Relationship Trading in OTC Markets*

TERRENCE HENDERSHOTT†
Haas School of Business
University of California, Berkeley
Berkeley, CA 94720

DAN LI‡
Board of Governors of the Federal Reserve System
Washington, DC 20551

DMITRY LIVDAN§
Haas School of Business
University of California, Berkeley
Berkeley, CA 94720

NORMAN SCHÜRHOFF¶
Faculty of Business and Economics and Swiss Finance Institute
Université de Lausanne
CH-1015 Lausanne, Switzerland

February 22, 2019

*We thank Philip Bond, the 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 Houston, University of Illinois-Chicago, University of Oklahoma, University of Piraeus, University of Wisconsin at Madison, University of Victoria, WU Vienna, Einaudi Institute for Economics and Finance, and the 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, 2017 OTC Markets and Their Reform Conference in Rigi for valuable comments and feedback.

†Tel.: (510) 643-0619; fax: (510) 643-1412; *email*: `hender@haas.berkeley.edu`

‡Tel.: (202) 452-2626; *email*: `dan.li@frb.gov`

§Tel.: (510) 642-4733; fax: (510) 643-1412; *email*: `livdan@haas.berkeley.edu`

¶Tel.: +41 (21) 692-3447; fax: +41 (21) 692-3435; *email*: `norman.schuerhoff@unil.ch`. Schürhoff is Research Affiliate of the CEPR. He 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.”

Relationship Trading in OTC Markets

Abstract

Relations are known to matter in banking, but not in trading. We examine the network of trading relations between insurers and dealers in the over-the-counter corporate bond market. Comprehensive regulatory data show that every third insurer uses a single dealer while a small fraction of insurers has networks of up to 40 dealers. Execution prices are non-monotone in network size, strongly declining for non-exclusive relations with more dealers but then increasing once networks exceed 20 dealers. To understand these facts we build a model of decentralized trade in which insurers trade off the benefits of repeat business against faster execution. The model quantitatively fits the distribution of insurers' network sizes and how prices depend on insurers' network size. Counterfactual analysis shows that proposed regulations to unbundle trade and non-trade dealer services may decrease welfare.

JEL Classification: G12, G14, G24

Key words: Over-the-counter market, corporate bond, trading relationship, trading cost, liquidity, financial network, unbundling

Insurance companies are vital for risk sharing. They compensate for loss, damage, injury, treatment or hardship in exchange for premium payments. To provide this coverage, insurance companies invest in a variety of financial assets. Corporate bonds comprise almost 70% of their investments with a total value close to $3.8 trillion.[1] To facilitate prompt compensation to policy holders insurers need to be able to liquidate their holdings fast without incurring large transaction costs (Koijen and Yogo, 2015; Chodorow-Reich, Ghent, and Haddad, 2016). However, corporate bonds trade on decentralized over-the-counter (OTC) markets. OTC markets 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 have to search for best execution across more than 400 active broker-dealers. It is an open question whether insurers and other market participants search randomly or build long-term relations with dealers to mitigate search frictions.

We study insurers' choice of trading networks and their execution prices in the OTC corporate bond market. Regulatory data provide information on the trading relations between more than 4,300 insurers and their dealers for all transactions from 2001 to 2014. We first empirically examine insurers' choice of trading networks and how these relate to transaction prices. Insurers form small, but persistent dealer networks. Figure 1 provides two examples of insurer-dealer trading relations over time. Panel A shows buys and sells for an insurer repeatedly trading exclusively with a single dealer. Panel B shows an insurer trading with multiple dealers over time. Every third insurer trades with a single dealer annually. A small fraction of insurers trades with up to 40 dealers in a year. The overall degree distribution follows a power law with exponential tail starting at about 10 dealers. We estimate trading costs as a function of network size N. Costs are non-monotone in N; costs decline with N for small networks and then increase once N exceeds 20 dealers.

Our evidence provides insights into which models of trade in OTC markets better describe the empirical evidence on client-dealer relations and trading costs. In random search models clients repeatedly search for best execution without forming a finite network of dealers (Duffie et al. 2005, 2007; Lagos and Rocheteau, 2007, 2009; Gavazza, 2016). The empirical 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 each additional dealer imposes a fixed cost on insurers. Insurers add dealers to improve prices up to the point where the marginal benefit equal the fixed cost. This leads to prices improving monotonically in

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

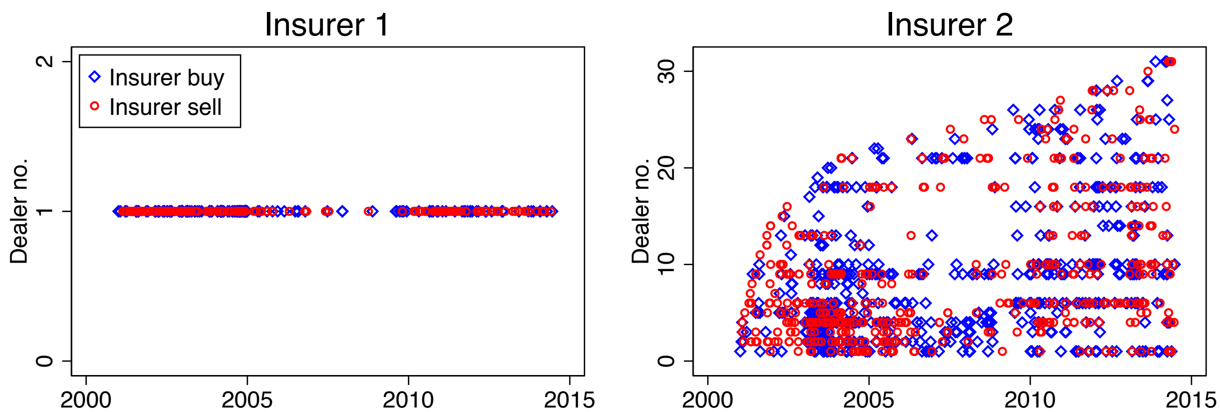

Figure 1: Example of dealer-client trading relations

The figure shows the buy (blue squares) and sell (red circles) trades of two insurance companies with different dealers. We sort the dealers on the vertical axis by the first time they trade with the corresponding insurance company.

network size, which is inconsistent with the empirical non-monotonicity of trading costs as a function of 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 against faster execution. We model the full relationship between clients and dealers as having two components. The first component is the repeat trading interactions between each client and the dealers. The client repeatedly buys and sells bonds from dealers in her network. This future repeat trading is taken into account by the client and the dealers when negotiating terms of each transaction. The second component of the relationship captures all other business between clients and dealers unrelated to bond trading. It includes transacting in other securities, the ability to purchase newly issued securities, as well as other soft dollar and non-monetary transfers such as investment research.

In our model a single console bond trades on an inter-dealer market which clients can only access through dealers. Dealers have search intensity λ and upon trading with a client transact at the competitive inter-dealer bid and ask prices. Clients initially start without a bond but stochastically receive trading shocks with intensity η which cause them to simultaneously contact N dealers to buy. Trading intensity η and, hence, N vary across clients. The 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. Thus, our trading mechanism is identical to repeated 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 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 either network size or prices, but not both. Random search models assume investors may contact every other counterparty. 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 chose to 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 with investors choosing between dealers with either large or small interdealer networks instead of asset markets. By contrast, both the network size and transaction prices are endogenously determined in the equilibrium of our model.

Our analysis delineates an important trade-off leading to optimal network size, N^*, and transaction prices. Clients trade off repeat relations with dealers against 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 losses of repeat business by charging higher spreads. Eventually the costs of having larger network outweigh the benefits and the dealers' spread starts to increase with network size. This corresponds to the empirical non-monotonicity in trading costs with respect to the 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 larger repeat business. Therefore, these larger clients use more dealers and get better execution as benefits from having larger repeat business outweigh the costs of having 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. The 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 by 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 enables 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 sizes, $p(N^*)$, to their empirical counterparts.

Using the structurally estimated parameters along with the distribution of trading intensities quantitatively reproduces both the empirical distribution of network sizes 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, while 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 the insurers' trading frequency. In contrast, dealers' bargaining power when buying declines significantly with the insurers' trading frequency.

The regulatory authorities in the U.S. and Europe follow different approaches to address OTC market frictions. Regulatory initiatives in the U.S., including the Dodd-Frank Act, encourage dealer competition through fostering of multilateral electronic trading platforms. These erode long-term client-dealer relations. In Europe, the Markets in Financial Instruments Directive (MiFID) II requires the unbundling of trading and non-trading 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 on dealers and their clients.

Our model enables the quantification of the regulatory impact on trading costs and the value of repeat relations. We perform counterfactual analysis by either reducing the probability of repeat trading with the same dealer or eliminating the non-trade value of relationships. All types of insurers incur higher transaction costs when the probability of repeat trading with the same dealer is reduced, e.g., due to required anonymous electronic trading venues/facilities.[2] The impact on various insurer types and sizes is different. Insurers that trade more frequently and, therefore, have larger networks see a smaller increase in transaction costs than insurers who trade less frequently and tend to trade repeatedly with 1 to 5 dealers.

Unbundling trade and non-trade insurer-dealer business, as stipulated by MiFID II regulations, reduces the optimal network sizes and decreases transaction costs for all insurers, except for the least active. The presence of non-trade relationship value causes trading networks to be larger than what is required to minimize trading costs. The loss of non-trade relationship value can cause infrequently trading insurers to cease trading completely, lowering market liquidity. Unless the non-trade relationship value can be captured efficiently in an unbundled market, insurer welfare declines.

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

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, et al. (2005), Bessembinder et al. (2006), Harris and Piwowar (2007), Green et al. (2007) document the magnitude and determinants of transaction costs for investors in OTC markets. Our paper deepens the understanding of OTC trading costs by using the identities of all insurers along with their trading networks and execution costs. These help explain the substantial heterogeneity in execution costs observed in these studies. O'Hara et al. (2015) and Harris (2015) examine best execution in OTC markets without formally studying investors' optimal network choice.

There exists an empirical literature on the value of relationships in financial markets. Similar to our findings, Bernhardt et al. (2004) show that on the London Stock Exchange broker-dealers offer greater price improvements to more regular customers.[3] Bernhardt et al. (2004) do not examine the client-dealer networks and in their centralized exchange setting quoted prices are observable. Afonso et al. (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. These have a significant impact on how liquidity shocks are transmitted across the market. Afonso et al. (2013) do not formally model the network and do not observe transaction prices. DiMaggio et al. (2015) study inter-dealer relationships on the OTC market for corporate bonds while our focus is on the client-dealer relations.[4]

The role of the interdealer market in price formation and liquidity provision are the focus in Hollifield et al. (2015) and Li and Schürhoff (2018). These studies explore the heterogeneity across dealers in their network centrality and how they provide liquidity and what prices they charge. By contrast, we focus on the heterogeneity across clients and how trading intensity affects their networks and transaction prices.

The search-and-matching literature is vast. Duffie et al. (2005, 2007) provide a prominent treatment of search frictions in OTC financial markets, while Weill (2007), Lagos and Rocheteau (2007, 2009), Feldhütter (2011), Neklyudov (2014), Hugonnier et al. (2015), Üslü (2015) generalize the economic setting. These papers do not focus on repeat relations and do not provide incentives to investors to have a finite size network. Gavazza (2016) structurally estimates a model of trading in decentralized markets with two-sided one-to-one search and

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

[4]Our paper also relates to a growing literature studying trading in a network, e.g., Gale and Kariv (2007), Gofman (2011), Condorelli and Galeotti (2012), Colliard and Demange (2014), Glode and Opp (2014), Chang and Zhang (2015), Atkeson et al. (2015), Babus (2016), Babus and Hu (2016), and Babus and Kondor (2016). These papers allow persistent one-to-one dealer-client relationships, while the main focus of our model is on clients' networks.

bilateral bargaining using aircraft transaction data. At the market-wide 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 relations 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 typically rely on a concept of competitive search equilibrium proposed by Moen (1997) for labor market. Examples include Guerrieri, Shimer, and Wright (2010), Guerrieri and Shimer (2014), and Lester et al. (2015). These papers explain assortative matching between clients and dealers and show how heterogeneity affects prices and liquidity. However the matching technology employed by these papers is one-to-one, thus limiting the network size to a single dealer.

The remainder of the paper is organized as follows. Section 1 describes the data. Sections 2 documents our findings on insurer trading activity, dealer networks, and trading costs. In Section 3 we develop an OTC model with repeat trade and dealer networks. We test its predictions and provide counterfactuals in Section 4. Section 5 concludes.

