The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXXV, NO. 2 • APRIL 2020

Relationship Trading in Over-the-Counter Markets

TERRENCE HENDERSHOTT, DAN LI, DMITRY LIVDAN,
and NORMAN SCHÜRHOFF*

ABSTRACT

We examine the network of trading relationships between insurers and dealers in the over-the-counter (OTC) corporate bond market. Regulatory data show that one-third of insurers use a single dealer, whereas other insurers have large dealer networks. Execution prices are nonmonotone in network size, initially declining with more dealers but increasing once networks exceed 20 dealers. A model of decentralized trade in which insurers trade off the benefits of repeat business and faster execution quantitatively fits the distribution of insurers' network size and explains the price–network size relationship. Counterfactual analysis shows that regulations to unbundle trade and nontrade services can decrease welfare.

INSURANCE COMPANIES ARE VITAL FOR risk sharing: In exchange for premium payments, they compensate for loss, damage, injury, treatment, or hardship. To provide this coverage, insurance companies invest in a variety of financial assets. However, corporate bonds dominate their investment portfolios, accounting for almost 70% of their investments with a total value close to

*Terrence Hendershott is with the Walter A. Haas School of Business, University of California. Dan Li is with the Board of Governors of the Federal Reserve System. Dmitry Livdan is with the Walter A. Haas School of Business, University of California. Norman Schürhoff is with the Faculty of Business and Economics and Swiss Finance Institute, University of Lausanne. He is also a Research Fellow of the CEPR. Hendershott provides expert witness services to a variety of clients. He gratefully acknowledges support from the Norwegian Finance Initiative. Li and Livdan have no conflict of interest to disclose. Schürhoff gratefully acknowledges research support from the Swiss Finance Institute and the Swiss National Science Foundation under Sinergia project CRSII1_154445/1 "The Empirics of Financial Stability" and has no conflict of interest to disclose. We thank Philip Bond; the Editor; an associate editor; two anonymous referees; Darrell Duffie; Jerôme Dugast; Burton Hollifield; Haoxiang Zhu; and seminar audiences at California Polytechnic University, Carnegie Mellon University, FED Board, Higher School of Economics Moscow, IDC Herzliya, Imperial College, Johns Hopkins University, Michigan State University, National Bank of Belgium, NY FED, Rice University, Texas A&M University, University of British Columbia, University of Houston, University of Illinois-Chicago, University of Oklahoma, University of Piraeus, University of Wisconsin at Madison, University of Toronto, University of Victoria, WU Vienna, Einaudi Institute for Economics and Finance, 2015 Toulouse Trading and Post-Trading Conference, 2016 AFA, 2016 NBER AP Meeting, 2016 SFI Research Day, 12th Annual Central Bank of France Workshop on the Microstructure of Financial Markets, 2016 LAEF Conference on OTC Markets, 2017 13th Csef-Igier Symposium on Economics and Institutions, and 2017 OTC Markets and their Reform Conference in Rigi for valuable comments and feedback.

DOI: 10.1111/jofi.12864

$4 trillion.[1] To facilitate prompt compensation to policy holders, insurers need to be able to liquidate their holdings quickly without incurring large transaction costs (Koijen and Yogo (2015), Chodorow-Reich, Ghent, and Haddad (2016)). Yet, corporate bonds trade on decentralized over-the-counter (OTC) markets, which are less liquid than centralized exchanges due to search frictions arising from fragmentation and limited transparency (Duffie, Garleanu, and Pedersen (2005, 2007), Weill (2007), Vayanos and Wang (2007)). Insurers therefore have to search for best execution across more than 400 active broker-dealers. In this paper, we investigate whether insurers and other market participants search randomly in the OTC corporate bond market, or whether they build long-term relations with dealers to mitigate search frictions.

Regulatory data provide information on the transactions between more than 4,300 insurers and their dealers over the period 2001–2014. Using these data, we begin by empirically examining insurers' choice of trading network and the relation between the network choice and transaction prices. Figure 1 provides examples of two different insurer-dealer trading networks. Panel A shows buys and sells for an insurer that trades exclusively with a single dealer, whereas Panel B shows buys and sells for an insurer that trades with multiple dealers over time. We find that insurers tend to form small but persistent dealer networks. At one extreme, approximately one-third of insurers trade with a single dealer annually. At the other extreme, a small fraction of insurers trades with up to 40 dealers each year. The overall degree distribution follows a power law with exponential tail starting at about 10 dealers. When we estimate trading costs as a function of network size N, we find that costs are nonmonotone in N—costs decline with N for small networks and then increase once N exceeds 20 dealers.

Our results provide insights into which models of trade in OTC markets better describe the empirical evidence on client-dealer networks and trading costs. In random search models, clients repeatedly search for best execution without forming a finite network of dealers (Duffie, Garleanu, and Pedersen (2005, 2007), Lagos and Rocheteau (2007, 2009), Gavazza (2016)). The fact that insurers form finite dealer networks suggests that adding dealers must be costly for insurers. Traditional models of strategic search (e.g., Stigler 1961) assume that each additional dealer imposes a fixed cost on insurers, in which case insurers add dealers to improve prices up to the point at which the marginal benefit of doing so equals the fixed cost. However, this leads prices to improve monotonically in network size, which is inconsistent with our finding that trading costs are nonmonotone in network size.

To rationalize the empirical evidence, we build a model of decentralized trade in which insurers—or more generally, clients—establish relationships with multiple dealers and trade off the benefits of repeat business and faster execution. We model the relationship between clients and dealers as having

[1] Insurers are major suppliers of capital to corporations. Schultz (2001) and Campbell and Taksler (2003) estimate that insurers hold between 30% and 40% of corporate bonds and account for about 12% of trading volume.

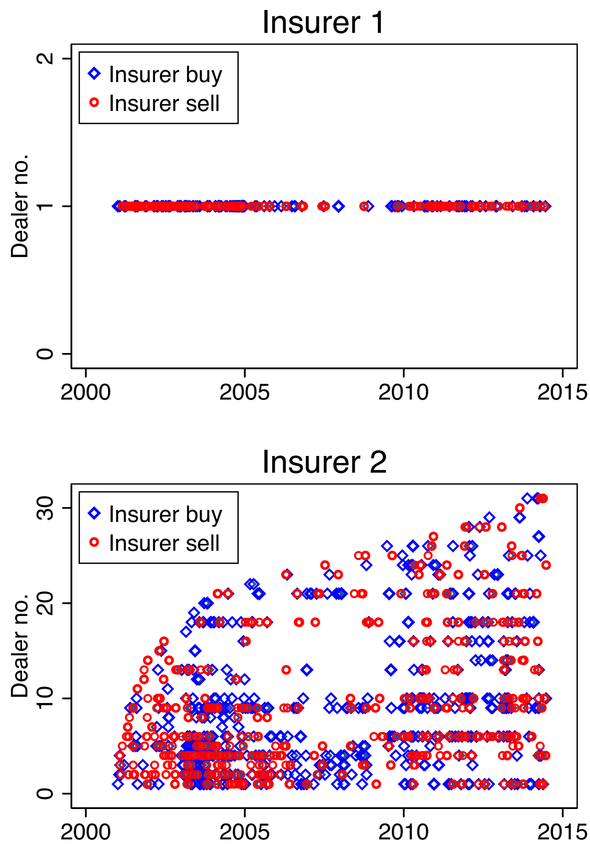

Figure 1. Two types of dealer-client trading networks. This figure shows the buy (squares) and sell (circles) trades of two insurance companies with different dealers. We sort dealers on the vertical axis by the first time they trade with the corresponding insurance company. (Color figure can be viewed at wileyonlinelibrary.com)

two components. The first component captures the extent of repeat trading between a client and a dealer in her network, which both the client and the dealer take into account when negotiating the terms of a transaction. The second component captures all client-dealer business unrelated to bond trading. The latter component includes transactions in other securities, the ability to purchase newly issued securities, as well as other soft dollar and nonmonetary transfers such as investment research.

In our model, a single console bond trades on an interdealer market that clients can access only through dealers. Dealers have search intensity λ and upon trading with a client transact at the competitive interdealer bid and ask prices. Clients initially start without a bond but stochastically receive trading shocks with intensity η that lead them to simultaneously contact N dealers to buy. Trading intensity η, and hence N, vary across clients. A client's effective search intensity, $\Lambda = N\lambda$, is increasing in the number of dealers. The first dealer to find the bond captures all benefits from the transaction. Our trading mechanism is therefore identical to repeat winner-takes-all races (Harris and Vickers (1985)). The bond's purchase price is set by Nash bargaining. Once an owner, the client stochastically receives a liquidity shock forcing her to sell the bond. The mechanics of the sell transaction are the same as those of the buy transaction. Both dealers and clients derive value from repeat transactions, leading to price improvement for more frequent clients in Nash bargaining.

Existing OTC models provide predictions about network size or prices, but not both. Random search models assume that investors may contact all other counterparties. Other models allow investors to choose specific networks or markets, but exogenously fix the structure of those networks. For instance, in Vayanos and Wang (2007), investors search for a counterparty between two markets for the same asset: a large market with faster execution but higher transaction costs, and a small market with slower execution but lower transaction costs. Neklyudov and Sambalaibat (2016) use a similar setup as in Vayanos and Wang (2007) but investors choose between dealers with either large or small interdealer networks instead of asset markets. By contrast, both network size and transaction prices are endogenously determined in the equilibrium of our model.

Our analysis highlights an important trade-off leading to optimal network size, N^*, and transaction prices. In particular, clients trade off repeat relations with dealers and the benefits of dealers competing to provide faster execution. More dealers lead to faster execution. However, the value of repeat relations declines in the number of dealers. Dealers compensate for a loss of repeat business by charging higher spreads. Eventually the costs of having a larger network outweigh the benefits and dealers' spread starts to increase with network size. This pattern corresponds to the empirical nonmonotonicity of trading costs with respect to network size. The value of repeat relations diminishes more slowly with the addition of dealers for clients with larger trading intensity as dealers compete for more repeat business. Therefore, these larger clients use more dealers and get better execution as benefits from having more repeat business outweigh the costs of having a larger network.

The model can quantitatively match the cross-section of insurers' observed network sizes and transaction prices. Doing so requires structural estimation of the model parameters not directly observable in the data, Θ. Clients' trading intensity, η, is the one parameter for which we observe its cross-sectional distribution, $p(\eta)$, in the data. Insurer i's trading shock intensity, η_i, can be estimated as the average number of bond purchases per year over the sample period. We estimate η_i separately for each insurer using multiple years of trade data, which enable us to construct $p(\eta)$. The model provides the optimal network size, $N^*(\Theta, \eta_i)$, for each client i as a function of its trading intensity η_i. The model predictions allow us to estimate the unobservable model parameters Θ by matching the model-implied distribution of network size, $p(N^*)$, to its empirical counterpart.

Using the structurally estimated parameters along with the distribution of trading intensities quantitatively reproduces both the empirical distribution of network size and the dependence of trading costs on network size found in the data. The model estimates reasonable unobserved parameters: Dealers can find the bond within a day or two, and insurers' average holding period ranges from two to four weeks. The sunk cost of obtaining quotes from a dealer is small, whereas insurers' estimated willingness to pay for immediacy is high. Allowing dealers' bargaining power to decrease with insurers' trading intensity improves the model's fit. Dealers' bargaining power when insurers are forced to sell is high and relatively insensitive to insurers' trading frequency. In contrast, dealers' bargaining power when insurers are buying declines significantly with insurers' trading frequency.

The regulatory authorities in the United States and Europe address OTC market frictions using different approaches. Regulatory initiatives in the United States, including the Dodd-Frank Act, encourage dealer competition by fostering multilateral electronic trading platforms. But these erode long-term client-dealer relationships. In Europe, the Markets in Financial Instruments Directive (MiFID) II requires the unbundling of trading and nontrading dealer services, which are now supposed to be priced and sold separately. Little evidence exists either empirically or theoretically on the welfare implications of the proposed regulations for dealers and their clients.

Our model enables us to quantify the regulatory impact on trading costs and the value of repeat interactions. We perform counterfactual analysis by reducing the probability of repeat trading with the same dealer or by eliminating the nontrade value of relationships. All types of insurers incur higher transaction costs when the probability of repeat trading with the same dealer is reduced, for example, due to mandatory use of anonymous electronic trading venues/facilities.[2] The impact on various insurer types and sizes is different. Insurers that trade more frequently and, as a result, have larger networks see a smaller increase in transaction costs than insurers that trade less frequently and tend to trade repeatedly with one to five dealers.

Unbundling trade and nontrade insurer-dealer business as stipulated by MiFID II regulations reduces the optimal network size and decreases transaction costs for all insurers but the least active. The presence of nontrade relationship value leads trading networks to be larger than what is required to minimize trading costs. The loss of nontrade relationship value can induce insurers that trade infrequently to cease trading completely, reducing market liquidity. Thus, unless the nontrade relationship value can be captured efficiently in an unbundled market, insurer welfare declines.

Relation to literature: Our paper complements the empirical literature on the microstructure of OTC markets and its implications for trading, price formation, and liquidity. Edwards, Harris, and Piwowar (2007), Bessembinder, Maxwell, and Venkataraman (2006), Harris and Piwowar (2006), Green, Hollifield, and Schürhoff (2007) document the magnitude and determinants of

[2] Our counterfactual analysis takes the current decentralized market structure as given. It cannot address what would happen if, say, the trading mechanism were centralized. The latter is an interesting and important topic left for future research.

transaction costs for investors in OTC markets. Our paper deepens our understanding of OTC trading costs by using the identities of all insurers along with their trading networks and execution costs. These characteristics help explain the substantial heterogeneity in execution costs observed in the studies above. Harris (2015) and O'Hara, Wang, and Zhou (2018) examine best execution in OTC markets without formally studying investors' optimal network choice.

Our study also relates to an empirical literature on the value of relationships in financial markets. Similar to our paper, Bernhardt et al. (2005) show that on the London Stock Exchange broker-dealers offer greater price improvement to more regular customers.[3] However, Bernhardt et al. (2005) do not examine client-dealer networks and in their centralized exchange setting quoted prices are observable. Also similar to our paper, Afonso, Kovner, and Schoar (2013) study the overnight interbank lending OTC market and find that a majority of banks in the interbank market form long-term, stable, and concentrated lending relationships that significantly affect the transmission of liquidity shocks across the market. But Afonso, Kovner, and Schoar (2013) do not formally model the network and do not observe transaction prices. DiMaggio, Kermani, and Song (2017) study interdealer relationships in the OTC market for corporate bonds. Our focus, in contrast, is on client-dealer relationships.[4]

The role of the interdealer market in price formation and liquidity provision are the focus in Hollifield, Neklyudov, and Spatt (2017) and Li and Schürhoff (2019). These studies examine the effect of heterogeneity in dealers' network centrality on the liquidity they provide and the prices they charge. By contrast, we focus on how heterogeneity in clients' trading intensity affects their network choices and transaction prices.

The search-and-matching literature is vast. Duffie, Garleanu, and Pedersen (2005, 2007) provide a prominent treatment of search frictions in OTC financial markets, whereas Lagos and Rocheteau (2007, 2009), Weill (2007), Feldhütter (2012), Hugonnier, Lester, and Weill (2018), Neklyudov (2019), and Üslü (2019) generalize the economic setting. However, these papers do not focus on repeat relationships, nor do they provide incentives to investors to have a finite network size. Gavazza (2016) structurally estimates a model of trading in decentralized markets with two-sided one-to-one search and bilateral bargaining using aircraft transaction data. At the market level, Gavazza (2016) quantifies the effects of market frictions on prices and allocations. We use the structural estimation of our one-to-many search-and-match model to quantify the effects of client-dealer relationships on execution quality in the OTC market for corporate bonds.

Directed search models allow for heterogeneous dealers and investors, as well as arbitrary trade quantities. These models typically rely on a concept of

[3] For a comprehensive theoretical model of loyalty, see Board (2011).

[4] Our paper also relates to a growing literature on trading in a network, for example, Gale and Kariv (2007), Gofman (2011), Condorelli and Galeotti (2012), Colliard and Demange (2014), Atkeson, Eisfeld, and Weill (2015), Chang and Zhang (2015), Babus (2016), Glode and Opp (2016), Babus and Hu (2017), and Babus and Kondor (2018). However, while these papers also consider persistent one-to-one dealer-client relationships, our model focuses largely on clients' networks.

competitive search equilibrium proposed by Moen (1997) in the context of labor markets; examples include Guerrieri, Shimer, and Wright (2010), Guerrieri and Shimer (2014), and Lester, Rocheteau, and Weill (2015). These papers explain assortative matching between clients and dealers and show that heterogeneity affects prices and liquidity. However, the matching technology employed by these papers is one-to-one, which limits the network size to a single dealer.

The remainder of the paper is organized as follows. Section I describes the data. Sections II documents our findings on insurer trading activity, dealer networks, and trading costs. In Section III, we develop an OTC model with repeat trading and dealer networks. We test the predictions of our model and provide counterfactuals in Section IV. Finally, Section V concludes.

I. Data

Insurance companies, which invest heavily in the corporate bond market, file quarterly reports of trades of stocks and long-term bonds to the National Association of Insurance Commissioners (NAIC). For each trade, the NAIC Schedule D data include the dollar amount, par value, and date of the transaction,[5] the insurer code, the counterparty dealer name, and the direction of the trade for both parties, whether the insurance company was buying from or selling to the dealer. NAIC started collecting the Schedule D data almost a decade before the Trade Reporting and Compliance Engine (TRACE) became available, which makes the NAIC transaction data an important data source for studies of the corporate bond market. For example, Schultz (2001) uses these data to provide one of the earlier estimates of transaction costs of corporate bonds. Ellul, Jotikasthira, and Lundblad (2011) provide a thorough discussion of the NAIC data. The NAIC data are uniquely suitable for our study as they allow us to observe, for each trade, the identities of both the customer (insurer) and the dealer.[6] The regulatory NAIC data differ from publicly available sources in this respect.

