocity based on the size of a dealer's local trading network, hypothesizing that smaller networks facilitate collusive behavior. Our findings confirm that dealer reciprocity positively correlates with transaction chain markups in the smallest quintile of local networks, while larger networks exhibit the opposite trend. Using machine learning techniques, such as the Louvain community detection algorithm, we identify dealer trading communities and find that when both buying and selling dealers belong to the same community, markups are significantly higher, ranging from 50 to 60 basis points. However, as the community size grows, the magnitude of these markups declines, reinforcing the idea that small, close-knit groups of dealers may engage in collusion to the detriment of customers. Further analysis reveals that reciprocity is a key factor in shaping these dealer communities.

Our paper contributes to the empirical literature on OTC markets, with a focus on municipal bond markets. While existing research often explores core-periphery network structures in OTC markets[12], our study unveils additional complexities within the municipal bond market network. Specifically, we demonstrate that factors like node reciprocity and trading communities contribute to price dispersion across the dealer network. Our findings suggest broader implications for OTC markets beyond municipal bonds. Notably, existing models have yet to explicitly incorporate the role of dealer reciprocity in dealer networks.[13]

In particular, our results highlight the importance of controlling for reciprocity when assessing the relationship between dealer centrality and markups. For high-reciprocity firms with low centrality, we observe a centrality discount, while elsewhere we find a centrality premium. This distinction may explain why previous studies have reported mixed results regarding the effect of dealer centrality on markups in various OTC markets. For instance, Hollifield et al. (2017) and Goldstein and Hotchkiss (2020) find a centrality discount in the

[12] For example, a core-periphery structure has been found in the corporate bond market (Di Maggio et al., 2017), the asset-backed securities market (Hollifield et al., 2017), the municipal bond market (Li and Schürhoff, 2019), inter-bank markets (Bech and Atalay, 2010; Roukny and Battiston, 2014; Afonso et al., 2013; Craig and Von Peter, 2014; Van Lelyveld and in't Veld, 2014), the credit default swaps market (Peltonen et al., 2014; Eisfeldt et al., 2023), and currency markets (King et al., 2012).

[13] Theoretical models include those based on search frictions which seek to endogenize the core-periphery structure (Munyan and Watugala, 2018; Üslü, 2019; Sambalaibat, 2022) or those that take the core-periphery structure as given (Gofman, 2011; Babus and Kondor, 2018; Malamud and Rostek, 2017; Collin-Dufresne et al., 2020). Other models examine how information frictions shape the market structure (Chang and Zhang, 2015).

asset-backed securities market and corporate bond market, while Di Maggio et al. (2017) and Li and Schürhoff (2019) document centrality premiums in the corporate bond and municipal bond markets. Examining the corporate bond market, Dick-Nielsen et al. (2021) find a centrality discount in client transaction costs but a centrality premium in interdealer bid-ask spreads. Our result also aligns with Schultz and Song (2019), which reports a shift toward a centrality premium following increased transparency in the asset-backed securities market.

Relatedly, our contribution extends to the investigation of OTC dealer characteristics and their impact on asset prices. These characteristics encompass dealer market share (Griffin et al., 2023), dealer relationships (Di Maggio et al., 2017), dealer centrality (Li and Schürhoff, 2019; Hollifield et al., 2017), and dealer specialization (Jotikasthira et al., 2023). Despite their interrelated nature, our study reveals that the influence of reciprocity on asset prices persists even after accounting for these dealer characteristics. In a related study, Hendershott et al. (2020) explore trading relationships between insurers and dealers in the corporate bond market. Their findings suggest that while some insurers exclusively engage with a single dealer, others maintain larger dealer networks. Notably, execution prices exhibit a non-monotonic relationship with network size, initially declining with more dealers but increasing once networks exceed 20 dealers. Similarly, our analysis of the inter-dealer market reveals trading costs decline with trading community size.

Finally, we contribute to the literature on bond market trading costs. Early studies by Harris and Piwowar (2006), Bessembinder et al. (2006), Green et al. (2007), and Edwards et al. (2007) highlight disparities in trading costs between small and large investors in municipal and corporate bond markets. These papers underscore the role of market opacity, allowing dealers to wield disproportionate power over less sophisticated retail investors. Green et al. (2007) and Schultz (2012) find significant markups in municipal bond transaction chains, while Griffin et al. (2023) observe substantial pricing variation for identical trades of the same bond on the same day from the same dealer, suggesting opportunistic behavior by dealers. Our study extends this literature by uncovering additional reasons for transaction chain markup variation and evidence of potential dealer malfeasance. We find that while reciprocity generally reduces transaction costs, a subset of dealers appears to exploit their local trading community to inflate costs for their own gain at the expense of customers.

Literature Review and Hypotheses

Literature Review

The literature on over-the-counter (OTC) markets is extensive, with notable contributions summarized by Weill (2020). Most theoretical analyses leverage search theory, network theory, or a combination of both. Search theory offers tractability, making it easier to model dynamic processes such as trading, pricing, and market reactions to aggregate shocks, while network theory embeds agent relationships into models, though at the cost of more complex mathematics.

The seminal work by Duffie et al. (2005) demonstrates how illiquidity premia arise from search frictions in OTC markets. Their model assumes independent, random matching between buyers and sellers, leading to transitory interdealer links and preventing natural emergence of persistent dealer networks or centrality without introducing additional assumptions such as dealer heterogeneity. These limitations motivate the incorporation of heterogeneity and network persistence in later models. Subsequent research builds on this search theory framework by adding dealer heterogeneity or extending economic settings, leading to emergent network features such as core-periphery structures, centrality premia, and dealer intermediation chains. For instance, differences in trader characteristics, such as their search technologies (Neklyudov, 2013), valuations (Hugonnier et al., 2014), meeting rates (Üslü, 2019), or trade frequencies (Sambalaibat, 2022), can generate core-periphery networks.

A key criticism of the search-theoretic framework is its inability to capture repeated counterparty trading, which empirical evidence suggests is prevalent in over-the-counter (OTC) markets. Studies such as Afonso et al. (2013) document stable trading relationships in the interbank lending market, while Li and Schürhoff (2019) find a 65% likelihood of municipal bond dealers trading with the same counterparties in consecutive months. Similarly, Hendershott et al. (2020) report persistent client-dealer relationships in corporate bond markets. To address this limitation, network-theoretic models offer an alternative or complementary approach by embedding agent relationships directly into the analysis. Some studies even adopt a hybrid approach, blending both search and network theories (Atkeson et al. (2015); Colliard et al. (2020); Chang and Zhang (2015); Dugast et al. (2022)). Network theory provides insights into how interconnectedness influences market stability, examining the effects of dealer position and centrality on liquidity, pricing, and efficiency. This sheds light on phenomena such as contagion, liquidity crises, and market resilience. For instance, more central dealers may have greater pricing power or access to liquidity, while shocks affecting one part of the network can propagate through these links. Neklyudov and Sambalaibat (2015) explore pricing in OTC markets by considering differences in dealers' search and matching efficiency due to network connections, predicting that bonds flow from fast, central dealers to slower, peripheral ones. Other models (Babus (2016); Wang (2016)) predict core-periphery network structures, highlighting the trade-offs between network efficiency and stability.[14]

Several network theories specifically examine how the structure of trading relationships affects the centrality premium that key dealers enjoy. Babus and Kondor (2018) shows that dealers' positions within the trading network influence their ability to acquire and disseminate information, which can enhance central dealers' advantages. Central dealers, by virtue of their network positions, can access information more quickly and trade on it, reinforcing their centrality premium. Chang and Zhang (2015) explore how market making and network formation occur endogenously in over-the-counter markets. They demonstrate that the network structure, shaped by the strategic formation of trading relationships, influences market efficiency and liquidity. The centrality premium enjoyed by key dealers arises from

[14] An incomplete list of network-based models includes Farboodi (2021), Gofman (2011), Gofman (2017), Malamud and Rostek (2017), Babus and Hu (2018), Babus and Kondor (2018), Wang (2016), Babus and Parlatore (2022), Babus and Farboodi (2020), Manea (2018), Eisfeldt et al. (2023), Aymanns et al. (2023), and Babus (2016).

their strategic positions within the network, underscoring how relationships affect market outcomes. Similarly, Gofman (2017) analyzes how the centralization of dealer networks impacts the efficiency and stability of financial systems. He shows that while a more centralized network can enhance overall efficiency by reducing transaction costs and improving liquidity, it may also concentrate risk and pricing power in the hands of central dealers, potentially increasing systemic risk. These findings relate to our result that dealer reciprocity can influence the advantages enjoyed by central dealers, although none of these models explicitly captures the role of reciprocity in detail.

The two models in this strand of literature most closely related to our findings are Sambalaibat (2022) and Hendershott et al. (2020). Sambalaibat (2022) models endogenous network formation between dealers, showing that core-periphery networks emerge as dealers specialize in clients with different trading needs. Her model assumes that customers must return to the same dealer to reverse trades, effectively embedding reciprocity into the customer-dealer relationship. Hendershott et al. (2020) highlight how repeated interactions between clients and dealers lower search costs, increase market efficiency, and stabilize liquidity through long-term relationships. While both models emphasize customer-dealer dynamics, they do not address how dealer reciprocity affects transaction chain markups.

While the literature on OTC markets has extensively explored both search and network theories, there is a noticeable gap in examining reciprocal relationships among dealers. Most models focus on customer-dealer dynamics or assume that dealer interactions are independent or random. This neglects the potential impact of long-term, repeated interactions between dealers themselves. Such reciprocal relationships could influence liquidity provision, pricing strategies, and centrality premiums, yet current models fail to fully capture how these dealer-to-dealer connections shape market outcomes, leaving a crucial aspect underexplored. To this end, we develop the following hypotheses which we test in our empirical analysis.

Hypotheses

Building upon the existing literature on over-the-counter (OTC) markets, dealer networks, and reciprocity, we formulate the following hypotheses for our empirical analysis.

H1: Impact of Dealer Reciprocity on Transaction Markups

Hypothesis 1A: *(Reciprocity Discount Hypothesis): Dealer reciprocity reduces transaction costs by lowering search frictions and incentivizing dealers to offer more competitive pricing in order to secure future business.*

Search frictions, as demonstrated by Duffie et al. (2005), can lead to higher transaction costs in decentralized over-the-counter markets. The lack of price transparency and the difficulty of finding counterparties in these fragmented markets further increase trading costs, particularly for investors without established networks. Dealer reciprocity, characterized by repeated trading relationships between counterparties, can serve as an effective mechanism to mitigate these frictions.

In the framework of Hendershott et al. (2020), reciprocal relationships between dealers help reduce the search intensity required to find counterparties and promote more efficient trading execution. In these relationships, dealers balance the benefits of maintaining repeat business with the need to compete for faster execution with other dealers. Reciprocity enhances intertemporal competition, where dealers aim to secure future transactions by offering better terms today, which ultimately lowers transaction costs. This dynamic fosters a more cooperative trading environment, leading to reduced markups as trust between reciprocal dealers develops.

Hypothesis 1B: *(Reciprocity Premium Hypothesis): In tightly knit dealer networks, reciprocity facilitates collusive behavior, leading to higher transaction markups.*

While generally we believe that reciprocity will reduce transaction costs and foster a more efficient trading environment, as described in Hypothesis 1A, it can also have unintended consequences in specific contexts. In smaller, tightly connected networks, the same reciprocal relationships that lower search frictions can enable collusion by providing dealers with opportunities to coordinate pricing strategies either implicitly or explicitly. In this hypothesis, we explore how the mutual expectation of future interactions, which enhances cooperation in large or competitive networks, can instead lead to collusive behavior in smaller, concentrated dealer networks.

Economic theory suggests that smaller, tightly knit networks are particularly conducive to building and sustaining collusive behavior due to several key mechanisms. In smaller groups, participants can more easily monitor each other's actions, reducing the likelihood of defection and enabling implicit coordination. Kranton and Minehart (2001) highlight that in strategic buyer-seller networks, the formation of links between participants can reduce competition by limiting the number of available trading partners. This insulation from competitive pressures allows network participants to implicitly coordinate on higher prices, as the network structure itself restricts competition. Similarly, Blume et al. (2011) emphasize that strategic complementarities—where one participant's actions reinforce those of others—make it easier for small groups to converge on collusive outcomes as an equilibrium, particularly when network connections are dense. In such networks, the mutual expectation of future interactions promotes long-term cooperation, allowing collusion to be sustained over time. The combination of repeated interactions, mutual dependence, and ease of monitoring in smaller networks creates a fertile environment for sustaining collusion, where participants can manipulate transaction prices with minimal external interference. Theories of collusion in oligopolistic markets further support this notion, as Stigler (1964) and Tirole (1988) demonstrate that smaller, cohesive groups are better able to maintain collusive arrangements, which can lead to inflated transaction markups.

However, regulatory oversight and the risk of detection may deter such behavior, limiting the extent to which reciprocity results in inflated prices. Empirical evidence from cartel studies (Levenstein and Suslow (2006); Harrington (2004)) shows that regulatory intervention and the threat of legal action often disrupt collusive behavior by increasing the risks and costs of maintaining price-fixing arrangements. Green and Porter (1984) further demonstrate that the risk of detection plays a critical role in destabilizing collusion, as firms are less likely to sustain inflated prices when there is a heightened chance of regulatory scrutiny. Furthermore, as networks grow larger or more competitive, the ability to maintain collusive agreements becomes more challenging. In such cases, reciprocity is less likely to lead to inflated markups, as increased competition dilutes the cohesiveness required for collusion. Ultimately, we propose that under certain conditions—such as smaller, more concentrated dealer networks with limited regulatory scrutiny—reciprocity can facilitate collusion and result in higher transaction markups, while in larger or more competitive networks, this effect is diminished.

H2: Dealer Reciprocity Predicts Future Quid Pro Quo Behavior and Longer Transaction Chains

Hypothesis 2: *Higher dealer reciprocity predicts future quid pro quo behavior among dealers, resulting in longer inter-dealer transaction chains and a lower likelihood of immediate placement with end customers.*

Reciprocity among dealers fosters mutual obligations and strengthens relationships, which can lead to quid pro quo behavior. We predict two ways this quid pro quo will manifest. First, reciprocal dealers may feel compelled to return favors by involving their reciprocal partners in transactions, even when direct trading with end customers is feasible. This practice can lengthen transaction chains as assets pass through additional intermediaries before reaching the final buyer or seller. Second, we expect reciprocal dealers to engage in more "reciprocal" or "reverse" trades than non-reciprocal dealers. For instance, if a reciprocal dealer sells a bond to another dealer, they are likely to purchase a bond from that dealer in the future as a way of returning the favor and maintaining the relationship.