1 Data

Insurance companies, who invest heavily in the corporate bond market, file quarterly reports of trades of long-term bonds and stocks to the National Association of Insurance Commissioners (NAIC). For each trade the NAIC Schedule D data report the dollar amount of transactions, par value of the transaction, insurer code, date of the transaction,[5] the counterparty dealer name, and the direction of the trade for both parties, e.g., whether the trade was an insurance company buying from a dealer or an insurance company selling to a dealer. NAIC started collecting the schedule D data almost a decade before the TRACE became available, which makes the NAIC transaction data an important data source for the study of the corporate bond market. For example, Schultz (2001) uses this data to provide one of the earlier estimates of transaction costs of corporate bonds. Ellul et al. (2011) also 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 (insurers) 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 "Time and Sales Data" from Mergent FISD available on WRDS, widely used in previous studies, such as Bessembinder et al. (2006). Most importantly, the public data do

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

not identify the insurers, which is crucial in identifying relationships in our 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. The annual reports suppress, however, 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.

We clean the raw NAIC data by keeping only corporate bond transactions. To focus on secondary trading we only include trades more than 60 days after issuance and trades more than 90 days to maturity. We further remove non-secondary transactions that are associated with redemptions and calls. One of the most challenging, yet important tasks of cleaning the NAIC data is the standardization of the counterparty (dealer) identities. The raw filings report hand-typed counterparty names, which contain typos and alternative names of the same dealer. We use fuzzy matching and human checking to standardize the dealer names. Details of the data cleaning and filtering are documented in the Appendix.

Our final 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 the Mergent Fixed Income Security Database (FISD) and Lipper eMAXX. Insurer financial information come 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 thousand insurer buys and 497 thousand insurer sells.

Table 1 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); (iii) 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. They 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. They 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. They trade on average with 4.81 dealers each year.

The distribution of trading activity is skewed with the top ten insurance companies accounting for 6.3% of trades and 14.3% of trading volume. They 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 as well as for 45.3% of trading volume. The 3,000 smallest insurers use on average 3.76 dealers.

Order splitting is not very common. On just 1.2% of all insurer-days (13K out of 1.1M trades) an insurer trades the same bond multiple times. This 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 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 trades are in investment grade while only 1% are in unrated with the remainder being high yield. Privately placed bond trades form a small minority of our sample at 8%.

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. Insurer size is reported assets. The cash-to-asset ratio is cash flow from the insurance business operations divided by assets.

Overall, there exists 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 the number of dealers ranging on average from one to as many as 40.

2 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, i.e., the number of trades, and intensive margin, i.e., the total dollar volume traded, of insurer trading in a given year. Both margins reveal that insurers have heterogeneous trading needs.

2.1 Insurer trading activity

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

Figure 2 shows the distribution in the average number of trades per year across insurers. A large fraction of insurers do not trade in a given month and we therefore 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 combined. The power law is $0.34 \times X^{-1.31}$ for insurer buys (depicted in Panel A) and

Table 1: Descriptive statistics

The 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)			
IG	0.74	0.76	0.72
HY	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
101-1000	75.66	46.22	11.56
1001+	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)		

$0.40 \times X^{-1.58}$ insurer sales (Panel B). Visually the two power law distributions for insurer buys and sales in Figure 2 look similar. This suggests insurers buy and sell at similar rates, even though these rates vary significantly across insurers.

We next examine what characteristics explain the heterogeneity in trading intensities. Table 2 documents the determinants of the intensive margins (trading volume in $bn, column (1)) and extensive margins (number of trades, column (2)) of the annual trading by insurance

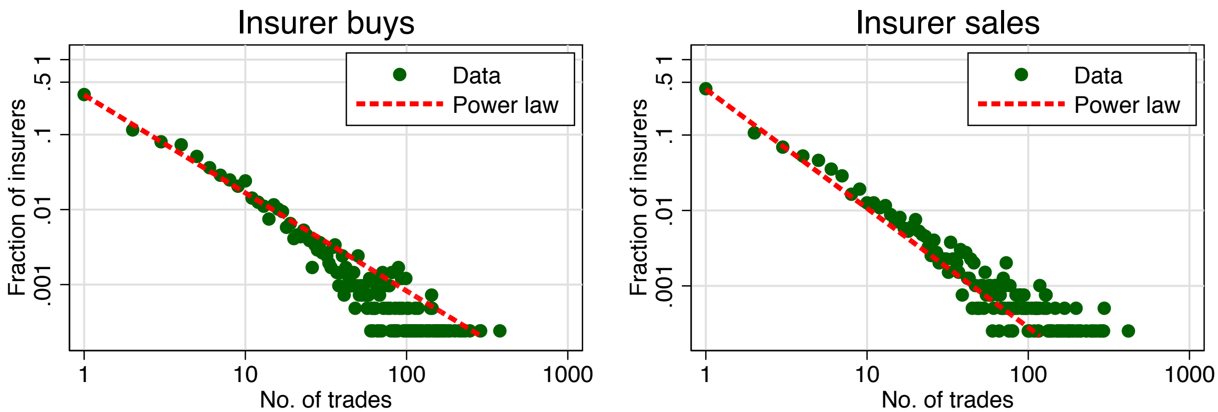

Figure 2: Insurer trading activity

The figure shows the distribution in the number of insurer buys per year (left) and insurer sales (right). We use a log-log scale.

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 the trade size and bond characteristics across all trades of the insurer during the year. The variables capturing the variation in characteristics capture complexity in insurers' portfolios and their need for more frequent rebalancing and for dealer specialization. Insurer characteristics include its 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, 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 the lagged log-volume is 0.67 and the coefficient on the lagged log-number-of-trades is 0.76. Both coefficients are statistically significant at 1% levels.

Insurer trading strongly correlates with insurer size, type, and quality, with bond types and bond varieties as these variables explain 79% of the variation in annual trading volume and 65% variation in annual number of trades. A two-fold increase in insurer's size increases trading volume by 22%. Larger insurance companies and insurers with higher cash-to-assets ratio also trade more often and submit larger orders. Insurers with higher RBC ratios trade less often than insurers with low RBC ratios. Both margins of trading increase with the insurer's rating, i.e., 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

Table 2: Insurers' trading activity

The determinants of insurance company trading activity are reported. We measure trading activity by 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 time period. Estimates are from pooled regressions with time fixed effects. Standard errors are adjusted for heteroskedasticity and clustering at the insurer and time level. Significance levels are indicated by * (10%), ** (5%), *** (1%).

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

traded more often and in larger quantities. Insurers tend to trade privately placed bonds less because potentially they just 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. Neither measure of trade intensity depends on bond issue size and remaining life as their 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 an opposite effect on the trading intensity. For instance, both measures of trading intensity increase with variation in bond rating and bond life. This is consistent with insurers increasing trading intensity when rebalancing their portfolios, i.e. shifting from high-yield to investment-grade bonds or from younger to older bonds.

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

2.2 Properties of insurer networks

The previous section's results demonstrate how insurer characteristics explain the intensity of their trading. There is a large degree of heterogeneity in insurers' trading intensity, with some insurers trading twice per day while others trade just once per year. This section studies how many dealers they trade with over time and how persistent are these networks. While insurers should have heterogeneous demands for their dealer network depending on their trading intensity, how concentrated and persistent their trading is reveals basic network formation mechanisms.

We start with the examples of the insurer-dealer relationship depicted in Figure 1. These 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 engage in long-term repeat, but non-exclusive, relations. We analyze how representative are the examples in Figure 1 and how insurer characteristics determine their network size.

Figure 3 plots the degree distribution across insurers by year, i.e., the fraction of insurers trading on average across all years with the given number of dealers, using a log-log scale. The figure shows insurers trade with up to 31 dealers every year, with some trading with as many as 40 but these represent less than 1/10,000 of the sample. Exclusive relations 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 is consistent with insurers building networks that they search randomly within. Fitting the degree distribution to a Gamma distribution by regressing the log of

[7] Because of the number of insurers that do not trade in a 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 1. For about half of the 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, while 14.8% of insurers trade just once in a year. Thus, many insurers who trade more than once in a year use a single dealer. Figure 4 further examines the relation between trading activity and network size.

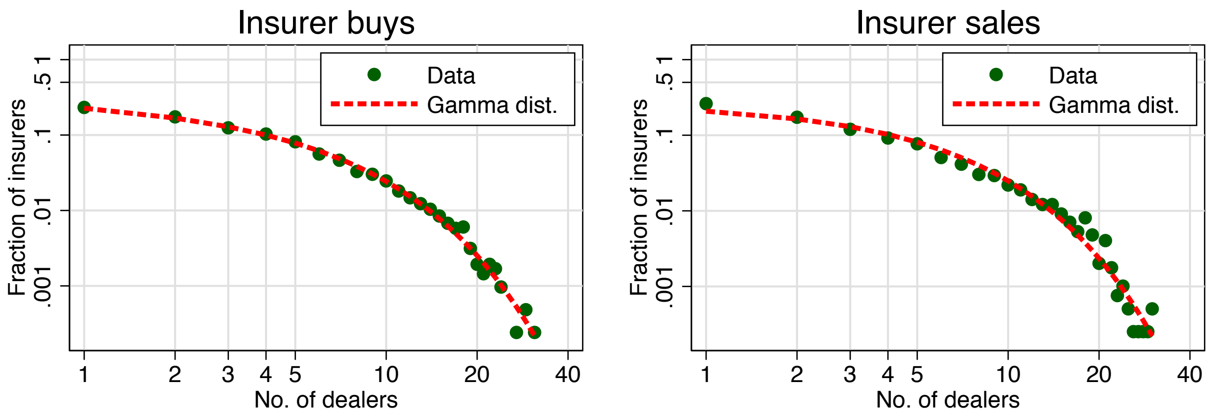

Figure 3: Size of insurer-dealer trading networks

The figure shows the degree distribution for insurer-dealer relations by year for insurer buys (left) and insurer sales (right). We use a log-log scale.

the probabilities of each N on a constant, the logarithm of N, and N yields the following coefficients:

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

Table 3 reports the determinants of insurer dealer network sizes using pooled regressions with time fixed effects. We measure the size of the trading network by the number of dealers that an insurance company trades with in a given year. We log-transform all dependent variables by $100 * log(1 + x)$ 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 whole sample (Column (1)) and, in order to examine how these vary with insurer's size, on sub-samples of small and large insurers based on asset size. We classify an insurer as small if it falls in the bottom three size quartiles and, respectively, 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 varieties matter for the size of the dealer network. Large insurers, which Table 2 shows have larger trading intensity, trade with more dealers. Insurers with demand for larger bond variety have larger networks even controlling for their size (column (2) and (3)). Higher quality insurers, i.e., insurers with higher cash-to-assets ratio and higher ratings, have larger networks, but this matters only for smaller insurers as column (2) indicates. This is potentially due to

Table 3: Size of insurers' trading network

The table reports the determinants of the size of insurers' trading network. We measure the size of the trading network by the number of different dealers that an insurance company trades with in a given year. See caption of Table 2 for additional details. Standard errors are adjusted for heteroskedasticity and clustering at the insurer and time level.

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

it being cheaper for a dealer to set up a credit account for higher quality insurers. These factors matter more for small insurers because they face larger adverse selection problems in forming permanent links with dealers. Insurers with greater variety in the bond they trade have larger networks.[8] Overall these findings suggest insurers' network choice is endogenous

[8]One potential explanation for this result could be that dealers tend to specialize in a few bonds and, as consequence, this relation arises mechanically for insurers demanding greater bond variety. Table IA.2 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.3 shows that for each of these same active bonds, the 10 most active insurers trade in each

and dependent on multiple factors.

Table 4: Persistence in insurers' trading network

The table reports switching probabilities, $p(\text{No. of dealers in } t+1|\text{No. of dealers in } t)$, for using a network size conditional on the insurer's past behavior for each 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	0.06	0.31	0.40	0.24
>10	0.01	0.07	0.17	0.75

Table 3 shows persistence in the size of the network with the coefficient on the lagged network size being 0.75 (column (1)). This result is mostly due to large insurers, because this coefficient equals only 0.62 for small insurers. Table 4 examines this 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 4. Trading relations are persistent from year to year. This is especially true for exclusive relations as the probability of staying with a single dealer each year is equal to 0.61.[9] 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 2 to 5 dealers and 0.06 for insurers with 6 to 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 the stable network sizes are difficult to reconcile with a "pure" random search model à la Duffie et al. (2005, 2007).

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, while 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 we plot in Figure 4 the number of insurer buys (left) and sales (right) per dealer and year by network size. The box and whisker plot provide information on the distribution of trading frequencies per dealer for each network size. While there is a fair amount of cross-section variation for small network

of them with more than one dealer on average. These statistics show that specialization does not appear to be a factor for the large dealers and large insurers which represent a substantial fraction of trading.

[9]For consecutive insurer-years with 1 dealer, 75% of the time the dealer is the same in both years.