There are several differences between the raw filing data used in this paper and, for instance, the widely used "Time and Sales Data" from Mergent Fixed Income Security Database (FISD) available on Wharton Research Data Services (WRDS) (e.g., Bessembinder, Maxwell, and Venkataraman (2006)). Most importantly, the public data do not identify the insurers, which is crucial to identifying relationships in our paper. Moreover, the Mergent data only use filings from the last quarter of each year, which report all purchases and sales made over the calendar year. However, the annual reports suppress some trades that are contained in the quarterly report. For completeness, we compile the trade data using quarterly reports for Q1–Q3 and all trades in Q4 from the annual report.

To clean the raw NAIC data, we first limit attention to corporate bond transactions. To focus on secondary trading, we further limit attention to

[5] The NAIC data preclude intraday analysis as they do not include trades' time stamps.
[6] Even the supervisory TRACE data contain only the identity of the dealers.

trades more than 60 days after issuance and more than 90 days to maturity. We additionally remove nonsecondary transactions that are associated with redemptions and calls. One of the most challenging yet important tasks of cleaning the NAIC data relates to standardization of counterparty (dealer) identities. The raw filings report hand-typed counterparty names, which contain typos and alternative names for a given dealer. We use fuzzy matching and human checking to standardize dealer names. Details on the data cleaning and filtering procedure are provided in the Appendix. Our final NAIC sample covers all corporate bond transactions between insurance companies and dealers reported in NAIC from January 2001 to June 2014.

We supplement the NAIC data with a number of additional sources. Bond and issuer characteristics come from Mergent FISD and Lipper eMAXX. Insurer financial information comes from A.M. Best and SNL Financial. For execution cost analysis, we incorporate daily corporate bond quotes (dealer bids) from Merrill Lynch. The final sample contains 506,000 insurer buys and 497,000 insurer sells.

Table I reports descriptive statistics for the corporate bond trades (Panel A) and insurers (Panel B) in our 2001–2014 sample. There are 4,324 insurance companies in our sample. Insurance companies fall into three groups based on their product types: (I) health, 617 companies (14% of the sample); (II) life, 1,023 companies (24% of the sample); and (III) property and casualty (P&C), 2,684 companies (62% of the sample). Health insurance companies account for 16.3% of trades and 4.4% of yearly trading volume, and trade on average with 6.59 dealers each year. Life insurance companies account for the majority (46.9%) of trades and 70.4% of yearly trading volume, and trade on average with 8.06 dealers each year. P&C insurance companies comprise 36.8% of trades and 25.2% of yearly trading volume, and trade on average with 4.81 dealers each year.

The distribution of trading activity is skewed, with the top 10 insurance companies accounting for 6.3% of trades and 14.3% of trading volume. These insurers use almost 30 dealers on average, which is much higher than the sample average of 5.83 dealers per insurer. The top 100 insurers account for 27.8% of trades and 45.3% of trading volume. The 3,000 smallest insurers use on average 3.76 dealers a year.

Order splitting is not common: An insurer trades the same bond multiple times on just 1.2% of all insurer-days (13,000 out of 1.1 million trades). This finding is consistent with larger transactions being cheaper to execute, which eliminates incentives for order splitting.

Insurers trade a variety of corporate bonds. The average issue size is quite large at $917 million and is similar across an insurer's buys and sells. The average maturity is nine years for insurer buys and eight years for insurer sells. Bonds are on average 2.88 years old, with sold bonds being a little older at 3.09 years. Finally, 75% of all bonds traded are investment grade, whereas only 1% are unrated with the remainder being high yield. Privately placed bond trades comprise a small part of our sample at 8%.

Table I

Descriptive Statistics

This table reports descriptive statistics for trades (Panel A) and insurers (Panel B) in our sample from 2001 to 2014. Panel A reports the average across all trades over the sample period. Panel B reports the yearly average across insurers.

Panel A: Trades			
	All Trades	Insurer Buys	Insurer Sells
No. of trades (k)	1,003	506	497
Trade par size ($mn)	1.80	1.73	1.87
Bond issue size ($mn)	916.66	921.37	911.87
Bond age (years)	2.88	2.67	3.09
Bond remaining life (years)	8.54	8.94	8.13
Private placement (%/100)	0.08	0.08	0.07
Rating (%/100)	–	–	–
Investment-grade	0.74	0.76	0.72
High-yield	0.25	0.23	0.28
Unrated	0.01	0.01	0.01

Panel B: Insurers ($N = 4{,}324$)			
	Volume ($mn)	No. of Trades	No. of Dealers
All insurers	17.32	9.52	5.83
Insurer type			
Health (617, 14%)	10.66	21.74	6.59
Life (1,023, 24%)	103.00	37.71	8.06
P&C (2,684, 62%)	14.08	11.29	4.81
Insurer activity			
Top 10	2,111.88	517.92	29.89
11–100	509.49	233.04	22.07
101–1,000	75.66	46.22	11.56
1,001+	3.80	4.24	3.76
Insurer characteristics:	Mean (*SD*)	–	–
Insurer size	4.97 (0.90)	–	–
Insurer RBC ratio	3.36 (0.35)	–	–
Insurer cash-to-assets	3.49 (10.79)	–	–
Life insurer	0.24 (0.42)	–	–
P&C insurer	0.62 (0.48)	–	–
Insurer rated A–B	0.37 (0.38)	–	–
Insurer rated C–F	0.01 (0.07)	–	–
Insurer unrated	0.53 (0.39)	–	–

The risk-based capital (RBC) ratio measures an insurer's capital relative to the riskiness of its business. The higher the RBC ratio, the better capitalized the firm is. Insurer size is captured by reported assets. The cash-to-assets ratio is cash flow from insurance business operations divided by assets.

Overall, there is a large degree of heterogeneity on the client side. Insurance companies buy and sell large quantities of different corporate bonds and execute these transactions with one to as many as 40 dealers.

II. Empirical Evidence on Insurer Trading Networks

This section empirically characterizes insurers' trading intensity and the size of their trading networks. We investigate the determinants of both the extensive margin, that is, the number of trades, and the intensive margin, that is, the total dollar volume traded, of insurer trading in a given year. Both margins reveal that insurers have heterogeneous trading needs.

A. Insurer Trading Activity

We start with univariate analysis. The majority of insurers do not trade often at an annual frequency: About 30% of insurers trade just once per year, whereas only 1% of insurers make at least 25 trades per year. This result is consistent with evidence in Table I indicating that while the top 100 insurers constitute just 0.23% of the insurer population, they account for as much as 32% of all trades in our sample. The mean number of trades per year is 19, and the median is 14, with several insurers making more than 1,000 trades in some years and up to 2,200 trades in a year.

Figure 2 shows the average number of trades per year across insurers. For the large fraction of insurers that do not trade in a given month, we report an annual figure. The annual distributions follow a power law with $p(X) \propto 0.27 \times X^{-1.21}$ for all insurer trades (buys and sells) combined. The power law is $0.34 \times X^{-1.31}$ for insurer buys (Panel A) and $0.40 \times X^{-1.58}$ for insurer sales (Panel B). Visually, the power law distributions for insurer buys and sales in Figure 2 look similar. This suggests that insurers buy and sell at similar rates, even though these rates vary significantly across insurers.

We next examine which characteristics explain the heterogeneity in trading intensity. Table II documents the determinants of the intensive margin (trading volume in \$bn, column (1)) and the extensive margin (number of trades, column (2)) of annual trading by insurance companies using pooled regressions with time fixed effects. The specification consists of the trade par size, insurer, and bond characteristics, as well as the variation in trade size and bond characteristics across all trades of the insurer during the year. The measures capturing variation in these characteristics capture complexity in insurers' portfolios and their need for more frequent rebalancing and dealer specialization. Insurer characteristics include the insurer's size, cash-to-assets ratio, type, RBC ratio, and rating. Bond and trade characteristics include size, age, maturity, rating, a private placement dummy, and trade size. Insurer size, the RBC ratio, and the dependent variables are log-transformed. All regressors are averaged across all trades of the insurer during the period and lagged by one year.

Our evidence is consistent with insurance companies having persistent portfolio rebalancing needs. Logarithms of both measures of trade intensity are persistent: the coefficient on lagged log-volume is 0.67 and the coefficient on lagged log-number-of-trades is 0.76. Both coefficients are statistically significant at the 1% level.

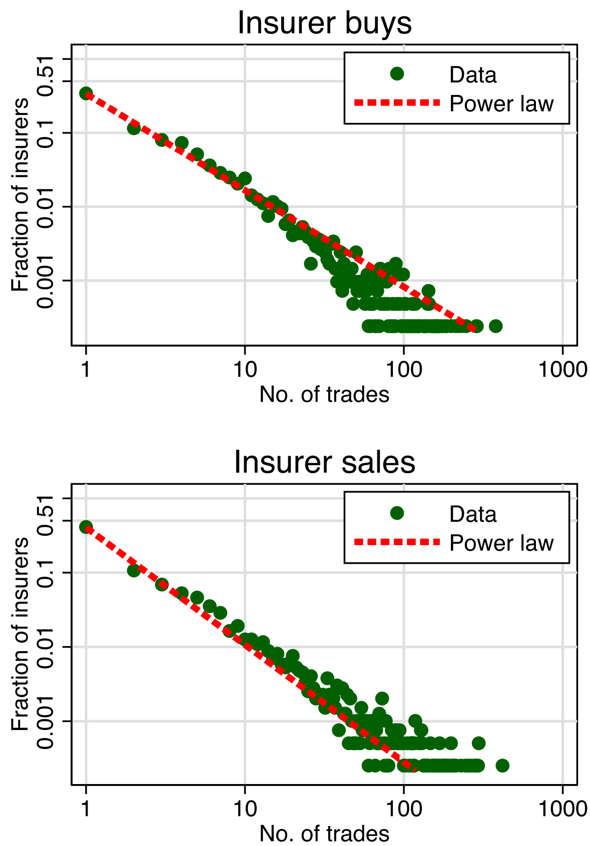

Figure 2. Insurer trading activity. This figure shows the distribution of the number of insurer buys (top) and sales (bottom) per year. We use a log-log scale. (Color figure can be viewed at wileyonlinelibrary.com)

Insurer trading strongly correlates with insurer size, type, and quality as well as with bond type and bond variety as these variables explain 79% of the variation in annual trading volume and 65% of the variation in annual number of trades. A twofold increase in insurer size increases trading volume by 22%. Larger insurance companies and insurers with a higher cash-to-assets ratio also trade more often and submit larger orders. Insurers with a higher RBC ratio trade less often than insurers with a low RBC ratio. Both margins of trading increase with the insurer's rating, that is, insurers with the lowest rating (C–F) trade less than higher rated insurers. Life insurers tend to submit larger orders.

Both margins of bond turnover increase as bond ratings decline: Lower rated bonds are traded more often and in larger quantities. Insurers tend to trade privately placed bonds less, potentially because they own fewer of them than publicly placed bonds. Both margins of bond turnover decline with par size and bond age, indicating that the majority of insurers are long-term investors.

Table II

Insurers' Trading Activity

This table presents results on the determinants of insurance company trading activity. We measure trading activity as the total dollar volume traded in a given year and, alternatively, by the number of trades over the same time horizon. All dependent variables are log-transformed by 100 × log(1 + x). All regressors are averaged across all trades of the insurer during the period and lagged by one period. Estimates are from pooled regressions with time fixed effects. Standard errors are adjusted for heteroskedasticity and clustering at the insurer and time levels. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Determinant	Volume ($mn) (1)	No. of Trades (2)
Lagged volume	0.67***	–
Lagged no. of trades	–	0.76***
Insurer size	21.95***	14.51***
Insurer RBC ratio	−1.53	−4.68***
Insurer cash-to-assets	0.27***	0.26***
Life insurer	4.96***	0.27
P&C insurer	−1.06	−3.89**
Insurer rated A–B	5.13**	5.80***
Insurer rated C–F	1.97	−0.43
Insurer unrated	6.39***	5.79**
Trade par size	−3.62***	−3.22***
Bond issue size	−0.00	0.00
Bond age	−0.79***	−1.25***
Bond remaining life	−0.04	−0.05
Bond high-yield rated	4.62***	4.65***
Bond unrated	−6.67	−12.81*
Bond privately placed	−5.37	−3.56
Variation in trade size	4.50***	1.52***
Variation in issue size	0.00	0.00
Variation in bond age	0.37	0.39
Variation in bond life	0.52***	0.65***
Variation in bond rating	−0.35	0.04
Variation in rated-unrated	39.60*	6.16
Variation in private-public	1.62	−1.60
No. varieties traded	9.23***	13.64***
Year fixed effects	Yes	Yes
R^2	0.789	0.646
N	30,029	30,029

Neither measure of trade intensity depends on bond issue size and remaining life as the respective coefficients are not statistically significant.

Finally, both trading volume and the number of trades decline if an insurer trades more bond varieties. However, a specific variety can have the opposite effect on trading intensity. For instance, both measures of trading intensity increase with variation in bond rating and bond life. This result is consistent with insurers increasing trading intensity when rebalancing their portfolios, that is, shifting from high-yield to investment-grade bonds or from younger to older bonds.

Overall, Table II reveals large heterogeneity in trading intensity across insurers. Trading intensity depends on the variety of bonds traded, bond-specific characteristics, and the insurer's type and quality. We now turn to how these characteristics affect insurers' choice of dealer network.

B. Properties of Insurer Networks

The results above document a large degree of heterogeneity in insurers' trading intensity, with some insurers trading twice per day while others trade just once per year. In this section, we study how many dealers they trade with over time and how persistent these networks are. While insurers should have heterogeneous demands for their dealer network depending on their trading intensity, the concentration and persistence of their trading reveals basic network formation mechanisms.

We start with the two illustrative insurer-dealer relationships depicted in Figure 1. These examples show that insurers do not trade with a dealer randomly picked from a large pool of corporate bond dealers. Instead, insurers buy from the same dealers that they sell bonds to and they build long-term repeat, but nonexclusive relationships. We examine how representative the examples in Figure 1 are and how insurer characteristics affect network size.

Figure 3 plots the degree distribution across insurers by year, that is, the fraction of insurers trading on average across all years with the given number of dealers, using a log-log scale. The figure shows that insurers trade with up to 31 dealers each year. We note that some insurers trade with as many as 40 dealers, but such instances represent less than 1/10,000 of the sample. Exclusive relationships are dominant, with almost 30% of insurers trading with a single dealer in a given year.[7] The degree distributions in Figure 3 follow a power law with exponential tail starting at about 10 dealers. This pattern is consistent with insurers building networks that they search randomly within. Fitting the degree distribution to a Gamma distribution by regressing the log of the probabilities of each N on a constant, the logarithm of N, and N yields the following coefficients:

$$\begin{aligned}&\text{For all insurer trades:} \quad p(N) \propto N^{0.15} e^{-0.20N},\\ &\text{For insurer buys:} \quad p(N) \propto N^{-0.12} e^{-0.22N},\\ &\text{For insurer sales:} \quad p(N) \propto N^{0.01} e^{-0.24N}. \qquad (1)\end{aligned}$$

[7] Because a large number of insurers do not trade in a given year. Figures 2 and 3 are not directly comparable. Figure 2 plots the average number of trades per year, rounded to the next integer and winsorized from below at one. For about half of insurer-years, there is no trade by an insurer. Figure 3 does not impute a zero for a given year if the insurer does not trade. Hence, 27.7% of insurers use a single dealer in a year when they trade, whereas 14.8% of insurers trade just once in a year. It follows that many insurers that trade more than once in a year use a single dealer. Figure 4 further examines the relation between trading activity and network size.

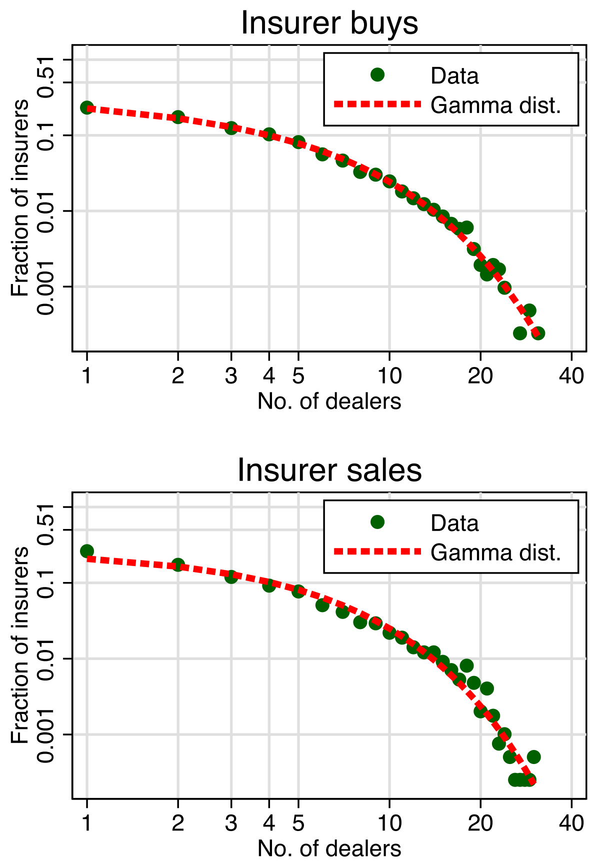

Figure 3. Size of insurer-dealer trading networks. This figure shows the degree distribution for insurer-dealer relationships by year for insurer buys (top) and sales (bottom). We use a log-log scale. (Color figure can be viewed at wileyonlinelibrary.com)

Table III reports the determinants of insurer-dealer network size using pooled regressions with time fixed effects.[8] We measure the size of the trading network using the number of dealers that an insurance company trades with in a given year. We log-transform all dependent variables by $100 \times \log(1 + x)$ and average all regressors across all trades by the same insurer during the year and lag them by one year. We perform the estimation on the full sample (column (1)) and, to examine how these vary with insurer size, on subsamples of small and large insurers based on asset size. We classify an insurer as small

[8] Table IA.I in the Internet Appendix, which is available in the online version of the article on *The Journal of Finance* website, reports correlations between insurer-dealer volume, network size, and insurer characteristics. The correlations between insurer-dealer volume, network size, and insurer size are all significantly positive, but less than 1. The correlations with other insurer characteristics are weak.