H3: Moderating Role of Dealer Centrality on the Effect of Reciprocity

Hypothesis 3: *The impact of dealer reciprocity on transaction markups is moderated by dealer centrality; reciprocity leads to lower markups for central dealers but may result in higher markups for peripheral dealers.*

Building on Hypotheses 1A and 1B, we propose that dealer centrality moderates the relationship between reciprocity and transaction markups. Central dealers, due to their extensive network connections and influence, are in a better position to leverage reciprocity to enhance market efficiency. Reciprocity facilitates cooperation, improves information flow, and builds trust, allowing for smoother trading execution. While Li and Schürhoff (2019) find that central dealers tend to charge a premium for liquidity and fast execution, reciprocity might reduce these premiums by lowering the central dealer's costs. Thus, central dealers with more reciprocal trading relationships may charge relatively lower markups compared to central dealers with fewer reciprocal relationships, as they can pass on the benefits of reduced costs to their counterparties. This comparison highlights the role of reciprocity in shaping pricing dynamics even among highly connected dealers (Hollifield et al. (2017); Wang (2016)).

In contrast, for peripheral dealers, the interaction between centrality and reciprocity may lead to different outcomes. Peripheral dealers, with their lower connectivity and reduced visibility, are less exposed to competitive pressures and may use reciprocal relationships to sustain collusive practices. Reciprocity among peripheral dealers could enable them to coordinate on inflated prices, avoiding the need for aggressive competition (Stigler (1964); Granovetter (1985)). Kranton and Minehart (2001) suggest that in networks with fewer links, participants can exploit these relationships to maintain higher markups, reinforcing the potential for collusive behavior as discussed in Hypothesis 1B. Thus, the interaction between peripheral dealers' weaker network position and reciprocity may contribute to sustained price inflation, as they face fewer market checks and balances.

H4: Reciprocity Premium Is More Pronounced in Smaller Dealer Networks

Hypothesis 4: *Dealer reciprocity is more likely to result in higher transaction markups in smaller dealer networks, where sustaining collusive arrangements is more feasible.*

In smaller dealer networks, coordination among dealers is more feasible, and the risk of detection by regulators or outsiders is lower, making collusion more sustainable (Stigler (1964); Tirole (1988)). Reciprocity in these contexts can facilitate collusive behavior, leading to higher markups. Theoretical models suggest that smaller groups are better able to enforce collusive agreements and punish deviations (Farrell and Maskin (1989); Green and Porter (1984)). Empirical studies, such as Levenstein and Suslow (2006), show that cartels tend to be more stable and successful in smaller, concentrated industries, which parallels the dynamics observed in tightly knit dealer networks.

Based on these arguments, we hypothesize that reciprocity in smaller dealer networks is more likely to lead to higher transaction markups through sustained collusion. The ability of dealers to easily monitor each other's actions, coupled with lower regulatory oversight, allows for more effective enforcement of collusive agreements. The close-knit structure of smaller networks ensures that deviations from collusive arrangements are quickly detected and punished, further strengthening the collusive behavior. Thus, dealer reciprocity in these networks is expected to result in higher markups due to the feasibility of sustaining collusion over time.

In contrast, larger networks make coordination more difficult and intensify competition among dealers, reducing the likelihood that reciprocity leads to inflated prices. In these more competitive settings, reciprocity may still play a role but primarily by promoting cooperation and efficiency, potentially leading to lower transaction markups. The increased scrutiny and number of participants in larger networks make it harder for collusion to persist, as the risks of defection and detection are higher.

Overall, by testing these hypotheses, we aim to clarify how dealer reciprocity can have dual effects in over-the-counter (OTC) markets: enhancing market efficiency through cooperation in larger networks while fostering collusive practices in smaller ones. Understanding the conditions that drive reciprocity toward one outcome versus the other is crucial for both regulators and market participants in their efforts to promote fair and efficient markets.

Data and Sample Construction

Our starting sample consists of over 40 million intra-day municipal bond trades between 2014 and 2018, taken from the Municipal Securities Rulemaking Board (MSRB) Electronic Municipal Market Access (EMMA) database, which records the universe of US Municipal Bond Market transactions. For every transaction, EMMA records the corresponding bond's CUSIP, state and locality of issuance, date of issuance, and maturity date. It also reports trade details such as whether the transaction was a customer purchase, customer sale, or inter-dealer transaction; the selling price of the bond; the par value of trade, and anonymized dealer identifiers (buy- and sell-side executing broker symbols (EBS)), with the latter being unavailable in the public version of the dataset. Unfortunately, this dataset does not contain customer identities. We augment this dataset with bond characteristics including issue size, and bond rating, along with information on bond complexity, which is dependent on whether the bond is insured, whether the bond is a general obligation bond, whether it has a call provision, whether it is qualified for tax exemptions for bank investors, whether the bond is subject to tax (state, federal, or alternative minimum tax), and whether the bond has a sinking fund feature.

For our paper, we focus on secondary market bond transactions and prepare our dataset following Li and Schürhoff (2019). To do this we limit our sample to seasoned bonds which are more than 90 days from issuance. This permits us to focus on customers' and dealers' typical bond trading experience.[15] We eliminate trades with par value under $5,000 and consolidate dealer IDs into dealer entities to account for some firms using multiple dealer IDs. We follow

[15]See Green et al. (2007) and Schultz (2012) for further discussion on the differences between the primary and secondary municipal bond market.

Li and Schürhoff (2019), Appendix B to construct roundtrip transaction chains, defined as series of transactions on the same bond with the same par value starting with a customer purchase and ending with a customer sale with none or many inter-dealer transactions in between. Finally, we remove transaction chains initiated by new dealers (i.e., dealers without a single transaction over the prior 30 trading days), as we cannot calculate their centrality and reciprocity statistics. Our final roundtrip transaction chain sample contains 1,536,551 roundtrip chains, with 739,411 chains having at least one inter-dealer transaction.

For our measure of dealer centrality, we use eigenvector centrality as it is 98% correlated with the Net (EW) Li and Schürhoff (2019), simpler to compute, and analogous to our node reciprocity measure, though our results hold using both their Net Equal-Weight and Net Value-Weight Centrality measures instead. For our measure of dealer reciprocity, we use the node reciprocity measure for each dealer, as defined in the seminal network paper Newman, Forrest, and Balthrop (2002). Node reciprocity in network analysis measures the tendency of nodes (entities within a network, in our case dealers) to reciprocate connections with each other. In simple terms, it assesses how likely two nodes are to form a bidirectional connection or relationship. High node reciprocity indicates that nodes in the network tend to have mutual connections with each other, while low node reciprocity suggests more one-sided connections. We calculate both dealer centrality and reciprocity on a rolling 30-day basis.

We further investigate the reciprocal relationships by documenting trade reversals (defined in Appendix A) and organizing dealers into communities using the Louvain community detection method (described in Appendix D), introduced by Blondel et al. (2008). Due to the random nature of the Louvain method we obtain a bootstrap sample of communities using 100 random seeds, ranging from 1 to 100, not only to ensure the replicability of our results but also to account for the inherent instability of any methods involving randomness, documented by Jain and Madhyastha (2019). Using a 100-seed bootstrap allows us to exclude community members that are assigned to a community in which they don't belong, as the Louvain algorithm is known for "forcing" nodes into communities regardless of how often they interact with the "true" members of the given community. All other variable definitions are provided in Appendix A.

Table 1 reports summary statistics for the variables used in our analysis. We measure the average round-trip transaction chain markup as 1.02%. The average markup is 1.10% for chains that trade bonds with small par sizes, and for 0.65% chains with medium and large par sizes, both of which are close to values found in Griffin et al. (2023), whose sample spans between 2011 and 2017. The average natural logarithm of days to maturity is 8.12, which roughly corresponds to an average of 11.87 years. The average time since issuance (seasoning) is approximately 5.72 years. Our centrality and reciprocity measures are normalized between 0 and 1, following Li and Schürhoff (2019), and are reported for the dealers that initiate the chain. The distribution of the eigenvector centrality measure is comparable to that of Li and Schürhoff (2019), while the average reciprocity sits at 0.41. The state and total market share variables are slightly lower than those reported in Griffin et al. (2023), but are generally within one standard deviation of their averages. The market share on specific bonds is somewhat higher than that reported in Griffin et al. (2023), with an average of 68%. The average local network size is approximately 20 dealers, suggesting that an average dealer trades a given state's bonds with an average of 20 other dealers. Finally, the average dealer inventory is somewhat higher than the average inventory reported in Griffin et al. (2023).

Finally, we estimate the correlation between dealer reciprocity and dealer centrality to be 63%. This is driven by the fact that most low reciprocity dealers have low centrality and most high centrality dealers have high reciprocity, but this is far from the whole story as much of the interesting effects lie in the low centrality, high reciprocity, or high centrality, low reciprocity dealers. We explore this more in the univariate analysis below.

Analysis

Univariate Analysis

We begin our analysis by examining the contrasting effects of centrality and reciprocity using univariate sorts. In Table 2 we show that centrality and reciprocity are indeed distinct features of a network's structure despite having a fairly high correlation as documented in Table 1. Before we do this though, it helps to discuss the differences between these two variables more in depth.

Network centrality reflects the importance and influence of individual dealers within a network, emphasizing their connectivity, role as bridges, or control over information flow. In contrast, network reciprocity highlights the mutuality of relationships, focusing on two-way interactions where both parties benefit or engage with each other in a similar manner. Centrality and reciprocity are distinct but not necessarily independent concepts. A central dealer may have numerous connections but relatively few reciprocal relationships, while a highly reciprocal dealer pair might not hold a central position within the overall network.[16]

Table 2 presents average values for key metrics, including the number of dealers, transactions, transaction chain markups, chain length, and community size, based on a double sorting of the sample into dealer centrality and dealer reciprocity terciles. Panel A focuses on the number of unique dealers contributing to our sample of transactions, providing an initial overview of the dealer network structure. In line with Li and Schürhoff (2019), our analysis reflects a core-periphery network structure, where a small number of central dealers account for a significant proportion of all transactions. Specifically, 25 dealers fall within the highest tercile of centrality, with the majority demonstrating medium to high levels of reciprocity. Similarly, 82 dealers fall within the highest tercile of reciprocity, with 51 dealers

[16]The simple analogy is a social network. A person with many friends and who is often involved in conversations might have high centrality. A pair of close friends who frequently interact and support each other would have high reciprocity. Some nodes might have both high centrality and reciprocity. For example, a popular and close-knit friend group could exhibit both. However, they can also be independent: a central node might have many connections but few reciprocal relationships, while a highly reciprocal pair might not be central to the overall network.

exhibiting the most contrasting levels of centrality and reciprocity (e.g., high centrality-low reciprocity or low centrality-high reciprocity).

Panel B provides the distribution of municipal bond transactions in categories defined by dealer centrality and reciprocity. Approximately half of all transactions originate from dealers classified as peripheral and non-reciprocal (lowest tercile of both centrality and reciprocity, 26.6%) or central and reciprocal (highest tercile of both centrality and reciprocity, 20.6%). This distribution reveals a concentration of transactions within the extreme terciles of dealer centrality, while the medium centrality terciles show a more even distribution across reciprocity levels.

Panel C of Table 2 presents the core findings of our analysis, highlighting the impact of dealer centrality and reciprocity on average transaction chain markups. Consistent with existing literature, our analysis reveals that transactions involving dealers with higher levels of centrality tend to be associated with larger markups. This pattern holds across different levels of dealer reciprocity, with one notable exception: Dealers with low centrality but high reciprocity impose some of the highest average markups, reaching 1.53%. A closer examination of the interaction between dealer centrality and reciprocity reveals nuanced patterns: markups for transactions with low-centrality dealers increase with reciprocity, whereas those with medium- or high-centrality dealers show the opposite trend, with markups decreasing as reciprocity grows.

These univariate results underscore the pivotal role of reciprocity in shaping dealer network behavior and transaction cost patterns. Reciprocal trading relationships can provide benefits such as improved information flow, liquidity sharing, and the fostering of trust within the network, all of which help lower transaction costs, as reflected in Hypothesis 1A. At the same time, the findings reveal a critical risk: in certain contexts, reciprocity may facilitate collusion, as suggested by Hypothesis 1B. For instance, in the low-centrality group, reciprocal relationships correlate with higher markups, likely due to easier coordination in smaller or less competitive networks. The findings in Panel C further support Hypothesis

3, showing that reciprocity moderates the effect of dealer centrality on markups: for low-centrality dealers, higher reciprocity corresponds to increased markups, while for medium- and high-centrality dealers, greater reciprocity is associated with lower markups.

We now examine the length of transaction chains. Longer chains generally result in higher markups, as each link in the chain adds to the cumulative markup. Panel D of Table 2 presents the average transaction chain length, segmented by terciles of dealer centrality and reciprocity. Across all transactions, chain lengths range from 2.74 to 4.18, with an overall average of 2.97 transactions. Focusing first on the effect of centrality, we confirm prior findings that central dealers are associated with shorter transaction chains. Specifically, chain lengths tend to decrease for dealers with medium and high levels of centrality. However, dealers with low reciprocity exhibit an exception, where chain lengths increase as centrality increases. Turning to the effect of reciprocity, we observe that among highly central dealers, chains initiated by more reciprocal dealers are shorter. In contrast, for dealers with low and medium centrality, chain lengths increase as reciprocity grows. These findings suggest that the interaction between dealer centrality and reciprocity plays a significant role in shaping the length of the transaction chain. Specifically, chains initiated by dealers with low centrality and high reciprocity stand out as the longest, averaging 4.18 transactions. This pattern aligns with Hypothesis 2, which posits that higher reciprocity among less central dealers fosters quid pro quo behavior, leading to extended transaction chains. Notably, these longer chains are also associated with one of the highest average markups, recorded at 1.53%, underscoring the dual impact of reciprocity on both chain structure and transaction costs.

