

18-05 | August 29, 2018
Revised May 24, 2021

OTC Intermediaries

Andrea L. Eisfeldt

Anderson School of Management, University of California, Los Angeles
andrea.eisfeldt@anderson.ucla.edu

Bernard Herskovic

Anderson School of Management, University of California, Los Angeles
bernard.herskovic@anderson.ucla.edu

Sriram Rajan

Office of Financial Research
sriram.rajan@ofr.treasury.gov

Emil Siriwardane

Harvard Business School, Harvard University
esiriwardane@hbs.edu

The Office of Financial Research (OFR) Working Paper Series allows members of the OFR staff and their coauthors to disseminate preliminary research findings in a format intended to generate discussion and critical comments. Papers in the OFR Working Paper Series are works in progress and subject to revision.

Views and opinions expressed are those of the authors and do not necessarily represent official positions or policy of the OFR or Treasury. Comments and suggestions for improvements are welcome and should be directed to the authors. OFR working papers may be quoted without additional permission.

OTC Intermediaries*

Andrea L. Eisfeldt[†], Bernard Herskovic[‡], Sriram Rajan[§], Emil Siriwardane[¶]

Current Draft: May 24, 2021

Abstract

We study the effect of dealer exit on prices and quantities in a model of an over-the-counter (OTC) market featuring a core-periphery network with bilateral trading costs. The model is calibrated using regulatory data on the entire U.S. credit default swap (CDS) market between 2010-2013. Prices depend crucially on the risk-bearing capacity of core dealers, yet unlike standard models featuring a dealer sector, we allow for heterogeneity in dealer risk-bearing capacity. This heterogeneity is quantitatively important. Depending on how well dealers share risk, the exit of a single dealer can cause credit spreads to rise by 8 to 24%.

Keywords: OTC markets, networks, intermediaries, dealers, systemic risk, credit default swaps

*The views expressed in this paper are those of the authors and do not reflect the position of Depository Trust & Clearing Corporation (DTCC) or the Office of Financial Research (OFR). This paper does not reveal any confidential DTCC information.

We are very grateful to the Macro Financial Modeling Project at the Becker Friedman Institute for generous grant funding. We thank participants at the following conferences: Laboratory for Aggregate Economics and Finance OTC, Laboratory for Aggregate Economics and Finance Women in Macroeconomics, Society for Economic Dynamics, Conference on OTC Markets and Their Reform by the SNSF, NBER Conference on Financial Market Regulation, Chicago Booth Recent Advances in Empirical Asset Pricing, the Macro Financial Modeling Annual Meeting, Rome Junior Finance, Lubrafin, Maryland Junior Finance, NBER 2018 Summer Asset Pricing, Junior Workshop in New Empirical Finance at Columbia University, FMA Conference on Derivatives and Volatility, Western Finance Association, and Search and Matching in Macro and Finance. We also thank seminar participants at the UCLA Macro Finance lunch, UCLA Finance Brown Bag, Pontifical Catholic University of Rio de Janeiro (PUC-Rio), University of Washington, Central Bank of Chile, Federal Reserve Bank of Dallas, Federal Reserve Board, the Office of Financial Research, and the University of Toronto for helpful comments and suggestions.

[†]Finance Area, Anderson School of Management, University of California, Los Angeles, and NBER, e-mail: andrea.eisfeldt@anderson.ucla.edu

[‡]Finance Area, Anderson School of Management, University of California, Los Angeles, e-mail: bernard.herskovic@anderson.ucla.edu

[§]Office of Financial Research, United States Department of the Treasury, e-mail: sriram.rajan@ofr.treasury.gov

[¶]Finance Unit, Harvard Business School, Harvard University, Boston, e-mail: esiriwardane@hbs.edu

1 Introduction

A substantial portion of global financial assets are traded in over-the-counter (OTC) markets, including virtually all corporate bonds, sovereign fixed-income products, and swaps (e.g., interest rate, currency, and credit). Trade in OTC markets occurs bilaterally between counterparties who are connected through a network, one that typically features a handful of central intermediaries or dealers. During the 2008 Global Financial Crisis, the distress (and rescue) of individual dealers seemingly led to large movements in prices in many OTC markets. Figure 1 provides an illustration of price dynamics in the credit default swap (CDS) market around events related to the financial soundness of specific dealers. These events raised larger questions about how dealers impact OTC market outcomes and whether their individual behavior should be more closely monitored by regulators.