Table III
Size of Insurers' Trading Network

This table presents results on the determinants of the size of insurers' trading network. We measure the size of the trading network as the number of dealers that an insurance company trades with in a given year. See the caption of Table II for additional details. Standard errors are adjusted for heteroskedasticity and clustering at the insurer and time levels. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Determinant	All Insurers (1)	Small Insurers (2)	Large Insurers (3)
Lagged no. of dealers	0.75***	0.62***	0.75***
Insurer size	9.95***	7.93***	5.04***
Insurer RBC ratio	−3.58***	−3.64***	0.11
Insurer cash-to-assets	0.15***	0.15***	0.11**
Life insurer	−0.13	0.75	−2.77**
P&C insurer	−2.60***	−1.20	−3.59***
Insurer rated A–B	4.12***	6.00***	0.13
Insurer rated C–F	−1.54	−0.18	−5.42***
Insurer unrated	3.05*	3.98*	−2.86**
Trade par size	−1.92***	−1.23***	−1.24***
Bond issue size	−0.00	−0.00	0.00
Bond age	−0.91***	−0.77***	−1.10***
Bond remaining life	−0.08	−0.11*	0.10
Bond high-yield rated	1.23	−2.93*	2.56
Bond unrated	−13.74***	−12.50***	−7.73
Bond privately placed	−3.85	1.10	−10.22
Variation in trade size	0.51	−0.10	0.81*
Variation in issue size	0.00	0.00	0.00
Variation in bond age	0.50**	0.56**	0.16
Variation in bond life	0.50***	0.41***	0.11
Variation in bond rating	0.05	0.19	0.36*
Variation in rated-unrated	3.60	−11.38	−24.50
Variation in private-public	1.20	−4.33	3.17
No varieties traded	5.88***	2.71***	12.50**
Year fixed effects	Yes	Yes	Yes
R^2	0.614	0.350	0.593
N	30,029	18,033	11,996

if it falls in the bottom three size quartiles and as large if it falls in the top quartile of the size distribution.

Column (1) indicates that insurer size and type, bond characteristics, and bond variety matter for the size of the dealer network. Large insurers, which Table II shows have larger trading intensity, trade with more dealers. Insurers with demand for more bond variety have larger networks even after controlling for their size (columns (2) and (3)). Higher quality insurers, that is, insurers with a higher cash-to-assets ratio and higher rating, have larger networks, potentially due to it being cheaper for a dealer to set up a credit account for higher quality insurers. However, this effect matters only for smaller insurers,

Table IV

Persistence in Insurers' Trading Network

This table reports switching probabilities, p(No. of dealers in $t+1$ | No. of dealers in t), for choosing a network size conditional on the insurer's behavior in the past year.

	No. of Dealers Next Year			
No. of Dealers This Year	1	2–5	6–10	>10
1	0.61	0.30	0.06	0.03
2–5	0.20	0.54	0.20	0.06
6–10	0.05	0.31	0.40	0.24
>10	0.01	0.07	0.17	0.75

as column (2) indicates, because smaller insurers face larger adverse selection problems in forming permanent links with dealers. Insurers with greater variety in the bonds they trade also have larger networks.[9] Overall, these findings suggest that insurers' network choice is endogenous and influenced by multiple factors.

Table III shows persistence in the size of the network, with the coefficient on lagged network size equal to 0.75 (column (1)). This result is mostly driven by large insurers, as this coefficient equals only 0.62 for small insurers. Table IV examines this result in more detail by reporting statistics for the frequency with which insurers adjust their network size. We compute the likelihood that an insurer uses a certain number of dealers in a given year and compare it to the corresponding number in the following year. The transition probabilities are reported in Table IV. Trading relationships are persistent from year to year. This is especially true for exclusive relationships as the probability of staying with a single dealer each year is equal to 0.61.[10] Insurers with more than one dealer are unlikely to switch to a single dealer as the annual switching probabilities are equal to 0.20 for insurers with two to five dealers and 0.06 for insurers with 6–10 dealers. Insurers with the largest networks (>10 dealers) tend to maintain large networks over time, with a 75% probability of staying with a large network. The distribution of insurers shown in Figure 3 together with stable network sizes are difficult to reconcile with a "pure" random search model à la Duffie, Garleanu, and Pedersen (2005, 2007).

[9] One potential explanation for this result could be that dealers tend to specialize in a few bonds and, as a consequence, this relation arises mechanically for insurers that demand greater bond variety. Table IA.II in the Internet Appendix examines whether some dealers trade all bonds and whether insurers use more than one dealer to trade the same bond. The 10 most active dealers trade virtually all active bonds (bonds with more than 200 trades for a bond in the sample)—the fraction of these bonds traded ranges from 95.9% to 99.8%. Table IA.III in the Internet Appendix shows that for each of these same active bonds, the 10 most active insurers trade in each of them with more than one dealer on average. These statistics show that specialization does not appear to be a factor for large dealers and large insurers, which represent a substantial fraction of trading.

[10] For consecutive insurer-years with one dealer, the dealer is the same in both years 75% of the time.

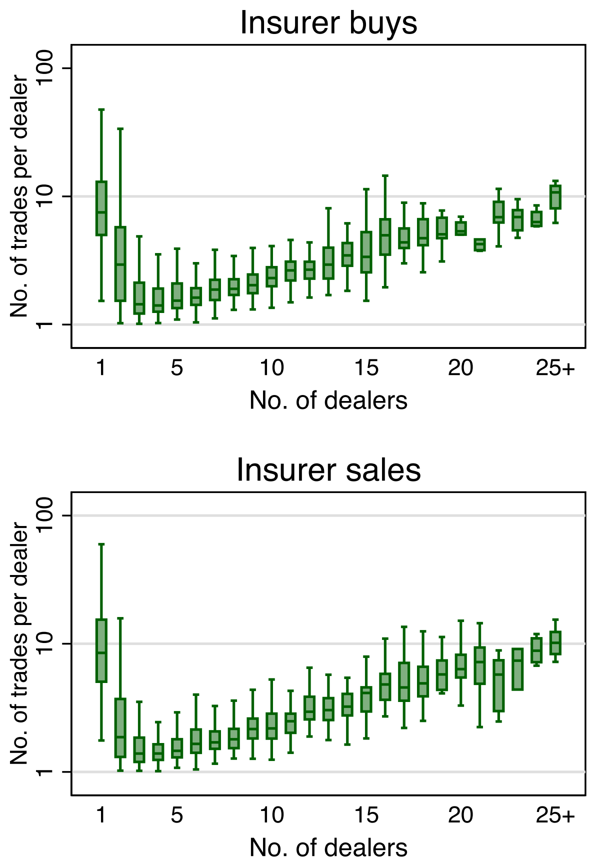

Figure 4. Insurer trading activity and network size. This figure shows the number of insurer buys per year per dealer (top) and insurer sales per year per dealer (bottom) for all insurers with more than one trade per year on average. Boxes correspond to the values at the lower and upper quartiles, with the medians indicated by the lines inside a box. Whiskers indicate the lower and upper adjacent Tukey values. We use a log-linear scale. (Color figure can be viewed at wileyonlinelibrary.com)

One potential concern is that the relationship between the size of the insurer and the size of its network arises mechanically. Small insurers trade once or twice per year and thus need only a single dealer, whereas large insurers need to execute many trades and therefore use many dealers. As a consequence, small exclusive networks are pervasive simply because they are used by small insurers. To examine this issue in Figure 4, we plot the number of insurer buys (left) and sales (right) per dealer and year by network size. The box and whisker plot provides information on the distribution of trading frequencies per dealer for each network size. Although there is a fair amount of cross-sectional variation for small networks, the average number of trades per dealer decreases when going from one to four dealers. Starting at five dealers, the number of trades per dealer is increasing in network size.

This nonmonotonicity of trades per dealer in network size is inconsistent with network size arising mechanically from trading needs.

In the next section, we study the relation between client-dealer networks and execution costs.

C. Insurer Trading Costs and Networks

Tables II and III suggest that bond characteristics impact insurers' trading intensity. To control for bond, time, and bond-time variation, we compare transaction prices to daily bond-specific Bank of America-Merrill Lynch (BAML) bid (sell) quotes. BAML is the largest corporate bond dealer, transacting with more than half of all insurers for almost 10% of trades and volume. The BAML (bid) quotes can be viewed as representative quotes for insurer sales and enable us to measure prices relative to a transparent benchmark price. The BAML quotes essentially provide bond-time fixed effects, which would be too numerous to estimate in our sample. Our measure of relative execution cost, given in basis points, is defined as

$$\text{Execution cost (bp)} = \frac{\text{BAML Quote} - \text{Trade Price}}{\text{BAML Quote}} \times (1 - 2 \times \mathbf{1}_{\text{Buy}}) \times 10^4, \qquad (2)$$

where $\mathbf{1}_{\text{Buy}}$ is a dummy variable indicating whether the insurer is buying or selling. Because some quotes may be stale or trades may be misreported, leading to extreme cost estimates, we winsorize the distribution at the 1% and 99% levels.

Execution costs depend on the bond being traded, time, whether the insurer buys or sells, the insurer's characteristics, the dealer's identity and characteristics, and the insurer's network size. To examine relationship-specific effects on execution costs, we control for bond and time fixed effects. In principle, if the BAML data perfectly controls for bond-time effects, the additional bond and time fixed effects are unnecessary.

The relationship component of transaction costs depends on the properties of insurers' networks. Figure 3 and equation (1) indicate that the degree distribution for insurer-dealer relations follows a Gamma distribution. Accordingly, we include both the size of the network, N, and its natural logarithm, $\ln(N)$, as explanatory variables. Both N and $\ln(N)$ are computed over the prior calendar year. We control for seasonality using time fixed effects, α_t, and for unobserved heterogeneity using either bond characteristics or bond fixed effects, α_i. The other explanatory variables consist of insurer characteristics, dealer characteristics, or both. We estimate the following panel regression for execution costs in bond i at time t:

$$\text{Execution cost}_{it} = \alpha_i + \alpha_t + \beta N_{it-1} + \gamma \ln N_{it-1} + \theta X_{it} + \epsilon_{it}. \qquad (3)$$

The set of explanatory variables X includes characteristics of the bond as well as of the insurer and dealer.

Table V

Execution Costs and Investor-Dealer Relationships

This table presents results on the determinants of execution costs. Execution costs are expressed in basis points relative to the Merrill Lynch quote at the time of the trade. Standard errors are adjusted for heteroskedasticity and clustered at the insurer, dealer, bond, and day levels. See the caption of Table II for additional details. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Determinant	(1)	(2)	(3)	(4)
Insurer no. of dealers	−0.22***	0.32***	0.32***	0.32***
ln(Insurer no. of dealers)	–	−6.29***	−6.51***	−6.55***
Insurer size	−3.72***	−3.59***	−3.52***	−3.95***
Insurer RBC ratio	−3.51***	−4.19***	−4.68***	−5.40***
Insurer cash-to-assets	−0.04**	−0.04**	−0.04**	−0.03*
Life insurer	4.43***	4.47***	5.73***	7.21***
P&C insurer	1.72**	1.73**	1.99**	2.82***
Insurer rated A–B	−0.47	0.01	−0.18	−0.50
Insurer rated C–F	11.29*	11.13*	10.90*	12.18*
Insurer unrated	0.74	0.62	0.55	0.14
Insurer buy	39.56***	39.27***	39.71***	40.17***
Trade size × Buy	−0.26***	−0.25***	−0.20**	−0.18*
Trade size × Sell	0.53***	0.50***	0.48***	0.52***
Bond issue size	–	–	−0.00***	−0.00***
Bond age	–	–	0.58***	0.63***
Bond remaining life	–	–	0.81***	0.82***
Bond high-yield rated	–	–	4.54***	4.16***
Bond unrated	–	–	−5.96***	−6.53***
Bond privately placed	–	–	3.38***	3.26***
Dealer size	–	–	–	−5.37***
NYC dealer	–	–	–	−6.66***
Primary dealer	–	–	–	2.43
Dealer leverage	–	–	–	−5.68**
Dealer diversity	–	–	–	0.41***
Dealer dispersion	–	–	–	0.07
Local dealer	–	–	–	0.23
Dealer distance	–	–	–	−0.20
Dealer leverage missing	–	–	–	−6.41**
Dealer dispersion missing	–	–	–	6.09
Bond fixed effects (16,823)	Yes	Yes	No	No
Dealer fixed effects (401)	Yes	Yes	Yes	No
Day fixed effects (3,375)	Yes	Yes	Yes	Yes
R^2	0.154	0.155	0.103	0.098
N	918,279	918,279	918,987	891,875

Table V provides trading cost estimates from panel regressions. We adjust standard errors for heteroskedasticity and cluster them at the insurer, dealer, bond, and day level. The coefficient on insurer buys captures the average bid-ask spread of roughly 40 basis points. Column (1) of Table V shows that execution costs decline with insurer network size. Adding a dealer decreases trading costs by 0.22 basis points. Large insurers pay on average lower execution costs. For instance, an insurer with 10 times as many assets has 3.72 basis points lower trading costs. Similar to the strength of the trading relationship mechanism in Bernhardt et al. (2005), the size coefficient suggests that an increase in the amount of trading between an insurer and a dealer can lead to better prices. However, Table IA.IV in the Internet Appendix shows that the inclusion of insurer-dealer trading volume does not impact the coefficients on network size in the specifications in Table V.

Column (2) adds the logarithm of N to the specification reported in column (1). The coefficient on N switches from -0.22 reported in column (1) to 0.32, whereas the coefficient on the logarithm of N is -6.29. Both coefficients are statistically significant at the 1% level. This result indicates that execution costs are nonmonotone in network size. Improvements in execution quality from having a larger dealer network are exhausted at $N = \frac{6.29}{0.32} \approx 20$. Clients with networks of 40 dealers and 10 dealers pay, on average, the same bid-ask spread of 40 basis points. This finding goes against the traditional wisdom that interdealer competition improves prices. It is also inconsistent with classic static strategic network formation models (e.g., Jackson and Wolinsky (1996)). In these models, a client trades off the fixed costs of adding an extra dealer and better execution due to increased dealer competition, thus making price a monotonically decreasing function of network size. In the next section, we use this and other network-related empirical evidence to motivate an alternative strategic model of finite network formation in which clients and the dealer share the benefits of repeat interactions.

Columns (3) and (4) replace bond and dealer fixed effects with bond and dealer characteristics. NYC-located dealers offer better prices to all insurers, and more diversified dealers charge, on average, higher prices. Bond characteristics matter for execution costs as insurers receive worse prices for special bonds and better prices for bonds with larger issue size.

In the next section, we use our evidence on networks and execution costs to motivate our model of OTC markets.

III. Model

The model is stylized but still rich enough to allow for the structural estimation of its primitives from regulatory NAIC data. We view the relationship between insurers and dealers as having two components. The first component corresponds to the repeat trading interactions between insurers and dealers. An insurer repeatedly buys and sells bonds from the dealers in her network. This future repeat trading is taken into account by both the insurer and dealers when negotiating the terms of a transaction. Because our data provide direct evidence on trading relations, we model this component explicitly. The second component captures all other business between insurers and dealers. In practice, it includes transactions in other securities, the ability to purchase newly issued securities, as well as other soft dollar and nonmonetary transfers such as investment research. Because we do not directly observe these nontrade relations in our data, we employ a reduced-form approach to model them. We quantify both components of the insurer-dealer relationship when structurally estimating the model from the NAIC data.

A. Setup and Solution

The economy has a single risk-free perpetual bond that pays a coupon flow C. The risk-free discount rate is constant and equal to r, so that the present value of the bond is $\frac{C}{r}$. To model client-dealer repeat trade interactions, we keep several attractive features of Duffie, Garleanu, and Pedersen (2005) such as liquidity supply/demand shocks on the client side and random search with constant intensity. Following Lester, Rocheteau, and Weill (2015), the bond trades on a competitive market accessible only to dealers.[11] Unlike the frictionless interdealer market in Lester, Rocheteau, and Weill (2015), in our model dealers face search frictions as in Duffie, Garleanu, and Pedersen (2005). We implicitly assume that the structure of insurer-dealer networks does not affect the interdealer market. Dealers therefore buy bonds at an exogenously given price M^{ask} from other dealers and sell bonds to other dealers at an exogenously given price M^{bid}. The bid-ask spread $M^{\text{ask}} - M^{\text{bid}} \geq 0$ reflects trading costs or the cost of carry.

In the model, we use the more generic term "client" instead of insurer. Clients act competitively with respect to other clients by not internalizing other clients' network choices in their own decisions.[12] Each client chooses a network of dealers, N, without knowledge of other clients' decisions. When a client wants to buy or sell a bond, she simultaneously contacts all N dealers in her network. Upon being contacted, each dealer starts searching the competitive dealer market for a seller or buyer with search intensity λ.[13] All dealers in the client's network search independently of each other. Therefore, the effective rate at which a client with N dealers in her network finds a counterparty equals λN.[14] When the client receives a subsequent trading shock, all dealers in the network are contacted to reverse the initial transaction.

Each client pays a fixed cost K per transaction.[15] The cost K corresponds to back-office costs of processing and clearing the transaction. Clients that trade more frequently incur K more often. Although K does not depend directly

[11] The interdealer market obviates the need to track where the entire stock of the asset is held at each moment in time.

[12] Although we do not have direct evidence on the information flows across insurers and between insurers and dealers regarding trades, it is unlikely that such information is shared even voluntary.

[13] Bessembinder et al. (2018) show that corporate bond dealers increasingly hold less inventory and facilitate trade via effectively acting as brokers by simultaneously buying and selling the same quantity of a given bond. We assume that the number of clients is large and therefore a fraction of them is hit by the buy (sell) liquidity shock at any time.

[14] We assume that there is no congestion in the dealers market, as in, for example, Afonso (2011).

[15] The costs of additional dealers could alternatively be modeled as per dealer or per dealer per transaction. Per dealer costs consist of the costs of forming a credit relationship and any other costs of maintaining the relationship independent of the number of trades. Such per dealer costs will immediately lead to clients with larger trading intensity using more dealers.

on network size, clients that choose a larger network trade faster and thus incur K more often. The cost K enters clients' value functions and thereby affects their reservation values in bargaining. However, K does not enter the bargaining directly. We postpone discussion of the impact of K on transaction prices and networks to the next section, where we estimate the model. Clients' search mechanism can be viewed as a winner-takes-all race with the first dealer to find the bond winning the race. The prize is the spread $P^b - M^{\text{ask}}$ when the client buys and $M^{\text{bid}} - P^s$ when the client sells, where P^b (P^s) is the price at which the client buys (sells) the bond from (to) the dealer.