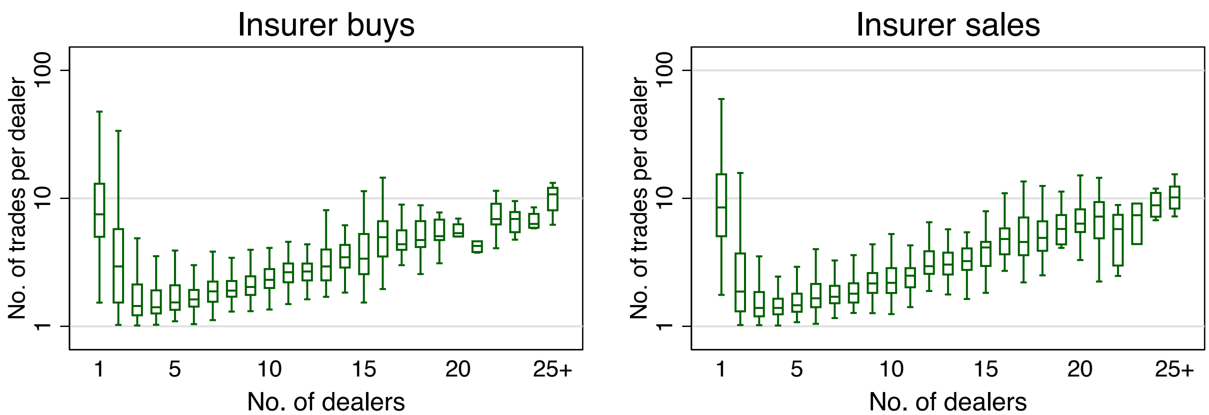

Figure 4: Insurer trading activity and network size

The figure shows the number of insurer buys per year per dealer (left) and insurer sales per year per dealer (right) for all insurers with more than one trade per year on average. The box corresponds to the values at the lower and upper quartile, with median indicated by the line inside the box. The whiskers indicate the lower and upper adjacent Tukey values. We use a log-linear scale.

sizes, 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 non-monotonicity of trades per dealer in network size is inconsistent with network size arising mechanically from trading needs.

The next section studies the relation between client-dealer networks and execution costs.

2.3 Insurer trading costs and networks

Tables 2 and 3 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 both the 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 relative execution cost measure in basis points is defined as

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

where $\mathbf{1}_{Buy}$ is an indicator for whether the insurer is buying or selling. Because some quotes may be stale or trades misreported, leading to extreme costs estimates, we winsorize the distribution at 1% and 99%.

Execution costs depend on the bond being traded, time, whether the insurer buys or sells, the insurer's characteristics, dealer identity and characteristics, and the insurer network size. To examine the relationship-specific effects on execution costs we control for bond and time fixed effects. In principle if the BAML 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 the insurers' networks. Figure 3 and equation (1) indicate that the degree distribution for insurer-dealer relations follows a Gamma distribution. Therefore, 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 either insurer or dealer characteristics, or both. We estimate the following panel regression for execution costs in bond i at time t:

$$\text{Execution cost}_{it} = \alpha_i + \alpha_t + \beta N_{it-1} + \gamma\, lnN_{it-1} + \theta X_{it} + \epsilon_{it}. \tag{3}$$

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

Table 5 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 buy captures the average bid-ask spread of roughly 40 basis points. Column (1) of Table 5 shows that execution costs decline with insurer network size. An insurer with an additional dealer has trading cost 0.22 basis point lower. Large insurers pay on average lower execution costs. An insurer with 10 times as many assets has trading cost 3.72 basis points lower. Similar to the strength of trading relationship mechanism in Bernhardt et. al (2004), the size coefficient suggests that the amount of trading between insurer and dealer can lead to better prices. However, Table IA.4 shows that the inclusion of insurer-dealer trading volume does not impact the coefficients on network size in the specifications in Table 5.

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, while the coefficient on the logarithm of N is -6.29. Both coefficients are statistically significant at 1%. This result indicates that the execution costs are non-monotone in the network size. Improvements in execution quality from having a larger dealer network are exhausted at $N = \frac{6.29}{.32} \approx 20$. Clients with networks of 40 dealers and 10 dealers pay, on average, the same bid-ask

Table 5: Execution costs and investor-dealer relations

The table reports 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 level. See caption of Table 2 for additional details.

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 HY 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

spread of 40 basis points. This finding goes against the traditional wisdom that inter-dealer 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 fixed costs of adding an extra dealer against better execution due to increased dealer competition thus making price a monotonically decreasing function of the 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 dealer share the benefits of repeated 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.

The next section uses our evidence on networks and execution costs to motivate our model of the OTC markets.

3 Model

The model is stylized but still rich enough to allow for the structural estimation of its primitives from the regulatory NAIC data. We view the full relationship between insurers and dealers as having two components. The first component is the repeat trading interactions between each insurer and the dealers. The insurer repeatedly buys and sells bonds from dealers in her network. This future repeat trading is taken into account by the insurer and the dealers when negotiating terms of each transaction. Because our data provides direct evidence on trading relations, we model this component explicitly. The second component of the relationship captures all other business between insurers and dealers unrelated to bond trading. In practice it includes transacting in other securities, the ability to purchase newly issued securities, as well as other soft dollar and non-monetary transfers such as investment research. Because we do not directly observe these non-trade relations in our data, we utilize 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.

3.1 Setup and Solution

The economy has a single risk-free perpetual bond paying 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 et al. (2005) such as liquidity supply/demand shocks on the client side and random search with constant intensity. Following Lester et al. (2015), the bond trades on a competitive market accessible only to dealers.[10] Unlike the frictionless inter-dealer market in Lester et

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

al. (2015), in our model dealers face search frictions as in Duffie et al. (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 it to other dealers at an exogenously given price M^{bid}. A bid-ask spread $M^{ask} - M^{bid} \geq 0$ reflects trading costs or cost of carry.

In the model we use a 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.[11] 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 a search intensity λ.[12] 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.[13] 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.[14] The cost K corresponds to back-office costs of processing and clearing the transaction. Clients trading more frequently incur K more often. While K does not depend directly on network size, clients who choose a larger network trade faster, thus, incurring K more often. The cost K enters clients' value functions thereby affecting their reservation values in bargaining. However, K does not enter the bargaining directly. We postpone the discussion of the impact of K on the transaction prices and networks until the next section where we present the model's solution. Clients' search mechanism can be viewed as a winner-takes-all race with the dealer first to find the bond winning the race. The prize is the spread $P^b - M^{ask}$ when the client buys and $M^{bid} - P^s$ 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

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

[12] Bessembinder et al. (2016) 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 the same bond. We assume the number of clients is large and, therefore, a fraction of them is hit by the buy(sell) liquidity shock at any time.

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

[14] The costs of additional dealers could alternatively be modeled as per dealer or per dealer per transaction. Per dealer costs consist of 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.

and transaction prices link valuations across the owner, non-owner, buyer, and seller states.

A client i starting as a non-owner with valuation $\widehat{V}^{no}$ 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 2.1 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 $\widehat{V}^b$. With $V^r(\eta, N)$ capturing exogenously given relation-specific non-trade flows to the client from her dealers, in steady state $\widehat{V}^{no}$ satisfies the Bellman equation linking it to $\widehat{V}^b$ and V^r:

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

which can be solved to yield

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

We are agnostic regarding whether $V^r(\eta, N)$ is increasing or decreasing in the 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 following 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 will 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 $\widehat{V}^o$:

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

yielding

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

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 forcing 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

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

yielding 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 $\widehat{V}^{no}$ are related by

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

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

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

This sequence of events continues in perpetuity and, therefore, we focus on the steady state of the model.

The above valuation equations depend upon the 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 a buyer and seller. Similarly, the dealers' reservation values arise from their transaction cycle. Each dealer acts competitively, i.e., without taking into account the effect of her actions on the actions of other dealers. In addition, we assume that each dealer internalizes only trade-specific value of her relation with each client.[15] When a client who wants to buy the bond contacts her dealer network, each dealer simultaneously starts looking for the bond at rate λ and expects to pay the inter-dealer ask price M^{ask} for the bond. The value to the

[15]We do not require the dealer to internalize the client's non-trade value V^r of the relationship and, therefore, omit it from dealer's valuations. This could happen, for instance, if the dealer is compensated to maximize trading profits and is not incentivized to maximize the enterprise value.

dealer searching for the bond satisfies

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

thus yielding

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

The last term in the expression for U^b captures the expected value of the future business with the same 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 the dealer to take ownership of the bond in the first place as well as the fact that the client is an owner. U^o satisfies the following Bellman equation

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

yielding 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 inter-dealer bid price M^{bid} and earn a markup of $M^{bid} - P^s$. Upon selling the bond the client becomes a non-owner but the dealer still anticipates future business with the same client which happens with frequency λN. This is because this client becomes a buyer with intensity η and approaches her in the future to buy the bond back thus once again generating value U^b to her. As a result, future buy-back business from the same client generates a total of U^{no} in value 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. U^{no} satisfies the following Bellman equation

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

yielding

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

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

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

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

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

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 client and dealer:

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

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 the existing models of OTC markets.

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

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

Substituting (22) into relations (23), (24), (25), and (26) yields the following relations for trade-specific client valuations:

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

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

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

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

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

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

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

The valuations and prices provide ten equations and ten unknowns. Proposition 1 in the Appendix provides the closed-form solutions (A.45) and (A.46) for transaction buy prices P^b and, respectively, sell prices P^s.

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

3.2 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 a 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

$$U^{no}_{\lambda\to\infty} = \frac{1}{N}\frac{\frac{\eta}{r+\eta}}{1-\frac{\kappa}{r+\kappa}\frac{\eta}{r+\eta}}\left[P^b_{\lambda\to\infty} - M^{ask} + \frac{\eta}{r+\eta}(M^{bid} - P^s_{\lambda\to\infty})\right], \tag{29}$$

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

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

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

Expressions (30) show that $U^o_{\lambda\to\infty}$ and $U^{no}_{\lambda\to\infty}$ represent the repeat relation buy discount and sell premium, respectively. Equations (29) show that both $U^o_{\lambda\to\infty}$ and $U^{no}_{\lambda\to\infty}$ 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 $N > 1$ is the improved transaction speed.

The buy and sell prices, given by expressions (A.45) and (A.46) in the Appendix, are nonlinear functions of the model primitives and N, complicating the analysis. However, we can verify that the prices are well-behaved functions of the network size, N, in the large network limit, $N \to \infty$. $N \to \infty$ implies $\frac{\lambda N}{r+\lambda N} \to 1$ and 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 only have a trivial solution $U^s_{N\to\infty} = U^b_{N\to\infty} = 0$, implying that dealers have no rents from future relations with clients. Clients have valuations $V^o_{N\to\infty} - V^b_{N\to\infty} = P^b_{N\to\infty} + K$ and $V^s_{N\to\infty} - V^{no}_{N\to\infty} = P^s_{N\to\infty} - K$ yielding the following expressions for transaction prices:

$$P^b_{N\to\infty} = M^{ask} + \frac{w}{1-w}K, \tag{31}$$
$$P^s_{N\to\infty} = M^{bid} - \frac{w}{1-w}K. \tag{32}$$

Dealers receive no relationship-based rents, but they charge clients a spread $\frac{w}{1-w}2K$ per roundtrip transaction over the inter-dealer spread. The coefficient $\frac{w}{1-w}$ indicates that the non-zero bargaining power enables dealers to extract some value from a client. This is because every dealer charges the same price thus making 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, $P^b - M^{ask}$ and $M^{bid} - P^s$, as the network size becomes larger. This result is similar to 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.

In order to find the optimal network size N^* in the general case, we maximize the total valuation of the "first-time" owner, i.e., 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. Then, we first solve for the optimal network size on the continuous grid $N^* \in \mathbb{R}_{\geq 0}$ and, finally, select the closest integer value to N^* maximizing $\widehat{V}^b(N)$. Proposition 2 given in the 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 our further discussion considers optimization on the continuous grid of network sizes.

Next we provide intuition for why the model yields finite size networks for a general set of model primitives $\{L, K, w, \kappa, \lambda, \eta\}$. Because equation (A.51) 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 the repeated trading, i.e., 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. $\frac{dV^r}{dN} > (<) 0$ implies that $N^{**} < (>) N^*$ thus making N^{**} the lower(upper) bound of N^*. As a result a sufficient condition for existence of a finite N^{**} also applies to N^*. In addition, our analysis helps to separate the trade-specific and the relation-specific network formation mechanisms.

To demonstrate that there exist a region of parameter space where 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 zero at $N = N^{**}$. $\frac{dV^b}{dN}$ consists of three terms in the parentheses. The first positive term is due to the direct effect of the larger network on the speed of execution as faster execution improves buyer's valuation. The second term reflects the marginal effect of the 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 the network size on the transaction buy price. We focus on the case where a buy transaction price, P^b, improves with the network size, $\frac{dP^b}{dN} < 0$, for which we later derive a sufficient condition. Because first and third terms in (33) are positive, $\frac{dV^b}{dN}$ can be zero for some N^{**} only if the second term is negative thus implying that 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 buy price are balanced by the 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 sale price is also decreasing with N:

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

where we have used that dealer's valuations, U^k, $k \in \{b, o, s, no\}$, decline with N. Using relation (26) the sufficiency condition for $\frac{dV^s}{dN} < 0$ can be written as:[16]

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

Therefore, V^s decreases with N when the marginal decline in the sale price exceeds the hazard rate of buying times the marginal increase in V^b sufficiently such 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:

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

where we have used the 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 model's primitives. While this conjecture cannot be proven analytically, we have verified numerically that parameters exist where 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 while the sell price declines with the network

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

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

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

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 marked down sell price.