We explore these issues using a model of OTC markets that is calibrated to match regulatory data on the U.S. CDS market between 2010-2013. There are two types of agents or “traders” in our model, dealers and customers. Dealers are located at the center of a core-periphery trading network and are connected to everyone, whereas customers sit on the periphery and are connected only to dealers. Traders are equally risk averse but differ in their initial exposure to aggregate default risk. They can adjust their initial exposures using CDS contracts, though bilateral positions are subject to trading costs. This cost is a reduced-form way of capturing risk-management position limits (Saita 2007), price impact concerns (Malamud and Rostek 2017), or information asymmetries (Kyle 1985). In equilibrium, traders simply weigh their desire to hedge (or bear) default risk via CDS at the best price against the cost of trading too much with a single trading partner.¹

Our baseline model is calibrated to match average CDS spreads in dealer-dealer trades, dealer-customer trades, and the distribution of net CDS positions across traders. This allows us to simulate how prices and quantities in the CDS market would respond to the exit of a core dealer. When a dealer exits, its risk-bearing capacity gets removed from the network and the remaining traders then re-trade in the CDS market. We therefore study how the CDS market would reequilibrate if the remaining traders cannot adjust any factors that determine their overall aversion to bearing credit risk (e.g., their ex-ante exposures). In this sense, our stress tests are designed to approximate the medium-run response of the market to dealer exit.

CDS spreads are very sensitive to dealer exit in the core-periphery model with trading costs. The reason is that the observed heterogeneity in dealers’ CDS positions implies large differences in their ex-ante risk bearing capacity. Removing the dealer with the largest risk-bearing capacity – equivalently, the one who sells the most CDS protection – causes spreads in dealer-dealer and customer-dealer trades to rise by 24% and 11%, respectively. When this dealer exits, the risk-bearing capacity of the core falls dramatically. The remaining dealers have higher implied initial risk exposures than the exiting dealer and are therefore less willing to sell CDS protection. In fact, the loss of the largest net provider of credit insurance drives the aggregate dealer sector to be a

¹We provide reduced-form evidence of this tradeoff in Section 3.2 by showing that traders offer price concessions to counterparties with whom they have relatively small existing bilateral positions and vice versa.

net demander, rather than a net seller, of CDS protection. As a result, CDS spreads rise sharply. In contrast, a dealer that is net-neutral in CDS does not contribute to the core’s overall ability to bear risk and so its failure has a minimal impact on spreads.

We extend our baseline model by relaxing the assumption that trading costs are the same for all counterparties. In particular, we study how dealer exit impacts CDS spreads when trading costs between dealers are smaller than those between dealers and customers. The limiting case in which dealer-dealer trades are costless is interesting because it implies that dealers can perfectly share risk with each other, as in the seminal work of Du e, Gârleanu, and Pedersen (2005). In this case, we show that exit by the largest net seller would still cause CDS spreads to rise by 8%. The fact that this dealer sells so much CDS protection means that even when interdealer trading costs are negligible the implied risk-bearing capacity of this dealer is substantially larger than that of the average dealer. As a result, the risk-bearing capacity of the core still declines by a material amount upon its exit. The resulting impact on spreads is lower than in our baseline model because the heterogeneity in dealer risk-bearing capacity that is implied by observed CDS positions is smaller when trading costs between dealers are low.² However, because dealer-customer trading costs prevent customers with high risk-bearing capacity from replacing the lost capacity in the core, dealer exit still has a sizable effect even when interdealer trading costs are negligible.