Clients transition through ownership and non-ownership based on liquidity shocks. At these transitions, clients act as buyers and sellers. The discounted transition probabilities and transaction prices link valuations across the owner, non-owner, buyer, and seller states.

A client i starting as a non-owner with valuation $\widehat{V}^{\text{no}}$ is hit by stochastic trading shocks to buy with intensity η_i. Intensity η_i varies across clients and its distribution can be directly inferred from the data. Section II.A characterizes the distribution of trading activity. For the sake of clarity, we omit the subscript i throughout our theoretical analysis.

The client contacts her network of N dealers, leading to her transiting to a buyer state with valuation $\widehat{V}^b$. Let $V^r(\eta, N)$ capture exogenously given relationship-specific nontrade flows to the client from her dealers. In steady state, $\widehat{V}^{\text{no}}$ satisfies the Bellman equation linking it to $\widehat{V}^b$ and V^r:

$$\widehat{V}^{\text{no}} = \frac{1}{1+rdt}[\lambda N dt \widehat{V}^b + (1-\lambda N dt)\widehat{V}^{\text{no}} + rV^r dt], \tag{4}$$

which can be solved to yield

$$\widehat{V}^{\text{no}} = \underbrace{\widehat{V}^b \frac{\eta}{r+\eta}}_{\text{Value from Trading}} + \underbrace{V^r \frac{r}{r+\eta}}_{\text{Nontrade Value}}. \tag{5}$$

We are agnostic regarding whether $V^r(\eta, N)$ is increasing or decreasing in trading intensity and network size, η and N; this will be determined by the data in our structural estimation. For our model, we assume that $V^r(\eta, N)$ is a monotonic function of N and satisfies the Inada condition $\lim_{N\to\infty} \frac{V^r(\eta,N)}{N} = 0$, which helps guarantee that the optimal network size is finite. Below we show that $V^r(\eta, N)$ does not directly affect transaction prices, but it does impact clients' choice of network size.

The buyer purchases the bond from her network at the expected price $\mathrm{E}[P^b]$ and transitions into being an owner with valuation $\widehat{V}^o$. In steady state, $\widehat{V}^b$ satisfies the Bellman equation linking it to $\widehat{V}^o$:

$$\widehat{V}^b = \frac{1}{1+rdt}[\lambda N dt(\widehat{V}^o - \mathrm{E}[P^b] - K) + (1-\lambda N dt)\widehat{V}^b + rV^r dt], \tag{6}$$

which yields

$$\widehat{V}^b = (\widehat{V}^o - \mathrm{E}[P^b] - K)\frac{\lambda N}{r + \lambda N} + V^r \frac{r}{r + \lambda N}. \tag{7}$$

While clients are owners, they receive a coupon flow C and have valuation $\widehat{V}^o$. Non-owners do not receive the coupon flow. With intensity κ, an owner receives a liquidity shock that forces her to become a seller with valuation $\widehat{V}^s$. In steady state, valuations in these two states are linked according to the Bellman equation

$$\widehat{V}^o = \frac{1}{1 + rdt}[dtC + \kappa dt V^s + (1 - \kappa dt)\widehat{V}^o + rV^r dt], \tag{8}$$

which yields the expression for $\widehat{V}^o$:

$$\widehat{V}^o = \frac{C}{r + \kappa} + \underbrace{\widehat{V}^s \frac{\kappa}{r + \kappa}}_{\text{Value From Future Sale}} + V^r \frac{r}{r + \kappa}, \tag{9}$$

where the second term captures the value from future sales. The liquidity shock received by the owner reduces the value of the coupon to $C(1 - L)$ until she sells the bond. After receiving the liquidity shock, she contacts her dealer network expecting to sell the bond for $\mathrm{E}[P^s]$. Upon selling she becomes a non-owner, completing the valuation cycle. Valuations $\widehat{V}^s$ and $\widehat{V}^{\mathrm{no}}$ are related by

$$\widehat{V}^s = \frac{1}{1 + rdt}[dtC(1 - L) + \lambda N dt(\mathrm{E}[P^s] + \widehat{V}^{\mathrm{no}} - K) + (1 - \lambda N dt)\widehat{V}^s + rV^r dt], \tag{10}$$

which can be solved for $\widehat{V}^s$ according to

$$\widehat{V}^s = \frac{C(1 - L)}{r + \lambda N} + (\mathrm{E}[P^s] + \widehat{V}^{\mathrm{no}} - K)\frac{\lambda N}{r + \lambda N} + V^r \frac{r}{r + \lambda N}. \tag{11}$$

This sequence of events continues in perpetuity. We therefore focus on the steady state of the model.

The above valuation equations depend on expected transaction prices. The realized transaction prices are determined by bilateral Nash bargaining. The clients' reservation values are determined by the differences in values between being an owner and non-owner and between being a buyer and seller. Similarly, dealers' reservation values arise from their transaction cycle. Each dealer acts competitively, that is, without taking into account the effect of her actions on the actions of other dealers. In addition, we assume that each dealer internalizes only the trade-specific value of her relation with each client.[16]

[16] We do not require that the dealer internalizes the client's nontrade value V^r of the relationship and therefore we omit it from dealers' valuations. This could happen, for instance, if the dealer is compensated to maximize trading profits and is not incentivized to maximize the enterprise value.

When a client who wants to buy the bond contacts her dealer network, each dealer starts looking for the bond at rate λ and expects to pay the interdealer ask price M^{ask} for the bond. The value to the dealer searching for the bond satisfies

$$U^b = \frac{1}{1+rdt}[\lambda dt(P^b - M^{\text{ask}}) + \lambda N dt U^o + (1-\lambda N dt)U^b], \tag{12}$$

which yields

$$U^b = \underbrace{(P^b - M^{\text{ask}})\frac{\lambda}{r+\lambda N}}_{\text{Transaction Profit/Loss}} + \underbrace{U^o\frac{\lambda N}{r+\lambda N}}_{\text{Value of Future Business}}. \tag{13}$$

The last term in the expression for U^b captures the expected value of future business with the client, which happens with frequency λN. This client, who is now the owner of the bond, becomes a seller with intensity κ and contacts dealers in her network to sell the bond. This generates a value U^o per dealer. The superscript "o" on the dealer's valuation U^o reflects the fact that delivering the bond to a client requires that the dealer takes ownership of the bond in the first place as well as the fact that the client is an owner. Dealer valuation U^o satisfies the Bellman equation

$$U^o = \frac{1}{1+rdt}[\kappa dt U^s + (1-\kappa dt)U^o], \tag{14}$$

which yields the following expression for U^o:

$$U^o = U^s\frac{\kappa}{r+\kappa}, \tag{15}$$

where U^s represents the valuation of the dealer searching to sell the bond. The dealer expects to resell the bond at rate λ for the interdealer bid price M^{bid} and earn a markup of $M^{\text{bid}} - P^s$. Upon selling the bond, the client becomes a non-owner but the dealer still anticipates future business with the client, which happens with frequency λN. This is because this client becomes a buyer with intensity η and approaches the dealer in the future to buy the bond back, thus again generating value U^b to the dealer. As a result, future buy-back business from a given client generates a total value of U^{no} to the dealer. The superscript "no" on the dealer's valuation U^{no} reflects the fact that both the dealer and the client are both non-owners at the time of the purchasing request. Dealer valuation U^{no} satisfies the Bellman equation

$$U^{\text{no}} = \frac{1}{1+rdt}[\eta dt U^b + (1-\eta dt)U^{\text{no}}], \tag{16}$$

which yields

$$U^{\text{no}} = U^b\frac{\eta}{r+\eta}. \tag{17}$$

Finally, the valuation of the dealer searching to sell the bond, U^s, satisfies

$$U^s = \frac{1}{1+rdt}[\lambda dt(M^{\text{bid}} - P^s) + \lambda N dt U^{\text{no}} + (1-\lambda N dt)U^s], \tag{18}$$

where the next-to-last term indicates that the dealer becomes a non-owner upon selling the bond on the interdealer market. Equation (18) can be solved to obtain

$$U^s = \underbrace{(M^{\text{bid}} - P^s)\frac{\lambda}{r+\lambda N}}_{\text{Transaction Profit/Loss}} + \underbrace{U^{\text{no}}\frac{\lambda N}{r+\lambda N}}_{\text{Value of Future Business}}. \tag{19}$$

Valuations U^{no} and U^o lead to price improvement for repeat business. Equations (13), (15), (17), and (19) fully characterize dealers' transaction cycle.

As in most OTC models, prices are set by Nash bargaining, resulting in prices that are the bargaining-power (w) weighted average of the reservation values of the client and the dealer:[17]

$$P^b = (\widehat{V}^o - \widehat{V}^b)w + (M^{\text{ask}} - U^o)(1-w), \tag{20}$$
$$P^s = (\widehat{V}^s - \widehat{V}^{\text{no}})w + (M^{\text{bid}} + U^{\text{no}})(1-w). \tag{21}$$

The above equations assume that the dealer loses all future business from the client if the bilateral negotiations fail. Upon dropping a dealer, the client maintains her optimal network size by forming a new link with another randomly picked identical dealer. Thus, by agreeing rather than not, the dealer receives U^o. As a consequence, dealers face intertemporal competition for future clients. This is a novel assumption missing from existing models of OTC markets.

Each client's valuation, $\widehat{V}^k$, $k \in \{b, o, s, \text{no}\}$, can be written as a sum of its trade-specific, V^k, and relationship-specific, V^r, values as follows:

$$\widehat{V}^k = V^k + V^r. \tag{22}$$

Substituting (22) into (23), (24), (27), and (26) yields the following relations for trade-specific client valuations:

$$V^{\text{no}} = V^b \frac{\eta}{r+\eta}, \tag{23}$$

$$V^b = (V^o - \text{E}[P^b] - K)\frac{\lambda N}{r+\lambda N}, \tag{24}$$

$$V^o = \frac{C}{r+\kappa} + V^s \frac{\kappa}{r+\kappa}, \tag{25}$$

$$V^s = \frac{C(1-L)}{r+\lambda N} + (\text{E}[P^s] + V^{\text{no}} - K)\frac{\lambda N}{r+\lambda N}. \tag{26}$$

[17] See Glode and Opp (2016) and others for models of OTC markets in which bargaining is not Nash. We use Nash bargaining because it offers a parsimonious way to model the value of repeat trading to the client.

Correspondingly, transaction prices depend only on clients' trade-specific valuations:

$$P^b = (V^o - V^b)w + (M^{\text{ask}} - U^o)(1-w), \tag{27}$$

$$P^s = (V^s - V^{\text{no}})w + (M^{\text{bid}} + U^{\text{no}})(1-w). \tag{28}$$

The valuations and prices provide 10 equations and 10 unknowns. Proposition 1 in the Internet Appendix provides the closed-form solutions (IA1) and (IA2) for transaction buy prices P^b and sell prices P^s, respectively.

Bargaining power could differ for buys and sells. For ease of exposition, we equate them here. In the structural estimation below, we allow for different bargaining powers when the insurer is buying a bond, w^b, versus selling a bond, w^s. In addition, we allow the bargaining power parameters to depend on trading intensity η.

B. Discussion

The main friction in the model motivating clients to select multiple dealers rather than a single dealer is the finite search intensity λ. The client always optimally selects a single dealer if this friction is removed. We illustrate this by considering the limiting case of large search intensity, $\lambda \to \infty$. In this case, a single dealer can instantaneously find the bond and the optimal size of the network is one. Taking limits in equations (13) and (19) yields

$$U^{\text{no}}_{\lambda\to\infty} = \frac{1}{N}\frac{\frac{\eta}{r+\eta}}{1-\frac{\kappa}{r+\kappa}\frac{\eta}{r+\eta}}\left[P^b_{\lambda\to\infty} - M^{\text{ask}} + \frac{\eta}{r+\eta}\left(M^{\text{bid}} - P^s_{\lambda\to\infty}\right)\right],$$

$$U^{o}_{\lambda\to\infty} = \frac{1}{N}\frac{\frac{\kappa}{r+\kappa}}{1-\frac{\kappa}{r+\kappa}\frac{\eta}{r+\eta}}\left[M^{\text{bid}} - P^s_{\lambda\to\infty} + \frac{\kappa}{r+\kappa}\left(P^b_{\lambda\to\infty} - M^{\text{ask}}\right)\right]. \tag{29}$$

Solving for transaction prices from (20) and (21), and using the fact that $(V^o - V^b)_{\lambda\to\infty} = P^b_{\lambda\to\infty} + K$ and $(V^s - V^{\text{no}})_{\lambda\to\infty} = P^s_{\lambda\to\infty} - K$, we obtain

$$P^b_{\lambda\to\infty} = M^{\text{ask}} + \frac{w}{1-w}K - U^o_{\lambda\to\infty},$$
$$P^s_{\lambda\to\infty} = M^{\text{bid}} - \frac{w}{1-w}K + U^{\text{no}}_{\lambda\to\infty}. \tag{30}$$

Expressions (30) show that $U^o_{\lambda\to\infty}$ and $U^{\text{no}}_{\lambda\to\infty}$ represent the repeat relation buy discount and sell premium, respectively. Equations (29) show that both $U^o_{\lambda\to\infty}$ and $U^{\text{no}}_{\lambda\to\infty}$ are strictly decreasing with the size of the network N. Thus, the client optimally chooses a single dealer. When λ is finite, clients' direct benefit from having a network where $N > 1$ is the improved transaction speed.

The buy and sell prices, given by expressions (IA1) and (IA2) in Proposition 1 in the Internet Appendix, are nonlinear functions of the model primitives and N, which complicates the analysis. However, we can verify that the prices are well-behaved functions of network size, N, in the large network limit, $N \to \infty$. Specifically, $N \to \infty$ implies $\frac{\lambda N}{r+\lambda N} \to 1$ and hence clients' search friction in terms of time is zero. Dealers' valuations, in this case denoted with subscript $N \to \infty$, satisfy the system of equations $U^b_{N\to\infty} = U^s_{N\to\infty}\frac{\kappa}{r+\kappa}$ and $U^s_{N\to\infty} = U^b_{N\to\infty}\frac{\eta}{r+\eta}$. These equations have only the trivial solution $U^s_{N\to\infty} = U^b_{N\to\infty} = 0$, which implies that dealers have no rents from future relationships with clients. Clients have valuations $V^o_{N\to\infty} - V^b_{N\to\infty} = P^b_{N\to\infty} + K$ and $V^s_{N\to\infty} - V^{\text{no}}_{N\to\infty} = P^s_{N\to\infty} - K$, which yields the following expressions for transaction prices:

$$P^b_{N\to\infty} = M^{\text{ask}} + \frac{w}{1-w}K, \tag{31}$$

$$P^s_{N\to\infty} = M^{\text{bid}} - \frac{w}{1-w}K. \tag{32}$$

Dealers receive no relationship-based rents, but they charge clients a spread of $\frac{w}{1-w}2K$ per roundtrip transaction over the interdealer spread. The coefficient $\frac{w}{1-w}$ indicates that the nonzero bargaining power enables dealers to extract some value from a client. This is because each dealer charges the same price, which makes clients' threat of ending the relationship and forming a new link not credible (Diamond's (1971) paradox).

Overall, dealers' surplus comes from both the immediate value of trade (the spread) and the value of future transactions. As the probability of transacting with the same client in the future declines with the size of the client's network, dealers may charge higher spreads, $P^b - M^{\text{ask}}$ and $M^{\text{bid}} - P^s$, as network size increases. This result is similar to that in Vayanos and Wang (2007) where an asset with more buyers and sellers has lower search times and worse prices relative to its identical-payoff counterpart with fewer buyers and sellers.

To find the optimal network size N^* in the general case, we maximize the total valuation of the "first-time" owner, that is, the client buying the asset, $\widehat{V}^b$. This is because the client has to take possession of the asset in the first place. Maximizing $\widehat{V}^b$ accounts for all trade and relation benefits the client receives from a larger network. We solve for the optimal network size on the grid of integers $N \in \mathbb{I}_{\geq 0}$. To do so, we start by assuming that all value functions are well-behaved functions of the network size, N. We then solve for the optimal network size on the continuous grid $N^* \in \mathbb{R}_{\geq 0}$. Finally, we select the closest integer value to N^* that maximizes $\widehat{V}^b(N)$. Proposition 2 given in the Internet Appendix demonstrates that, for a given client type described by the model primitives $\{L, K, w, \kappa, \lambda, \eta\}$, there may exist an optimal network size $N^* = N(L, K, w, \kappa, \lambda, \eta)$. For the sake of clarity in our discussion below, we consider optimization on the continuous grid of network sizes.

Next, we provide intuition for why the model yields finite-sized networks for a general set of model primitives $\{L, K, w, \kappa, \lambda, \eta\}$. Because the equation describing the optimal size of the client's dealer network, given by condition (IA24) in Proposition 2 in the Internet Appendix, is nonlinear in N, the existence of the solution is not guaranteed for an arbitrary parameter set $\{L, K, w, \kappa, \lambda, \eta\}$. Given the complexity of the problem, it is convenient to consider the optimal network resulting solely from repeat trading, N^{**} maximizing V^b from $\frac{dV^b}{dN} = 0$. At this point, we do not take a stand on whether V^r is an increasing or decreasing function of N. The derivative $\frac{dV^r}{dN} > (<)0$ implies that $N^{**} < (>)N^*$, thus making N^{**} the lower (upper) bound of N^*. As a result, a sufficient condition for the existence of a finite N^{**} also applies to N^*. In addition, our analysis helps separate the trade-specific and the relationship-specific network formation mechanisms.