The effect of N on buy and sell transaction prices is not, however, symmetric due to the time lag between each buy and sell as well as the different effect of relations 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 $SP \equiv P^b - P^s$:

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

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

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

For the 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 $\frac{dV^b}{dN}|_{N=N^{**}} = 0$ thus violating inequality (39). This implies that $\frac{dSP}{dN}$ switches its 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 who 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 client's "buyer" and "seller" states thus creating 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 the 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 the bid-ask spread widens. At this point the client starts trading off the benefit due to increased execution speed against the cost due to wider bid-ask spread when deciding to add another dealer to the network. The optimal size of the network is set when the benefits from transacting with the larger number of dealers equal the costs. When the relation-specific non-traded value, $V^r(\eta, N)$, is added, and if it is increasing with the network size for a given value of η, the client is willing to accept even wider bid-ask spreads in exchange for greater non-trade relation-specific benefits. If the marginal non-traded value improvement from a larger network declines with the size of the network, the overall optimal network size N^* is reached when relation specific benefits together with benefits from transacting faster are balanced by the wider bid-ask spread.

3.3 Empirical predictions

While 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 2.1 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 the 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 channel for price improvement is that active insurers have more repeat trade business with the dealers who internalize the benefits from future business and grant price improvements. A second, complementary explanation is that large insurers have higher bargaining power than small insurers vis-a-vis the dealers. A third explanation is that the trading activity in corporate bonds is correlated with the value of other businesses they conduct with the dealers.

The intuition for the price improvement channel from future business is as follows. It is convenient to discuss the effect of η on prices while keeping the size of the network fixed. When $\eta = 0$ a client owning the bond will never be an owner again after she sells. Therefore, all of the dealers' surplus comes from the price at which the sale occurs and none from future trades. In this case, a dealer's buy reservation value is lower, resulting in a wider bid-ask spread. When η is large, dealers derive significant value from repeated trades, leading to smaller bid-ask spreads. This implies that, conditional on network size, clients with greater trading intensity (higher η) get better execution than clients with smaller trading intensity (lower η). This is consistent with the empirical findings in Table 2 that insurers' trading activity increases in insurer size and that trading costs decline in insurer size in Table 5.

Network size: Increasing the network size increases buyers' valuation, $\widehat{V}^b$, through improved speed of execution, greater inter-temporal dealer competition,[17] and increased relation-specific non-trade value. The optimal network size, N^*, as a function of η balances these

[17]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.

gains against a lower buyer valuation due to the intra-temporal dealer competition leading to the loss of repeat trading business.

The optimal network size is increasing in η for a non-empty parameter range over which the marginal gains from a larger network increase in η, $\frac{\partial^2 \widehat{V}^b}{\partial N \partial \eta}|_{N=N^*} > 0$. Intuitively, dealers' profit from repeated trade improves with the increasing trading intensity and dealers offer better execution. Clients with greater trading intensity benefit from a larger network which improves the speed of execution and may generate larger non-trading relation-specific value. Therefore, consistent with Table 3, the model can generate larger insurers with more frequent trading needs having larger dealer networks.

Ultimately, the model must be structurally estimated to test its ability to fit the cross-sectional distribution of network sizes from Figure 3 as well as the cross-sectional distribution of execution costs as a function of the network size from Table 5.

4 Model Estimation and Policy Analysis

The empirical findings in equation (1) and Tables 5 and IA.4 suggest that in the cross-section trading networks follow a power law with exponential tail and that trading costs are non-monotone in network size. While the model can produce such non-monotonicity, 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 these stylized facts.

4.1 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, Poste-Vinay and Robin (2002) and Eckstein and Van den Berg (2007) for labor search and Hitsch et al. (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), the identification of unobserved parameters relies on key moments in the data. But unlike the aggregate moments in Gavazza (2016), we utilize the granular nature of our data and the heterogeneity in insurers' trading intensities and networks to facilitate our estimation.

Figures 2 and 3 present insurers' heterogeneous trading intensities and network sizes. To estimate the model's distribution of network sizes we infer the distribution of trading intensities, η_i, $i = 1, ..., \mathcal{I}$, across insurers i. Section 2.1 characterizes the distribution of trading activity. If trading shocks occur at Poisson times, the intensity η_i of the shocks can be estimated by the expected number of buy trades per year. This 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 utilize the multiple years of trade data, we perform the estimation separately for each insurer. This yields a cross-sectional distribution of trading intensities, which we index by $p(\eta)$. Section 2.1 and Figure 2 show 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 L and K are constant. The competitive interdealer spread is set to match the average execution costs of insurers by choosing $C = 1$, $r = 0.1$, $M^{ask} = (1 + 0.0018)\frac{C}{2r}$, and $M^{bid} = (1 - 0.0018)\frac{C}{2r}$. 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 they are estimated structurally. To ensure that the estimated parameters L, K, κ, and λ remain in their natural domains, i.e., positive or between zero and one, 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.[18] The evidence in Table 5 shows trading costs depend on insurer size even after controlling for network size. The most straightforward way to accommodate this is to allow dealer bargaining power to vary across insurers, e.g., large active insurers 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:

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

Specification 1 with constant bargaining power sets both $\theta^1_{w^s}$ and $\theta^1_{w^b}$ to zero.

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

The functional form of the non-trade value is Cobb-Douglas:

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

with a constant "technology" $e^{\theta_V^0}$ and elasticities to trading intensity and network size given by θ_V^1 and θ_V^2 respectively. The Cobb-Douglas functional form captures the potential complementarity between trading intensity, measured by η, and network size, while precluding infinitely large non-trade value. We do not impose any specific restrictions on the elasticities beyond that their sum must be less than one. This allows the non-trade value to be an increasing or decreasing function of either N or η or both. By construction, the non-trade value is non-negative. 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 2.2 and Figure 3 show the degree distribution of client-dealer relations follows a mixed power-exponential law. The model also predicts 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 5 as $c(N) = 51+0.32N-6.29\,lnN$. The distance between the empirical and the model-generated relations depends on the parameters Θ. We fit them by minimizing the sum of two probability-weighted distances between the data and model:

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

where ω is the weight put on matching the network sizes compared to trading cost, $p(N)$ 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 sizes and trading cost, $\omega = 0.5$.[19]

Fitting the model based on network sizes (N) rather than insurer trading needs (η) reduces the computational requirements. Inverting the power law distribution for trading intensities 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 the 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) non-smooth. To accommodate this we optimize using simulated annealing with fast decay and

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

fast reannealing.

Overall we estimate a total of 9 model parameters in specification 1 and 11 model parameters in specification 2 by fitting a total of 80 empirical moments. 40 moments are the individual data points from the degree distribution of client-dealer buy relations presented in the left panel of Figure 3. The remaining 40 moments are from the empirical trading costs and network size relation, $c(N) = 51 + 0.32N - 6.29\,lnN$, evaluated on the 40-point grid $N = 1, 2, \ldots, 40$.

4.2 Estimated model parameters

The model can be solved assuming the client maximizes her value as a buyer or as a seller. To simplify exposition we use her value as a buyer. Panel A of Table 6 reports the estimates for the transformed model parameters from (40–42). Panel B provides 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 provides statistics on the 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, 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), a holding period from three weeks to a month. Longer holding periods, 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.

The 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. The 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.

Dealers' bargaining power on the buy and the sell sides in specification 1 is large, 98% and 94%, suggesting that dealers capture most of the trade surplus in this specification. In specification 2, when bargaining power depends on insurers' type, dealers' average bargaining power remains large and fairly symmetric across buys, 86%, and sells, 88%. Dealers' bargaining power on sales is relatively insensitive to insurer trading frequency. In contrast,

Table 6: Estimated model parameters

The table reports the estimated model parameters $\theta = (\theta_L, \theta_K, \theta_\kappa, \theta_\lambda, \theta^j_{w^i}, \theta_V)$ in Panel A. Estimates are from the minimum-distance estimation. Standard errors are computed using the sandwich estimator and reported in parenthesis. 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 non-trade value $V^r = e^{\theta^0_V} \eta^{\theta^1_V} N^{\theta^2_V}$. Panel C provides statistics on the 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 $(*10^4)$	0.06	15.82
κ	15.36	21.19
λ	290.99	145.56
w^b (S.D.)	0.98 (0.00)	0.86 (0.05)
w^s (S.D.)	0.94 (0.00)	0.88 (0.01)
Non-trade value V^r (S.D.)	0.01 (0.00)	0.32 (0.03)
Panel C: Model fit		
Minimum distance $(\times 10^2)$	145.35	2.82
S.D. residuals	1.21	0.17
S.D. residuals network	1.09	0.19
S.D. residuals prices	1.31	0.14

$\theta^1_{w^b} = -1.02$ indicates that dealers' bargaining power when buying declines significantly with insurers trading intensity η.[20] 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 whereby a dealer can more easily locate a seller than a buyer. The heterogeneity in bargaining across insurers and

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

across buy and sell transactions suggests richer modeling of the price setting process may enable deeper understanding of OTC trading.

The estimated non-trade value, V^r, is quite small in specification 1. The estimated non-trade value, V^r, is 25 times larger in specification 2 and it is increasing with 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.5 in the Internet Appendix reports the parameter estimates when we vary the mininum-distance weight ω between 0.1 and 0.9. Across columns, we re-estimate the model for specification 2. The estimates are generally consistent with the ones reported in Table IA.5. The variability of the estimates across columns is small. In column 2 with $\omega = 0.5$, the residuals for network sizes 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 sizes, while the reverse is the case in column 3 with $\omega = 0.9$. We conclude that our results are robust to the relative weight we put on matching network sizes compared to trading cost.

4.3 Sensitivity to model parameters and identification

In order to understand the intuition behind the parameter identification we estimate the sensitivity of model implied moments to baseline parameters using the following exercise. We pick four sets of moments related to: the dealer network sizes N across insurers, the cross-sectional distribution of network sizes $Pr(N)$, trading costs $c(N)$, and the cost premium across different network sizes N relative to the minimum trading costs. For the dealer network we calculate the average network size, minimum, N_{min}, and maximum, N_{max}, network sizes, and the smallest-cost network size, $N_{min-cost}$. We also characterize the network size distribution by 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, $c_{min} = c(N_{min-cost})$, and maximum, $c_{max} = c(N_{min})$.[21] We also calculate the cost for the largest network, $c(N_{max})$. Finally, we calculate the min-N cost premium as $c_{max} - c_{min}$ and the max-N cost premium as $c(N_{max}) - c_{min}$. Trading costs and cost premia are expressed in basis points.

Our baseline is the four sets of moments described above using specification 2 from Table 6. 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, resolve the model, and recalculate all four sets of moments. The procedure is then repeated using a 50% reduction in each parameter. Because

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

we extensively discuss the role of dealer's bargaining power later in this section, we exclude w^b and w^s from the sensitivity analysis.

Table 7 illustrates the sensitivity of the model fit to the parameters in Θ_S. In order to understand the mechanics of the comparative statics exercise reported in Table 7 we need to consider that a client with a trading intensity η selects her equilibrium network of size N by balancing gains from increased transaction speed against losses from wider spreads. Changing any parameter from the set Θ_S directly affects either the transaction speed or the spread or both thus inducing this client to change the size of her network. Because clients with larger trading needs and, therefore, larger networks, are more sensitive to changes in the transaction speed and spread, their networks are affected more by the parameter changes. Clients with low trading needs are less sensitive to such changes and are much less likely to change their networks.

Reducing the liquidity shock parameter L has two direct effects. First, it reduces seller's need for immediacy thus decreasing the benefits of higher transaction speed provided by larger networks. As a result, some of the network distribution mass shifts towards smaller networks. Table 7, 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 sale 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 2 dealers, changes in the cost premium are mostly due 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 7 shows that c_{min} declines by 2% while c_{max} remains unchanged thus leading to an increase of 1% (6%) in $c_{max} - c_{min}$. $c(N_{max})$ remains unchanged (decreases by 1%) thus leading to an increase of 5% (15%) in $c(N_{max}) - c_{min}$.

Reducing the fixed costs for 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 induce even clients with very low trading needs to increase the size of their dealer network thus making $N_{min} > 1$. Indeed, Table 7, columns 3 and 4, confirm this 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 1 to 2 and the maximum network size more than doubles. Overall, reducing K leads to significant changes in the network size distribution as a sizable fraction of its mass shifts towards larger network sizes.