The extended model also sheds light on how interdealer vs. customer-dealer trading costs uniquely impact equilibrium prices and quantities. The cost of trading between dealers does not impact their total CDS position as a group, but does alter the distribution of net CDS positions within the core. By contrast, trading costs between customers and dealers, along with the relative risk-bearing capacity of the two groups, are a central determinant of how much total protection is sold by dealers and how spreads differ between dealer-dealer and dealer-customer trades on average. Together, these properties allow us to estimate dealer-customer trading costs and the difference in total risk-bearing capacity between the two groups. To do so, we match the observed difference between average spreads in dealer-dealer and dealer-customer trades, as well as the net position of the dealer sector.

Even in the absence of any trading costs, the exit of a dealer could in theory impact CDS spreads by changing aggregate risk-bearing capacity. We gauge the strength of this channel by simulating dealer failure when the network is complete (or when trading costs are zero) and find almost no resulting spread impact, regardless of which dealer we remove. When the network is complete, all of the remaining dealers and customers can step in to replace the lost dealer’s trading capacity, which is small relative to the market.

Our sample period of 2010-2013 covers a relatively calm period in U.S. credit markets. However, in the last part of the paper, we show broader economic turmoil like the Fall 2008 or March 2020 could amplify the effects of dealer exit even further. In particular, we show that CDS spreads would rise by an additional factor of up to two or three if dealer exit coincides with an increase in risk

²In our baseline model, heterogeneity in risk-bearing capacity arises from differences in pre-trade risk exposures, however we show that our model produces quantitatively similar results with differences in risk aversion.

aversion, bilateral trading costs, or both. This magnification result holds true in both our baseline model and the extension with lower interdealer trading costs, though in the complete-network case there is no interaction between these structural parameters and dealer exit.

These insights about the effects of dealer exit inform the question of whether regulators should actively monitor the activity of individual dealers. Suppose that dealers as a group have no net position in CDS. It would be tempting in this case to conclude that individual dealers need not be monitored because they collectively act as a pure intermediary. Our results suggest that this intuition is incomplete without more data on the distribution of risk-bearing capacity within the core. If one dealer sells a large amount of CDS protection to another, then the sector will be net neutral, yet our stress tests indicate that the failure of either dealer could still substantially impact spreads. This situation resembles what we observe empirically, since a few dealers account for most of the CDS protection sold within the core during our sample.

Another important practical element of our baseline model is that equilibrium prices and quantities can be solved in closed form. Thus, there is transparent mapping between its parameters (e.g., trading costs) and two properties that regulators could easily monitor in real time: (i) the net CDS position of individual dealers and (ii) price dispersion, defined as the difference between CDS spreads in dealer-dealer transactions and dealer-customer transactions. Dealers are net sellers of CDS protection during our sample, providing an average of \$0.05 of default insurance for each dollar of their equity. As mentioned, a handful of dealers account for the bulk of these protection sales. In terms of price dispersion, we estimate that CDS spreads in dealer-dealer transactions are five percent lower than spreads in dealer-customer transactions.

Our paper contributes to several strands of research in asset pricing and systemic risk measurement. The specific objective function and endowment structure that we use builds on Atkeson, Eisfeldt, and Weill (2015), though there are important differences in our respective approaches. Network shape is not relevant in their search-based model because traders are fully connected, while the core-periphery structure plays a critical role in our model's equilibrium.³ In addition, our model is closely related to Malamud and Rostek (2017), who explore the implications of risk-sharing in a network where traders have quadratic utility. The setup in Malamud and Rostek (2017) is quite general and is useful for understanding the theoretical implications of endogenous price impact. In comparison, our model is tailored so that we can more easily estimate the empirical impact of dealer removal.