To demonstrate that there exists a region of the parameter space in which the optimal network size, N^{**}, is finite, we examine different terms of the first-order condition for maximizing V^b with respect to N. The derivative

$$\frac{dV^b}{dN} = \frac{\lambda N}{r + \lambda N}\left(\frac{rV^b}{\lambda N^2} + \frac{dV^o}{dN} - \frac{dP^b}{dN}\right) \tag{33}$$

must be positive on $[1, N^{**})$ and equal to 0 at $N = N^{**}$. As can be seen, $\frac{dV^b}{dN}$ contains three terms in the parentheses. The first positive term captures the direct effect of network size on the speed of execution as faster execution improves the buyer's valuation. The second term reflects the marginal effect of network size on the owner's value and can be written using relation (25) as $\frac{dV^o}{dN} = \frac{\kappa}{r+\kappa}\frac{dV^s}{dN}$. The third term represents the effect of network size on the transaction price. We focus on the case in which the buy price, P^b, improves with the network size, $\frac{dP^b}{dN} < 0$, for which we derive a sufficient condition below. Because the first and third terms in (33) are positive, $\frac{dV^b}{dN}$ can be 0 for some N^{**} only if the second term is negative, thus implying that the seller's valuation must be decreasing with N, $\frac{dV^s}{dN} < 0$. In summary, in equilibrium the marginal improvements in the buyer's valuations due to faster transaction speed and the buy price are balanced by a marginal decline in the seller's valuation adjusted by the hazard rate of selling.

The next step is to derive a sufficient condition for $\frac{dV^s}{dN} < 0$. We start by noting that as long as V^s is decreasing with N, the sell price is also decreasing with N:

$$\frac{dP^s}{dN} = \left(\underbrace{\frac{dV^s}{dN}}_{<0} - \frac{\eta}{r+\eta}\underbrace{\frac{dV^b}{dN}}_{\geq 0}\right)w + \underbrace{\frac{dU^{\text{no}}}{dN}}_{<0}(1-w) < 0, \quad N \in [1, N^{**}], \tag{34}$$

where we use the fact that dealers' valuations, U^k, $k \in \{b, o, s, \text{no}\}$, decline with N. Using relation (26), the sufficiency condition for $\frac{dV^s}{dN} < 0$ can be written as[18]

[18] Multiplying (26) by $r + \lambda N$, differentiating it with respect to N, and multiplying the resulting expression by N yields

$$\lambda N^2 \frac{dV^s}{dN} = -\lambda N V^s + \lambda N(P^s + V^{\text{no}} - K) + \lambda N^2\left(\frac{dP^s}{dN} + \frac{\eta}{r+\eta}\frac{dV^b}{dN}\right) \leq 0, \tag{35}$$

where we use the relation (23). After multiplying both sides of (26) by $r + \lambda N$, it follows that $\lambda N V^s = -rV^s + C(1-L) + \lambda N(P^s + V^{\text{no}} - K)$. Substituting this expression back into (35) yields (36) after some algebra.

$$V^s + \frac{\lambda N^2}{r}\left(\underbrace{\frac{dP^s}{dN}}_{<0} + \frac{\eta}{r+\eta}\underbrace{\frac{dV^b}{dN}}_{\geq 0}\right) \leq \frac{C(1-L)}{r}, \quad N \in [1, N^{**}]. \tag{36}$$

Therefore, V^s decreases with N when the marginal decline in the sell price exceeds the hazard rate of buying times the marginal increase in V^b sufficiently that the seller prefers to hold the discounted asset forever to obtain $\frac{C(1-L)}{r}$.

Finally, we establish the sufficient condition for $\frac{dP^b}{dN} < 0$. Differentiating equation (20) with respect to N, we obtain

$$-\underbrace{\frac{dU^o}{dN}}_{<0} < \frac{w}{1-w}\left(\underbrace{\frac{dV^b}{dN}}_{\geq 0} - \frac{\kappa}{r+\kappa}\underbrace{\frac{dV^s}{dN}}_{<0}\right), \quad N \in [1, N^{**}], \tag{37}$$

where we use relation (26). We conjecture that there exists an interval of network sizes belonging to $[1, N^{**}]$ on which inequalities (36) and (37) are simultaneously satisfied for a range of values of the model's primitives. Although this conjecture cannot be proven analytically, we verify numerically that there exist parameters over which these inequalities are satisfied. In particular, they are satisfied for the structurally estimated parameters below. In the equilibrium, the buy price P^b improves with N, whereas the sell price declines with network size. As a result, the buyer's value is maximized by trading off more frequent buys at a discounted buy price against more frequent sells at a marked down sell price.

The effect of N on buy and sell transaction prices is not symmetric, however, due to the time lag between each buy and sell as well as the different effect of repeat relationships on P^b and P^s. Therefore, instead of focusing on the individual transaction prices, we investigate the sign of the marginal effect of increasing the network size on the bid-ask spread $SP \equiv P^b - P^s$:

$$\begin{aligned}\frac{dSP}{dN} &= -w\left(\frac{r}{r+\kappa}\frac{dV^s}{dN} + \frac{r}{r+\eta}\frac{dV^b}{dN}\right) \\ &\quad - (1-w)\left(\frac{dU^o}{dN} + \frac{dU^{\text{no}}}{dN}\right), \quad N \in [1, N^{**}].\end{aligned} \tag{38}$$

If inequalities (36) and (37) are satisfied simultaneously, the bid-ask spread improves with the size of the network, $\frac{dSP}{dN} < 0$, as long as the following inequality holds:

$$0 < -\frac{r}{r+\kappa}\frac{dV^s}{dN} - \frac{1-w}{w}\left(\frac{dU^o}{dN} + \frac{dU^{\text{no}}}{dN}\right) < \frac{r}{r+\eta}\frac{dV^b}{dN}, \quad N \in [1, N^{**}]. \tag{39}$$

For network sizes below the trade-specific optimal network size N^{**}, we must have that $\frac{dV^b}{dN} > 0$, whereas at $N = N^{**}$, we must have that $\frac{dV^b}{dN} \mid_{N=N^{**}} = 0$, which violates inequality (39). This implies that $\frac{dSP}{dN}$ switches sign from negative to positive at some $\widetilde{N} < N^{**}$.

A finite optimal size of the trade-specific network follows from several trade-offs. Starting with a client that has a small dealer network, when the client adds another dealer to her network, the speed of execution, measured by $\frac{\lambda N}{r+\lambda N}$, increases. Improved execution speed implies faster transitions between the client's "buyer" and "seller" states and thus creates more future repeat trading for dealers. Some of the surplus from increased repeat trading is then passed by dealers back to the client via price improvements, leading to a narrower bid-ask spread. However, adding an additional dealer to the existing network implies sharing the value of the relationship with one more dealer. Consequently, the extra per-dealer value from the repeat trading created through the increased execution speed declines with network size.

As the size of the network becomes larger, the per-dealer benefits of faster execution are outweighed by losses from the additional rent-sharing with other dealers. Dealers respond by reducing price improvements and hence the bid-ask spread widens. At this point, the client starts trading off the benefit due to increased execution speed and the cost due to a wider bid-ask spread when deciding whether to add another dealer to her network. The optimal size of the network is achieved when the benefits from transacting with the increased number of dealers equal the costs. When the relationship-specific nontraded value, $V^r(\eta, N)$, is added, and if it is increasing with network size for a given value of η, the client is willing to accept even wider bid-ask spreads in exchange for greater nontrade relationship-specific benefits. If the marginal nontraded value improvement from a larger network decreases with the size of the network, the overall optimal network size N^* is reached when relationship-specific benefits together with benefits from transacting faster are balanced by the wider bid-ask spread.

C. Empirical Predictions

Although there is no client heterogeneity in the model, each client creates her own access to the interdealer market and the interdealer market links clients together. Therefore, the model's empirical predictions can be considered via comparative statics with respect to the client's trading intensity, η. Section II.A shows that η follows a power law with exponential tail in the cross-section of insurers. Based on these empirical findings, we next discuss how larger trading intensity impacts clients' choice of network size and the prices they receive. The predictions are complex as network size and transaction prices are interrelated.

Transaction prices: There exist several competing channels through which more active insurers receive better prices than less active insurers. The first explanation for price improvement is that active insurers have more repeat trade business with those dealers that internalize the benefits from future business and grant price improvements. A second, complementary, explanation is that large insurers have greater bargaining power than small insurers vis-à-vis dealers. A third explanation is that insurers' trading activity in corporate bonds is correlated with the value of other business they conduct with the dealers.

To shed light on the intuition for price improvement from future business, it is convenient to discuss the effect of η on prices while keeping the size of the network fixed. When $\eta = 0$, a client that owns the bond will never be an owner again after she sells. Therefore, all of a dealer's surplus comes from the price at which the sale occurs and none from future trades. In this case, the dealer's buy reservation value is lower, resulting in a wider bid-ask spread. When η is large, dealers derive significant value from repeat trades, leading to smaller bid-ask spreads. This implies that, conditional on network size, clients with larger trading intensity (higher η) get better execution than clients with smaller trading intensity (lower η). This is consistent with our empirical findings above that insurers' trading activity increases in insurer size (Table II) and that trading costs decline in insurer size (Table V).

Network size: Increasing the network size increases buyers' valuation, $\widehat{V}^b$, through improved speed of execution, greater intertemporal dealer competition,[19] and increased relationship-specific nontrade value. The optimal network size, N^*, as a function of η balances these gains against a lower buyer valuation due to intratemporal dealer competition leading to a loss of repeat trading business.

The optimal network size is increasing in η for a nonempty parameter range over which the marginal gains from a larger network increase in η, $\frac{\partial^2 \widehat{V}^b}{\partial N \partial \eta} \mid_{N=N^*} > 0$. Intuitively, dealers' profit from repeat trade improves with increasing trading intensity and thus dealers offer better execution. Clients with larger trading intensity benefit from a larger network, which improves the speed of execution and may generate larger nontrading relationship-specific value. Therefore, consistent with Table III, the model can generate larger insurers with more frequent trading needs that have larger dealer networks.

Ultimately, the model must be structurally estimated to test its ability to fit the cross-sectional distribution of network size from Figure 3 as well as the cross-sectional distribution of execution costs as a function of network size from Table V.

IV. Model Estimation and Policy Analysis

The empirical findings in equation (1) and Table V and Table IA.IV in the Internet Appendix suggest that in the cross-section, trading networks follow a power law with exponential tail and trading costs are nonmonotone in network size. Although the model can produce such nonmonotonicity, it is unclear if it can quantitatively match the empirical relationship between network size and trading costs. In this section, we test the model's ability to quantitatively match

[19] The buy transaction price may or may not be improving with N, $\frac{dP^b}{dN} \gtrless 0$. However, the cross derivative $\frac{\partial^2 P^b}{\partial N \partial \eta}$ is positive.

these stylized facts. Given the model's complexity, readers more interested in the policy implications can focus their attention on Section IV.E.

A. Estimation Procedure

Settings such as labor markets and marriage markets involve one-to-one matching. These literatures structurally estimate models to examine their quantitative fit to the data (for example, Postel-Vinay and Robin (2002) and Eckstein and Van den Berg (2007) for labor search, and Hitsch, Hortaçsu, and Ariely (2010) and Choo (2015) for marriage search). Gavazza (2016) estimates a one-to-one search-and-bargaining model of a decentralized market using aircraft transaction data. The structural identification of our model shares several similarities with Gavazza (2016). As in Gavazza (2016), identification of unobserved parameters relies on key moments in the data. But unlike the aggregate moments in Gavazza (2016), we use the granular nature of our data and the heterogeneity in insurers' trading intensity and network size to facilitate our estimation.

Figures 2 and 3 highlight the heterogeneity in insurers' trading intensity and network size. To estimate the model's network size distribution, we infer the distribution of trading intensities, η_i, $i = 1, \ldots, \mathcal{I}$, across insurers i. Section II.A characterizes the distribution of trading activity. If trading shocks occur at Poisson times, the intensity η_i of the shocks can be captured by the expected number of buy trades per year. This fact yields the maximum likelihood estimator $\widehat{\eta_i} = \frac{1}{T}\sum_{t=1}^{T} X_{it}$, where X_{it} is the number of bond purchases by insurer i in year t. To use the multiple years of trade data, we perform the estimation separately for each insurer. This procedure yields a cross-sectional distribution of trading intensities, which we index by $p(\eta)$. Section II.A and Figure 2 show that the distribution of the insurer trades per year follows approximately a power law. For insurer buys, this distribution is best described by $p(\eta) = 0.34 \times \eta^{-1.31}$.

As discussed above, insurers' trading shocks are heterogeneous and directly observable in the data. In the model, κ measures clients' selling intensity after having bought. Because the empirical number of buys and sells are equal, and buys and sells follow a similar power law, there is no evidence of heterogeneity in κ. Therefore, κ does not vary across insurers. Because we fit the model to the two observable outcomes of network size and trading costs, we limit the number of model parameters that vary across insurers by assuming that L and K are constant. The competitive interdealer spread is set to match the average execution costs of insurers by choosing $C = 1$, $r = 0.1$, $M^{\text{ask}} = (1 + 0.0018)\frac{C}{2r}$, and $M^{\text{bid}} = (1 - 0.0018)\frac{C}{2r}$. Because insurers can potentially access all dealers, λ is constant across insurers.

Model parameters $\Theta = (L, K, \kappa, \lambda, w^b, w^s, V^r)$ are not directly observable in the data, so we estimate them structurally. To ensure that the estimated parameters L, K, κ, and λ remain in their natural domains (i.e., positive or between 0 and 1) in the estimation, we transform them as follows:

$$L = \Phi(\theta_L) \in (0, 1), \quad K = e^{\theta_K} \geq 0, \quad \kappa = e^{\theta_\kappa} \geq 0, \quad \lambda = e^{\theta_\lambda} \geq 0, \tag{40}$$

where $\Phi(x)$ is the normal cdf.[20] The evidence in Table V shows that trading costs depend on insurer size even after controlling for network size. The most straightforward way to accommodate this fact is to allow dealer bargaining power to vary across insurers, as large active insurers likely have higher effective bargaining power than small inactive insurers. We report the model fit under both constant (specification 1) and variable (specification 2) bargaining power by holding w^b and w^s constant and allowing w^b and w^s to be functions of η as follows:

$$w^s = \Phi\left(\theta^0_{w^s} + \theta^1_{w^s}\ln\eta\right) \in (0,1), \quad w^b = \Phi\left(\theta^0_{w^b} + \theta^1_{w^b}\ln\eta\right) \in (0,1). \tag{41}$$

Specification 1 with constant bargaining power sets both $\theta^1_{w^s}$ and $\theta^1_{w^b}$ to 0. The functional form of the nontrade value is Cobb-Douglas,

$$V^r = e^{\theta^0_V}\eta^{\theta^1_V}N^{\theta^2_V},\ \theta^1_V + \theta^2_V < 1, \tag{42}$$

with a constant "technology" $e^{\theta^0_V}$ and elasticities to trading intensity and network size given by θ^1_V and θ^2_V, respectively. The Cobb-Douglas functional form captures the potential complementarity between trading intensity, measured by η, and network size, while precluding infinitely large nontrade value. We do not impose specific restrictions on the elasticities beyond requiring that their sum be less than 1. This specification allows the nontrade value to be an increasing or decreasing function of N, η, or both. By construction, the nontrade value is nonnegative. Taking into account transformations (40) and (41) and parametrization (42), the estimated parameter set takes the form $\Theta = (\theta_L, \theta_K, \theta_\kappa, \theta_\lambda, \theta^0_{w^b}, \theta^1_{w^b}, \theta^0_{w^s}, \theta^1_{w^b}, \theta^0_V, \theta^1_V, \theta^2_V)$.

Section II.B and Figure 3 show that the degree distribution of client-dealer relationships follows a mixed power-exponential law. The model also has implications for how percentage trading costs, $(P^b - P^s)/(0.5(P^b + P^s))$, depend on the parameters Θ and η. Empirically, the estimated relation between trading costs and network size is in column (4) of Table V as $c(N) = 51 + 0.32N - 6.29\ln N$. The distance between the empirical and model-generated relationships depends on the parameters Θ. We fit them by minimizing the sum of two probability-weighted distances between the data and model,

$$\underbrace{\omega\sum_{N=1}^{\overline{N}} p(N)[N - N^*(\Theta, \eta(p(N)))]^2}_{\text{Size Distribution Criterion}} + \underbrace{(1-\omega)\sum_{N=1}^{\overline{N}} \frac{1}{\overline{N}}[c(N) - c^*(\Theta, \eta(p(N)))]^2}_{\text{Trading Cost Criterion}}, \tag{43}$$

where ω is the weight put on matching the network size compared to trading costs, $p(N)$ is the empirical network size distribution, $\overline{N}$ is the largest network size, and $N^*(\Theta, \eta)$ and $c^*(\Theta, \eta)$ are the model-implied network size and trading cost, respectively. In our estimation, we choose equal weights on matching network size and trading costs, $\omega = 0.5$.[21]

[20] Our results are robust to alternative transformations in place of Φ and e.

[21] The Internet Appendix reports similar estimates with various asymmetric weights.

Fitting the model based on network size (N) rather than insurer trading needs (η) reduces the computational requirements. Inverting the power law distribution for trading intensity from Panel A of Figure 2 yields $\eta(p) = (\frac{p}{0.34})^{-\frac{1}{1.31}}$. It then follows from equation (1) that for insurer buys $p(N) = 0.28\,e^{-0.22N}N^{-0.12}$. Therefore, by substituting $p(N)$ into the expression for $\eta(p)$, we obtain the mapping of trading intensities into network sizes, $\eta(p(N)) = (\frac{0.28}{0.34})^{-\frac{1}{1.31}}e^{\frac{0.22}{1.31}N}N^{\frac{0.12}{1.31}}$. We use this relation in (43). To fit the data, the model's optimal N^* is constrained to be an integer. This makes the objective function (43) nonsmooth. To accommodate this, we optimize using simulated annealing with fast decay and fast reannealing.

Overall, we estimate a total of nine model parameters in specification 1 and 11 model parameters in specification 2 by fitting a total of 80 empirical moments. Of these, 40 moments are the individual data points from the degree distribution of client-dealer buy relationships presented in the left panel of Figure 3. The remaining 40 moments are from the empirical trading costs and network size relation, $c(N) = 51 + 0.32N - 6.29\ln N$, evaluated on the 40-point grid $N = 1, 2, \ldots, 40$.