In Panel E, we analyze the size of local networks within the municipal bond market, defining it as the number of unique counterparties that a dealer traded with (bought from or sold to) during the past 30 trading days, calculated for each state of issuance. The results reveal consistent trends: the average size of the local network increases for more central dealers across all terciles of reciprocity and similarly increases for more reciprocal dealers. These findings suggest that centrality and reciprocity reinforce one another in expanding a dealer's local network. However, an exception arises among dealers in the low centrality-high reciprocity tercile, who have the smallest local networks. Notably, within the low centrality group, local network size increases from the lowest to the medium reciprocity tercile but sharply declines in the highest tercile.

These findings align with Hypothesis 4, which posits that smaller dealer networks, particularly those involving low-centrality dealers, are more conducive to sustaining collusive arrangements. The observed decline in local network size among low-centrality, high-reciprocity dealers may indicate that these dealers rely more heavily on a few reciprocal relationships to coordinate transactions, a hallmark of tightly knit, collusive arrangements. In contrast, the strengthening effects of centrality and reciprocity in expanding local networks highlight the broader cooperative dynamics observed in larger or more competitive networks. This dual role underscores the importance of network structure in mediating the impact of reciprocity on market behavior and transaction costs.

Centrality Premium and Reciprocity Discount

Our results in the previous section reveal interesting patterns, but these may be driven by factors such as the types of bonds traded by different dealers or other dealer-specific characteristics. To account for these possibilities, we now turn to our results from multivariate analysis.

The initial analyses explore the determinants of transaction costs, focusing on the effects of dealer centrality and reciprocity on the markups of transaction chains, measured in percent. The markup of a chain is calculated as the pricing differential between the final transaction of the chain (the sale to a customer) and the initial transaction of the chain (the purchase from the customer). We estimate the following regression specification:

$$\text{Markup}_i = \beta_1 \text{Centrality}_i + \beta_2 \text{Reciprocity}_i + \beta_3 \text{MarketShare}_i + \mathbf{X}_i \boldsymbol{\beta} + \alpha_s + \alpha_{my} + \epsilon_i \quad (1)$$

Centrality, reciprocity, and market shares are measured for the initiating dealer, specifically the dealer that purchases the bond from the customer. In this analysis, we standardize all explanatory variables to facilitate comparisons of their effects on markups across different covariates. $\mathbf{X}_i$ represents a vector of additional control variables, including market share, chain length, and bond characteristics such as time to maturity, time since issuance, insurance, tax status, provisions, rating, and whether the bond is a general obligation bond. Trade characteristics, including par size (modeled as small, medium, or large binary variables interacted with the natural logarithm of par size) and dealer inventory levels, are also included. The terms α_s and α_{ym} denote state fixed effects and month-year fixed effects, respectively, accounting for regional and temporal variations in markups. ϵ_i is the error term. Results are reported in Table 3.

In Column 1, we present results for dealer centrality and reciprocity controlling for bond and trade characteristics, but not dealer market share or chain length. Consistent with findings from Li and Schürhoff (2019), we observe a positive impact of dealer centrality on transaction costs, though the results are economically small. Specifically, in our sample, a one standard deviation increase in dealer centrality is associated with a 2.4 basis points increase in markups. The effect of a one standard deviation increase in dealer reciprocity is similarly sized, at 2.1 basis points.

Upon introducing measures of market share in Column 2, including overall dealer market share, dealer market share in the bond's domiciled state, and dealer market share in the bond, we observe a reversal in the estimated effect of dealer centrality. We find a reduction in markups by 3.5 basis points for a one standard deviation increase in dealer centrality. Notably, the estimated effects for dealer reciprocity remain unchanged in terms of both sign and magnitude.

These results are consistent with the evidence presented by Griffin et al. (2023) that shows that dealer market share measures attenuate the effect of dealer centrality on transaction chain markups. The estimated magnitudes of dealer market share, particularly for the overall dealer market share, are significantly larger when compared to the effects of dealer centrality and dealer reciprocity on markups. We estimate a one standard deviation increase in dealer market share leads to a 12.2 basis points increase in markups.

Moving to Column 3, where we control for chain length and its interaction with centrality and reciprocity, we observe a reversion of the baseline coefficient on centrality to a positive sign, akin to our estimate in Column 1. However, dealer reciprocity now exhibits a negative baseline coefficient. Additionally, chain length itself shows a positive association with markups. Although a one standard deviation increase in either centrality or reciprocity yields coefficient estimates for their respective interaction with chain length, producing offsetting effects, these effects are too small to counterbalance the sign of the baseline effects described earlier.[17] Column 4 introduces bond-month fixed effects, so the reciprocity estimate now compares trades in the same bond in the same month; the coefficients remain virtually unchanged, confirming the robustness of our results.

In Columns 5 - 8, we re-estimate the specifications used by columns 1 - 4 this time using a sample of transactions excluding those involving dealers in the low-centrality, high-reciprocity category identified in our univariate analysis. These excluded transactions, comprising 2.31% of all transactions in our sample, are clear outliers to the general patterns observed in terms of higher markups, longer transaction chains, and smaller local networks in our univariate analysis.

Here, two main observations emerge: First, dealer centrality consistently exhibits a positive association with transaction costs, with estimated magnitudes ranging from 7.2 to 13.3 basis points. These magnitudes are not only much larger than those estimated using the full sample but are also of similar magnitude to the estimated effects of dealer market share, which remain relatively unchanged. Second, we consistently estimate a negative effect of dealer reciprocity on markups. The magnitude of these estimated effects is 4-5 times larger

[17]Note that the median of our standardized measure of chain length is 0.0322, leading to coefficient estimates of 0.018 and -0.039 for dealer centrality and dealer reciprocity, respectively, at the median value of chain length. Both estimates are statistically significant at the 1 percent level.

than those previously estimated.

Thus, we find that excluding a specific set of transactions restores one of Li and Schürhoff (2019)'s insights: more central dealers charge higher markups, even when controlling for dealer market share, effectively reconciling the tension between Li and Schürhoff (2019) and Griffin et al. (2023) regarding the effect of centrality on markups. Additionally, our results show that transactions involving more reciprocal dealers, when controlling for dealer centrality and market share, are executed with lower markups, consistent with Hypothesis 1A. These results are also important because they help explain why dealer centrality leads to a premium in some markets, such as the municipal and corporate bond markets (Di Maggio et al., 2017; Li and Schürhoff, 2019), and a discount in others, such as ABS/MBS markets (Hollifield et al., 2017).

Finally, we implement a two-step approach to further explore the effects of centrality and reciprocity on markups, isolating their impact from bond and trade characteristics. In the first step, we regress markups on bond and trade variables, as well as fixed effects, to calculate residuals, which we term abnormal markups. In the second step, we double-sort these residuals into terciles based on measures of centrality and reciprocity. The results, presented in Panel F of Table 2, corroborate our findings from the univariate analysis and those presented in Table 3. Specifically, we observe an increase in abnormal markups as reciprocity increases within the lowest centrality tercile. Conversely, contrasting results emerge for the medium and high centrality terciles, indicating a different relationship between reciprocity and markups in these categories. Notably, the highest abnormal markups are identified in the category characterized by the lowest centrality and highest reciprocity, highlighting again the unique dynamics present within this subgroup.

Reciprocity and customer search costs

In this section, we examine how dealers find counterparties by analyzing their likelihood of directly selling bonds to customers versus involving other dealers. Consistent with

Hypothesis 2, we propose that reciprocal dealers are more likely to include intermediaries in transactions to return favors and maintain relationships, even when direct sales to customers are feasible. This behavior is expected to be most pronounced among low-centrality, high-reciprocity dealers, who may rely heavily on inter-dealer trades, potentially inflating transaction costs.

Following the approach of Li and Schürhoff (2019), we employ a panel Probit model to analyze dealer-to-customer (DC) trades and inter-dealer transactions. Our explanatory variables, measured at the initiating dealer of a transaction chain, include dealer centrality, dealer reciprocity, dealer market shares, and the complete set of control variables outlined in Table 3. To improve interpretability, we report marginal effects instead of probit coefficients, showing the change in the probability of the dependent variable given a one-unit increase in the explanatory variable. Thus, since continuous explanatory variables are standardized, each coefficient can be directly interpreted as the change in the probability of a direct customer trade due to a one standard deviation increase in the explanatory variable.

In Column 1 of Table 4, we find that dealer centrality positively influences the likelihood of the next trade being with a customer, while reciprocity exerts a negative impact. In our sample, the probability of the next trade being with a customer is 51.74% when all explanatory variables are set to their mean values. The marginal effect of centrality at these means translates to an increase of 6.24 percentage points, constituting 12.1% of the baseline probability. The corresponding effect of reciprocity is a decrease of 6.22 percentage points, equivalent to 12% of the baseline probability. Notably, the effects of these two factors counterbalance each other.

In Column 2, we use a categorical representation of dealer centrality and reciprocity tercile affiliation, replacing the continuous measures. Here, the medium category serves as the reference point. High centrality maintains its positive impact on the likelihood of subsequent trades with a customer, although there is variation depending on reciprocity: low reciprocity leads to a lower impact of high centrality. However, overall, for high centrality, variation in reciprocity does not fully offset the positive baseline effect. What stands out most from these results is the estimate for the low-centrality, high-reciprocity group. Here we find a 36% lower probability of placing the deal with a customer compared to the middle centrality, middle reciprocity group, which aligns with Hypothesis 2.

Overall, these results highlight that while dealer centrality invariably enhances the probability of direct customer transactions, reciprocity presents a nuanced impact, depending on its interplay with centrality. An increase in the initiating dealer's degree of reciprocity, especially when paired with low centrality, markedly lowers the probability of a dealer-customer trade, either because reciprocal in-network dealers agree to purchase a bond without access to an adequate customer base to market the bond to, or because those dealers engage in some form of collusive behavior to increase markups. In either case, we expect these instances to lead to longer transaction chains and ultimately increased markups, issues we explore in the following analyses.

Reciprocity and chain length

In Columns 3 and 4 of Table 4, we estimate a Poisson model to analyze the length of transaction chains and its relationship with dealers' network reciprocity. Our analysis in Column 3, utilizing continuous measures of dealer centrality and dealer reciprocity, reveals a negative impact of centrality on chain length, with a magnitude of 8.24 percent. This aligns with previous research by Li and Schürhoff (2019), suggesting that more central dealers are associated with shorter transaction chains. Conversely, reciprocity has a positive effect on chain length, with a magnitude of 9.64 percent. This implies that transaction chains initiated by reciprocal dealers tend to be longer, consistent with Hypothesis 2.

Using categorical centrality and reciprocity terciles in Column 4, the results remain consistent with previous findings. Specifically, high centrality and high reciprocity are linked to shorter and longer transaction chains, respectively. Conditional on high centrality, the shortening effect persists in the medium and high terciles of reciprocity, albeit somewhat muted for low levels of reciprocity. Conversely, the lengthening effect of reciprocity remains consistent across all terciles of centrality, with the most pronounced impact observed for low levels of centrality. Once again, the most pronounced deviation occurs in the low-centrality, high-reciprocity group, where transaction chain length is estimated to be 44% higher than in the middle centrality, middle reciprocity group.

Overall this set of results aligns closely with our second hypotheses regarding customer search costs. Specifically, our analysis indicates that while dealer centrality tends to decrease the length of transaction chains, suggesting more efficient market transactions facilitated by central dealers, the presence of reciprocity among dealers has the opposite effect, resulting in longer transaction chains. This elongation implies that dealers involved in reciprocal relationships may prioritize trading within their network, potentially compromising transaction efficiency and contributing to increased market fragmentation.

Corroborating evidence is presented in Columns 5 and 6 of Table 4, where we examine the probability of completing a transaction chain within one minute, referred to as prearranged trades. Two key observations emerge. First, in Column 5, we find that both centrality and reciprocity have a positive impact on the probability of prearranging a chain: a one standard deviation increase in centrality and reciprocity increases this probability by 3.1% and 2%, respectively. Second, in Column 6, using categorical centrality and reciprocity terciles, we observe that the groups with both low levels of centrality and reciprocity and high levels of centrality and reciprocity drive these results. However, we find a negative impact on the probability of prearranging, relative to the middle terciles, for groups characterized by high centrality and low reciprocity or low centrality and high reciprocity. The finding of a lower probability of quick customer-to-customer pass-throughs in the low centrality and high reciprocity group, combined with their longer transaction chains and lower probabilities of dealer-customer trades, reinforces their unique role in this market.

Reciprocal dealers and reciprocal behavior

Next, we investigate evidence of reciprocity among dealers, specifically focusing on those characterized by high levels of reciprocity. To do this, we estimate the probability of a reverse inter-dealer trade, such as dealer 2 selling to dealer 1, occurring shortly after the original trade, where dealer 1 sold to dealer 2. This analysis is conducted within pre-specified time frames of a day, a week, or a month. The first three columns of Table 5 present our findings, where we estimate the reverse trade propensities over the next 1, 7, and 30 days, respectively. As in previous analyses, we control for the initiating dealer's centrality and market shares, along with a comprehensive set of control variables and fixed effects adopted from Table 3. To make the results easier to interpret and compare, we report marginal effects instead of probit coefficients, as in Table 4.

Across all three specifications, we obtain negative and consistent coefficients, indicating that most of the activity occurs within the first day following the original trade. Over the one-day horizon, shown in Column 1, we obtain a negative and significant coefficient on dealer centrality, with a one standard deviation increase reducing the probability of a reverse trade by 2.7 percent. In contrast, a one standard deviation increase in dealer reciprocity raises the probability of a reverse trade by 7.1 percent. Over a 30-day period, the effect of dealer centrality on reversals, while still negative, diminishes to 0.9 percent. In contrast, the positive effect of reciprocity remains pronounced, increasing the probability of a reversal by 6.4 percent over the same timeframe.

Holding all control variables at their mean values, the baseline probability of observing a reverse trade after one day is 24.85 percent. Accordingly, a one standard deviation increase in dealer centrality represents a 10.95 percent decrease from the baseline probability, while a one standard deviation increase in dealer reciprocity leads to a substantial 28.37 percent increase.

The observed impact of reciprocity on reversal probabilities not only validates our measure of reciprocity out of sample but also underscores the significance of within-network trades for more reciprocal dealers. This finding is particularly pronounced when high reciprocity is paired with low centrality, a combination we previously emphasized. In Column 4 of Table 5, we once again employ a model utilizing categorical representation of centrality and reciprocity terciles, focusing on estimating the likelihood of 1-day reversals. Once again we find the low-centrality, high-reciprocity dealers acting as outliers with a substantially larger probability of reverse trades than all other dealer groups. For this dealer set we estimate the probability of the reverse deal occurring within one day as 34% higher than the middle centrality, middle reciprocity group.