Spreads in our setup are heavily influenced by the risk-bearing capacity of core dealers. For example, we show that the net position of dealers determines whether observed spreads will be lower or higher than the Walrasian spread. Intuitively, when the dealer sector becomes a net demander of credit insurance, spreads widen considerably and are above the Walrasian spread. The importance of core dealers for market outcomes in our model is similar in spirit to theories

³The role of networks has been studied in several other related settings, including financial intermediation (Farboodi 2017), OTC markets (Golub and Livne 2010, Hendershott, Li, Livdan, and Schürhoff 2017, Babus and Kondor 2018), investor networks (Akerlof and Holden 2016), decentralized exchanges (Malamud and Rostek 2017), firm volatility (Herskovic, Kelly, Lustig, and Van Nieuwerburgh 2017), and input-output (Herskovic 2018).

of intermediary-based asset pricing (He and Krishnamurthy 2013). Consistent with our approach, Haddad and Muir (2018) also show that return predictability from an intermediary-based factor is greater for OTC-traded assets like foreign exchange, mortgage-backed securities, and credit default swaps compared to exchange traded assets such as equities.

While our results support the idea that intermediary risk-bearing capacity is a crucial determinant of asset prices, our work highlights some key directions in which these models might be extended in future work. First, there is mounting evidence (Siriwardane (2018), Lewis, Longstaff, and Petrusek (2017)) that the financial health of individual dealers can significantly impact asset prices. From the perspective of theory, this suggests a benefit to relaxing the standard assumption of a representative intermediary. Our model is one such example. We emphasize that heterogeneity within the core is important for understanding the consequences of dealer failure, especially when dealers do not perfectly share risks with one another.

Second, our findings also imply that net worth may not accurately measure the risk-bearing capacity of dealers, as is commonly assumed in applied work on intermediary-based asset pricing (Adrian, Etula, and Muir 2014). In our baseline model, the dealers who sell the most protection do so because they start with low initial exposures to aggregate default risk, which in principal could be driven by factors unrelated to size or net worth.⁴ Our model allows for other sources of heterogeneity. For example, differences in dealers' business models, risk-management abilities, or risk aversion could all lead to different pre-trade risk exposures. This observation means that policy makers may need to look further than simple measures of financial soundness and connectivity in order to assess the systemic importance of a specific dealer for market outcomes.

The broad motivation for this paper naturally draws on a large body of work that examines systemic risk in financial networks (see Allen and Gale (2000), Eisenberg and Noe (2001), Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015), and Denbee, Julliard, Li, and Yuan (2014), among many others).⁵ That literature largely focuses on the propagation of a liquidity or default shock through a lending network, whereas we study how dealers in a core-periphery trading network could disrupt capital market outcomes. Importantly, we show that dealer failure can materially impact CDS spreads even without realized corporate defaults or contagion of any kind.

Another salient feature of our OTC network model is price dispersion, which in our context means that dealer-dealer trades occur at lower CDS spreads than customer-dealer trades. The prevailing explanation for price dispersion in OTC markets centers around strategic pricing on the part of core dealers (Colliard, Foucault, and Hoffmann (2018), Di Maggio, Kermani, and Song (2017)). In our setting, prices are set competitively and dispersion still occurs in equilibrium because trading costs prevent dealers and customers from perfectly sharing risk. In equilibrium, prices between dealers reflect their willingness to pay for insurance, which may differ from customers' willingness to pay and therefore generate price dispersion. Consistent with this theoretical result,

⁴Indeed, there is a weak cross-sectional correlation between dealer CDS positions and leverage during our sample.

⁵In related work in mathematical finance, Cont and Minca (2016) use publicly available aggregate CDS data and a random graph framework to ask how liquidity shocks cascade through derivatives networks. The fixed point they study is the number of nodes that fail after a liquidity shock, whereas we focus on equilibrium asset pricing effects.

we also document sizable price dispersion in CDS transactions where competition is more likely to be high.