Table 7: Sensitivity to model parameters and identification

The 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 sizes N across insurers, the cross-sectional distribution of network sizes $\Pr(N)$, trading costs $c(N)$, and the cost premium across different network sizes N relative to the 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	
		-10%	-50%	-10%	-50%	-10%	-50%	-10%	-50%	-10%	-50%
	Base	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(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%*
$N_{min-cost}$	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%*	*2339%*	*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 independently of their trading needs. Shifting the mass of the network size distribution towards larger networks reduces the spread for networks smaller than $N_{min-cost}$, while it increases the spread for networks greater than $N_{min-cost}$. Because a substantial mass of the network size distribution shifts to the interval $N > N_{min-cost}$, the overall effect of decreasing K on the trading costs $c(N)$ is modest. Indeed, Table 7 (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 smallest networks, c_{max} decreases and, as a result, the cost curve flattens as $c_{max} - c_{min}$ falls. Conversely, $c(N_{max}) - c_{min}$ is non-monotonic in K. Table 7 shows that the cost premium declines 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 thus decreasing the speed of transition between client's owner and seller states. Dealers widen the spread across all types of clients to compensate for the lost repeat trading. Clients with very high trading needs, 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 to them outweighs extra trading costs. Clients with medium trading needs, with network sizes in the range of $N_{min-cost}$ 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, while clients with $N = N_{min-cost}$ reduce their networks from 21 to 18 (17) dealers. Overall, the 50% reduction in κ leads to quite drastic changes in the right tail of the network size distribution. Column 6 of Table 7 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%) while 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 types of clients to compensate for the lost repeat trading, while clients either increase or decrease their network size depending on whether the benefit from transacting faster outweighs higher trading costs for them. However, in the case of λ, precisely which clients end up increasing or decreasing their network size depends on the magnitude of the reduction in λ. Column 7 of Table 7 shows that when λ is reduced by 10%, clients with very high trading needs increase their networks from 44 to 45 dealers. 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 $N_{min-cost}$ 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 quite drastic when λ is halved. Column 8 of Table 7 shows that $\Pr(N > 20)$ increases by 149% in this case. For clients with medium to high trading needs, $N \geq N_{min-cost}$, benefits from faster execution outweigh losses from higher trading costs and, as a result, $N_{min-cost}$ increases from 21 to 23 dealers while N_{max} increases from 44 to 48 dealers. Consequently, 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 lowers the cost premium at N_{min}, while the reduction in κ increases the cost premium at N_{min}, we can identify λ and κ separately.

Reducing V^r reduces the non-trade benefits from having larger network size for clients with both low and high trading needs, with latter being affected more. Clients respond by scaling down their networks. When V^r is reduced by 10% (50%), by reducing 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 non-trade 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%), while the cost premium at N_{max}, $c(N_{max}) - c_{min}$, declines by 13% (65%).

Overall, Table 7 clearly demonstrates that there exists a substantial variation in how each structural parameter affects the size distribution and transaction costs thus allowing for the parameter identification.

4.4 Model fit and discussion

Figures 5 and 6 visually examine the quality of the model's fit. They each provide four plots illustrating the fits for their respective specification. The top left graph plots 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 graph plots the degree distribution for insurer-dealer relations 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 graph plots the degree distribution for insurer-dealer relations 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-degree line measures model fit.

The specification with uniform dealer bargaining power fits the network distribution between 2 and 19 dealers. In the bottom left graph in Figure 5 the blue circles are on the 45-degree line for almost all values of N less than 20. However, the goodness of fit deteriorates for larger networks and the model does not generate networks larger than 22 dealers.

The bottom right corner graph in Figure 5 plots the empirical relation between trading

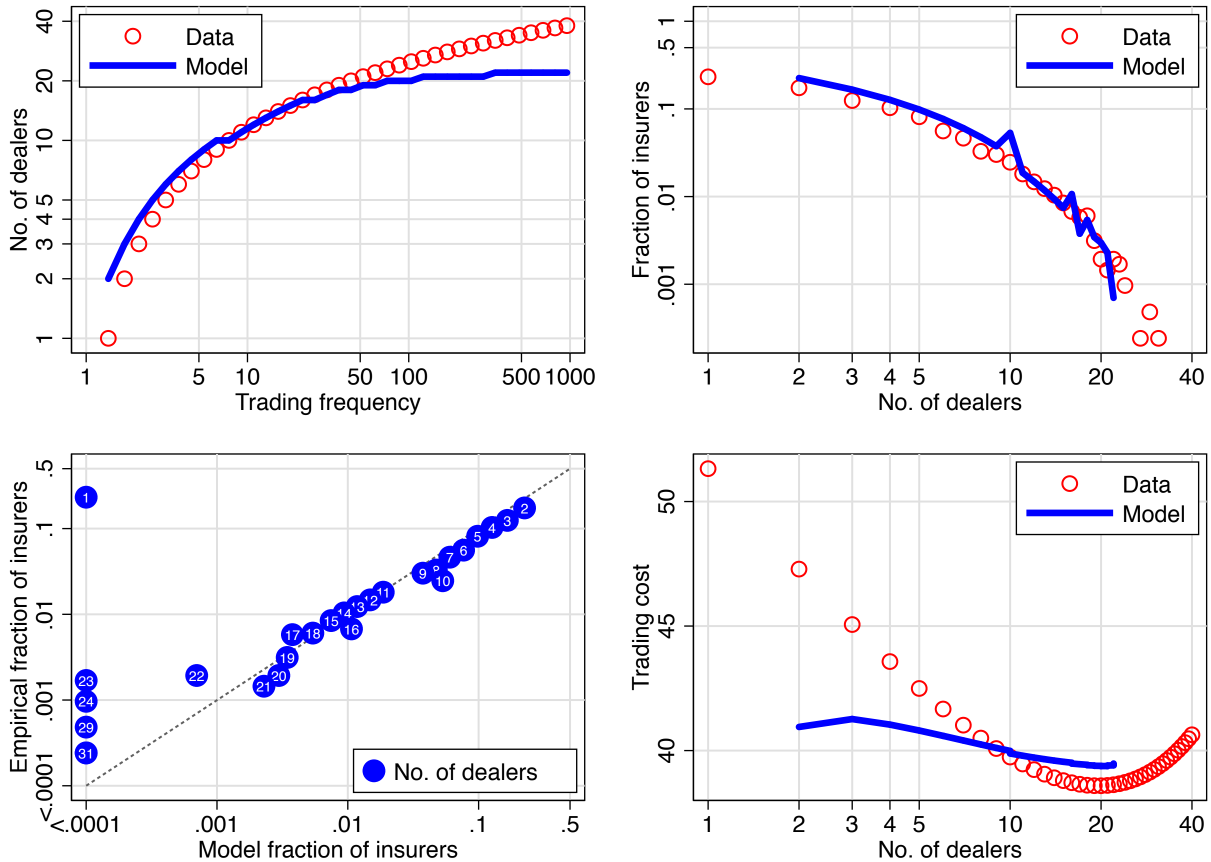

Figure 5: Model fit with constant bargaining power w^b, w^s

The figure shows the model fit for the base specification of constant dealers' bargaining power w. The top left plot (log-log scale) shows the network size of insurer-dealer relations 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 plot (log-log scale) shows 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 plot (log-log scale) compares the empirical probability of a given network size to the model-implied probability. The bottom right plot (linear-log scale) shows the bid-ask spread in the data as a function of the insurers' network size in the data (circles) compared to the model-implied bid-ask spread under the estimated parameters (solid line).

costs and network size from column (4) of Table 5 (circles), $c(N^*) = 51 + .32N^* - 6.29\, lnN^*$, against its model-implied counterpart (dashed line).[22] The distance between the two lines is a measure of model fit, showing the parameter estimates from specification 1 do not well describe the relationship between network sizes and trading costs in the data. The model's

[22]Because values reported in Table 5 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.

relation is too weak as the line is too flat. In addition, the non-monotonicity in the relation is barely visible. This suggests that uniform bargaining power limits the variation in the benefits and costs of having a larger network.

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

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

The top left graph in Figure 5 is the inverse of $\frac{d\eta}{dN^*}$. The top right graph in Figure 5 weighs this 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 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 non-decreasing function of η, $\frac{d\eta}{dN^*} \geq 0$, so the second term in (44) is always non-positive. The opposing signs of the two terms explain how the model produces the non-monotonic relationship between network size and trading costs in the data. To fit the trading cost-network size relationship in Table 5 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 in neither initially negative enough nor does it increase noticeably enough at $N = 20$ to fit the data.

The log-log plot in the top left 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 dealer's bargaining power $w^{b,s}$ close to one, but 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. On the other hand, 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, clients with smallest networks pay wider spreads, with the widest spread paid by clients with a single dealer, and the spread monotonically narrows with the network size until N = 20 and then widens. The 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, to match the widest spreads, and to make the spread to 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 single 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 for 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 data. The graphs show network sizes greater than 22 and the U-shape in the bid-ask spread starting at 20 dealers. In this specification dealers' bargaining power is greater than 0.8 for insurers with trading intensity less than five.[23] 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 network sizes for insurers with large η as the value 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 larger bargaining power, lowering their trading costs. This strengthens the relationship between trading costs and network sizes.

Two things of note in Figures 5 and 6 are the "stair-like" appearance of the top left model-implied graph and the "spiky" appearance of the top right model-implied graph. Both graphs also show that larger changes in η are required to add an additional dealer to the optimal network as η increases. When we numerically solve for the optimal network size, N^*, we use a grid of integer values for N, while the grid for η has a much more refined spacing. Therefore, as η increases on the grid, the optimal network size remains constant until η generates sufficient enough benefits to increase N^* by one dealer at which point the plot jumps up by 1. Since $\frac{dN^*}{d\eta}$ declines with η, the length of flat intervals of N^* increases with η. These "stairs" are completely smoothed out when we allow for a grid of non-integer values for N. Using log-log scale also smooths these "stairs" but does not completely eliminate them. The top right plot in both Figures is a nonlinear transformation of the top left plot yielding spikes at values of $\log(\eta)$ where $\log(N^*)$ becomes flat.

To further examine the ability of the model to fit the data we compute the model-implied

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

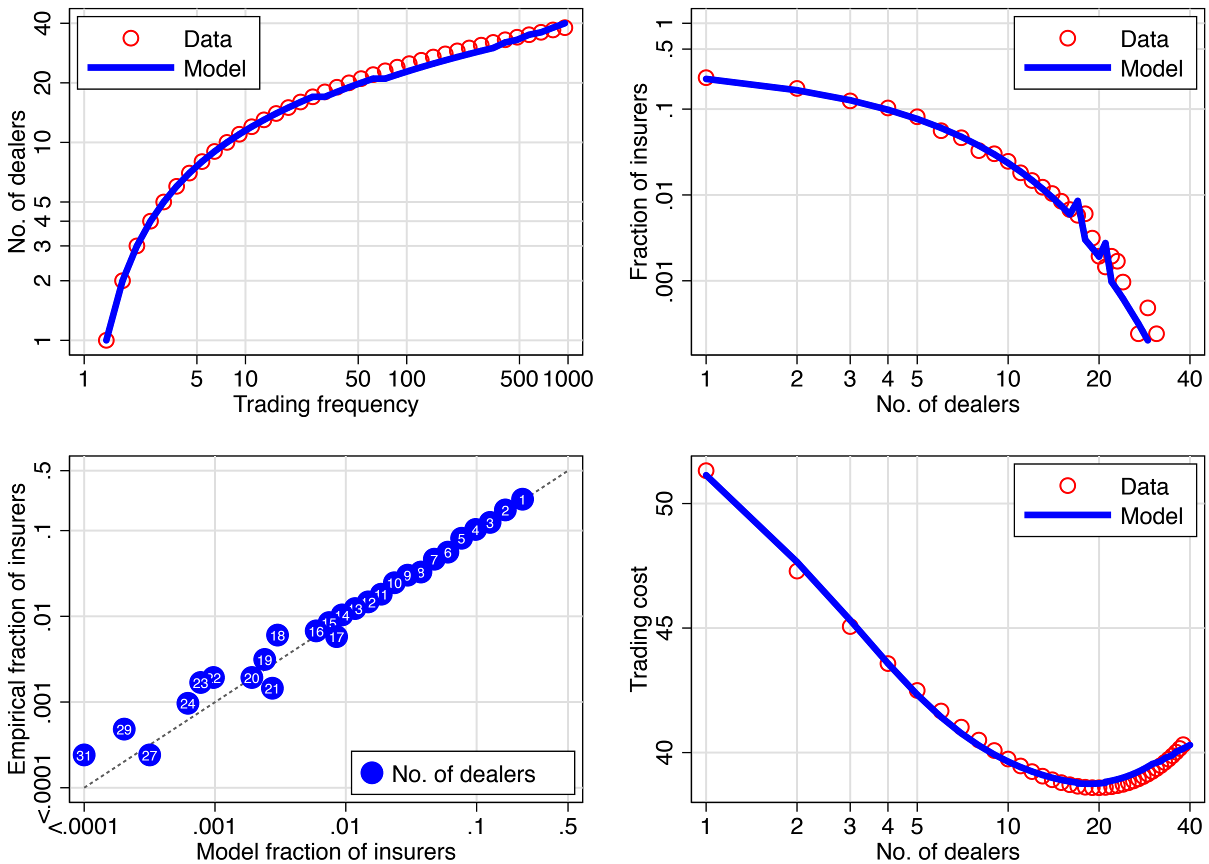

Figure 6: Model fit with client-specific bargaining power $w^b(\eta), w^s(\eta)$

The figure shows the model fit when dealers' bargaining power is a function of trading intensity, $w^b(\eta)$ and $w^s(\eta)$. Please see caption of Figure 5 for further details.

switching probabilities, $p(\text{No. of dealers in } t+1 | \text{No. of dealers in } t)$, between different network sizes and compare them to their empirical counter-parties reported in Table 4.[24] To do this we simulate panels of realized network sizes N_{it} for insurer i in year t using the actual realized number of buy transactions by the insurer i in that year. To do so, for each insurer in the sample we use the trading intensity η_i, as computed earlier in this section, and the estimated model parameters from specification 2 of Table 6, as inputs to determine the insurer's optimal network size N_i^*. We then simulate for each trade in the actual data set which of the N_i^* dealers in the insurer's network is matched to a given trade according to model's random search across dealers in the network. We record the identity of the dealer

[24]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.

handling each such trade and use these identities to compute the realized dealer network, N_{it}, as the unique number of dealers from the client'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 very stable from simulation to simulation, so that 10 simulation rounds suffice. Table 8 shows that the model can match the persistence in networks, as the model implied transition matrix is close to the actual.