By demonstrating that dealers characterized by high reciprocity are more likely to engage in reverse trades and that such reserve trades occur relatively quickly, our analysis potentially edges towards uncovering evidence suggestive of collusive behavior. While the observed behavior may reflect practices of trust and liquidity among reciprocal local dealer networks, the rapidity of these reversals also opens the door to interpretations that may imply strategic, coordinated actions aimed at inflating prices to the dealers' mutual benefit.

Trade characteristics by dealer reciprocity

In this section we deepen our analysis of dealer behavior by relating network characteristics to the attributes of the trades they arrange. We examine five dimensions in Table 6 that jointly characterize the extent to which dealers (i) transact with retail customers, (ii) favor complex bonds, (iii) generate anomalous "daisy–chain" transaction paths, (iv) quote coarse prices, and (v) suppress informative fine–yield quotations. These metrics together sharpen the picture of how low-centrality, high-reciprocity dealers appear to exploit opacity.

We hypothesize that if dealers use reciprocal relationships to extract rents (e.g., the low-centrality, high-reciprocity dealer group), they will focus more on trades with retail customers, who are more prone to cognitive biases compared to institutional customers. They are also more likely to trade more complicated bonds, which provide an advantage over less sophisticated retail investors and enable them to conceal price inflation more effectively than with vanilla bonds, where such behavior is easier to detect. Moreover, we expect their trading to display anomalous patterns, such as selling a bond to another dealer and then repurchasing it at a much higher price after multiple inter-dealer trades, before finally offloading it to a customer. We also anticipate them to rely heavily on coarse prices (quarter and odd eighths) rather than fine price increments, signaling a large discretionary element in their pricing.

We begin by considering trade size, which serves as a proxy for retail activity, with smaller trades (par $\leq$ \$100,000) typically linked to less sophisticated investors. Sorting dealers into centrality and reciprocity terciles reveals that among low-centrality dealers, the share of small-par trades rises sharply with reciprocity, from 42.36% in the lowest reciprocity tercile to 51.07% in the highest. This 8.71-percentage-point increase is both statistically and economically significant. By contrast, among medium and high-centrality dealers, the relationship reverses: higher reciprocity correlates with fewer small-par trades. This cross-sectional contrast supports the notion that low-centrality, high-reciprocity dealers target retail flows as part of a broader rent-extraction strategy.

We next examine whether the same group of dealers prefers to trade bonds with more complex contractual features, which exacerbate information asymmetries and may obscure markups. Complexity is quantified by counting optional features such as callability, sinking-fund provisions, and credit enhancement mechanisms. We find that within the low-centrality tier, complexity increases with reciprocity. In contrast, this association weakens or reverses among more central dealers. These findings align with Brancaccio et al. (2022), who argue that complex structures can facilitate surplus extraction from less-informed investors.

To explore potentially anomalous trading behavior, we analyze the incidence of daisy chains, i.e. round-trip paths in which the initiating dealer repurchases the same bond after a sequence of inter-dealer trades and eventually sells it to a customer. These patterns artificially inflate prices without any observable inventory benefit. Daisy chain probability increases monotonically with reciprocity across all centrality groups and peaks at 19.17% for the medium-centrality, high-reciprocity group. The low-centrality, high-reciprocity cell also exhibits a high rate of 11.31%, far exceeding the sub-2% rates observed for low-reciprocity or high-centrality dealers. These patterns are consistent with the use of favor-trading to camouflage price markups passed on to customers.

Supporting evidence of opacity also arises from dealers' price-quoting behavior. We measure the frequency of coarse pricing, e.g., quotes in eighths, quarters, or whole-dollar increments, which tend to obscure true economic concessions. Among low-centrality dealers, the share of trades executed at coarse price points increases markedly with reciprocity, from 7.15% to 11.14%. This 3.99-percentage-point rise contrasts sharply with medium- and high-centrality groups, where higher reciprocity is associated with more granular quoting. A complementary measure is the frequency of fine yield quotes, where precision is expressed in basis points. For low-centrality dealers, reciprocity is associated with less frequent use of such informative quotes (falling from 22.29% to 15.57%), while for high-centrality dealers, the pattern reverses, with transparency increasing alongside reciprocity. The low-centrality, high-reciprocity group thus not only relies more on coarse price grids but also provides the least yield information, reinforcing the hypothesis of deliberate opacity.

Overall, these results present consistent evidence that dealers who are both peripheral and highly reciprocal disproportionately engage in retail-oriented trades, prefer complex bond structures, orchestrate circular transaction chains, avoid transparent execution, and routinely obscure pricing information. This constellation of behaviors supports a rent-extraction hypothesis, wherein reciprocity is not merely a marker of cooperative trading but a strategic tool to exploit less-informed market participants.

The impact of reciprocity on markups and trading community size

In this section we examine our fourth hypothesis about how the effect of dealer reciprocity on markups changes will vary with dealer community size. Intuitively, collusion becomes increasingly difficult with more participants since coordinating actions and maintaining confidentiality become more complex, while the risk of detection rises, incentivizing individual members to cheat and increasing the likelihood of internal conflicts. Thus, we hypothesize that when local dealer networks are small, reciprocity will lead to inflated markups; and when networks are large, reciprocity will lead to reduced markups.

To test this hypothesis we attempt to measure dealer network size in two ways. First, we try a direct measurement and calculate local network size as a dealer's number of interdealer trading partners over the past 30 days. Second, we apply machine learning techniques to classify dealers into local trading communities using the Louvain community detection algorithm (Blondel et al., 2008). Overall, we find results consistent with our hypothesis.

Local Dealer Networks

To examine the impact of reciprocity conditional on dealer local network size, we re-run the regression specification from Column 4 of Table 3, stratifying the sample into quintiles based on each dealer's local network size, or the number of their inter-dealer trading partners over the preceding 30 days. Table 7 presents these results.

Consistent with our hypothesis, we observe that reciprocity increases markups for dealers operating within smaller networks, specifically those in the lowest quintile. For transactions within this subgroup, a one standard deviation increase in dealer reciprocity is associated with a 9.5 basis point increase in markup. These findings suggest that dealers in smaller networks may exploit reciprocal relationships to raise markups, potentially engaging in quid pro quo trading practices.

Conversely, our analysis reveals that reciprocity leads to a reduction in markups for dealers with larger networks across all other quintiles. Specifically, a one standard deviation increase in reciprocity corresponds to a decline in markup ranging from 12.7 to 26.3 basis points for these dealers. All coefficient estimates are significant at the 1 percent level. Overall, this consistent trend suggests that in larger networks, the positive effects of liquidity provision facilitated by reciprocal relationships, such as reduced search costs, outweigh any propensity toward markup-increasing behavior.

Louvain communities

We next apply the Louvain algorithm for community detection (Blondel et al., 2008) to cluster dealers into trading communities.[18] This clustering allows us to analyze how community affiliation impacts transaction chain markups. Specifically, we focus on the effect on markups when both dealers belong to the same community. Table 8 presents these results. We employ essentially the same regression specification as in Column 4 of Table 3, omitting centrality and reciprocity measures, and incorporating our Louvain community classification. We introduce the variable 'Same Community' to identify transactions where both dealers are within the same Louvain community, which occurs in 2.09% of cases. This analysis uses dynamic Louvain communities defined at a resolution parameter of 40. We also account for the sizes of the seller's and buyer's communities and their interactions with the 'Same Community' variable.

In the first three columns of Table 8, we observe a large, positive coefficient for the 'Same Community' indicator. These estimates indicate that markups are approximately 60 to 70 basis points higher when both the seller and buyer belong to the same community. Furthermore, the magnitude of this effect diminishes with the size of the buyer's and seller's communities, as shown by the negative coefficients on the interaction terms between the 'Same Community' indicator and the sizes of the respective communities. Notably, since the average community consists of 3.56 dealers, we consistently observe a positive impact of intra-community trades on markups across a typical community. These results remain robust across all three specifications, regardless of whether we control for seller community size or buyer community size. In the last three columns of Table 8, we exclude transactions involving dealers assigned to the low-reciprocity, high-centrality group. This attenuates the

[18] To enhance the reliability of this non-deterministic algorithm, we run it multiple times with different random seeds (from 1 to 100) and aggregate the outcomes. This method mitigates bias from any single random state and ensures consistent dealer community assignments. Dealers assigned to multiple communities are excluded from the analysis. The assignment process is detailed in Appendix D.

effect of the same community dummy to around 50 basis points, despite representing only 0.7% of the transactions in the sample. We continue to observe the attenuating effects of community size.

Overall, these findings closely relate to our earlier results on dealer reciprocity. The elevated markups within the same community suggest that dealers may exploit close-knit relationships to increase profits, similar to how reciprocity can lead to higher markups in smaller networks. The attenuation of this effect with larger community sizes parallels the observation that reciprocity reduces markups in larger networks due to enhanced liquidity and reduced search costs. We conclude that community affiliation and dealer reciprocity both play significant roles in shaping transaction costs, with their impacts conditioned by the size and structure of the trading network.

Determinants of dealer communities

To provide a better understanding of the Louvain community measure employed in our previous analysis, we now explore the primary factors that determine the community affiliation of dealers within the network structure. Table 9 outlines the significance of the top 20 determinants affecting whether trade counterparties are part of the same dynamic Louvain community. Our analysis employs machine learning models, focusing on two types of tree-based classifier models: LightGBM, a gradient-boosting machine learning framework, introduced by Ke et al. (2017), and Random Forest (RF) classifier, introduced by Breiman (2001). Despite some disparities between the models, LightGBM emerges as generally more accurate, boasting over 90% correct predictions. Nonetheless, the RF classifier shows a significantly lower rate of false negatives, demonstrating its effectiveness in accurately identifying trades within the same community. Feature importances reported in Table 9 are averaged across 100 seeds to avoid dependence on random states, documented by Jain and Madhyastha (2019).

Some of the results provided in Table 9 are quite intuitive - a larger local network size of either of the counterparties will likely result in the dealers being from the same, larger community. A larger fraction of purchases and sales from one another (Di Maggio et al., 2017) is another somewhat intuitive result - the more the two counterparties trade with each other, the more likely it is that they belong to the same community. We also find that the reciprocity measures and, to a smaller degree, centrality measures, tend to be major predictors of whether the two counterparties belong to the same community, with the reciprocity also being a prominent predictor in the RF classifier models. The importance of the buyer's and seller's returns also hints at potential quid pro quo situations: A buyer (seller) who has been financially well-off in the prior month wants to help out his community members. The total value of the reverse trades and the number of such trades also seem to accurately predict whether the two counterparties belong to the same community.

Additionally, the machine learning models reveal noteworthy findings regarding the dealers' retail trade market shares. This metric, defined as the percentage of sales a dealer makes to customers, appears to influence whether two trading counterparties are from the same community. It hints at potential quid pro quo dynamics, where a seller with a larger retail market share might facilitate a transaction for a community member. Collectively, these models indicate that the trades within the same community are primarily driven by the reciprocity and centrality of the counterparties, their financial standing, and their history of interactions. Our analysis underscores the complexity of factors influencing intra-community trading behaviors.

Conclusion

The dynamics of dealer networks hold significant implications for asset pricing and market efficiency. Our study illuminates the complex interplay between dealer reciprocity and transaction chain markups, revealing both the beneficial and detrimental impacts of such relationships on market outcomes.

Our research indicates that overall, reciprocity among dealers tends to foster a more cohesive and efficient trading environment. This mutual exchange of favors, grounded in the expectation of future reciprocity, generally contributes to lower transaction costs for participants. Notably, our findings underscore that higher levels of dealer reciprocity correlate with reduced bond markups, particularly among central dealers with significant network connections. This phenomenon suggests that reciprocal networks can enhance market efficiency, streamline the flow of capital, and ultimately, benefit the broader financial ecosystem through more competitive pricing.

Conversely, our study also reveals the problematic aspects of reciprocity, highlighting the risks of collusion and market manipulation within these networks. Specifically, we identify a particular group of peripheral dealers characterized by high levels of reciprocity, who maintain significant markups, suggesting a form of tacit collusion to artificially inflate transaction chain markups.

References

Afonso, G., Kovner, A., Schoar, A., 2013. Trading partners in the interbank lending market. FRB of New York Staff Report .

Atkeson, A. G., Eisfeldt, A. L., Weill, P.-O., 2015. Entry and exit in otc derivatives markets. Econometrica 83, 2231–2292.

Aymanns, C., Georg, C.-P., Golub, B., 2023. Exit spirals in coupled networked markets. Operations Research 71, 1441–1924.

Babus, A., 2016. The formation of financial networks. The RAND Journal of Economics 47, 239–272.

Babus, A., Farboodi, M., 2020. The hidden costs of strategic opacity. Working Paper 27471, National Bureau of Economic Research.

Babus, A., Hu, T.-W., 2018. Endogenous intermediation in over-the-counter markets. Journal of Financial Economics 125, 200–215.

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

Babus, A., Parlatore, C., 2022. Strategic fragmented markets. Journal of Financial Economics 145, 876–908.

Bech, M. L., Atalay, E., 2010. The topology of the federal funds market. Physica A: Statistical mechanics and its applications 389, 5223–5246.

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

Blau, P., 1964. Exchange and power in social life .

Blondel, V. D., Guillaume, J.-L., Lambiotte, R., Lefebvre, E., 2008. Fast unfolding of communities in large networks. Journal of statistical mechanics: theory and experiment 2008, P10008.

Blume, L. E., Brock, W. A., Durlauf, S. N., Ioannides, Y. M., 2011. Identification of social interactions. North-Holland, vol. 1 of *Handbook of Social Economics*, pp. 853–964.

Brancaccio, G., Kang, K., 2022. Search frictions and product design in the municipal bond market. Tech. rep., National Bureau of Economic Research.

Brancaccio, G., Li, D., Schürhoff, N., 2017. Learning by Trading: The Case of the US Market for Municipal Bonds. Unpublished Paper. Princeton University .

Breiman, L., 2001. Random Forests. Machine Learning 45.

Chang, B., Zhang, S., 2015. Endogenous market making and network formation .

Colliard, J.-E., Foucault, T., Hoffmann, P., 2020. Inventory management, dealers' connections, and prices in otc markets. The Journal of Finance, Forthcoming .

Collin-Dufresne, P., Junge, B., Trolle, A. B., 2020. Market structure and transaction costs of index cdss. The Journal of Finance 75, 2719–2763.