B. Estimated Model Parameters

The model can be solved assuming that the client maximizes her value as a buyer or a seller. To simplify exposition, we use her value as a buyer. Panel A of Table VI reports the estimates for the transformed model parameters from (40) to (42). Panel B reports the original model parameters $\Theta = (L, K, \kappa, \lambda, w^b, w^s, V^r)$. All parameters are significantly different from zero and appear well identified, except for θ_L in specification 1. Panel C reports statistics on model fit. Specification 2 allows bargaining power to vary across insurers and generates a substantially lower minimum distance between the model and empirical distributions than does the uniform bargaining power in specification 1. The standard deviations from the size distribution criterion and the trading cost criterion are small.

The estimated liquidity shock parameter L in both specifications is close to 100% of the flow income from the bond, suggesting a high willingness to pay for immediacy. The cost K of processing each transaction is small, at about 0.1 (specification 1) and 16 (specification 2) basis points. The estimated selling shock intensity κ is between 15 (specification 1) and 21 (specification 2), which amounts to a holding period from three weeks to a month. Longer holding periods (i.e., smaller κ) reduce the value of repeat relations and increase network size. Without heterogeneity in κ, the model needs relatively short holding periods to reproduce the large fraction of insurers using a single dealer.

Dealers' search efficiency λ is estimated to be 291 in specification 1, corresponding to dealers taking less than a day ($250/\lambda$) to locate a bond. Dealers' search efficiency λ is estimated to be 146 (1.7 days) in specification 2. Given overall corporate bond trading frequencies, these estimates seem reasonable.

Table VI

Estimated Model Parameters

Panel A reports the estimated model parameters $\theta = (\theta_L, \theta_K, \theta_\kappa, \theta_\lambda, \theta^j_{w^i}, \theta_V)$. Estimates are based on the minimum-distance estimation. Standard errors are computed using the sandwich estimator and reported in parentheses. Panel B reports the implied values for the model parameters $L = \Phi(\theta_L)$, $K = e^{10*\theta_K}$, $\kappa = e^{10*\theta_\kappa}$, $\lambda = e^{10*\theta_\lambda}$, $w^i = \Phi(\theta^0_{w^i} + \theta^1_{w^i} \ln \eta)$ for $i = s, b$, where $\Phi \in (0, 1)$ is the normal cdf and nontrade value $V^r = e^{\theta^0_V} \eta^{\theta^1_V} N^{\theta^2_V}$. Panel C reports statistics on model fit.

	w^b, w^s Constant (1)	$w^b(\eta)$, $w^s(\eta)$ (2)
Panel A: Parameter Estimates		
θ_L	7.46 (17.31)	2.51 (0.01)
θ_K	−1.20 (0.01)	−0.64 (0.00)
θ_κ	0.27 (0.00)	0.31 (0.00)
θ_λ	0.57 (0.01)	0.50 (0.00)
$\theta^0_{w^b}$	2.93 (0.12)	2.52 (0.00)
$\theta^1_{w^b}$	–	−1.02 (0.00)
$\theta^0_{w^s}$	1.76 (0.07)	1.24 (0.00)
$\theta^1_{w^s}$	–	0.04 (0.00)
θ^0_V	−3.82 (0.00)	−1.70 (0.00)
θ^1_V	0.12 (0.01)	0.33 (0.00)
θ^2_V	−0.46 (0.00)	0.13 (0.00)
Panel B: Implied Model Parameters		
L	1.00	0.99
K ($\times 10^4$)	0.06	15.82
κ	15.36	21.19
λ	290.99	145.56
w^b (*SD*)	0.98 (0.00)	0.86 (0.05)
w^s (*SD*)	0.94 (0.00)	0.88 (0.01)
Nontrade value V^r (*SD*)	0.01 (0.00)	0.32 (0.03)
Panel C: Model Fit		
Minimum distance ($\times 10^2$)	145.35	2.82
SD residuals	1.21	0.17
SD residuals network	1.09	0.19
SD residuals prices	1.31	0.14

Dealers' bargaining power on the buy and sell sides in specification 1 is large, at 98% and 94%, suggesting that dealers capture most of the trade surplus in this specification. In specification 2, when bargaining power depends on the insurer's type, dealers' average bargaining power remains large and fairly symmetric across buys and sells, at 86% and 88%. Dealers' bargaining power on sales is relatively insensitive to insurer trading frequency. In contrast, $\theta^1_{w^b} = -1.02$ indicates that dealers' bargaining power when buying declines significantly with insurers' trading intensity η.[22] Dealers' bargaining power with insurers with trading intensity $\eta = 1$ is close to one. As insurer trading intensity increases to five, dealers' bargaining power falls to about half. Dealers have almost zero bargaining power when the largest insurers are buying. This could arise from insurers being buy-and-hold investors, in which case a dealer can more easily locate a seller than a buyer. The heterogeneity in bargaining across insurers and across buy and sell transactions suggests that richer modeling of the price-setting process may enable deeper understanding of OTC trading.

The estimated nontrade value, V^r, is quite small in specification 1. The estimated nontrade value, V^r, is 25 times larger in specification 2, and it is increasing in both the size of the network and trading intensity. The decreasing returns to scale parameter restriction is satisfied. The trade intensity elasticity is more than twice the network size elasticity.

Table IA.V in the Internet Appendix reports the parameter estimates when we vary the minimum-distance weight ω between 0.1 and 0.9. Across columns, we reestimate the model for specification 2. The estimates are generally consistent with those reported in Table VI. The variation in the estimates across columns is small. In column (2) with $\omega = 0.5$, the residuals for network size have the same order of magnitude as for prices. Naturally, the specification in column (1) with $\omega = 0.1$ yields smaller residuals for trading costs than network size, while the reverse is the case in column (3) with $\omega = 0.9$. We conclude that our results are robust to the relative weight put on matching network size compared to trading costs.

C. Sensitivity to Model Parameters and Identification

To provide intuition on the parameter identification, we estimate the sensitivity of the model-implied moments to the baseline parameters using the following exercise. We pick four sets of moments related to dealer network size N across insurers, the cross-sectional distribution of network size $\Pr(N)$, trading costs $c(N)$, and the cost premium across different network sizes relative to minimum trading costs. For the dealer network, we calculate the average network size, the minimum network size $N_{\min}$, the maximum network size $N_{\max}$, and the smallest cost network size $N_{\min-\text{cost}}$. We also characterize the distribution of network size using the probability of having a network smaller than three dealers, $\Pr(N \leq 2)$, and the probability of having a network greater than 20 dealers, $\Pr(N > 20)$. For trading costs $c(N)$, we calculate its average, minimum $c_{\min} = c(N_{\min-\text{cost}})$, and maximum $c_{\max} = c(N_{\min})$.[23] We also calculate the cost for the largest network $c(N_{\max})$. Finally, we calculate the min-N cost premium $c_{\max} - c_{\min}$ and the max-N cost premium $c(N_{\max}) - c_{\min}$. Trading costs and cost premia are expressed in basis points.

[22] The asymmetry in buy-sell bargaining power's sensitivity to η arises from our choice to maximize clients' valuation as a buyer.

[23] Although it is not true in general, this is true for all reported parameter specifications.

Our baseline comprises the four sets of moments described above using specification 2 from Table VI. We then take each parameter from the set $\Theta_S = (L, K, \kappa, \lambda, V^r)$, reduce it by 10% while keeping all other parameters fixed, solve the model, and recalculate all four sets of moments. Next, we repeat this procedure using a 50% reduction in each parameter. Because we discuss the role of dealers' bargaining power later in this section, we exclude w^b and w^s from the sensitivity analysis.

Table VII illustrates the sensitivity of the model fit to the parameters in Θ_S. To understand the mechanics of the comparative statics exercise reported in Table VII, recall that a client with trading intensity η selects her equilibrium network size N by balancing the gains from increased transaction speed against the losses from wider spreads. Changing any parameter from the set Θ_S directly affects the transaction speed, the spread, or both, thus inducing the client to change the size of her network. Because clients with larger trading needs and, in turn, larger networks are more sensitive to changes in transaction speed and spread, their networks are more affected by changes in the parameters. Clients with low trading needs are less sensitive to such changes and thus are much less likely to change their network size.

Reducing the liquidity shock parameter L has two direct effects. First, it reduces the seller's need for immediacy and thus decreases the benefits of higher transaction speed provided by larger networks. As a result, some of the network distribution mass shifts toward smaller networks. Table VII, columns (1) and (2), show that when L is reduced by 10% (50%), the average network size declines by 7% (30%), $\Pr(N \leq 2)$ increases by 32% (58%), and $\Pr(N > 20)$ declines by 20% (68%). Second, reducing L increases the reservation utility of the seller thus improving the sell price and narrowing the spread across all network sizes. The average $c(N)$ declines by 2% (5%) while the maximum trading costs remain unchanged. Overall, $c(N)$ is either reduced or remains unchanged.

Finally, because the minimum network size remains unchanged while the largest network size is reduced by only two dealers, changes in the cost premium are due mostly to changes in $c_{\min}$. For example, if $c_{\max}$ is fixed while $c_{\min}$ declines, the cost premium at $N_{\min}$, $c_{\max} - c_{\min}$, must increase. Table VII shows that $c_{\min}$ declines by 2% while $c_{\max}$ remains unchanged, thus leading to an increase of 1% (6%) in $c_{\max} - c_{\min}$, while $c(N_{\max})$ remains unchanged (decreases by 1%), thus leading to an increase of 5% (15%) in $c(N_{\max}) - c_{\min}$.

Reducing the fixed costs of each trade K also has two direct effects. First, it makes it cheaper to add extra dealers to the network. A large reduction in K may even induce clients with very low trading needs to increase the size of their dealer network, thus making $N_{\min} > 1$. Table VII, columns (3) and (4), confirm this conjecture by showing that when K is reduced by 10% (50%), the average network size increases by 10% (123%), $\Pr(N \leq 2)$ remains unchanged (declines by 42%), and $\Pr(N > 20)$ increases by 58% (2,339%). When K is halved, the minimum network size increases from one to two and the maximum network size more than doubles. Overall, reducing K leads to significant changes in the distribution of network size, as a sizable fraction of its mass shifts toward larger networks.

Table VII

Sensitivity to Model Parameters and Identification

This table reports the sensitivity of different moments describing the predictions of the model to the main model parameters. The four sets of moments we consider are related to the dealer network size N across insurers, the cross-sectional distribution of network size $\Pr(N)$, trading costs $c(N)$, and the cost premium across different network sizes N relative to minimum trading costs. The smallest cost network size is the network size corresponding to minimum trading costs. Trading costs and cost premiums are expressed in basis points.

		L		K		κ		λ		V^r	
	Base	−10% (1)	−50% (2)	−10% (3)	−50% (4)	−10% (5)	−50% (6)	−10% (7)	−50% (8)	−10% (9)	−50% (10)
Network size N:											
mean	4.6	4.3	3.2	5.1	10.3	4.7	5.0	4.8	5.5	4.2	3.3
		−7%	−30%	10%	123%	2%	7%	3%	19%	−10%	−28%
$N_{\min}$	1	1	1	1	2	1	1	1	1	1	1
		0%	0%	0%	100%	0%	0%	0%	0%	0%	0%
$N_{\max}$	44	44	42	52	>100	46	65	45	48	39	21
		0%	−5%	18%	–	5%	48%	2%	9%	−11%	−52%
$N_{\min-\text{cost}}$	21	17	13	21	45	18	17	19	23	17	13
		−19%	−38%	0%	114%	−14%	−19%	−10%	10%	−19%	−38%
Network size distribution $\Pr(N)$:											
$\Pr(N \leq 2)$	0.40	0.53	0.63	0.40	0.23	0.40	0.40	0.40	0.40	0.53	0.53
		32%	58%	0%	−42%	0%	0%	0%	0%	32%	32%
$\Pr(N > 20)$	0.01	0.01	0.00	0.01	0.19	0.01	0.01	0.01	0.02	0.00	0.00
		−20%	−68%	58%	2,339%	0%	58%	0%	149%	−50%	−100%
Trading cost $c(N)$:											
mean	45.1	44.4	42.8	44.4	43.7	45.3	46.7	45.2	46.0	44.5	44.2
		−2%	−5%	−2%	−3%	0%	4%	0%	2%	−1%	−2%
$c_{\min}$	38.6	38.5	37.9	38.6	38.8	38.8	40.5	38.7	39.2	38.4	37.8
		0%	−2%	0%	0%	1%	5%	0%	2%	−1%	−2%
$c_{\max}$	51.1	51.1	51.1	49.7	48.3	51.4	53.5	51.1	50.4	51.1	51.1
		0%	0%	−3%	−6%	1%	5%	0%	−1%	0%	0%
$c(N_{\max})$	41.0	40.9	40.6	41.2	41.0	41.7	48.1	41.1	41.6	40.5	38.7
		0%	−1%	0%	0%	2%	17%	0%	2%	−1%	−6%
Cost premium:											
min-N	12.5	12.7	13.3	11.1	9.5	12.6	13.0	12.4	11.1	12.7	13.3
		1%	6%	−12%	−24%	1%	4%	−1%	−11%	2%	7%
max-N	2.3	2.5	2.7	2.6	2.2	2.9	7.6	2.4	2.3	2.1	0.9
		5%	15%	9%	−8%	21%	224%	3%	0%	−11%	−61%

Second, reducing K narrows the spread across all clients independent of their trading needs. Shifting the mass of the network size distribution toward larger networks reduces the spread for networks smaller than $N_{\text{min-cost}}$, whereas it increases the spread for networks greater than $N_{\text{min-cost}}$. Because a substantial mass of the network size distribution shifts to the interval $N > N_{\text{min-cost}}$, the overall effect of decreasing K on trading costs $c(N)$ is modest. Indeed, Table VII (columns (3) and (4)) shows that when K is reduced by 10%

(50%), the average $c(N)$ declines only by 2% (3%). However, reducing K has a strong effect on the cost premia mainly through its effect on $c_{\max}$, as $c_{\min}$ remains unchanged. Because the spread narrows for the smallest networks, $c_{\max}$ decreases and, as a result, the cost curve flattens as $c_{\max} - c_{\min}$ declines. Conversely, $c(N_{\max}) - c_{\min}$ is nonmonotonic in K. Table VII shows that the cost premium decreases for $N = N_{\min}$ by 12% (24%) while it first increases (9%) and then falls (−8%) for $N = N_{\max}$.

Reducing κ increases bond holding periods. This decreases the speed of transition between a client's owner and seller states. To compensate for the loss of repeat trading, dealers widen the spread for all client types. Clients with the highest trading needs, that is, with network sizes in the range of $N_{\max}$ for which the spread increases with N, scale their networks up to compensate for the reduction in speed, the benefit of which outweighs the extra trading costs. Clients with medium trading needs, that is, with network sizes in the range of $N_{\text{min-cost}}$ for which the spread also increases with N, may scale their networks down to avoid paying extra trading costs. When κ is reduced by 10% (50%), clients with very high trading needs increase their networks from 44 to 46 (65) dealers, whereas clients with $N = N_{\text{min-cost}}$ reduce their networks from 21 to 18 (17) dealers. Overall, the 50% reduction in κ leads to pronounced changes in the right tail of the network size distribution. Column (6) of Table VII shows that when κ is halved, $\Pr(N > 20)$ increases by 58%. Reducing κ by 10% (50%) widens spreads for all network sizes, as both $c_{\max}$ and $c_{\min}$ increase by 1% (5%) while $c(N_{\max})$ increases by 2% (17%). Therefore, the cost premium at $N_{\min}$, $c_{\max} - c_{\min}$, increases by 1% (4%), whereas the cost premium at $N_{\max}$, $c(N_{\max}) - c_{\min}$, increases by 21% (224%).

Reducing search intensity λ decreases the speed of transition between the buyer, owner, seller, and non-owner states for both clients and dealers. Similar to the case of κ, dealers widen the spread across all client types to compensate for the loss in repeat trading, whereas clients either increase or decrease their network size depending on whether the benefit from transacting faster outweighs the higher trading costs. However, in the case of λ, which clients end up increasing or decreasing their network size depends on the magnitude of the reduction in λ. Column (7) of Table VII shows that when λ is reduced by 10%, clients with very high trading needs increase their networks from 44 to 45 dealers. In contrast, clients with low trading needs are not affected by this change in λ, as $\Pr(N \leq 2)$ remains unchanged. Clients with medium trading needs reduce their networks, as $N_{\text{min-cost}}$ falls from 21 to 19 dealers. Overall, changes in the network size distribution and trading costs are minor in this case. However, changes in the network size distribution are pronounced when λ is halved—column (8) of Table VII shows that $\Pr(N > 20)$ increases by 149% in this case. For clients with medium to high trading needs, $N \geq N_{\text{min-cost}}$, benefits from faster execution outweigh losses from higher trading costs and, thus, $N_{\text{min-cost}}$ increases from 21 to 23 dealers while $N_{\max}$ increases from 44 to 48 dealers. As a consequence, $c_{\min}$, $c(N_{\max})$, and the average trading cost all increase by 2%. Therefore, the cost premium at $N_{\max}$, $c(N_{\max}) - c_{\min}$, stays flat while the cost premium at $N_{\min}$, $c_{\max} - c_{\min}$, declines by 11%. Since the reduction in λ decreases the cost premium at $N_{\min}$ while the reduction in κ increases the cost premium at $N_{\min}$, the model can identify λ and κ separately.

Finally, reducing V^r decreases the nontrade benefits from having a larger network size for clients with both low and high trading needs, with the latter being more affected. Clients respond by scaling down their networks. When V^r is reduced by 10% (50%), by decreasing θ_0^v, the average network size declines by 10% (28%), the maximum network size declines by 11% (52%), $\Pr(N \leq 2)$ increases by 32% (32%), and $\Pr(N > 20)$ declines by 50% (100%). Because clients trade off the nontrade benefits against wider spreads, reducing network size should lead to a narrower spread, except for $N_{\min} = 1$. Indeed, the average $c(N)$ declines by 1% (2%), $c_{\min}$ is unchanged (declines by 52%), $c(N_{\max})$ declines by 1% (5%), and $c_{\max}$ is unchanged. Consequently, the cost premium at $N_{\min}$, $c_{\max} - c_{\min}$, increases by 1% (6%), whereas the cost premium at $N_{\max}$, $c(N_{\max}) - c_{\min}$, declines by 13% (65%).