Table 8: Model-implied persistence in insurers' trading network: Evidence from simulation

The table reports model-implied switching probabilities, p(No. of dealers in $t+1$|No. of dealers in t), for using a network size conditional on the insurer's past behavior for each year. Simulation uses estimated model parameters from Table 6. The 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*	0.03	*0.31*	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*	0.83

To summarize, the theoretical model with the dealers' bargaining power declining with a trading intensity η fits the data well. In particular, it is important to have the dealer's bargaining power on the ask side to depend on η. This leads to small inactive insurers having very weak outside options both when they buy and sell. By contrast, large active insurers have strong outside options when they buy and weak bargaining power when they sell.

4.5 Impact of policy changes

The parameter estimates can help quantify the value of both the repeat client-dealer trading and non-trade value V^r. 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 compute counterfactual networks and bid-ask prices under the assumption that the dealers give only partial price improvement for the future repeat trading from the same client. In this case, the search and bargaining proceeds in the same way, but the dealers in the network are implicitly re-chosen with probability $\frac{1}{2}$ after every trade. This is counterfactual scenario 1. In 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 two.

To capture the impact of V^r on networks and transaction prices we compute counterfactual networks and bid-ask prices under the assumption that clients fully benefit from the repeat trading but non-trade relationships with the 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 non-trade value is zero. Setting the non-trade value to zero also directly reduces client valuations. An interpretation of the direct effect on clients' valuation is that unbundling of trading from research and other activities means that the dealers now no longer provide the non-trade benefits to clients or dealers charge clients for the services that provide the benefits. Our expectation is that clients will lose (some of) these non-trade benefits. We focus on the indirect impact on client valuations by adding the non-trade value back in when we recompute insurer value. An alternative policy experiment would be to assume clients lose a fraction or all of the direct non-trade benefits, in which case the drop in value would be substantially larger.

The upper left plot in Figure 7 corresponds to the top left plot in Figure 6 and shows the optimal network sizes, N^*. Because the network sizes change under the counterfactuals, the upper 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 estimated parameters from specification 2 in Table 6. Each graph shows the results from specification 2 (solid line), counterfactual 1 (crosses), and counterfactual 2 (circles).

Figure 7 highlights the fundamental trade-offs 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 lower trading intensities have smaller networks. Consequently, the widening in bid-ask spread is slightly larger for clients with lower trading intensities than for clients who trade very frequently. Both large (high η) and small (low η) clients respond to wider bid-ask spread by reducing the size of their networks. The magnitude of this effect is much smaller for clients with low η as they already have small networks.

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

The U-shape in the bid-ask spread is not present without the non-trade relationship

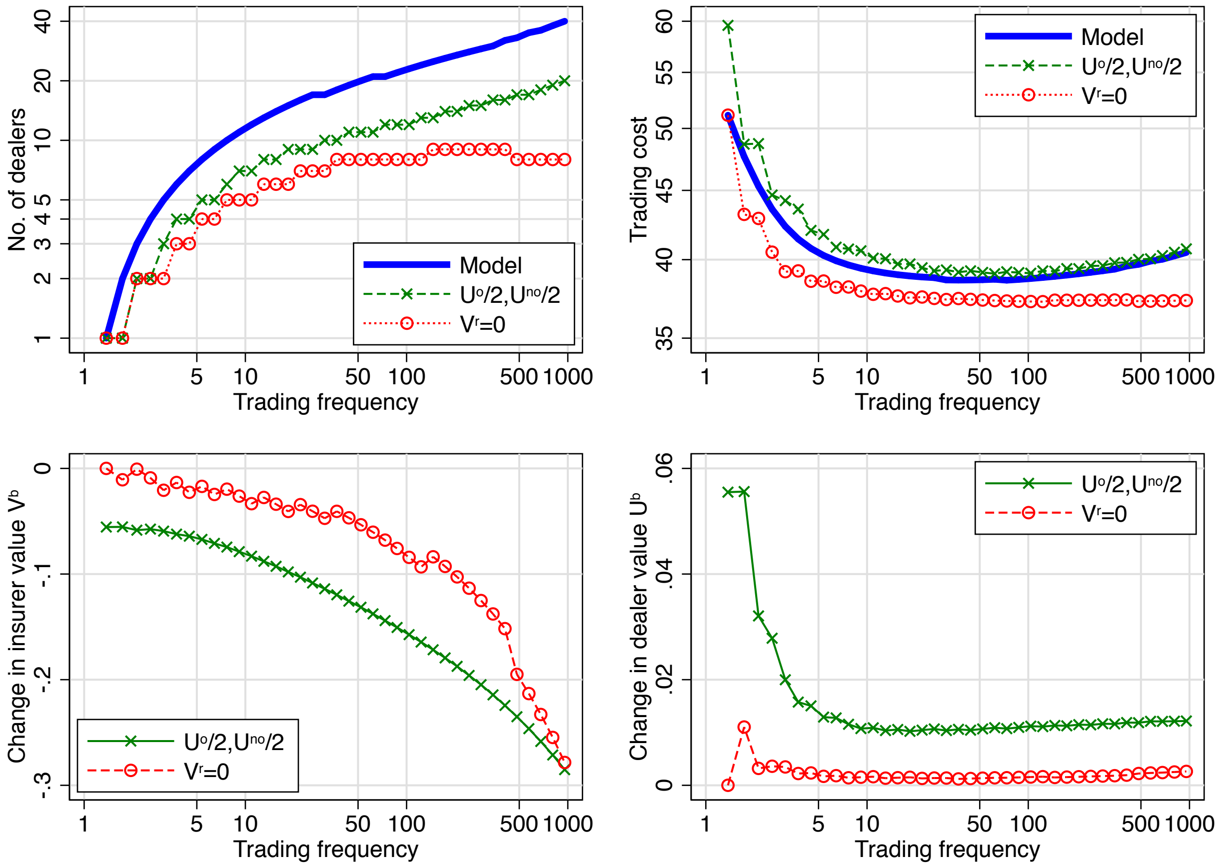

Figure 7: Counterfactual analysis

The left plot shows the client-dealer network size as a function of the clients' trading frequency, η. The right plot shows the bid-ask spread as a function of η. The bottom left plot shows the change in clients' value from trading, V^b, as a function of the clients' trading frequency, η. The bottom right plot shows the change in dealers' value from trading, U^b, as a function of the clients' trading frequency, η. Each plot shows the corresponding variable in under the estimated parameters from specification 2 in Table 6 (solid line) compared to the counterfactuals that would arise if: (i) dealers have suffered a permanent loss of 50% of the repeat trade business (crosses); (ii) the non-trade relationship value is zero (circles). We use log-log scale.

value V^r. This highlights the role V^r increasing in N^* plays in the non-monotonicity of relationship between network size and the bid-ask spread. As V^r increases with N^*, clients have an incentive to choose larger networks than in the absence of the non-trade value because larger transaction costs are offset by the non-trade value. Because V^r increases with η as well, this incentive is especially strong for clients with very large trading intensities and large networks, 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 non-trading 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, TruMid, and others. These initiatives likely impact the transaction costs incurred by clients, but very little empirical and/or theoretical evidence exists to quantify these effects. Our counterfactuals provide quantitative insights regarding both: Unbundling trading and non-trading dealer services (counterfactual 2) and moving the bond trading from OTC to more centralized electronic platforms where the probability of repeat trading with the same dealer is reduced (counterfactual 1).

The upper 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 who trade more frequently and, therefore, have larger networks will see much smaller increase in the transaction costs than clients who trade less frequently and tend to trade repeatedly with 1 to 5 dealers. For instance, clients with 100 annual trades have on average 11 dealers in their network and will see less than 1 bps increase in transaction costs 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 5 dealers and will on average pay 3 to 7 bps more per bond per transaction.

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

The bottom left plot of Figure 7 shows that the proposed regulations would decrease the value from trading, V^b, for all types of clients. Larger, more active clients would be affected differently than smaller, less active clients. In the model, higher-η clients incur the largest drop in V^b. These clients would also reduce the size of their trading networks the most. If unbundling only impacts network size and clients still receive non-trade relationship value from their dealers, the clients' welfare loss is significant but does not turn the value from trading negative. However, clients with lower-η already extract a small 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 non-trade benefits, they may exit the market all together. Thus, the proposed regulations can potentially reduce trading activity and decrease clients' welfare. Especially in the case that unbundling, via MiFID II or other regulation, eliminates both trade and non-trade relationship value the decline in client welfare is substantial.

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

5 Conclusion

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

We use comprehensive regulatory corporate bond trading data for all U.S. insurance companies to study how investors choose the size of their trading networks and how this impacts the transaction prices they receive in the current decentralized OTC market setting. We document that 30% of insurers trade with a single dealer. There exist few insurers with networks of up to 40 dealers. The cross-sectional distribution in trading activity is a power law, while the network sizes follow a hybrid of power law with exponential tail. Trading costs decline with network size up to 20 dealers and increase for larger network sizes.

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

We use the structurally estimated model parameters to counter factually assess the impact of regulations that erode client-dealer relations. We find that clients will incur higher transaction costs when repeat trading is reduced through, e.g., anonymous trading. Unbundling trade and non-trade services will reduce optimal network sizes and decrease transaction costs for all insurers, except for the least active ones. If clients cease to receive non-trading benefits from dealers, clients that trade infrequently may cease trading, which can reduce market liquidity. Unless non-trading services can be provided in a significantly more efficient way in an unbundled setting, unbundling decreases welfare.

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," Federal Reserve Bank of New York Staff Report.

Atkeson, Andrew, Andrea Eisfeld, and Pierre-Olivier Weill, 2015, "Entry and exit in OTC Derivatives Markets," *Econometrica* 83, 2231-2292.

Babus, Ana and Tai-Wei Hu, 2016, "Endogenous Intermediation in Over-the-Counter Markets," *Journal of Financial Economics*, Forthcoming.

Babus, Ana, 2016, "The Formation of Financial Networks," *RAND Journal of Economics* 47, 239-272.

Babus, Ana and Peter Kondor, 2013, "Trading and Information Diffusion in Over-The-Counter Markets," Working Paper, Federal Reserve Bank of Chicago.

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, 2016, "Capital Commitment and Illiquidity in Corporate Bonds," working paper.

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.

Chaney, Thomas, 2018, "The gravity equation in international trade: An explanation," *Journal of Political Economy* 126, 150-77.

Chang, Briana and Shengxing Zhang, 2015, "Endogenous market making and network formation," Mimeo.

Chodorow-Reich, Gabriel, Andra Ghent, and Valentin Haddad, 2016, "Asset Insulators," Mimeo.

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," Paris School of Economics Working Paper No. 2014-8.

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, 2015, "The Value of Trading Relationships in Turbulent Times ," Mimeo.

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

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

Edwards, Amy, Lawrence Harris, and Michael Piwowar, 2005, "Corporate bond markets transparency and transaction costs," *Journal of Finance* 62, 1421-1451.

Eckstein, Zvi, and Gerard J. Van Den Berg, 2007, "Empirical Labor Search: A Survey," *Journal of Econometrics*, 136, 531564.

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

Feldhutter, 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, 2014, "Adverse Selection and Intermediation Chains," Working Paper, University of Pennsylvania.

Gofman, Michael, 2011, "A Network-Based Analysis of Over-the-Counter Markets," Working Paper.

Green, Richard C., Burton Hollifield, and Norman Schürhoff, 2007, "Financial Intermediation and the Costs of Trading in an Opaque Market," *Review of Financial Studies* 20(2), 275-314.

Green, Richard C., Dan Li, and Norman Schürhoff, 2010, "Price Discovery in Illiquid Markets: Do Financial Asset Prices Rise Faster Than They Fall?," *The Journal of Finance* 65(5), 1669-1702.