Craig, B., Von Peter, G., 2014. Interbank tiering and money center banks. Journal of Financial Intermediation 23, 322–347.

Di Maggio, M., Kermani, A., Song, Z., 2017. The value of trading relationships in turbulent times. Journal of Financial Economics 124, 266–284.

Dick-Nielsen, J., Kjær Poulsen, T., Rehman, O., 2021. Dealer networks and the cost of immediacy. The 19th International Paris December Finance Meeting, Paris December 2021 ; Conference date: 16-12-2021 Through 16-12-2021.

Duffie, D., Gârleanu, N., Pedersen, L. H., 2005. Over-the-counter markets. Econometrica 73, 1815–1847.

Dugast, J., Üslü, S., Weill, P.-O., 2022. A theory of participation in otc and centralized markets. The Review of Economic Studies 89, 3223–3266.

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

Eisfeldt, A. L., Herskovic, B., Rajan, S., Siriwardane, E., 2023. Otc intermediaries. The Review of Financial Studies 36, 615–677.

Farboodi, M., 2021. Intermediation and voluntary exposure to counterparty risk. Working Paper 29467, National Bureau of Economic Research.

Farrell, J., Maskin, E., 1989. Renegotiation in repeated games. Games and Economic Behavior 1, 327–360.

Gofman, M., 2011. A network-based analysis of over-the-counter markets. In: *AFA 2012 Chicago meetings paper*.

Gofman, M., 2017. Efficiency and stability of a financial architecture with too-interconnected-to-fail institutions. In: *Journal of Financial Economics*, vol. 124, pp. 113–146.

Goldstein, M. A., Hotchkiss, E. S., 2020. Providing liquidity in an illiquid market: Dealer behavior in us corporate bonds. Journal of Financial Economics 135, 16–40.

Gouldner, A. W., 1960. The norm of reciprocity: a preliminary statement. American Sociological Review 25.

Granovetter, M., 1985. Economic action and social structure: The problem of embeddedness. American Journal of Sociology 91, 481–510.

Green, E. J., Porter, R. H., 1984. Noncooperative collusion under imperfect price information. Econometrica 52, 87–100.

Green, R. C., Hollifield, B., Schürhoff, N., 2007. Financial intermediation and the costs of trading in an opaque market. The Review of Financial Studies 20, 275–314.

Griffin, J. M., Hirschey, N., Kruger, S., 2023. Do municipal bond dealers give their customers "fair and reasonable" pricing? The Journal of Finance 78, 887–934.

Harrington, J. E., 2004. Cartel pricing dynamics in the presence of an antitrust authority. Tech. Rep. 4.

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

Hendershott, T., Li, D., Livdan, D., Schürhoff, N., 2020. Relationship trading in otc markets. Journal of Finance 75, 683–734.

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

Hugonnier, J., Lester, B., Weill, P.-O., 2014. Heterogeneity in decentralized asset markets. Working Paper 20746, National Bureau of Economic Research.

Jain, R., Madhyastha, P., 2019. On Model Stability as a Function of Random Seed .

Jotikasthira, C., Lundblad, C. T., Xue, J., 2023. Dealer specialization and market segmentation. Available at SSRN 4636821 .

Kamada, T., Kawai, S., 1989. An algorithm for drawing general undirected graphs. Information Processing Letters 31.

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., Liu, T.-Y., 2017. LightGBM: A Highly Efficient Gradient Boosting Decision Tree. Advances in Neural Information Processing Systems 30.

King, M. R., Osler, C., Rime, D., 2012. Foreign exchange market structure, players, and evolution. Handbook of exchange rates pp. 1–44.

Kranton, R. E., Minehart, D. F., 2001. A theory of buyer-seller networks. American Economic Review 91, 485–508.

Levenstein, M. C., Suslow, V. Y., 2006. What determines cartel success? Journal of Economic Literature 44, 43–95.

Li, D., Schürhoff, N., 2019. Dealer networks. The Journal of Finance 74, 91–144.

Malamud, S., Rostek, M., 2017. Decentralized exchange. American Economic Review 107, 3320–3362.

Manea, M., 2018. Intermediation and resale in networks. Journal of Political Economy 126, 1250–1301.

Munyan, B., Watugala, S. W., 2018. What makes dealers central? evidence from credit interdealer networks. Working Paper .

Neklyudov, A., Sambalaibat, B., 2015. Endogenous specialization and dealer networks. Working Paper .

Neklyudov, A. V., 2013. Bid-ask spreads and the decentralized interdealer markets: Core and peripheral dealers. Tech. rep.

Newman, M. E., Forrest, S., Balthrop, J., 2002. Email networks and the spread of computer viruses. Physical Review E 66, 035101.

Peltonen, T. A., Scheicher, M., Vuillemey, G., 2014. The network structure of the cds market and its determinants. Journal of Financial Stability 13, 118–133.

Piwowar, M. S., 2007. Corporate and Municipal Bonds. Securities Litigation and Consulting Group, Inc .

Roukny, T., Battiston, S., 2014. A network analysis of the evolution of the german interbank market. Bundesbank Discussion Paper .

Sambalaibat, B., 2022. Endogenous specialization and dealer networks. Available at SSRN 2676116 .

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

Schultz, P., Song, Z., 2019. Transparency and dealer networks: Evidence from the initiation of post-trade reporting in the mortgage backed security market. Journal of Financial Economics 133, 113–133.

Sobel, J., 2005. Interdependent preferences and reciprocity. Journal of economic literature 43, 392–436.

Stigler, G. J., 1964. A theory of oligopoly. Journal of Political Economy 72, 44–61.

Tirole, J., 1988. The theory of industrial organization. MIT press.

Üslü, S., 2019. Pricing and liquidity in decentralized asset markets. Econometrica 87, 2079–2140.

Van Lelyveld, I., in't Veld, D., 2014. Finding the core: Network structure in interbank markets. Journal of Banking & Finance 49, 27–40.

Wang, C., 2016. Core-periphery trading networks. University of Pennsylvania Wharton School of Business Working Paper .

Weill, P.-O., 2020. The search theory of otc markets. Working Paper 27354, National Bureau of Economic Research.

Table 1: Univariate statistics for transaction chains.

This table contains summary statistics for key variables used in our analysis using our sample of 1,536,551 transaction chains. We report the mean, median, standard deviation, 25^{th}, and 75^{th} percentiles for all variables. The last column of the table reports the correlation of the given variable with the dealer reciprocity. Transaction characteristics reflect all trades in a transaction chain, while dealer characteristics are based on the initial dealer only. All variable definitions are provided in Appendix A.

Variable	Mean	St. Dev.	25%	Median	75%	Corr.
Transaction Characteristics						
Markup (%)	1.02	2.00	0.17	0.68	1.94	0.04
Par Size (in $000s)	68.93	616.94	10.00	25.00	40.00	-0.03
Log(par) × Small	9.02	2.90	9.21	9.90	10.13	0.10
Log(par) × Medium	1.00	3.30	0.00	0.00	0.00	-0.09
Log(par) × Large	0.09	1.19	0.00	0.00	0.00	-0.04
Dealer Characteristics						
Dealer centrality	0.13	0.07	0.07	0.14	0.18	0.63
Dealer reciprocity	0.41	0.17	0.29	0.44	0.52	1.00
Dealer market share	0.02	0.02	0.01	0.02	0.03	0.33
Dealer market share (State)	0.03	0.04	0.01	0.02	0.04	0.17
Dealer market share (Bond)	0.68	0.35	0.33	0.87	1.00	-0.00
Dealer network size	19.68	18.09	5.00	15.00	29.00	0.50
Dealer inventory	0.92	33.64	-1.47	0.00	1.54	-0.04
Bond Characteristics						
Maturity	8.12	0.75	7.63	8.20	8.74	0.07
Seasoning	7.30	0.91	7.97	7.56	6.77	0.02
Issue size	18.78	1.39	17.81	18.86	19.81	-0.02
Rating	2.24	2.56	0.00	2.00	4.00	0.02
Junk	0.01	0.09	0.00	0.00	0.00	-0.03
Unrated	0.40	0.49	0.00	0.00	1.00	-0.04
Callable	0.65	0.48	0.00	1.00	1.00	0.09
Insured	0.00	0.05	0.00	0.00	0.00	-0.00
General obligation	0.33	0.47	0.00	0.00	1.00	0.01
Taxable	0.08	0.26	0.00	0.00	0.00	-0.03
Bank qualified	0.02	0.13	0.00	0.00	0.00	0.01
Subject to AMT	0.02	0.12	0.00	0.00	0.00	-0.02
Sinking fund	0.55	0.50	0.00	1.00	1.00	0.03

Table 2: Transaction chain characteristics across dealer centrality and reciprocity terciles.

This table presents the distribution and characteristics of municipal bond transaction chains segmented by dealer centrality and reciprocity terciles. Panels A and B report the number of unique initial dealers and transaction chains, respectively. Panels C through F provide averages for key metrics, including transaction chain markups (Panel C), chain length (Panel D), local network size (Panel E), and abnormal markups (Panel F). Local network size is defined as the number of unique counterparties a dealer has bought from or sold to in the past 30 trading days, and abnormal markup is the difference between actual markup and predicted markup based on a regression model controlling for trade, dealer, and bond characteristics. The data is based on 1,536,661 transaction chains. Statistical significance for differences between high and low centrality (reciprocity) terciles is indicated at the 10% (*), 5% (**), and 1% (***) levels. Definitions for all variables are provided in the appendix.

Panel A. Dealer entity counts

		Reciprocity		
		Lowest	Medium	Highest
Centrality	Lowest	636	110	49
	Medium	25	34	21
	Highest	2	11	12

Panel B. Number of C(N)DC transaction chains

		Reciprocity		
		Lowest	Medium	Highest
Centrality	Lowest	408,043 (26.56%)	68,640 (4.47%)	35,506 (2.31%)
	Medium	93,485 (6.08%)	258,015 (16.79%)	160,710 (10.46%)
	Highest	10,658 (0.69%)	185,578 (12.08%)	315,916 (20.56%)

Panel C. Average markups

		Reciprocity				
		Lowest	Medium	Highest	**Total**	**H-L**
Centrality	Lowest	0.81	1.00	1.53	0.88	0.72***
	Medium	1.41	0.97	0.79	0.99	-0.62***
	Highest	1.99	1.13	1.19	1.18	-0.80***
	Total	0.94	1.03	1.09		
	H-L	1.18***	0.13***	-0.34***		

Panel D. Average chain length

		Reciprocity				
		Lowest	Medium	Highest	**Total**	**H-L**
	Lowest	2.74	3.55	4.18	2.94	1.44***
	Medium	3.00	3.11	3.39	3.18	0.39***
Centrality	Highest	3.20	2.81	2.77	2.79	-0.43***
	Total	2.79	3.06	3.06		
	H-L	0.46***	-0.74***	-1.41***		

Panel E. Average local network size

		Reciprocity				
		Lowest	Medium	Highest	**Total**	**H-L**
	Lowest	6.46	8.86	1.34	6.43	-5.12***
	Medium	16.33	20.06	24.11	20.65	7.78***
Centrality	Highest	20.91	22.85	37.43	31.80	16.52***
	Total	8.59	19.57	30.75		
	H-L	14.45***	13.99***	36.09***		

Panel F. Average abnormal markup

		Reciprocity				
		Lowest	Medium	Highest	**Total**	**H-L**
	Lowest	-0.06	0.02	0.33	-0.03	0.39***
	Medium	0.15	-0.06	-0.25	-0.08	-0.40***
Centrality	Highest	0.29	0.03	0.15	0.11	-0.14***
	Total	-0.02	-0.02	0.03		
	H-L	0.35***	0.01	-0.18***		

Table 3: Baseline regressions of markups on dealer centrality and reciprocity.

This table presents the results of baseline regressions analyzing the relationship between transaction markups and dealer network characteristics, transaction attributes, and bond features. The dependent variable is the total markup charged to the customer in the final sale of the transaction chain. Independent variables include dealer eigenvector centrality and node reciprocity (see Appendix A for definitions), along with controls for chain length, dealer market share, bond characteristics, and trade attributes. Bond controls include time to maturity, time since issuance, insurance status, tax status, provisions, ratings, whether the bond is a general obligation bond, bond issuer state fixed effects, and in columns 4 and 8 bond-trading month fixed effects. Trade controls include par size (categorized as small, medium, or large and interacted with the natural logarithm of par size), dealer inventory levels, and trade month-year fixed effects. Columns 1 through 4 use our entire transaction chain sample, while columns 5 through 8 exclude the transaction chains where the initiating dealer is in the lowest centrality-highest reciprocity tercile group to illustrate the impact of this small, but influential set of dealers. All independent variables are standardized to have a mean of 0 and a variance of 1. Standard errors, clustered by bond and time, are reported in parentheses below the parameter estimates. Statistical significance is denoted at the 10% (*), 5% (**), and 1% (***) levels.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sample:	All	All	All	All	Ex. low c., high. r	Ex. low c., high. r	Ex. low c., high. r	Ex. low c., high. r
Dep. Variable:	Total Markup	Total Markup	Total Markup	Total Markup	Total Markup	Total Markup	Total Markup	Total Markup
Dealer centrality	0.024***	-0.035***	0.022***	0.018***	0.133***	0.072***	0.094***	0.090***
	(0.007)	(0.006)	(0.006)	(0.005)	(0.014)	(0.012)	(0.010)	(0.008)
Dealer reciprocity	0.021***	0.024***	-0.040***	-0.033***	-0.095***	-0.087***	-0.114***	-0.107***
	(0.007)	(0.007)	(0.006)	(0.005)	(0.014)	(0.013)	(0.011)	(0.009)
Chain length			0.020***	0.002			0.017**	0.001
			(0.007)	(0.005)			(0.007)	(0.006)
Dealer centrality × Chain length			-0.110***	-0.101***			-0.084***	-0.073***
			(0.007)	(0.006)			(0.018)	(0.013)
Dealer reciprocity × Chain length			0.046***	0.037***			0.012	-0.000
			(0.009)	(0.007)			(0.022)	(0.016)
Dealer market share		0.122***	0.112***	0.112***		0.121***	0.114***	0.114***
		(0.008)	(0.007)	(0.006)		(0.008)	(0.007)	(0.005)
Dealer market share (State)		-0.036**	-0.033**	-0.030***		-0.039**	-0.036**	-0.033***
		(0.015)	(0.015)	(0.010)		(0.016)	(0.015)	(0.010)
Dealer market share (Bond)		0.052***	0.051***	0.043***		0.051***	0.051***	0.043***
		(0.006)	(0.006)	(0.005)		(0.006)	(0.007)	(0.005)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Bond-Month FE	No	No	No	Yes	No	No	No	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,536,551	1,536,551	1,536,551	1,536,551	1,501,045	1,501,045	1,501,045	1,501,045
Adjusted R^2	0.04	0.04	0.05	0.18	0.04	0.04	0.05	0.18

Table 4: Regressions analyzing the impact of centrality and reciprocity on customer search costs, chain size and chain type

This table presents regression results examining three key transaction outcomes: (1) the probability of a transaction being a chain-concluding customer sale (columns 1 and 2), (2) the number of unique dealers involved in the transaction chain (columns 3 and 4), and (3) the probability of the chain being a prearranged trade (columns 5 and 6). Independent variables include dealer centrality, reciprocity, their terciles, dealer market share, and bond- and trade-level controls. Bond controls comprise attributes such as time to maturity, time since issuance, insurance status, tax status, provisions, ratings, and whether the bond is a general obligation bond. Trade controls include par size (categorized as small, medium, or large and interacted with the natural logarithm of par size) and dealer inventory levels. Columns 1 and 2 use the full sample of all transactions and report marginal effects for probit regressions to enhance interpretability, reflecting the change in the dependent variable's probability given a unit increase in the explanatory variable. Columns 3 through 6 are restricted to transaction chain-level data. All independent variables are standardized to have a mean of 0 and unit variance. Standard errors are shown in parentheses, with statistical significance denoted at the 10% (*), 5% (**), and 1% (***) levels.