The remainder of the paper is organized as follows. We start by outlining our baseline model in Section 2. In Section 3, we discuss our data, provide reduced-form evidence supporting our assumption of bilateral execution costs, and calibrate the model. We use the calibrated model to study the impact of dealer failure in Section 4 and extend our baseline setup to accommodate lower execution costs within the core. Section 5 contains concluding remarks.

2 Model

In this section, we lay out the baseline version of the model that we use as the basis for our empirical work. One attractive feature of the baseline model is that it can be solved in closed form and is thus a parsimonious workhorse with which to develop intuition. The transparent nature of our model allows us to easily calibrate its key parameters and to develop intuition for how the OTC network impacts market prices and stability. We consider extensions of the baseline model which incorporate preference heterogeneity, along with the associated quantitative implications related to dealer exit, in Section 4.2.

2.1 Investor Preferences and the Trading Network

There are n agents in the economy and one asset with random payoff equal to $(1 - D)$. D can be interpreted as aggregate default risk and it has a mean μ and variance σ^2 . Agent i is initially endowed with an exposure ω_i to the underlying asset. Agents can trade in the CDS market before aggregate default risk is realized, which allows them to hedge or take on additional default risk. A CDS contract between agents i and j specifies that agent i promises to pay D to agent j , and in exchange, agent j makes a payment of R_{ij} to agent i . Formally, ψ_{ij} is the amount of CDS sold by agent i to agent j . The bilateral position ψ_{ij} is positive when agent i sells insurance against aggregate default to j , and ψ_{ij} is negative when i purchases insurance from j . The CDS spreads (R_{ij}) and the bilateral positions (ψ_{ij}) are determined in equilibrium.

There is a network of trade connections that determines which agents can trade CDS with each other. The network is characterized by an $n \times n$ matrix G of zeros and ones, where each entry (i, j) of G is denoted by $g_{ij} \in \{0, 1\}$. If agents i and j can trade, then $g_{ij} = 1$, and, if they cannot trade, then $g_{ij} = 0$. Furthermore, the network is symmetric, i.e., $g_{ij} = g_{ji}$ for every i and j , which means that if agent i can trade with agent j , then j can also trade with i . Without loss of generality, we assume that $g_{ii} = 1$ for every i .⁶ We treat the trading network G as exogenously determined.⁷

Agents have mean-variance preferences over default risk. They must also pay a trading cost that increases in the size of their bilateral positions. Overall, their optimization problem is given

⁶The equilibrium allocation is identical whether we set $g_{ii} = 0$ or $g_{ii} = 1$.

⁷As we argue in the Internet Appendix, this approach is supported by the fact that new connections in the CDS market are rarely formed and existing connections are rarely broken: conditional on no connection in week t , two counterparties have a 0.01% chance of making a new connection in the following week.

by:

$$\max_{\{g_{ij}\}_{j=1}^n, z_i} \omega_i(1 - \mu) + \sum_{j=1}^n g_{ij}(R_{ij} - \mu) - \frac{\alpha\psi}{2}(\omega_i + z_i)^2 - \frac{\gamma}{2} \sum_{j=1}^n g_{ij}^2 \quad (1)$$

subject to $g_{ij} = 0$ if $g_{ji} = 0$, and

$$z_i = \sum_{j=1}^n g_{ij} \quad (2)$$

where z_i is agent i 's net position in the CDS market, $\alpha\psi > 0$ is a risk aversion parameter and $\gamma > 0$ is the trading cost parameter. The first restriction guarantees that agent i can trade with agent j if they are connected. The parameters $\alpha\psi$ and γ play distinct roles in the model. Risk aversion represented by the parameter $\alpha\psi$ measures aversion to total post-trade exposure to the aggregate default risk, measured by $\omega_i + z_i$.