Overall, Table VII clearly demonstrates that the effect of each structural parameter on the network size distribution and transaction costs exhibits substantial variation, thus allowing for the parameter identification.

D. Model Fit and Discussion

Figures 5 and 6 provide visual evidence on the quality of the model's fit. In particular, both figures provide four plots illustrating the fit for their respective specification. The top left graphs plot the number of dealers as a function of their trading intensity both in the data (circles) and in the model (dashed line). Similar to Figure 3, the top right graphs plot the degree distribution for insurer-dealer relationships in the data (circles) compared to the model-implied distribution under the estimated parameters (dashed line). In both cases, the distance between the two lines captures the model fit.

To further visualize the goodness of fit, the bottom left graphs plot the degree distribution for insurer-dealer relationships in the data against the model-implied distribution under the estimated parameters. Each circle is labeled with the corresponding network size. The deviation from the 45° line measures model fit.

The specification with uniform dealer bargaining power fits the network distribution between two and 19 dealers. In the bottom left graph in Figure 5, the circles are on the 45° line for almost all values of N less than 20. The goodness of fit deteriorates, however, for larger networks, and the model does not generate networks larger than 22 dealers.

The bottom right graph in Figure 5 plots the empirical relation between trading costs and network size from column (4) of Table V (circles), $c(N^*) = 51 + 0.32N^* - 6.29 \ln N^*$, against its model-implied counterpart (dashed line).[24] The distance between the two lines is a measure of model fit

[24] Because values reported in Table V are estimated on individual transactions and contain many control variables, it is simpler to plot the functional relationship rather than some other transformation of the underlying data.

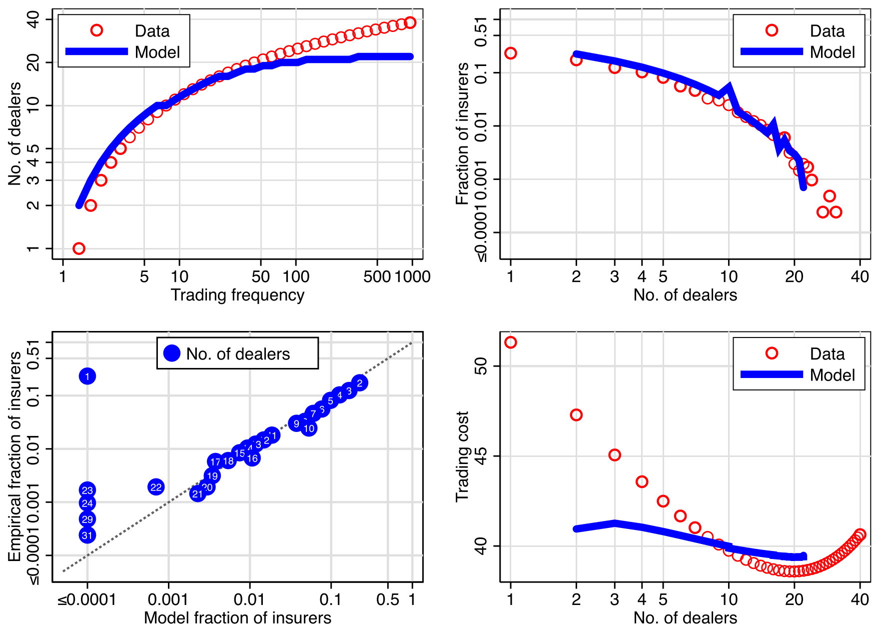

Figure 5. Model fit with constant bargaining power w^b, w^s. This figure shows the model fit for the base specification of constant dealers' bargaining power, w. The top left graph (log-log scale) plots the network size of insurer-dealer relationships for insurer buys as a function of the insurers' trading frequency in the data (circles) compared to the model-implied distribution under the estimated parameters (solid line). The top right graph (log-log scale) plots the degree distribution for network size in the data (circles) compared to the model-implied distribution under the estimated parameters (solid line). The bottom left graph (log-log scale) compares the empirical probability of a given network size to the model-implied probability. The bottom right graph (linear-log scale) plots the bid-ask spread in the data as a function of insurers' network size in the data (circles) compared to the model-implied bid-ask spread under the estimated parameters (solid line). (Color figure can be viewed at wileyonlinelibrary.com)

and shows that the parameter estimates from specification 1 do not well describe the relationship between network size and trading costs in the data. The model's relation is too weak as the line is too flat. In addition, the nonmonotonicity in the relation is barely visible. This result suggests that uniform bargaining power limits the variation in the benefits and costs of having a larger network.

In the model, network size affects trading costs, insurers with different η choose different $N^*(\eta)$, and η impacts trading costs independent of the network size. Understanding how these effects interact in the model provides insight into why the model with uniform bargaining power weakly fits the relationship between network size and trading costs. The spread can be written as $SP(\eta(N^*), N^*)$. Its full derivative with respect to N^* is

$$\frac{dSP}{dN^*} = \underbrace{\frac{\partial SP}{\partial N^*}}_{>0} + \underbrace{\frac{\partial SP}{\partial \eta}}_{<0} \cdot \underbrace{\frac{d\eta}{dN^*}}_{\geq 0}. \tag{44}$$

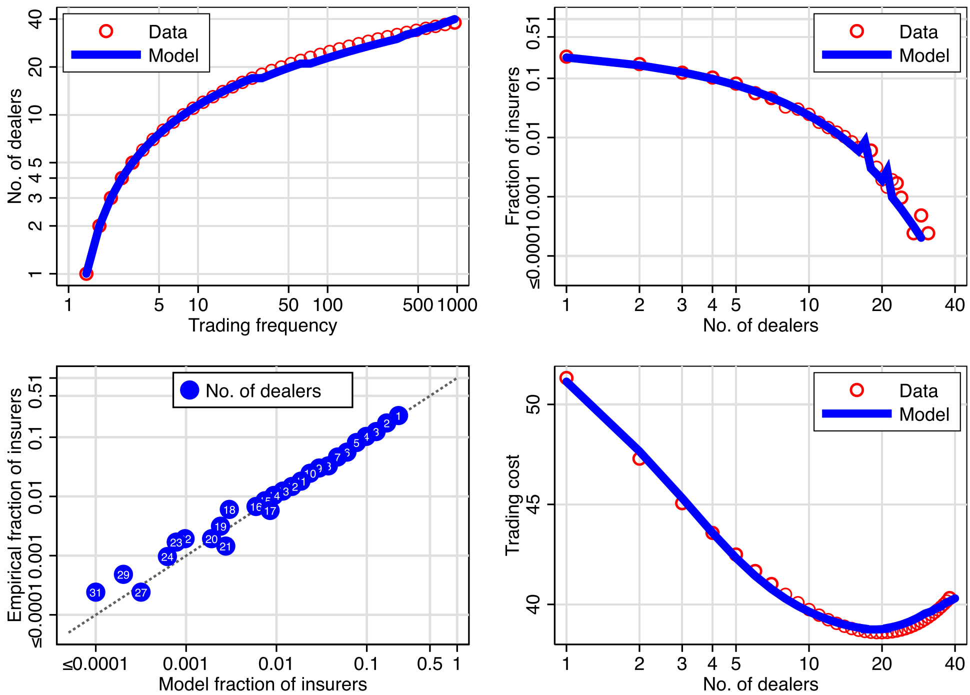

Figure 6. Model fit with client-specific bargaining power $w^b(\eta)$, $w^s(\eta)$. This figure shows the model fit when dealers' bargaining power is a function of trading intensity, $w^b(\eta)$ and $w^s(\eta)$. Please see the caption of Figure 5 for further details. (Color figure can be viewed at wileyonlinelibrary.com)

The top left graph in Figure 5 is the inverse of $\frac{d\eta}{dN^*}$. The top right graph in Figure 5 weighs this term by the observed empirical distribution of trading intensities. Therefore, fitting the model requires fitting $\frac{d\eta}{dN^*}$ and $\frac{dSP}{dN^*}$, and $\frac{dSP}{dN^*}$ depends on $\frac{d\eta}{dN^*}$. The model determines $\frac{\partial SP}{\partial N^*}$ and $\frac{\partial SP}{\partial \eta}$.

The first term in (44) corresponds to the derivative of the spread with respect to optimal network size in the model. As discussed earlier, at the optimum this derivative is positive. The second term in (44) corresponds to the effect of insurers with different η choosing different N^* and η impacting trading costs independent of the network size. The value of future business causes the bid-ask spread to improve with η, $\frac{\partial SP}{\partial \eta} < 0$, and the optimal network size is a nondecreasing function of η, $\frac{d\eta}{dN^*} \geq 0$, so the second term in (44) is always nonpositive. The opposing signs of the two terms explain how the model produces a nonmonotonic relationship between network size and trading costs in the data. To fit the trading cost-network size relationship in Table V, the full derivative in (44) must be initially negative and then become positive at 20 dealers. The slope of the line for the model with uniform bargaining power is not negative enough initially, nor does it increase noticeably enough at $N = 20$, to fit the data.

The log-log plot in the top left graph of Figure 5 shows that fitting $\frac{d\eta}{dN^*}$ requires that $N^* = 1$ for a range of η's. The estimation can achieve this for a wide range of parameters by setting dealers' bargaining power $w^{b,s}$ close to one, but it does not do so in specification 1. This is because increasing dealers' bargaining power has two effects for a fixed η. First, it widens the spread, and second by shifting benefits of repeat trading from a client to a dealer, it reduces the optimal network size. In contrast, increasing trading frequency η while keeping $w^{b,s}$ fixed increases the total value of repeat trading thus narrowing the spread and increasing the optimal network size. In the data, the clients with the smallest networks pay wider spreads, with the widest spread paid by clients with a single dealer, and the spread narrows monotonically with network size until $N = 20$ and then widens. Bargaining power is fixed in specification 1 of the model, while η varies over a very wide range of values. Setting $w^{b,s}$ close to one allows the model to fit small networks (from $N = 2$ to $N = 9$) for a wide range of η, match the widest spreads, and make the spread decline with N for $N < 20$. Unfortunately, this same mechanism makes it costly to have larger dealer networks, so the model produces no insurer with a single dealer and no networks larger than 22. Finally, high dealer bargaining power weakens the relationship between spreads and network size. If dealer bargaining power declines with insurer trading intensity, then larger networks may be optimal and the relationship between network size and trading costs may be larger. Specification 2 allows bargaining power to vary with η.

The model fit improves significantly when dealers' bargaining power depends on η. All four plots in Figure 6 show a close correspondence between the model and the data. The graphs show network sizes greater than 22, and the U-shape in the bid-ask spread starts at 20 dealers. In this specification, dealers' bargaining power is greater than 0.8 for insurers with trading intensity less than 5.[25] Dealer bargaining power for insurer buys is about 0.5 when trading intensity approaches 10. When η exceeds 30, dealer bargaining power for insurer buys is small. This enables the model to produce large networks for insurers with large η, as the value that dealers place on repeat business increases with η. In addition, bargaining power varying with η has a direct effect on trading costs. Larger insurers have greater bargaining power, which reduces their trading costs. This strengthens the relationship between trading costs and network size.

Consistent with the treatment above, the top right graphs of Figures 5 and 6, which depict the fraction of insurers with different network size, contain several nonmonotonic spikes. These spikes arise from the fact that the number of dealers in the model are constrained to be integers. This constraint is also visible in the top left graphs of Figures 5 and 6, which depict the network size for different trading frequencies. The model line in the top left graphs has flat regions since a range of trading frequencies η lead to the same network size. These flat regions cause the fraction of dealers to be nonmonotonic. The locations of the flat regions correspond to the spikes in the top right graphs.

[25] Figure 2 shows that the majority of insurers have trading intensity less than 5.

Table VIII
Model-Implied Persistence in Insurers' Trading Network: Simulation Evidence

This table reports model-implied switching probabilities, p(No. of dealers in $t+1$ | No. of dealers in t), for choosing a network size conditional on the insurer's behavior in the past year. The simulation uses model parameter estimates from Table VI. Details of the simulation are reported in the text.

	No. of Dealers Next Year							
	1		2–5		6–10		>10	
No. of Dealers this Year	Data	Model	Data	Model	Data	Model	Data	Model
1	0.61	0.75	0.30	0.21	0.06	0.04	0.03	0.01
2–5	0.20	0.13	0.54	0.67	0.20	0.17	0.06	0.03
6–10	0.06	0.04	0.30	0.26	0.40	0.53	0.24	0.17
>10	0.01	0.00	0.07	0.03	0.17	0.13	0.75	0.84

To further examine the ability of the model to fit the data, we compute the model-implied switching probabilities, p(No. of dealers in $t+1$ | No. of dealers in t), between different network sizes and compare them to their empirical counterpart reported in Table IV.[26] To do so, we simulate panels of realized network sizes N_{it} for insurer i in year t using the actual number of buy transactions by insurer i in year t. Specifically, for each insurer in the sample, we use trading intensity η_i, as computed earlier in this section, and the estimated model parameters from specification 2 of Table VI as inputs to determine the insurer's optimal network size N_i^*. Then, for each trade in the actual data set, we simulate which of the N_i^* dealers in the insurer's network is matched to a given trade according to random search across the dealers in the network. We record the identity of the dealer handling each such trade and use these identities to compute the realized dealer network, N_{it}, as the unique number of dealers from the insurer's network matched to a trade with the insurer in year t. The resulting panel of realized network sizes N_{it} for all insurers and all years yields model-implied switching probabilities. We then bootstrap the switching probabilities by repeating the trade-by-trade simulation of the insurer-dealer matching several times and averaging over simulations. The switching probabilities are highly stable from simulation to simulation, so 10 simulation rounds suffice. Table VIII shows that the model can match the persistence in networks, as the model-implied transition matrix is close to the actual transition matrix.

To summarize, the theoretical model in which dealers' bargaining power declines with trading intensity η fits the data well. In particular, it is important that dealers' bargaining power on the ask side depend on η. This leads to small inactive insurers having very weak outside options both when they buy and

[26] This assumes that the network is in steady state. In our model, all dealers are identical. How clients dynamically choose new dealers when dealers are heterogeneous is an important question that our model does not address. See Chaney (2018) for an example of a dynamic network formation model with heterogeneous counterparties.

when they sell. By contrast, large active insurers have strong outside options when they buy and weak bargaining power when they sell.

E. Impact of Policy Changes

The parameter estimates help us quantify both the value of repeat client-dealer trading and nontrade value V^r. In this section, we use the model and estimated parameters to construct counterfactuals that capture the impact of repeat trade business on client-dealer networks and prices. We first compute counterfactual networks and bid-ask prices under the assumption that dealers extend only partial price improvement for future repeat trading from a given client. In this case, search and bargaining proceed as above, but the dealers in an insurer's network are implicitly rechosen with probability one-half after every trade. This is counterfactual scenario 1. Under this counterfactual scenario, the bid and ask price equations in the model are adjusted by dividing both U^o and U^{no} in (20) and (21) by 2.

To capture the impact of V^r on networks and transaction prices, we also compute counterfactual networks and bid-ask prices under the assumption that clients fully benefit from repeat trading but nontrade relationships with dealers are segregated, which we achieve by setting $V^r = 0$ in the optimization and then adding it back at the new optimum. This is counterfactual scenario 2. Setting $V^r = 0$ in the optimization has the indirect effect that clients choose smaller networks when the nontrade value is zero. Setting the nontrade value to zero also reduces client valuations. An interpretation of the direct effect on clients' valuation is that unbundling trading from research and other activities means that either dealers now no longer provide nontrade benefits to clients or dealers charge clients for the services that provide these benefits. Our expectation is that clients will lose (some of) these nontrade benefits. We focus on the indirect impact on client valuations by adding the nontrade value back in when we recompute client value. An alternative policy experiment would be to assume that clients lose a fraction or all of their direct nontrade benefits, in which case the decrease in value would be substantially larger.

The top left plot in Figure 7 corresponds to the top left plot in Figure 6 and depicts the optimal network size, N^*. Because network size changes under the counterfactuals, the top right plot in Figure 7 shows the bid-ask spread as a function of trading intensity, η. All four graphs in Figure 7 use the parameters estimated from specification 2 in Table VI. Each graph plots the results from specification 2 (solid line), counterfactual 1 (crosses), and counterfactual 2 (circles).