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Variable:	Pr(DC trade)	Pr(DC trade)	N of dealers	N of dealers	Pr(Prearranged)	Pr(Prearranged)
Dealer centrality	0.062***		-0.086***		0.031***	
	(0.001)		(0.001)		(0.001)	
Dealer reciprocity	-0.062***		0.092***		0.020***	
	(0.001)		(0.001)		(0.001)	
Low reciprocity		0.152***		-0.015**		-0.310***
		(0.008)		(0.006)		(0.09)
High reciprocity		-0.114***		0.093***		-0.051***
		(0.003)		(0.002)		(0.007)
Low centrality		-0.199***		0.188***		-0.175***
		(0.004)		(0.003)		(0.009)
High centrality		0.318***		-0.095***		-0.245***
		(0.005)		(0.004)		(0.007)
Low centrality × Low reciprocity		0.300***		-0.294***		0.468***
		(0.009)		(0.007)		(0.013)
High centrality × Low reciprocity		-0.333***		0.159***		-0.718***
		(0.016)		(0.017)		(0.059)
Low centrality × High reciprocity		-0.045***		0.029***		-0.293***
		(0.005)		(0.004)		(0.016)
High centrality × High reciprocity		0.088***		-0.115***		0.511***
		(0.006)		(0.004)		(0.010)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Market share controls	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,169,705	4,169,705	1,536,551	1,536,551	1,536,551	1,536,551
Pseudo R^2	0.01	0.01	0.01	0.01	0.10	0.10

Table 5: Analysis of reverse transaction probabilities and dealer network characteristics

This table contains probit regressions of the probability of the reverse transactions (a buyer and a seller in a previous transaction becoming the seller and buyer in a later transaction) occurring in the next 1, 7, and 30 trading days. Columns (1)-(3) report probit regressions of the respective reversal probabilities on centrality, reciprocity, their terciles, dealer market share controls as well as bond and trade controls. Bond controls include the bond characteristics, such as time to maturity, time since issuance, insurance, tax status, provisions, rating, and whether the bond is a general obligation bond, with state and month fixed effects, clustered by bond, and time. Trade controls include par size (small, medium, or large binary variables interacted with the natural logarithm of par size) controls and dealer inventory levels. All of the independent variables are standardized to have a mean of 0 and unit variance. Column (4) reports the average baseline probabilities of each dealer category relative to the dealer of average centrality (reciprocity). To aid in interpretability, rather than report probit coefficient estimates all columns report their marginal effects, i.e. the change in the probability of the dependent variable conditional on a unit increase in the explanatory variable. Standard errors, clustered by bond and time, are reported in parentheses below the parameter estimates. Statistical significance is denoted at the 10% (*), 5% (**), and 1% (***) levels.

	(1)	(2)	(3)	(4)
Dep. Variable:	Reversal 1 day	Reversal 7 day	Reversal 30 day	Baseline Prob.
Dealer centrality	-0.027***	-0.016***	-0.009***	
	(0.001)	(0.001)	(0.001)	
Dealer reciprocity	0.071***	0.076***	0.064***	
	(0.001)	(0.001)	(0.001)	
Low reciprocity				-0.520***
				(0.013)
High reciprocity				-0.037***
				(0.009)
Low centrality				-0.180***
				(0.011)
High centrality				-0.311***
				(0.011)
Low reciprocity × Low centrality				0.291***
				(0.017)
Low reciprocity × High centrality				0.095***
				(0.024)
High reciprocity × Low centrality				0.556***
				(0.015)
High reciprocity × High centrality				0.242***
				(0.013)
Controls	Yes	Yes	Yes	Yes
Market share controls	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes
Observations	739,411	739,411	739,411	739,411
Pseudo R^2	0.06	0.06	0.02	0.07

Table 6: Comparison of trade size, bond complexity, daisy chains, and prearranged trades across dealer centrality and reciprocity terciles.
This table reports the percentage of transaction chains with par size less than \$100K (Panel A), the percentage of transaction chains involving complex bonds –i.e., bonds with above-median complexity score (Panel B), the percentage of daisy chains–i.e., chains initiated and ended by the same dealer entity (Panel C), the percentage of transactions that use coarse dollar prices (Panel D), and the percentage of transactions that use fine yields (Panel E), by dealer centrality and reciprocity tercile. Percentages are reported as the total number of transaction chains initiated by each dealer centrality and reciprocity tercile-sorted group that fall under a given category (e.g. 42.36% of the transaction chains initiated by the dealers with the lowest centrality and reciprocity involve the par amount of less than \$100K.)

Panel A. Percentage of chains involving par size <\$100K

		Reciprocity				
		Lowest	Medium	Highest	**Total**	**H-L**
	Lowest	42.36%	49.96%	51.07%	43.98%	8.71%***
	Medium	62.21%	47.93%	55.83%	53.02%	-6.38%***
Centrality	Highest	84.16%	45.18%	47.70%	47.54%	-36.46%***
	Total	46.85%	47.21%	50.49%		
	H-L	41.80%***	-4.78%***	-3.37%***		

Panel B. Percentage of chains involving complex bonds

		Reciprocity				
		Lowest	Medium	Highest	**Total**	**H-L**
	Lowest	8.91%	10.48%	14.62%	9.51%	5.71%***
	Medium	16.36%	12.48%	12.81%	13.29%	-3.55%***
Centrality	Highest	21.53%	15.11%	13.50%	14.25%	-8.03%***
	Total	10.53%	13.16%	13.36%		
	H-L	12.62%***	4.63%***	-1.12%***		

Panel C: Percentage of chains with same initiating and ending dealer

		Reciprocity				
		Lowest	Medium	Highest	**Total**	**H-L**
	Lowest	0.56%	4.95%	11.31%	1.90%	10.7%***
	Medium	0.33%	12.12%	19.17%	12.19%	18.84%***
Centrality	Highest	0.37%	1.05%	1.92%	1.56%	1.55%***
	Total	0.52%	7.15%	7.98%		
	H-L	-0.19%***	-3.80%***	-9.39%***		

Panel D. Percentage of coarse-price transactions

		Reciprocity				
		Lowest	Medium	Highest	**Total**	**H-L**
	Lowest	7.15%	5.41%	11.14%	7.44%	3.99%***
	Medium	4.94%	7.12%	4.07%	5.78%	-0.87%***
Centrality	Highest	6.36%	6.46%	5.18%	5.74%	-1.18%***
	Total	6.68%	6.65%	5.62%		
	H-L	-0.79%***	1.05%***	-5.96%***		

Panel E. Percentage of fine-yield transactions

		Reciprocity				
		Lowest	Medium	Highest	**Total**	**H-L**
	Lowest	22.29%	20.28%	15.57%	21.13%	-6.72%***
	Medium	21.47%	29.98%	22.73%	26.19%	1.26%***
Centrality	Highest	16.52%	18.59%	24.22%	21.54%	7.70%***
	Total	21.63%	24.61%	22.62%		
	H-L	-5.77%***	-1.69%***	8.65%***		

Table 7: Regressions of markups on dealer centrality and reciprocity by dealer network size

This table presents estimates of the same specification used in column 4 of Table 3 on transaction chain samples sorted by the initiating dealers local network size quintile, where local network size is defined as the number of unique counterparties a dealer has bought from or sold to in the past 30 trading days. Column (1) reports results for dealers with the smallest networks and column (5) for dealers with the largest networks. All samples are limited to dealers within local networks containing at least two members. Independent variables are standardized to have a mean of 0 and unit variance. Standard errors are clustered by bond and time and are shown in parentheses, with statistical significance denoted at the 10% (*), 5% (**), and 1% (***) levels.

	(1)	(2)	(3)	(4)	(5)
Local network size percentile:	0-20	21-40	41-60	61-80	81-100
Dep. Variable:	Total Markup	Total Markup	Total Markup	Total Markup	Total Markup
Dealer Centrality	-0.074***	0.092***	0.193***	0.151***	0.194***
	(0.006)	(0.010)	(0.011)	(0.012)	(0.015)
Reciprocity	0.095***	-0.135***	-0.263***	-0.230***	-0.127***
	(0.006)	(0.010)	(0.011)	(0.013)	(0.022)
Controls	Yes	Yes	Yes	Yes	Yes
Market share controls	Yes	Yes	Yes	Yes	Yes
Bond FE	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes	Yes
Observations	267,917	289,157	258,065	272,531	272,630
Adjusted R^2	0.27	0.24	0.24	0.23	0.25

Table 8: Regressions of markups on community membership and size

This table presents estimates of the same specification used in column 4 of Table 3 substituting dealer centrality and reciprocity variables with a dummy variable, *Same community*, which equals one if both dealers belong to the same Louvain community and is zero otherwise. Seller and buyer community sizes reflect the sizes of the respective dealers' Louvain communities. All of the independent variables are standardized to have a mean of 0 and unit variance. Columns (1)-(3) report regression results on the entire sample of dealers that are assigned to a community, and columns (4)-(6) report regression results for the sample of dealers assigned to a community that does not belong to the group of dealers with high reciprocity and low centrality. Independent variables are standardized to have a mean of 0 and unit variance. Standard errors are clustered by bond and time and are shown in parentheses, with statistical significance denoted at the 10% (*), 5% (**), and 1% (***) levels.

	(1)	(2)	(3)	(4)	(5)	(6)
Dealers	All	All	All	Ex. low c., high. r	Ex. low c., high. r	Ex. low c., high. r
Dep. Variable	Total Markup	Total Markup	Total Markup	Total Markup	Total Markup	Total Markup
Same community	0.494***	0.500***	0.507***	0.343***	0.343***	0.351***
	(0.020)	(0.020)	(0.020)	(0.038)	(0.038)	(0.038)
Seller community size	0.054***		0.055***	0.056***	0.056***	0.057***
	(0.004)		(0.004)	(0.005)	(0.005)	(0.005)
Same community ×	-0.120***			-0.079***	-0.079***	
Seller community size	(0.014)			(0.016)	(0.016)	
Buyer community siz		0.043***	0.044***			0.042***
		(0.003)	(0.003)			(0.003)
Same community ×		-0.087***	-0.125***			-0.096***
Buyer community size		(0.010)	(0.010)			(0.011)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Market share controls	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Bond FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	691,298	691,311	691,297	659,126	659,126	659,125
Adjusted R^2	0.17	0.17	0.17	0.16	0.16	0.16

Table 9: Feature importance of the determinants of the trades within the same community.

This table documents the feature importance for 20 of the most important features used by LightGBM (Ke et al., 2017) and Random Forest (Breiman, 2001) classifier models, bootstrapped across 100 random seeds ranging from 1 to 100. The feature importance score in LightGBM models is determined by the total number of improvements in all decision trees. In Random Forest Models, the feature importance is defined as the decrease in node impurity weighted by the probability of reaching that node. The training sample uses 70% of the total sample of the inter-dealer trades, while the test sample uses the remaining 30%. The forecast accuracy statistics are reported for the test sample.

Variable	LightGBM	Random Forest
Buyer local network size	88.00	0.04
Seller local network size	81.64	0.07
Seller's retail market share	70.12	0.01
Purchase fraction	61.21	0.03
Seller's reciprocity	58.52	0.08
Buyer's retail market share	56.05	0.01
Buyer's reciprocity	54.18	0.06
Seller's par-weighted reciprocity	46.96	0.07
Seller's 30-day return volatility	45.83	0.00
Buyer's EW centrality	41.85	0.02
Buyer's 30-day aggregate return	41.79	0.00
Buyer's VW centrality	41.26	0.03
Seller's EW centrality	40.25	0.02
Seller's VW centrality	36.96	0.05
Reverse purchase fraction	34.74	0.02
Total value traded in reverse trades	34.52	0.01
Buyer's eigenvector centrality	34.52	0.02
Total number of reversals in past 30 days	34.24	0.01
Seller's 30-day aggregate return	32.08	0.00
Average dealer centrality in the given chain	31.69	0.01
Correct Predictions, %	90.53	71.80
Predicted same when different, %	0.04	29.12
Predicted different when same, %	9.43	0.08

Appendix A. Variable Definitions

Below are the definitions of the variables used in our analysis:

1. Dependent Variables

Total Markup: Par-weighted difference of selling prices in the dealer-customer trades and purchasing prices in initial customer-dealer trades divided by purchasing price in initial customer-dealer trade, defined by Li and Schürhoff (2019).

Abnormal Markup: Difference between actual markup and predicted markup. Predicted markup is estimated using the following specification that excludes dealer centrality and dealer reciprocity:

$$\widehat{Markup}_{i,t} = Trade\ controls + Dealer\ Inventory + Bond\ Controls + State\ FE + Month\ FE, \quad (1)$$

where *Bond Controls* are bond characteristics, and *Trade Controls* include natural logarithms of par value for small, medium, and large trades.