The parameter γ is new to our model and it captures bilateral trading costs. It represents agents' aversion to hold large bilateral positions as well as their willingness to smooth trades out across counterparties. In practice, this desire could arise for several reasons, such as efforts to minimize information leakage, reduce hold-up problems, ensure the ability to trade in case of a counterparty failure, or diversify relationship capital.⁸ Later, in Section 3.2, we provide empirical evidence that institutions in the CDS market do indeed prefer to spread their trades out across multiple counterparties, i.e. $\gamma > 0$.

There are several ways to model agents' desire to spread out their trades. One advantage of our exact functional form is its simplicity—as we will derive below, the model's equilibrium delivers a closed-form mapping between easily observable quantities in the data and the main structural parameters. This makes it easier to calibrate the model later in Section 3.4 and to develop intuition for the resulting equilibrium. One critique of our functional form might be that it penalizes high-volume traders, even if they perfectly spread out their trades across the counterparties in their network. In the Internet Appendix I.3.1, we develop a variant of the model without this feature. While this variant cannot be solved in closed form, its equilibrium carries the same intuition of our benchmark model. More importantly, both versions deliver similar findings on the impact of dealer removal.

2.2 Equilibrium

Having defined each individual agent's optimization problem, let us now define the market clearing conditions. In equilibrium, how much exposure agent i wants to sell to agent j has to be equal to how much agent j wants to buy from agent i . Hence, bilateral clearing conditions for any two

⁸See Balasubramaniam, Gomes, and Lee (2019) for an explicit model of the net benefit of access to multiple counterparties in times of market stress.

counterparties in our model are given by:

$$i_j + j_i = 0 \quad \forall i, j = 1, \dots, n. \quad (3)$$

Finally, we assume no transaction costs between counterparties, which means that a payment agent i receives from selling to agent j is exactly the amount agent j pays for such contract. In addition, there is no strategic pricing in our model, so dealers do not explicitly have monopolist or oligopolistic power over customers. While market power is certainly an interesting and realistic feature of OTC markets (e.g. Hau, Hoffmann, Langfield, and Timmer (2018)), we do not model it in order to focus on how network frictions impact equilibrium pricing.⁹ Formally, a no-transaction-cost assumption means that prices satisfy the following condition:

$$R_{ij} = R_{ji} \quad \forall i, j = 1, \dots, n. \quad (4)$$

We solve this model for a competitive equilibrium, in which agents optimize taking prices as given, and all markets clear, using the following equilibrium concept.

Definition. *An economy consists of a finite number of agents n , a trading network G , preferences described in Equation (1), and pre-trade exposures given by $\{\omega_i\}_{i=1}^n$. A competitive equilibrium with no transaction costs consists of spot market prices $\{R_{ij}\}_{i,j=1,\dots,n}$ and traded quantities $\{i_j\}_{i,j=1,\dots,n}$ such that: (i) agents optimize, taking the network of trading connections and prices as given (Equation 1); (ii) markets clear (Equation 3); and (iii) there are no transaction costs (Equation 4).*

In this subsection, we fully characterize the equilibrium of the model. If agent i can trade with agent j , i.e., $g_{ij} = 1$, then agent i 's first-order condition with respect to i_j is:

$$\underbrace{R_{ij}}_{\text{MB of selling}} \underbrace{\mu}_{\text{MC of bilateral concentration}} = \underbrace{i_j}_{\text{Shadow cost of risk bearing}} + \underbrace{\hat{z}_i}_{\text{Shadow cost of risk bearing}}, \quad (5)$$

where

$$\hat{z}_i = (\omega_i + z_i)\alpha\sigma^2. \quad (6)$$

Equation (5) specifies agent i 's optimal exposure to aggregate default risk as a function of the contract premium, $R_{ij} \mu$, along with an additional term, \hat{z}_i . We interpret this last term as the shadow price of aggregate default risk for agent i , since it is the Lagrange multiplier on the constraint given by equation (2). Hence, \hat{z}_i is agent i 's willingness to pay to insure against one additional unit of exposure to aggregate default.¹⁰

⁹In the Internet Appendix [I.3.2](#), we present a version of our theoretical model allowing agents to take into account their impact on equilibrium prices, inspired by the more detailed analysis by Malamud and Rostek (2017). We show that our analysis holds in an environment with price impact as well. In addition, we address the potential for strategic pricing in the empirical estimates we use in our calibration.