Figure 7 highlights the fundamental trade-offs that the client faces in our model. When dealers' repeat trade benefits are reduced (counterfactual 1), dealers charge wider bid-ask spreads. Dealers in smaller networks lose more repeat trade surplus than dealers in large networks because each dealer's per-trade loss is scaled by the network size. Smaller clients with a lower trading intensity have smaller networks. Consequently, the widening in bid-ask spread is slightly larger for clients with a lower trading intensity than for clients who trade very

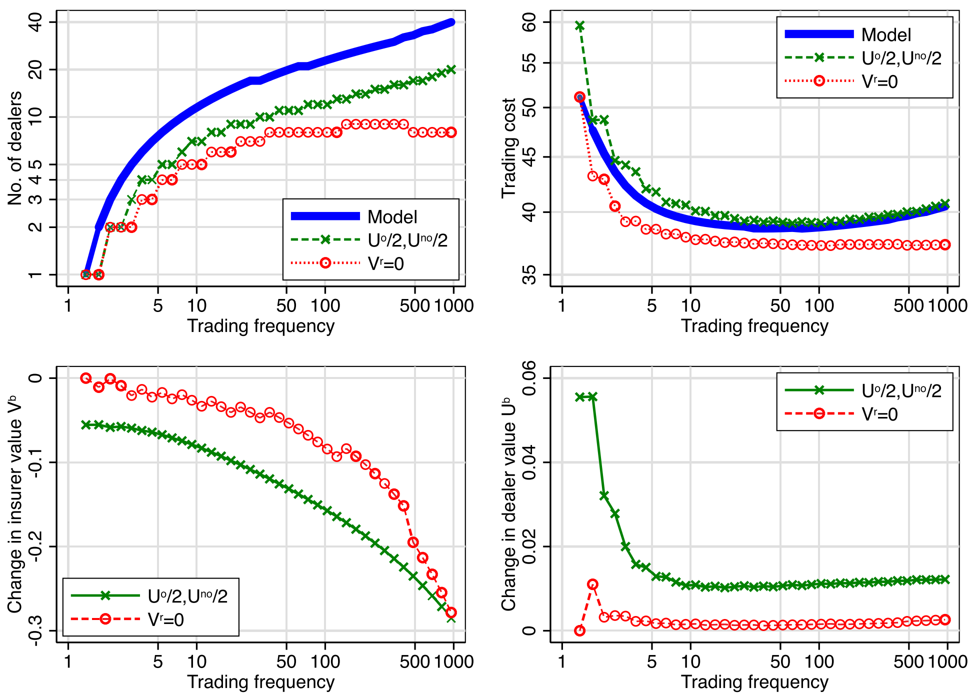

Figure 7. Counterfactual analysis. The top left graph plots the client-dealer network size as a function of clients' trading frequency, η. The top right graph plots the bid-ask spread as a function of η. The bottom left graph plots the change in clients' value from trading, V^b, as a function of clients' trading frequency, η. The bottom right graph plots the change in dealers' value from trading, U^b, as a function of clients' trading frequency, η. Each plot shows the corresponding variable under the parameters estimated from specification 2 in Table VI (solid line) compared to the counterfactuals that would arise if (I) dealers suffer a permanent loss of 50% of repeat trade business (crosses) or (II) the nontrade relationship value is zero (circles). We use a log-log scale. (Color figure can be viewed at wileyonlinelibrary.com)

frequently. Both large (high η) and small (low η) clients respond to wider a bid-ask spread by reducing the size of their network. The magnitude of this effect is much smaller for clients with low η as they already have small networks.

Clients reduce their network size when the nontrade value V^r is segregated from trading in counterfactual 2. Because V^r is larger for clients with a larger trading intensity, they reduce the size of their network more than clients with low η. Dealers respond to the reduction in network size by charging better transaction prices, thus leading to narrower bid-ask spreads for all η-type clients. The magnitude of the bid-ask spread improvement is larger (smaller) for clients with high (low) η because they reduce the size of their network more (less).

The U-shape in the bid-ask spread is not present without the nontrade relationship value V^r. This highlights the role that V^r increasing in N^* plays in the nonmonotonicity of the relationship between network size and the bid-ask spread. As V^r increases with N^*, clients have an incentive to choose a larger network than in the absence of the nontrade value, because larger transaction costs are offset by the nontrade value. Because V^r increases with η as well, this incentive is especially strong for clients with a very large trading intensity and a large network, thus leading to the U-shape in the bid-ask spread for large networks.

MiFID II attempts to unbundle various aspects of relationships between clients and dealers. Trading and nontrading services are supposed to be priced and sold separately. In addition, U.S. and European regulators seem to favor shifting corporate bond trading to centralized electronic platforms like MarketAxess, Tradeweb, and TruMid. These initiatives likely impact the transaction costs incurred by clients, but very little empirical and/or theoretical evidence exists that quantifies these effects. Our counterfactuals provide quantitative insights regarding both unbundling trading and nontrading dealer services (counterfactual 2) and moving bond trading from OTC to more centralized electronic platforms where the probability of repeat trading with the same dealer is reduced (counterfactual 1).

The top right panel of Figure 7 shows that when the probability of repeat trading with the same dealer is reduced (green crosses), all clients will incur higher transaction costs. Insurers that trade more frequently and, in turn, have a larger network will see a much smaller increase in transaction costs than clients that trade less frequently and tend to trade repeatedly with one to five dealers. For instance, clients with 100 annual trades have on average 11 dealers in their network and see an increase in transaction costs of less than 1 basis point per bond per transaction. These clients tend to be larger with more market power. In contrast, smaller clients with less than 10 trades annually have networks with up to five dealers and pay on average 3–7 basis points more per bond per transaction.

With unbundling (circles), transaction costs improve for all clients, with the largest clients trading most frequently realizing the largest improvement. This is because with bundling, these clients obtain the largest nontrade benefits from dealers at the expense of transaction prices. The largest clients reduce their network size the most to benefit from trade-related relationship value. Overall, our results indicate that decoupling trade and nontrade client-dealer business are likely to decrease transaction costs for all but a few clients.

The bottom left plot of Figure 7 shows that the proposed regulations are expected to decrease the value from trading, V^b, for all client types, with larger and more active clients affected differently than smaller and less active clients. In the model, higher η clients incur the largest reduction in V^b. These clients also reduce the size of their trading network the most. If unbundling only impacts network size and clients still receive nontrade relationship value from their dealers, clients' welfare loss is significant but does not turn the value from trading negative. However, clients with lower η already extract little value from trading. As a consequence, under the new regulations the value they generate from trading may become negative. If small clients lose both trade and nontrade benefits, they may exit the market altogether. Thus, the proposed regulations can potentially reduce trading activity and decrease clients' welfare. Indeed, if MiFID II or other regulation eliminate both trade and nontrade relationship value, the decline in client welfare may be substantial.

As for the dealers, the bottom right graph of Figure 7 shows that the reduction in network sizes by clients raises the value of trading for those dealers that still remain in the network. The profits from trading, U^b, also rise for the remaining dealers across all types of investors, which is somewhat contrary to the intent of the proposed regulations, as our results suggest that profits rise the most for low η clients, for which U^b is already the highest.

V. Conclusion

OTC markets, which are pervasive across asset classes, are often considered as functioning poorly due to search frictions arising from a lack of transparency and fragmented trading. Regulators have attempted to address these concerns by unbundling trade and nontrade services, increasing transparency, and encouraging the use of electronic trading platforms. However, while changes in regulations and market structure impact heterogeneous investors differently, there is limited theory closely linked to empirical work to guide these decisions.

In this paper, we use comprehensive regulatory corporate bond trading data for all U.S. insurance companies to examine investors' choice of trading network size and the relation between network size and the transaction prices investors receive in the current decentralized OTC market setting. We document that 30% of insurers trade with a single dealer, whereas few insurers have networks of up to 40 dealers. The cross-sectional distribution in trading activity follows a power law, whereas network size follows a power law with exponential tail. Trading costs decline with network size up to 20 dealers and then increase for larger networks.

A parsimonious model of OTC trading in which insurers choose one-to-many dealer relationships can match the empirical regularities in our sample. In the model, insurers trade off the value of repeat relationships with fewer dealers and the benefits of faster execution with more dealers. The value of repeat relationships declines in the number of dealers as an increase in competition erodes the probability of transacting. Dealers compensate for a loss of repeat business by charging higher spreads. The value of repeat relationships diminishes more slowly with the addition of dealers for clients with larger trading intensity as dealers compete for more repeat business. Thus, larger clients use more dealers and get better execution as the benefits of more repeat business swamp the costs of a larger network. Dealers provide better prices to these larger clients because their repeat business is more valuable. Eventually, the costs of a larger network outweigh the benefits and the spread starts to increase with network size.

We use the structurally estimated model parameters to counterfactually assess the impact of regulations that erode client-dealer relationships. We find that clients incur higher transaction costs when repeat trading is reduced through, for example, anonymous trading. Unbundling trade and nontrade services reduce the optimal network size and decrease transaction costs for all insurers except the least active ones. If clients cease to receive nontrading benefits from dealers, clients that trade infrequently may cease trading altogether, which can reduce market liquidity. Our analysis thus suggests that, unless nontrading services can be provided in a significantly more efficient way in an unbundled setting, unbundling decreases welfare.

Initial submission: March 15, 2018; Accepted: January 9, 2019
Editors: Stefan Nagel, Philip Bond, Amit Seru, and Wei Xiong

Appendix A: Data Filters

Our insurance company Schedule D raw filing data are downloaded from SNL financial, which obtains the data from the insurance company regulator, NAIC. Parts 3 and 4 of Schedule D filed with the NAIC contain purchases and sales made during the quarter, except for the last quarter. In the last quarter of each year, insurers file an annual report in which all transactions during the year are reported. Part 3 of Schedule D reports all long-term bonds and stocks acquired during the year but not disposed of, whereas Part 4 of Schedule D reports all long-term bonds and stocks disposed of. In addition, all long-term bonds and stocks acquired during the year and fully disposed of during the current year are reported in the special Part 5 of Schedule D. We compile the information in Parts 3–5 of Schedule D to obtain a comprehensive set of corporate bond transactions by all insurance companies regulated by NAIC.

There are several differences between the raw filing data used in this paper and the "Time and Sales Data" from Mergent FISD available on WRDS and widely used in previous studies such as Bessembinder, Maxwell, and Venkataraman (2006). Most importantly, the Mergent data do not identify the insurers, which is crucial in identifying relationships in this paper. Moreover, the Mergent data use only filings from the last quarter of each year, which report all purchases and sales made in the calendar year. However, we find that annual reports appear to be missing many trades that show up in the quarterly report. Thus, for completeness, we compile the trade data using the quarterly reports from Q1 to Q3 and all trades that happened in Q4 using the Q4 annual report.

NAIC's counterparty field reports names in text, which can sometimes be mistakenly typed. The bank with the most variation in spelling is DEUTSCHE BANK. We manually clean this field to account for different spellings of broker-dealer names.

We apply various FISD-based data filters based on Ellul, Jotikasthira, and Lundblad (2011) to eliminate outliers and establish a corporate bond universe with complete data. The data filters are described in Table A1, which summarizes the number of observations that are affected by each step. We exclude a bond if it is exchangeable, preferred, convertible, MTN, foreign currency denominated, puttable, or has a sinking fund. We also exclude CDEB (U.S. Corporate Debentures) bonds, CZ (Corporate Zero) bonds, and all government bonds (including municipal bonds) based on the reported industry group,

Table A1
Data Filters

Filter	Full sample	Corporate bonds
1. All trades from original filings (includes all markets, all trades since 2001).	19.1	4.5
2. Remove all trades that do not involve a dealer (e.g., paydown, redemption, mature, and correction).	6.6	3.1
3. Remove duplicates, aggregate all trades of the same insurance company in the same CUSIP on the same day with the same dealer.	6.5	3.1
4. Map dealer names to SEC CRD number, drop trades without a name match, and drop trades with a dealer that trades fewer than 10 times over the sample period.	6.1	2.9
5. Drop if not fixed coupon (based on eMaxx data), drop if information on outstanding amount information is missing in both eMaxx and FISD.	5.3	2.5
6. Drop if trade is on a holiday or weekend.	5.2	2.5
7. Drop if counterparty is "various."	4.1	2.1
8. Drop trades less than 90 days to maturity or less than 60 days since issuance (i.e., primary market trade).	2.8	1.5
9. Merge with FISD data, keep only those securities that are not exchangeable, preferred, convertible, issued by domestic issuer, taxable muni, or missing the offering date, offering amount, or maturity, and for which offering amount is not less than $100K.	1.0	1.0

as well as bonds for which any of the following fields are missing: offering date, offering amount, and maturity. We further restrict our sample to bonds with offering amount greater than $10 million, as issues smaller than this amount are highly illiquid and hence rarely traded. Ellul, Jotikasthira, and Lundblad (2011) use a cutoff of $50 million, which we find restrictive for our purpose. We winsorize the cash-to-assets ratio at 99% to remove extreme values.

REFERENCES

Afonso, Gara, 2011, Liquidity and congestion, *Journal of Financial Intermediation* 66, 1109–1139.

Afonso, Gara, Anna Kovner, and Antoinette Schoar, 2013, Trading partners in the interbank lending market, Staff report, Federal Reserve Bank of New York.

Atkeson, Andrew, Andrea Eisfeld, and Pierre-Olivier Weill, 2015, Entry and exit in OTC derivatives markets, *Econometrica* 83, 2231–2292.

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

Babus, Ana, and Tai-Wei Hu, 2017, Endogenous intermediation in over-the-counter markets, *Journal of Financial Economics* 125, 200–215.

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

Bernhardt, Dan, Vladimir Dvoracek, Eric Hughson, and Ingrid Werner, 2005, Why do large orders receive discounts on the London Stock Exchange? *Review of Financial Studies* 18, 1343–1368.

Bessembinder, Hendrik, Stacey Jacobsen, William Maxwell, and Kumar Venkataraman, 2018, Capital commitment and illiquidity in corporate bonds, *Journal of Finance* 73, 1615–1661.

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

Board, Simon, 2011, Relational contracts and the value of loyalty, *American Economic Review* 101, 3349–3367.

Campbell, John Y., and Glen B. Taksler, 2003, Equity volatility and corporate bond yields, *Journal of Finance* 58, 2321–2350.

Chaney, Thomas, 2018, The gravity equation in international trade: An explanation, *Journal of Political Economy* 126, 150–177.

Chang, Briana, and Shengxing Zhang, 2015, Endogenous market making and network formation, Working paper, University of Wisconsin–Madison.

Chodorow-Reich, Gabriel, Andra Ghent, and Valentin Haddad, 2016, Asset insulators, Working paper, Harvard University.

Choo, Eugene, 2015, Dynamic marriage matching: An empirical framework, *Econometrica* 83, 1373–1423.

Colliard, Jean-Edouard, and Gabrielle Demange, 2014, Cash providers: Asset dissemination over intermediation chains, Working paper, Paris School of Economics.

Condorelli, Daniele, and Andrea Galeotti, 2012, Endogenous trading networks, Working paper, University of Essex.

Diamond, Peter, 1971, A model of price adjustment, *Journal of Economic Theory* 3, 156–168.

DiMaggio, Marco, Amir Kermani, and Zhaogang Song, 2017, The value of trading relationships in turbulent times, *Journal of Financial Economics* 124, 266–284.

Duffie, Darrell, Nicolae Garleanu, and Lasse Pedersen, 2005, Over-the-counter markets, *Econometrica* 73, 1815–1847.

Duffie, Darrell, Nicolae Garleanu, and Lasse Pedersen, 2007, Valuation in over-the-counter markets, *Review of Financial Studies* 20, 1865–1900.

Eckstein, Zvi, and Gerard J. Van Den Berg, 2007, Empirical labor search: A survey, *Journal of Econometrics*, 136, 531–564.

Edwards, Amy, Lawrence Harris, and Michael Piwowar, 2007, Corporate bond market transaction costs and transparency, *Journal of Finance* 62, 1421–1451.

Ellul, Andrew, Chotibhak Jotikasthira, and Christian Lundblad, 2011, Regulatory pressure and fire sales in the corporate bond market, *Journal of Financial Economics* 101, 596–620.

Feldhütter, Peter, 2012, The same bond at different prices: Identifying search frictions and selling pressures, *Review of Financial Studies* 25, 1155–1206.

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

Gavazza, Alessandro, 2016, An empirical equilibrium model of a decentralized asset market, *Econometrica* 84, 1755–1798.

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

Gofman, Michael, 2011, A network-based analysis of over-the-counter markets, Working paper, University of Rochester.

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

Guerrieri, Veronica, and Robert Shimer, 2014, Dynamic adverse selection: A theory of illiquidity, fire sales, and flight to quality, *American Economic Review* 104, 1875–1908.

Guerrieri, Veronica, Robert Shimer, and Randall Wright, 2010, Adverse selection in competitive search equilibrium, *Econometrica* 78, 1823–1862.

Harris, Christopher, and John Vickers, 1985, Perfect equilibrium in a model of a race, *Review of Economic Studies* 52, 193–209.

Harris, Larry, 2015, Transaction costs, trade throughs, and riskless principal trading in corporate bond markets, Working paper, University of Southern California.

Harris, Lawrence E., and Michael S. Piwowar, 2006, Municipal bond liquidity, *Journal of Finance* 61, 1330–1366.

Hitsch, Gunter J., Ali Hortaçsu, and Dan Ariely, 2010, Matching and sorting in online dating markets, *American Economic Review* 100, 130–163.

Hollifield, Burton, Artem Neklyudov, and Chester Spatt, 2017, Bid-ask spreads, trading networks, and the pricing of securitizations, *Review of Financial Studies* 30, 3048–3085.

Hugonnier, Julien, Benjamin Lester, and Pierre-Olivier Weill, 2018, Frictional intermediation in over-the-counter markets, NBER Working paper No. 24956.

Jackson, Matthew O., and Asher Wolinsky, 1996, A strategic model of social and economic networks, *Journal of Economic Theory* 71, 44–74.

Koijen, Ralph S.J., and Motohiro Yogo, 2015, The cost of financial frictions for life insurers, *American Economic Review* 105, 445–475.

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

Lagos, Ricardo, and Guillaume Rocheteau, 2009, Liquidity in asset markets with search frictions, *Econometrica* 77, 403–426.

Lester, Benjamin, Guillaume Rocheteau, and Pierre-Olivier Weill, 2015, Competing for order flow in OTC markets, *Journal of Money, Credit and Banking* 47, 77–126.

Li, Dan, and Norman Schürhoff, 2019, Dealer networks, *Journal of Finance* 74, 91–144.

Moen, Espen R., 1997, Competitive search equilibrium, *Journal of Political Economy* 105, 385–411.

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

Neklyudov, Artem, and Batchimeg Sambalaibat, 2016, Endogenous specialization and dealer networks, Working paper, Swiss Finance Institute.

O'Hara, Maureen, Yihui Wang, and Xing Zhou, 2018, The best execution of corporate bonds, *Journal of Financial Economics* 130, 308–326.

Postel-Vinay, Fabien, and Jean-Marc Robin, 2002, Equilibrium wage dispersion with worker and employer heterogeneity, *Econometrica* 70, 2295–2350.

Schultz, Paul, 2001, Corporate bond trading costs: A peek behind the curtain, *Journal of Finance* 56, 677–698.

Stigler, George J., 1961, The economics of information, *Journal of Political Economy* 69, 213–225.

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

Vayanos, Dimitri, and Tan Wang, 2007, Search and endogenous concentration of liquidity in asset markets, *Journal of Economic Theory* 136, 66–104.

Weill, Pierre-Olivier,2007, Leaning against the wind, *Review of Economic Studies* 74, 1329–1354.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.
Replication code.