DC Trade: Binary variable set to 1 if the given transaction is a customer sale.

N of Dealers: Total number of different dealers in a given transaction chain.

1-day Reversal: Binary variable set to 1 if the buyer and seller in a given transaction were on the opposite ends of the trade in the previous trading day.

7-day Reversal: Binary variable set to 1 if the buyer and seller in a given transaction were on the opposite ends of the trade in the previous 7 trading days.

30-day Reversal: Binary variable set to 1 if the buyer and seller in a given transaction were on the opposite ends of the trade in the previous 30 trading days.

Same Community: Binary variable set to 1 if the buyer and seller belong to the same Louvain community Blondel et al. (2008). For a definition of communities, see Appendix D.

2. Explanatory Variables

Chain Length: Total number of transactions composing the transaction chain.

Dealer Inventory: Aggregate dealer inventory, calculated over the past 30 trading days and standardized by subtracting its mean and dividing by its daily standard deviation, defined by Li and Schürhoff (2019).

Dealer Market Share: The dealer's national market share of par traded in customer purchases, calculated on a rolling basis using trades from the past 30 trading days.

Dealer Market Share (State): The dealer's market share of par traded in customer purchases in the state where the bond is issued, calculated on a rolling basis using the past 30 trading days.

Dealer Market Share (Bond): The dealer's market share of par traded in customer purchases for the specific bond, calculated on a rolling basis using the past 30 trading days.

Retail Market Share: The dealer's national market share of par traded in customer sales, calculated on a rolling basis using trades from the past 30 trading days.

Small: Binary variable set to 1 if the par volume is less than \$100K.

Medium: Binary variable set to 1 if the par volume is between \$100K and \$1M.

Large: Binary variable set to 1 if the par volume is greater than \$1M.

log(par): Natural logarithm of the par volume of the trade.

Maturity: Natural logarithm of the number of days left until the bond matures.

Seasoning: Natural logarithm of the number of days since the bond was issued.

Issue Size: Natural logarithm of the bond's issue size.

Rating: Bond's credit rating based issued by Standard and Poor's, arranged in a descending manner of credit quality (e.g. 1: AAA, 2: AAA-, ..., 20: D).

Junk: Binary variable set to 1 if bond's credit rating is BB+ and worse.

Unrated: Binary variable set to 1 if the bond is not rated.

Insured: Binary variable set to 1 if the bond has a measure of default protection.

General Obligation: Binary variable set to 1 if the bond is backed by credit and taxing power of the municipality it is issued by.

Callable: Binary variable set to 1 if the bond has a call provision.

Sinking Fund: Binary variable set to 1 if the bond has a sinking fund feature.

Bank Qualified: Binary variable set to 1 if the bond is qualified for tax exemptions for bank investors.

Taxable: Binary variable set to 1 if the bond's interest income is subject to federal or state taxes.

Subject to AMT: Binary variable set to 1 if the bond is a "private activity" bond and the interest may be subject to Alternative Minimum Tax (AMT).

Low Centrality: Binary variable set to 1 if the dealer belongs to the first centrality tercile (the tercile with the lowest centrality).

High Centrality: Binary variable set to 1 if the dealer belongs to the third centrality tercile (the tercile with the highest centrality).

Low Reciprocity: Binary variable set to 1 if the dealer belongs to the first reciprocity tercile (the tercile with the lowest reciprocity).

High Reciprocity: Binary variable set to 1 if the dealer belongs to the third reciprocity tercile (the tercile with the highest reciprocity).

30-day aggregate return: Total percentage of markup charged by dealer i to all other dealers and customers they sold to, net of the markup they paid to other dealers and percent losses in the first inter-dealer transaction after customer purchase, aggregated over prior 30 trading days.

30-day aggregate return volatility: Standard deviation of the returns for dealer i, aggregated over different bonds they traded over the prior 30 trading days.

Bond Complexity: Sum of Insured, General Obligation, Callable, Sinking Fund, Bank Qualified, and Subject to AMT binary variables.

3. Centrality, Reciprocity, and Community variables

Eigenvector (Dealer Centrality): Relative score based on the number of connections of dealers with central and peripheral dealers, calculated over the prior 30 trading days.

Node reciprocity: see Appendix C.

Static (Dynamic) Community: Louvain community identifier. See Appendix D for the definition of Louvain communities. Calculated over the entire sample (past 30 trading days).

Static (Dynamic) Community Length: Number of dealers in dealer i's community. Size is set to 1 if a dealer is not a part of any community. Calculated over the entire sample (past 30 trading days).

Local Network Size: Number of unique counterparties dealer i bought from or sold to in the past 30 trading days. Calculated on a dealer level and identified separately for the buyer, seller, and chain initiator.

Appendix B. Data Filters, Roundtrip Transaction Chains

Our sample was prepared using the data filters from Li and Schürhoff (2019). These filters include removing transactions on bonds with a par volume of less than 5,000, winsorizing the sample using price filters for coupon and zero-coupon bonds separately, and restricting to seasoned issues (issued over 90 days prior) with more than 1 year until maturity. We began with 45.62M trades. After leaving only bonds with a par volume of over 5,000, we were left with 42.27M transactions. Once we left bonds with at least one year until maturity and 90 days since issuance, the number of trades decreased to 32.31M. Finally, removing the trades with prices more than 3 standard deviations from the mean in either direction, our final sample before chain construction was 31.20M.

The table below illustrates the progression of our sample after applying each of the filters.

Filter	Customer Purchases	Inter-dealer Transactions	Customer Sales
Initial Sample	10,170,036	17,678,982	17,778,700
Keep trades with a par value of at least 5000	9,340,373	16,457,010	16,476,605
Keep trades at least 90 days since issuance	8,888,977	13,166,118	12,024,561
Keep trades at least 1 year until maturity	8,285,387	12,637,347	11,389,753
Remove trades with prices more than 3 SD away from mean	7,862,082	12,468,011	10,915,560

After applying the filters from Li and Schürhoff (2019), we construct roundtrip transaction chains, defined as series of transactions on the same bond with the same par value starting with a customer purchase and ending with a customer sale with none or many inter-dealer transactions in between. The roundtrip transaction chains must also be performed on the same bond with the same par size. The full algorithm of roundtrip chain construction is described in Appendix B of Li and Schürhoff (2019). Our final roundtrip transaction chain sample contains 1,536,551 roundtrip chains, with 739,411 chains having at least one inter-dealer transaction.

Appendix C. Node Reciprocity Calculation

We define our node reciprocity measure following Newman et al. (2002), as the total number of dealers with whom dealer i has had two-way interactions (buying and selling) divided by the total number of other dealers dealer i has interacted with over the previous 30 trading days:

$$r_{i,t} = \frac{L^{\leftrightarrow}_{i,t}}{max(L_{i,t},1)}, \tag{2}$$

where $L_{i,t}$ is the total number of interactions in both ways (in the context of municipal bond markets, the number of both purchases and sales) for dealer i for the 30-day period t, and $L^{\leftrightarrow}_{i,t}$ is the number of interactions in both ways for dealer i with counterparties who both bought from and sold to dealer i, calculated over the 30-day period t.[19] Period t is defined as a 30-trading-day-long period that ends on the trading day before the transaction. For example, for February 15, 2014, the reciprocity measure for dealer i will be calculated over the period between January 3, 2014 and February 14, 2014. For the dealer i that does not have any two-way interactions, node reciprocity for dealer i is 0.

A visualization of node reciprocity can be seen in the figure below, which uses Kamada and Kawai (1989) drawing method. The figure also uses various core-periphery thresholds based on the number of two-way interactions, similar to Li and Schürhoff (2019), who use 10,000 trades as a threshold. Dealers with above-median reciprocity levels are highlighted in red.

[19] This is a slight variation of the measure in Newman et al. (2002) in that we divide by $max(L_{i,t},1)$, rather than $L_{i,t}$, to avoid a potential division by zero error due to dealers which only have both way interactions.

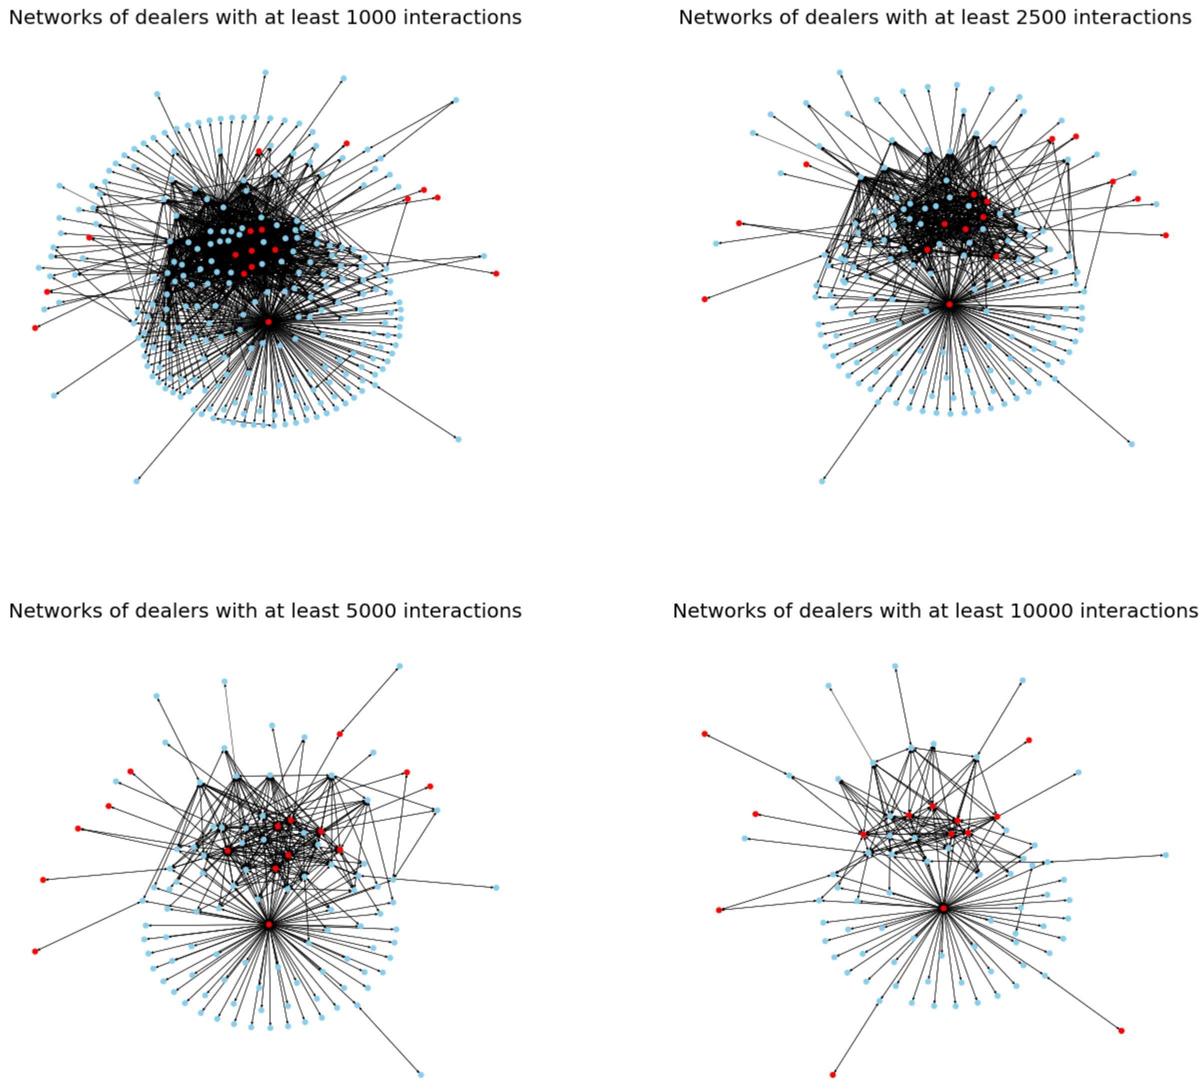

Fig. A1. **Order flow in the network with different core-periphery thresholds**

Appendix D. Community Detection Algorithm

We use the Louvain (Blondel et al., 2008) algorithm of community detection. Since the algorithm is random, we fix the seed and make sure we obtain results from 1 different seed (ranging from 1 to 100) to alleviate several concerns. First, we ensure that our results are not driven by a singular random state (Jain and Madhyastha, 2019). Second, it reduces the likelihood of dealers being arbitrarily "forced" into communities they do not naturally belong to, as varying the random seeds can otherwise result in inconsistent community assignments for these dealers across iterations. By aggregating the outcomes from multiple runs, we improve the reliability and consistency of our community classifications.

After obtaining the 100 community samples, we group them and delete observations with members that belong to a different distinct community, which is best demonstrated in an example below:

Dealer ID	Seed 1 community	Seed 2 community	Seed 3 community	...	Seed 100 community
A	10	19	13	...	4
B	10	19	13	...	4
C	10	14	11	...	4
...	...	...	...	...	...

Dealers A and B belong to the same community, while dealer C does not, even though in some cases, such as in seeds 1 and 100, dealer C was in the same community as dealers A and B. As a result, dealer C will be either in a different community or unassigned (in their community). For this example, we assume that dealers A and B have the same community assignment in the other 96 seeds.

Community detection is applied to both static (full 4 years) and dynamic (30-day rolling windows) datasets, with fairly similar results. We use the dynamic community assignments in reported results.

We also select different resolution parameters, which determine the size of the communities, explained in Blondel et al. (2008). The higher resolution parameter results in a larger community size. We also confirm the robustness of our findings across different resolution parameters (1, 20, 40, and 60) and types of communities (dynamic and static). The subsequent discussion will focus on the outcomes for dynamic Louvain communities, derived using a resolution parameter of 40. The figures below show the number of communities for a given size for different resolution parameters in both static and dynamic communities.