¹⁰In equilibrium, agents trade because they have different pre-trade exposures and want to share risk, consistent with risk management considerations. In the Internet Appendix [I.3.3](#), we argue how a model where agents differ in their beliefs about aggregate default risk μ_i delivers the same equilibrium allocation.

Agent i 's first-order condition equalizes the marginal benefit of selling insurance to its own shadow cost of risk bearing combined with the marginal cost associated with bilateral trading costs. The risk aversion parameter, α , determines how much agent i values net positions through the shadow cost of risk bearing, while the trading cost parameter, ψ , determines how much agent i values bilateral contracts individually. In other words, α drives total net positions, while ψ defines how much agents sell to and buy from each counterparty. Agent i will buy CDS from agent j at a higher spread than from agent k in order to avoid holding a large position with agent k , implying that connected agents only exploit price dispersion up to the extent that the benefits from attractive spreads outweigh the costs of concentrated bilateral positions.

By combining the first-order condition in Equation (5) with the counterparty clearing conditions in Equation (3) and the no-transaction cost assumption in Equation (4), we can write equilibrium prices as a linear combination of counterparties' shadow prices of insurance:

$$\underbrace{R_{ij}}_{\text{contract premium}} - \mu = \frac{\hat{z}_i + \hat{z}_j}{2}, \quad (7)$$

for every i and j who can trade, i.e., $g_{ij} = g_{ji} = 1$.

The contract premium, which is the contract price in excess of the expected default in the underlining asset, depends on agents' shadow prices of insurance. As a result, whenever there are differences in agents' shadow prices of insurance, there is price dispersion in the cross section of agents in equilibrium, even if agents have identical preferences.

The CDS premium in Equation (7) is a function of the shadow prices of risk, which are determined in equilibrium. We can use Equations (2), (5), (6), and (7) to solve for equilibrium net positions as a linear combination of initial exposures and the net positions of other agents:

$$z_i + \omega_i = (1 - \gamma_i) \omega_i + \gamma_i \sum_{j=1}^n \tilde{g}_{ij} (z_j + \omega_j) \quad \forall i = 1, \dots, n, \quad (8)$$

where $\tilde{g}_{ij} = \frac{g_{ij}}{K_i}$, $K_i = \sum_{j=1}^n g_{ij}$, and $\gamma_i = \frac{K_i \alpha \sigma \psi}{K_i \alpha \sigma \psi + 2} \in (0, 1)$.

Agent i 's post-trade exposure to aggregate default risk is given by $z_i + \omega_i$. In equilibrium, i 's post-trade exposure is a convex combination of her pre-trade exposure, i.e. ω_i , and a network-weighted average of agent i 's neighbors' equilibrium post-trade exposures. The weight γ_i defines how close agent i is to the average of her neighbors' post-trade exposures, and it makes a clear distinction between risk aversion, α , and aversion to bilateral concentration, ψ . Notice that γ_i depends on the ratio $\frac{2}{K_i \alpha \sigma \psi}$ and parameters α and ψ have opposite effects on γ_i . Risk aversion increases agents' willingness to diversify risk away and makes agent i 's post-trade exposures closer to her neighbors, by increasing γ_i . Larger trading costs, however, make agents less willing to hold larger positions with their counterparties at the expense of lower risk sharing. As a result, γ_i decreases, which leads to *lower* risk diversification.¹¹

¹¹The limiting cases of the equilibrium are also instructive. As ψ goes to infinity, we have that γ_i goes to zero, and

2.3 Risk-sharing Benchmark: Complete Network

In this subsection, we consider the model when the trading network is complete, i.e., $g_{ij} = 1$ for every i and j . We consider this risk-sharing benchmark because the trading network itself does not impose any additional trading frictions. Under the complete network benchmark, the average prices in equilibrium would be:

$$\bar{R}_{\text{Complete Network}} \equiv \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n R_{ij} = \alpha \bar{\omega} + \mu, \quad (9)$$

where $\bar{\omega} \equiv \frac{1}{n} \sum_{i=1}^n \omega_i$.