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

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.

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," Mimeo.

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, 2015, "Bid-Ask Spreads and the Pricing of Securitizations: 144a vs. Registered Securitizations," Working Paper, Carnegie Mellon University.

Hugonnier, Julien, Benjamin Lester, and Pierre-Olivier Weill, 2015, "Heterogeneity in Decentralized Asset Markets," Working Paper, Swiss Finance Institute.

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, 2018, "Dealer Networks," *Journal of Finance*, Forthcoming.

Moen, Espen R., 1997, "Competitive Search Equilibrium," *Journal of Political Economy* 105, 385-411.

Neklyudov, Artem, 2014, "Bid-Ask Spreads and the Over-the-Counter Interdealer Markets: Core and Peripheral Dealers," Working Paper, Swiss Finance Institute.

Neklyudov, Artem and Batchimeg Sambalaibat, 2016, "Endogenous Specialization and Dealer Networks," Working Paper, Swiss Finance Institute.

O'Hara, Maureen, Yihui Wang, and Xing Zhou, 2015, "The Best Execution of Corporate Bonds," Working paper, Cornell University.

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.

Üslü, Semih, 2015, "Pricing and Liquidity in Decentralized Asset Markets," Working Paper, UCLA.

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.

Appendix

Data Filters

Our insurance company Schedule D raw filing data are downloaded from SNL financial, which receives the data from the insurance company regulator, NAIC (National Association of Insurance Companies). Parts 3 and 4 of Schedule D filed with the NAIC contains 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, while 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, 4, and 5 of Schedule D to obtain a comprehensive set of corporate bond transactions by all insurance companies regulated by NAIC.

Table A.1: Data filters

Filter	Full sample	Corp. 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, 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, drop trades with a dealer that trades less than 10 times in total over the sample period	6.1	2.9
5. Drop if not fixed coupon (based on eMaxx data), drop if outstanding amount information is in neither eMaxx nor 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 securities that are not exchangeable, preferred, convertible, issued by domestic issuer, taxable muni, missing the offering date, offering amount, or maturity, and offering amount is not less than 100K	1.0	1.0

There are a couple of differences between the raw filing data used in this paper and the "Time and Sales Data" from Mergent FISD available on WRDS, widely used in previous studies, such as Bessembinder et al. (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, our experience is that annual reports appear to be missing many trades that show up in the quarterly report. For completeness, we compile the trade data using quarterly reports from Q1-Q3 and all trades that happened in the 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 the field to account for different spellings of broker-dealer names.

We apply various FISD-based data filters based on Ellul et al. (2011) to eliminate outliers and establish a corporate bond universe with complete data. The data filters are in Table A.1 which summarizes the number of observations that is 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 (US Corporate Debentures) bonds, CZ (Corporate Zero) bonds, and all government bond (including municipal bonds) based on the reported industry group. Finally, we also drop a bond if any of the following fields is missing: offering date, offering amount, and maturity. We restrict our sample to bonds with the offering amount greater than \$10 million, as issues smaller than this amount are very illiquid and hence are rarely traded. Ellul et al. (2011) have used \$50 million which we find restrictive for our purpose. We winsorize the cash-to-asset ratio at 99 percent to remove extreme values.

Proofs

PROPOSITION 1: *Bid and ask prices are*

$$\begin{aligned} P^b &= \left[\frac{C(1-L)}{\lambda N}\Psi_1 + \frac{C}{r+\kappa}\Psi_2 + K\Pi_1(\kappa,\eta)\right] w \qquad \text{(A.45)}\\ &+ \left[M^{ask}\Pi_2 + M^{bid}\Pi_3 \frac{\kappa}{r+\kappa}\right](1-w), \\ P^s &= \left[\frac{C(1-L)}{\lambda N}\Psi_3 - \frac{C}{r+\kappa}\Psi_4 - K\Pi_1(\eta,\kappa)\right] w \qquad \text{(A.46)}\\ &+ \left[M^{bid}\Pi_2 + M^{ask}\Pi_3 \frac{\eta}{r+\eta}\right](1-w). \end{aligned}$$

where the coefficients Ψ_1, Ψ_2, Ψ_3, Ψ_4, $\Pi_1(x,y)$, Π_2, *and* Π_3 *are given below in the proof.*

Proof: Define $\Lambda = \lambda N$, $\widetilde{\lambda} = \frac{\lambda}{r+\Lambda}$, $\widetilde{\Lambda} = \frac{\lambda N}{r+\lambda N}$, $\widetilde{\eta} = \frac{\eta}{r+\eta}$, and $\widetilde{\kappa} = \frac{\kappa}{r+\kappa}$. We can rewrite both the client's and the dealer's valuations as

$$\begin{aligned} V^b &= \frac{C}{r+\kappa}\widetilde{\Lambda} - (P^b+k)\widetilde{\Lambda} + V^s\widetilde{\kappa}\widetilde{\Lambda} = \frac{\widetilde{\Lambda}}{1-\widetilde{\kappa}\widetilde{\eta}(\widetilde{\Lambda})^2}\left[\frac{C(1-l)}{r+\Lambda}\widetilde{\kappa} + \frac{C}{r+\kappa} - (P^b+K) + (P^s-k)\widetilde{\kappa}\widetilde{\Lambda}\right], \\ V^s &= \frac{C(1-l)}{r+\Lambda} + (P^s-k)\widetilde{\Lambda} + V^b\widetilde{\eta}\widetilde{\Lambda} = \frac{\widetilde{\Lambda}}{1-\widetilde{\kappa}\widetilde{\eta}(\widetilde{\Lambda})^2}\left[\frac{C(1-l)}{\Lambda} + \frac{C}{r+\kappa}\widetilde{\eta}\widetilde{\Lambda} + (P^s-K) - (P^b+K)\widetilde{\eta}\widetilde{\Lambda}\right], \end{aligned}$$

and

$$
\begin{aligned}
U^s &= (M^{bid} - P^s)\widetilde{\lambda} + U^b\widetilde{\eta}\widetilde{\Lambda} = \frac{\widetilde{\lambda}}{1 - \widetilde{\kappa}\widetilde{\eta}(\widetilde{\Lambda})^2}\left[(M^{bid} - P^s) + (P^b - M^{ask})\widetilde{\eta}\widetilde{\Lambda}\right], \\
U^b &= (P^b - M^{ask})\widetilde{\lambda} + U^s\widetilde{\kappa}\widetilde{\Lambda} = \frac{\widetilde{\lambda}}{1 - \widetilde{\kappa}\widetilde{\eta}(\widetilde{\Lambda})^2}\left[(P^b - M^{ask}) + (M^{bid} - P^s)\widetilde{\kappa}\widetilde{\Lambda}\right].
\end{aligned}
$$

After some algebra one obtains

$$
\begin{aligned}
V^o - V^b &= \frac{C}{r+\kappa} + V^s\widetilde{\kappa} - V^b \\
&= \frac{C}{r+\kappa} + \frac{\widetilde{\Lambda}}{1 - \widetilde{\kappa}\widetilde{\eta}(\widetilde{\Lambda})^2}\left[\frac{C(1-l)}{\Lambda}\widetilde{\kappa} - \frac{C(1-l)}{r+\Lambda}\widetilde{\kappa} - \frac{C}{r+\kappa}(1 - \widetilde{\kappa}\widetilde{\eta}\widetilde{\Lambda}) + \right. \\
&\quad \left. + (P^b + K)(1 - \widetilde{\kappa}\widetilde{\eta}\widetilde{\Lambda}) + (P^s - K)\widetilde{\kappa}(1 - \widetilde{\Lambda})\right], \\
V^s - V^{no} &= V^s - V^b\widetilde{\eta} \\
&= \frac{\widetilde{\Lambda}}{1 - \widetilde{\kappa}\widetilde{\eta}(\widetilde{\Lambda})^2}\left[\frac{C(1-l)}{\Lambda}(1 - \widetilde{\kappa}\widetilde{\eta}\widetilde{\Lambda}) - \frac{C}{r+\kappa}\widetilde{\eta}(1 - \widetilde{\Lambda}) + \right. \\
&\quad \left. + (P^b + K)\widetilde{\eta}(1 - \widetilde{\Lambda}) + (P^s - K)(1 - \widetilde{\kappa}\widetilde{\eta}\widetilde{\Lambda})\right].
\end{aligned}
$$

Substituting these expressions into (20) and (21) yields

$$
\begin{aligned}
P^b &= \frac{C}{r+\kappa}w \\
&\quad + \frac{\widetilde{\Lambda}}{1 - \widetilde{\kappa}\widetilde{\eta}(\widetilde{\Lambda})^2}\left[\frac{C(1-l)}{\Lambda}\widetilde{\kappa} - \frac{C(1-l)}{r+\Lambda}\widetilde{\kappa} - \frac{C}{r+\kappa}(1 - \widetilde{\kappa}\widetilde{\eta}\widetilde{\Lambda}) + (P^b + K)(1 - \widetilde{\kappa}\widetilde{\eta}\widetilde{\Lambda}) + (P^s - K)\widetilde{\kappa}(1 - \widetilde{\Lambda})\right] w \\
&\quad + M^{ask}(1 - w) - \frac{\widetilde{\lambda}}{1 - \widetilde{\kappa}\widetilde{\eta}(\widetilde{\Lambda})^2}\left[(M^{bid} - P^s) + (P^b - M^{ask})\widetilde{\eta}\widetilde{\Lambda}\right]\widetilde{\kappa}(1 - w), \\
P^s &= \frac{\widetilde{\Lambda}}{1 - \widetilde{\kappa}\widetilde{\eta}(\widetilde{\Lambda})^2}\left[\frac{C(1-l)}{\Lambda}(1 - \widetilde{\kappa}\widetilde{\eta}\widetilde{\Lambda}) - \frac{C}{r+\kappa}\widetilde{\eta}(1 - \widetilde{\Lambda}) + (P^b + K)\widetilde{\eta}(1 - \widetilde{\Lambda}) + (P^s - K)(1 - \widetilde{\kappa}\widetilde{\eta}\widetilde{\Lambda})\right] w \\
&\quad + M^{bid}(1 - w) + \frac{\widetilde{\lambda}}{1 - \widetilde{\kappa}\widetilde{\eta}(\widetilde{\Lambda})^2}\left[(P^b - M^{ask}) + (M^{bid} - P^s)\widetilde{\kappa}\widetilde{\Lambda}\right]\widetilde{\eta}(1 - w).
\end{aligned}
$$

These expressions can be rewritten as

$$
\begin{aligned}
P^b A_1(w) &= \left[\frac{C(1-l)}{\Lambda}\widetilde{\kappa}\widetilde{\Lambda} + \frac{C}{r+\kappa}\right](1 - \widetilde{\Lambda})w + KA_2(\widetilde{\kappa}, w) \\
&\quad + M^{ask}A_3(w) - M^{bid}\widetilde{\kappa}\widetilde{\lambda}(1 - w) + P^s\widetilde{\kappa}A_4(w) \\
P^s A_1(w) &= \left[\frac{C(1-l)}{\Lambda}(1 - \widetilde{\kappa}\widetilde{\eta}\widetilde{\Lambda}) - \frac{C}{r+\kappa}\widetilde{\eta}(1 - \widetilde{\Lambda})\right]\widetilde{\Lambda}w - KA_2(\widetilde{\eta}, w) \\
&\quad - M^{ask}\widetilde{\eta}\widetilde{\lambda}(1 - w) + M^{bid}A_3(w) + P^b\widetilde{\eta}A_4(w),
\end{aligned}
\tag{A.47}
$$

where we have defined

$$
\begin{aligned}
A_1(w) &\equiv 1 - \widetilde{\Lambda}w - \widetilde{\kappa}\widetilde{\eta}(\widetilde{\Lambda})^2(1 - w) + \widetilde{\kappa}\widetilde{\eta}\widetilde{\Lambda}\widetilde{\lambda}(1 - w), \\
A_2(x, w) &\equiv [1 - \widetilde{\kappa}\widetilde{\eta}\widetilde{\Lambda} - x(1 - \widetilde{\Lambda})]\widetilde{\Lambda}w, \\
A_3(w) &\equiv [1 - \widetilde{\kappa}\widetilde{\eta}(\widetilde{\Lambda})^2 + \widetilde{\kappa}\widetilde{\eta}\widetilde{\Lambda}\widetilde{\lambda}](1 - w), \\
A_4(w) &\equiv (1 - \widetilde{\Lambda})\widetilde{\Lambda}w + \widetilde{\lambda}(1 - w), \\
A_5(w) &\equiv A_1(w) - \widetilde{\kappa}\widetilde{\eta}\frac{A_4(w)^2}{A_1(w)}.
\end{aligned}
\tag{A.48}
$$