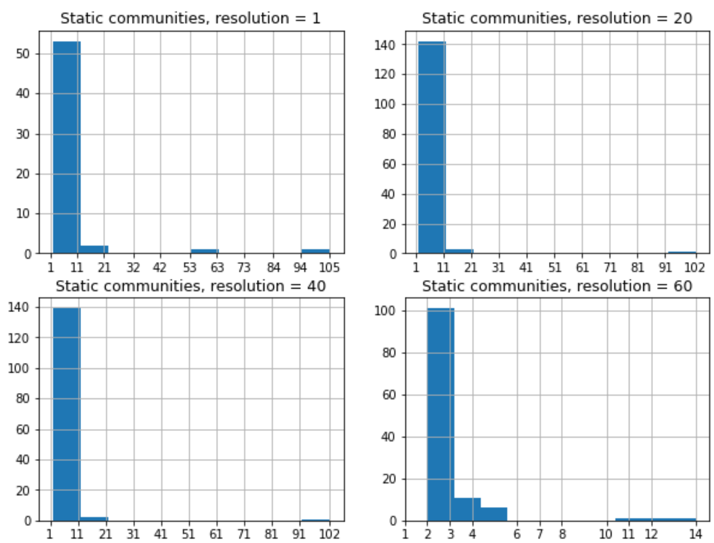

Fig. A2. **Number of static communities of a given length, by resolution parameter**

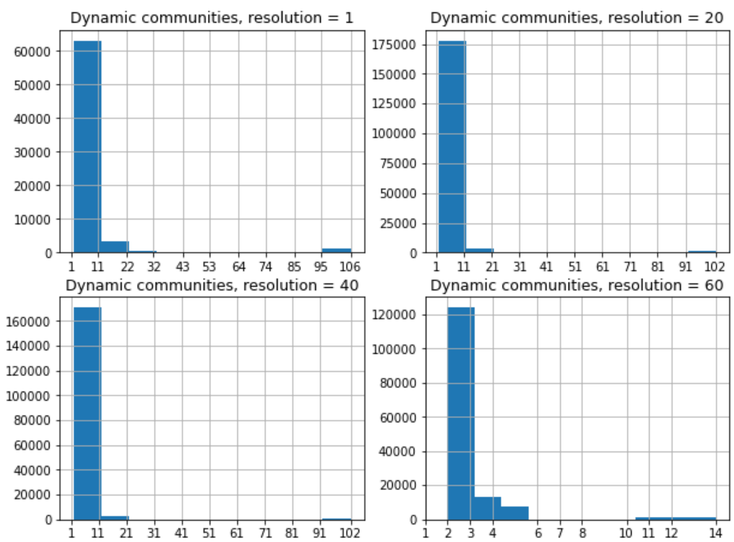

Fig. A3. **Number of dynamic communities of a given length, by resolution parameter**

The table on the next page shows the number and the average size of Louvain communities, the percentage of dealers assigned to a community, as well the size of the largest Louvain community, and the number of dealers that did not belong to any community after we obtained our bootstrap sample. The summary statistics are reported for the default resolution parameter of 1, as well as parameters of 20, 40, and 60. In addition to different resolution parameters, we report statistics for two different types of communities - dynamic (recovered from the entire 4-year sample) and static (constructed over the previous 30 trading days, consistent with centrality calculations in Li and Schürhoff (2019)). On average, the default parameter of 1 returns the largest communities, but also results in a large percentage of dealers being unassigned to a community. Resolution parameters of 20 and 40 tend to return much smaller communities, even though with either parameter, there is one large community of 102 dealers. Finally, if we set the resolution parameter to 60, the proportion of dealers in a community is roughly the same as when we set the parameter to 1, but the size of the communities reduces radically as compared with when the resolution parameter is set to 40, In regressions and machine learning models that use communities as dependent or outcome variable, we find that our results are robust to different resolution parameters.

Table A1: Louvain Communities descriptive statistics.
This table documents the summary statistics for Louvain communities, reported for 4 different resolution parameters: 1, 20, 40, and 60. The statistics are calculated for both dynamic (formed over the prior 30 trading days) and static (for the full, 4-year sample) communities. The statistics reported are the total number of static and dynamic communities, the average number of dealers in the communities, the percent of dealers that belong to a community, the size of the largest Louvain community, and the number of dealers unassigned to a community at any point. For the definition of the Louvain communities, see Appendix.

Resolution	Statistic	Static	Dynamic
	Number of communities	427	68,801
	Average size (N of dealers)	6.05	6.07
1	% of dealers in a community	48.25	47.59
	Size of largest community	105	106
	N of dealers w/o community	370	212
	Number of communities	340	182,615
	Average size (N of dealers)	3.57	3.54
20	% of dealers in a community	72.87	73.66
	Size of largest community	102	102
	N of dealers w/o community	194	162
	Number of communities	502	174,468
	Average size (N of dealers)	3.57	3.56
40	% of dealers in a community	70.62	70.67
	Size of largest community	102	102
	N of dealers w/o community	210	208
	Number of communities	427	148,589
	Average size (N of dealers)	2.76	2.76
60	% of dealers in a community	46.71	46.71
	Size of largest community	14	14
	N of dealers w/o community	381	381

Table A2: Markups and trading relationships. This table reports the results of the linear regressions of markups on centrality, reciprocity, and purchase/sale fractions (Di Maggio et al., 2017) for the first inter-dealer transactions of each transaction chain. Standard errors are calculated using bond and time clusters and are reported in parentheses below the parameter estimates, and the level of statistical significance is reported for parameters significant at 10 (*), 5 (**), and 1 (***) percent significance levels.

	(1)	(2)
Dealers	All	Excl. low c., high r.
Dep. Variable	Total Markup	Total Markup
Dealer Centrality	-0.113***	-0.002
	(0.007)	(0.018)
Reciprocity	0.054***	-0.084***
	(0.010)	(0.025)
Purchase Fraction	0.138***	0.194***
	(0.050)	(0.058)
Sale Fraction	-0.257***	-0.281***
	(0.026)	(0.029)
Controls	Yes	Yes
Market share controls	Yes	Yes
State FE	Yes	Yes
Month-Year FE	Yes	Yes
Observations	715,151	683,133
Pseudo R^2	0.10	0.09

Table A3: Tercile comparisons of dealer trade participation and market shares. This table contains the total number of trades each dealer type participated in and provides a comparison of sizes of dealer market shares, expressed in percentage of unique bonds traded by the dealer. Panel A shows the total number of trades each dealer group has initiated. Panel B shows the average market shares within the overall municipal bond market in the prior 30 trading days, panel C depicts the total average market share of the dealers on bonds issued in specific states, and panel D shows the average dealer market share on trade of specific bonds.

Panel A. Number of Trades Initiated by a Given Dealer Type

		Reciprocity		
		Lowest	Medium	Highest
Centrality	Lowest	1,158,670 (25.84%)	200,566 (4.47%)	135,686 (3.03%)
	Medium	285,525 (6.37%)	730,124 (16.28%)	479,302 (10.69%)
	Highest	50,724 (1.13%)	564,586 (12.59%)	879,497 (19.61%)

Panel B. Average Dealer Market Share (Overall)

		Reciprocity				
		Lowest	Medium	Highest	**Total**	**H-L**
Centrality	Lowest	0.66%	1.02%	0.34%	0.68%	-0.32***
	Medium	1.50%	2.96%	2.68%	2.61%	1.18%***
	Highest	3.04%	5.13%	3.42%	4.03%	0.38***
	Total	0.86%	3.49%	2.98%		
	H-L	2.38%***	4.11***	3.08%***		

Panel C. Average Dealer Market Share (Within State)

		Reciprocity				
		Lowest	Medium	Highest	**Total**	**H-L**
Centrality	Lowest	1.29%	1.44%	0.62%	1.26%	-0.67***
	Medium	3.47%	3.79%	3.48%	3.63%	0.01%
	Highest	8.06%	5.64%	4.22%	4.82%	3.84***
	Total	1.83%	4.15%	3.74%		
	H-L	6.77%***	4.20***	3.60%***		

Panel D. Average Dealer Market Share (Within Specific Bonds)

		Reciprocity				
		Lowest	Medium	Highest	**Total**	**H-L**
Centrality	Lowest	68.37%	65.86%	63.06%	67.66%	-5.30***
	Medium	69.83%	68.25%	68.63%	68.66%	-1.19%
	Highest	71.94%	66.25%	68.72%	67.90%	-3.21***
	Total	68.71%	67.21%	68.30%		
	H-L	3.57%***	0.39%**	5.66%***		

Table A4: Tercile comparisons second dealer relationships in the transaction chain.
This table contains the average values of centrality and reciprocity of the second dealer in a given transaction chain, based on the centrality and reciprocity terciles of the initial dealer. Panel A shows the average centrality values. Panel B shows the average reciprocity values. Panel C reports the probability of the initiating the reversal trade.

Panel A. Centrality of Second Dealers

		Reciprocity				
		Lowest	Medium	Highest	**Total**	**H-L**
	Lowest	0.05	0.06	0.01	0.05	0.04***
	Medium	0.14	0.14	0.15	0.14	0.01***
Centrality	Highest	0.18	0.19	0.22	0.21	0.04***
	Total	0.07	0.15	0.18		
	H-L	0.13***	0.12***	0.21***		

Panel B. Reciprocity of Second Dealers

		Reciprocity				
		Lowest	Medium	Highest	**Total**	**H-L**
	Lowest	0.22	0.39	0.77	0.28	0.55***
	Medium	0.25	0.44	0.53	0.43	0.28***
Centrality	Highest	0.20	0.44	0.59	0.53	0.39***
	Total	0.07	0.15	0.18		
	H-L	-0.02***	0.05**	0.18***		

Panel C. Percent of Chains Involving 1-day reversals

		Reciprocity				
		Lowest	Medium	Highest	**Total**	**H-L**
	Lowest	16.36	22.98	38.79	20.22	22.43%***
	Medium	15.06	36.80	31.52	31.24	16.46%***
Centrality	Highest	12.82	27.69	30.79	28.10	17.97%***
	Total	15.81	31.68	32.07		
	H-L	-3.54%***	4.71%***	-8.00%***		

Table A5: Feature importance of the determinants of the trades within the same community (outside of top 20)

This table documents the feature importance for the features used by LightGBM (Ke et al., 2017) and Random Forest (Breiman, 2001) classifier models, bootstrapped across 100 random seeds ranging from 1 to 100. The feature importance score in LightGBM models is determined by the total number of improvements in all decision trees. In Random Forest Models, the feature importance is defined as the decrease in node impurity weighted by the probability of reaching that node. The training sample uses 70% of the total sample of the inter-dealer trades, while the test sample uses the remaining 30%. The forecast accuracy statistics are reported for the test sample.

Variable	LightGBM	Random Forest
Buyer bond 30-day return volatility	30.59	0.00
Average dealer centrality in the chain	29.01	0.01
Seller eigenvector centrality	24.47	0.04
Seasoning	24.30	0.00
Log(transaction size)	24.10	0.00
Total par traded in the previous month's reverse r.	23.72	0.01
Dealer's sale fraction to the counterparty	20.82	0.04
Dealer total market share	20.23	0.00
Reverse sale fraction to the counterparty	18.26	0.01
Chain length * reverse purchase fraction	18.06	0.01
Maturity	17.28	0.00
Average VW reciprocity in the chain	17.19	0.03
Chain length * backward number of chains	17.01	0.00
Median EW centrality in the chain	15.89	0.00
Initiator interactions	15.53	0.02
Chain length * purchase fraction	15.47	0.01
Average reverse purchase fraction in the chain	15.16	0.01
Median EW reciprocity in the chain	15.13	0.00
Seller's state market share	14.57	0.00
Chain length * dealer's sale fraction to the counterparty	12.98	0.01
Average VW centrality in the chain	12.52	0.01
Chain length * average purchase fraction in the chain	12.38	0.00
Chain length * average sale fraction in the chain	12.09	0.00
Average EW centrality in the chain	11.19	0.01
Chain length * average reverse purchase fraction in the chain	11.05	0.01
Issue size	10.82	0.00
Initiator's state market share	10.81	0.00
Chain length * forward number of chains	10.78	0.00
Median VW reciprocity in the chain	10.76	0.00
Chain length	10.34	0.00

Variable	LightGBM	Random Forest
Chain length * reverse sale fraction	9.91	0.00
WA Year FE	9.59	0.00
Median VW centrality in the chain	9.24	0.00
Average sale fraction in the chain	8.86	0.00
Median eigenvector centrality in the chain	8.77	0.00
Forward number of chains	8.73	0.00
Average EW reciprocity in the chain	7.77	0.03
Average purchase fraction in the chain	7.52	0.01
Backward number of chains	7.44	0.00
Chain length * VW centrality	7.10	0.01
Chain length * average reverse sale fraction in the chain	6.92	0.00
Initiator Inventory	6.77	0.00
Chain length * EW reciprocity	6.75	0.01
Buyer state market share	6.32	0.00
Chain length * VW reciprocity	5.87	0.01
Average forward number of chains in the chain	5.54	0.00
Forward Number of daisy chains	5.24	0.01
Average forward number of chains in the chain	5.22	0.00
Initiator market share (State)	4.94	0.00
Chain length * eigenvector centrality	4.55	0.01
Average backward number of daisy chains in the chain	4.35	0.00
Chain length * average backward number of chains in the chain	4.21	0.00
WE Rating FE	4.19	0.00
Initiator local network size	3.82	0.00
Average backward number of chains in the chain	3.71	0.00
Chain length * EW centrality	3.50	0.00
Chain length * average forward number of chains in the chain	3.48	0.00
Initiator market share (bond)	3.23	0.00
Number of reversals (7 days)	3.17	0.00
Log(par)*small	2.50	0.00
Is a reversal within 30 days	2.40	0.00
Backward number of daisy chains	2.25	0.01
Log(par)*medium	1.98	0.00
Subject to AMT	1.71	0.00
Log(par)*large	1.42	0.00
Sinking fund provision	1.37	0.00
Callable	1.33	0.00
Average forward number of daisy chains	1.18	0.00
Chain length * average backward number of daisy chains	1.00	0.00

Variable	LightGBM	Random Forest
Chain Length * forward number of daisy chains	0.97	0.00
Chain Length * forward number of daisy chains	0.84	0.01
Number of reversals (1 day)	0.76	0.00
Bank qualified	0.63	0.00
Chain Length * backward number of daisy chains	0.62	0.00
WA Month FE	0.53	0.00
Is a reversal within 1 day	0.40	0.00
WA State FE	0.35	0.00
Insured	0.34	0.00
Taxable	0.33	0.00
General obligation	0.33	0.00
Unrated	0.10	0.00
Junk	0.08	0.00
Medium	0.07	0.00
Reversal 7 day	0.06	0.00
Small	0.03	0.00
Large	0.01	0.00
Chain length * average eigenvector centrality in the chain	0.00	0.01