The system of equations (A.47) can be solved to yield expressions (A.45) and (A.46), where we have defined

$$\Psi_1 \equiv \frac{\widetilde{\kappa}\widetilde{\Lambda}}{A_5(w)}\left[1-\widetilde{\Lambda}+(1-\widetilde{\kappa}\widetilde{\eta}\widetilde{\Lambda})\frac{A_4(w)}{A_1(w)}\right], \tag{A.49}$$

$$\Psi_2 \equiv \frac{(1-\widetilde{\Lambda})}{A_5(w)}\left[1-\widetilde{\kappa}\widetilde{\eta}\frac{A_4(w)}{A_1(w)}\widetilde{\Lambda}\right],$$

$$\Psi_3 \equiv \frac{\widetilde{\Lambda}}{A_5(w)}\left[1-\widetilde{\kappa}\widetilde{\eta}\widetilde{\Lambda}+\widetilde{\kappa}\widetilde{\eta}(1-\widetilde{\Lambda})\frac{A_4(w)}{A_1(w)}\right],$$

$$\Psi_4 \equiv \frac{\widetilde{\eta}(1-\widetilde{\Lambda})}{A_5(w)}\left[\widetilde{\Lambda}-\frac{A_4(w)}{A_1(w)}\right],$$

$$\Pi_1(x,y) \equiv \frac{1}{A_5(w)}\left[A_2(x)-xA_2(y)\frac{A_4(w)}{A_1(w)}\right],$$

$$\Pi_2 \equiv \frac{1}{A_5(w)}\left[A_3-\widetilde{\kappa}\widetilde{\eta}\frac{A_4(w)}{A_1(w)}\widetilde{\lambda}\right],$$

$$\Pi_3 \equiv \frac{1}{A_5(w)}\left[A_3\frac{A_4(w)}{A_1(w)}-\widetilde{\lambda}\right].$$

PROPOSITION 2: *The optimal size of the client's dealer network solves*

$$\underset{N\in\{\lfloor N^*\rfloor,\lceil N^*\rceil\}}{\operatorname{argmax}} \widehat{V}^b(N), \tag{A.50}$$

where N^* is given by the following condition

$$\frac{rV^b}{(\lambda N^*)^2}\left(1+\frac{\kappa}{r+\kappa}\frac{\eta}{r+\eta}\left(\frac{\lambda N^*}{r+\lambda N^*}\right)^2\right) = -\left(\frac{r+\lambda N^*}{\lambda N^*}-\frac{\eta}{r+\eta}\frac{\kappa}{r+\kappa}\frac{\lambda N^*}{r+\lambda N^*}\right)\frac{dV^r}{d\lambda N^*}+ \tag{A.51}$$
$$+\frac{\kappa}{r+\kappa}\frac{r}{(r+\lambda N^*)^2}\left(\frac{C(1-L)}{r}-P^s+K\right)+\frac{dP^s}{d\lambda N^*}\frac{\kappa}{r+\kappa}\frac{\lambda N^*}{r+\lambda N^*}-\frac{dP^b}{d\lambda N^*}.$$

Proof: We are solving for the optimal network size on the grid of integers $N \in \mathbb{I}_{\geq 0}$. To do so we assume that all value functions are well-behaved as functions of the network size, N. Then, we first solve for the optimal network size on the continuous grid $N^* \in \mathbb{R}_{\geq 0}$ and, finally, select the closest integer value to N^* maximizing $\widehat{V}^b(N)$. Equation (A.50) reflects this exercise.

Since $\frac{d\widehat{V}^b}{dN} = \frac{dV^b}{dN}+\frac{dV^r}{dN}$ we need to calculate the derivative of V^b with respect to N. We start by rewriting the expression (24) as

$$V^b = \left(\frac{C}{r+\kappa}+\frac{C(1-L)}{r+\lambda N}\frac{\kappa}{r+\kappa}+\left(E[P^s]+V^b\frac{\eta}{r+\eta}-K\right)\frac{\kappa}{r+\kappa}\frac{\lambda N}{r+\lambda N}-P^b-K\right)\frac{\lambda N}{r+\lambda N}, \tag{A.52}$$

and solve it for V^b to obtain

$$V^b = \frac{\frac{\lambda N}{r+\lambda N}}{1-\frac{\kappa}{r+\kappa}\frac{\eta}{r+\eta}(\frac{\lambda N}{r+\lambda N})^2}\left[\frac{C}{r+\kappa}+\frac{C(1-L)}{r+\lambda N}\frac{\kappa}{r+\kappa}+(P^s-K)\frac{\kappa}{r+\kappa}\frac{\lambda N}{r+\lambda N}-(P^b+K)\right]. \tag{A.53}$$

Taking into account that

$$\frac{d\left(\frac{\lambda N}{r+\lambda N}\right)}{d\lambda N} = \frac{r}{(r+\lambda N)^2}, \tag{A.54}$$

we obtain the following expression for the derivative in question

$$\frac{dV^b}{d\lambda N} = \left(\frac{r+\lambda N}{\lambda N} - \frac{\eta}{r+\eta}\frac{\kappa}{r+\kappa}\frac{\lambda N}{r+\lambda N}\right)^{-1}\left[\frac{rV^b}{(\lambda N)^2}\left(1+\frac{\kappa}{r+\kappa}\frac{\eta}{r+\eta}\left(\frac{\lambda N}{r+\lambda N}\right)^2\right) - \right. \tag{A.55}$$

$$\left. -\frac{\kappa}{r+\kappa}\frac{r}{(r+\lambda N)^2}\left(\frac{C(1-L)}{r} - P^s + K\right) + \frac{dP^s}{d\lambda N}\frac{\kappa}{r+\kappa}\frac{\lambda N}{r+\lambda N} - \frac{dP^b}{d\lambda N}\right], \tag{A.56}$$

which after setting $\frac{d\widehat{V}^b}{dN} = 0$ leads to (A.51) after some algebra.

Internet Appendix

Table IA.1: Correlation matrix

The table reports correlations between insurer-dealer volume, network size, and insurer characteristics.

	Insurer-dealer volume, year t	N	$ln(N)$	Size	RBC	Cash	Life	P&C	A-B
Insurer no. of dealers N	0.27	1.00							
$ln(N)$	0.20	0.90	1.00						
Insurer size	0.54	0.61	0.56	1.00					
Insurer RBC ratio	-0.08	-0.09	-0.11	-0.20	1.00				
Insurer cash-to-assets	-0.04	-0.04	-0.03	-0.09	0.03	1.00			
Life insurer	0.36	0.27	0.26	0.47	0.02	-0.04	1.00		
P&C insurer	-0.24	-0.27	-0.24	-0.34	0.03	-0.03	-0.72	1.00	
Insurer rated A-B	-0.14	0.01	0.05	-0.15	0.01	0.03	-0.06	0.06	1.00
Insurer rated C-F	-0.03	-0.05	-0.05	-0.07	-0.05	-0.07	-0.02	0.03	-0.07

Table IA.2: Scope of bond coverage by dealers

The table documents the scope of bond coverage by top dealers. We report the fraction of active bonds traded by each of the top 10 dealers in each category. We define a bond as active if it has at least a certain number of trades during the sample period, with threshold indicated in the column header. We vary the activity threshold across columns.

	Fraction of active bonds traded by each dealer				
Dealer rank	No. of bond trades≥200 (No. of bonds=992) (1)	≥100 (3,135) (2)	≥50 (6,821) (3)	≥10 (15,398) (4)	≥1 (21,007) (5)
1	0.998	0.983	0.923	0.731	0.587
2	0.997	0.981	0.919	0.707	0.564
3	0.997	0.965	0.887	0.680	0.544
4	0.996	0.965	0.868	0.651	0.519
5	0.995	0.956	0.863	0.645	0.512
6	0.994	0.948	0.853	0.631	0.505
7	0.994	0.939	0.812	0.576	0.453
8	0.985	0.919	0.799	0.576	0.451
9	0.981	0.913	0.795	0.568	0.449
10	0.959	0.884	0.752	0.533	0.419

Table IA.3: Dealer concentration by insurers

The table documents the dealer concentration by top insurers. We report descriptive statistics for the number of dealers used to trade any given bond across all active bonds. We define a bond as active if it has 200 or more trades during the sample period.

	Number of dealers used to trade an active bond				
Insurer rank	Mean	S.D.	Median	95%	Max
1	2.83	1.73	2	6	11
2	2.67	1.64	2	6	9
3	2.42	1.38	2	5	8
4	2.01	1.22	2	4	7
5	2.18	1.28	2	5	8
6	2.66	1.73	2	6	10
7	2.13	1.44	2	5	8
8	2.42	1.48	2	5	9
9	2.57	1.69	2	6	9
10	2.65	1.70	2	6	13
Large insurers	1.71	0.99	1	4	13
Small insurers	1.29	0.52	1	2	6

Table IA.4: Investor-dealer relations and execution costs

The table reports 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 level. See caption of Table 2 for additional details.

Determinant	(1)	(2)	(3)	(4)	(5)	(6)
Insurer-dealer volume, year t	-15.62**			14.08*		
Insurer-dealer volume, year t-1		-35.31***			-8.29	
Insurer-dealer volume, years t & t-1			-19.10***			3.16
Insurer no. of dealers				0.31***	0.33***	0.32***
ln(Insurer no. of dealers)				-6.12***	-6.37***	-6.23***
Insurer size				-3.80***	-3.49***	-3.67***
Insurer RBC ratio	1.04	0.88	0.87	-4.26***	-4.16***	-4.21***
Insurer cash-to-assets	-0.02	-0.02	-0.02	-0.04**	-0.04**	-0.04**
Life insurer	3.27***	3.51***	3.52***	4.36***	4.54***	4.42***
P&C insurer	3.54***	3.53***	3.51***	1.70**	1.75**	1.71**
Insurer rated A-B	2.95***	2.71***	2.72***	-0.00	-0.00	0.01
Insurer rated C-F	18.12***	17.81***	17.82***	11.05*	11.15*	11.10*
Insurer unrated	6.29***	6.00***	6.02***	0.55	0.63	0.60
Insurer buy	40.17***	40.15***	40.15***	39.28***	39.27***	39.27***
Trade size × Buy	-0.69***	-0.67***	-0.64***	-0.28***	-0.24***	-0.26***
Trade size × Sell	0.16	0.19*	0.21*	0.46***	0.51***	0.48***
Bond fixed effects (16,823)	Yes	Yes	Yes	Yes	Yes	Yes
Dealer fixed effects (401)	Yes	Yes	Yes	Yes	Yes	Yes
Day fixed effects (3,375)	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.152	0.153	0.152	0.155	0.155	0.155
N	918,279	918,279	918,279	918,279	918,279	918,279

Table IA.5: Estimated model parameters—robustness to minimum-distance weight ω

The table reports the estimated model parameters $\theta = (\theta_L, \theta_K, \theta_\kappa, \theta_\lambda, \theta^j_{w^i}, \theta_V)$ in Panel A. Estimates are from the minimum-distance estimation with weight ω indicated in the column header. Standard errors are computed using the sandwich estimator and reported in parenthesis. 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 non-trade value $V^r = e^{\theta^0_V}\eta^{\theta^1_V}N^{\theta^2_V}$. Panel C provides statistics on the model fit.

	$\omega = .1$ (1)	$\omega = .5$ (2)	$\omega = .9$ (3)
Panel A: Parameter estimates			
θ_L	2.61 (0.07)	2.51 (0.01)	2.57 (0.00)
θ_K	-0.64 (0.00)	-0.64 (0.00)	-0.64 (0.00)
θ_κ	0.31 (0.00)	0.31 (0.00)	0.30 (0.00)
θ_λ	0.50 (0.00)	0.50 (0.00)	0.51 (0.00)
$\theta^0_{w^b}$	2.51 (0.00)	2.52 (0.00)	2.51 (0.00)
$\theta^1_{w^b}$	-1.03 (0.00)	-1.02 (0.00)	-1.03 (0.00)
$\theta^0_{w^s}$	1.23 (0.00)	1.24 (0.00)	1.25 (0.00)
$\theta^1_{w^s}$	0.04 (0.00)	0.04 (0.00)	0.04 (0.00)
θ^0_V	-1.70 (0.00)	-1.70 (0.00)	-1.70 (0.00)
θ^1_V	0.32 (0.00)	0.33 (0.00)	0.32 (0.00)
θ^2_V	0.14 (0.00)	0.13 (0.00)	0.14 (0.00)
Panel B: Implied model parameters			
L	1.00	0.99	0.99
K $(*10^4)$	16.65	15.82	16.28
κ	22.46	21.19	20.14
λ	147.07	145.56	158.10
w^b (S.D.)	0.86 (0.05)	0.86 (0.05)	0.86 (0.05)
w^s (S.D.)	0.88 (0.01)	0.88 (0.01)	0.89 (0.01)
Non-trade value V^r (S.D.)	0.32 (0.03)	0.32 (0.03)	0.32 (0.03)
Panel C: Model fit			
Minimum distance $(\times 10^2)$	1.97	2.82	2.10
S.D. residuals	0.14	0.17	0.15
S.D. residuals network	0.24	0.19	0.14
S.D. residuals prices	0.13	0.14	0.20