In Appendix A.1, we provide detailed derivation and we also show that when $\alpha = 0$, there is perfect risk sharing and *all* equilibrium spreads are equal to $\alpha \bar{\omega} + \mu$, whenever any two agents are connected through a sequence of links. Hence, the complete network economy features an average equilibrium price that is the same as the one in an economy with perfect risk sharing. The difference between these two cases is that the complete network still has price dispersion in the cross section whenever $\alpha \neq 0$ and pre-trade exposures are heterogeneous.

2.4 Core-Periphery Equilibrium

We now further characterize the market equilibrium taking into account that the trading network is core-periphery. We analyze a core-periphery trading network because it is well-established that many OTC trading networks are characterized by a core-periphery structure in which a central set of dealers trades with a periphery set of customers (e.g., Li and Schürhoff (2018), Peltonen, Scheicher, and Vuillemeys (2014), or Hollifield, Neklyudov, and Spatt (2017)). In the Internet Appendix, we confirm that this is the case in the U.S. CDS market as well.¹² Appendix A.2 contains the detailed derivation of the model under the core-periphery trading network. Formally, we assume that G has a core-periphery shape, defined as:

Definition. A *core-periphery trading network* consists of two groups of agents, a core and a periphery, in which (i) all agents in the core are connected to all other core agents, and all peripheral agents, and (ii) peripheral agents are connected to all core agents, and no other peripheral agents.

Specifically, let n_d be the number of members in the core (“dealers”), and $n_c = n - n_d$ the number of agents in the periphery (“customers”). We use subscripts d and c to denote dealers and customers. Without loss of generality we set agents $i = 1, \dots, n_d$ to be core agents. Hence, a autarky allocation is achieved in equilibrium (i.e. $z_i = 0$ and $\omega_{ij} = 0 \quad \forall i, j$). Alternatively, as α goes to zero, ω_i goes to one, and neighboring agents are able to perfectly share risks. In Appendix A.1, we further discuss these limiting cases as well as the model solution at $\alpha = 0$.

¹²In Internet Appendix I.2.1, we consider an example with three agents in order to provide intuition and to highlight key features of our framework. In addition, in the Internet Appendix I.2.2, we provide a core-periphery example with the smallest possible number of agents (five).

core-periphery trading network is defined as:

$$G_{\text{core-periphery}} = \begin{bmatrix} \mathbb{1}_{n_d} \mathbb{1}'_{n_d} & \mathbb{1}_{n_d} \mathbb{1}'_{n_c} \\ \mathbb{1}_{n_c} \mathbb{1}'_{n_d} & \mathbb{I}_{n_c} \end{bmatrix}, \quad (10)$$

where $\mathbb{1}_{n_d}$ is a column vector of ones with n_d elements and \mathbb{I}_{n_c} is an $n_c \times n_c$ identity matrix. The core-periphery G matrix with n_d dealers consists of blocks of ones along the top n_d rows, and left-most n_d columns, representing the complete connections of core agents. While the upper right square of ones represents connections between dealers, the upper right and lower left squares of ones represent the full connections between dealers and customers. Finally, the bottom right identity matrix represents that customers are not connected with other customers.

In the dealer market, the average post-trade exposures of dealers is given by:

$$\bar{z}_d + \bar{\omega}_d \equiv \frac{1}{n_d} \sum_{i=1}^{n_d} (z_i + \omega_i) = (1 - \alpha) \bar{\omega}_d + \alpha \bar{\omega}. \quad (11)$$

where $\bar{\omega}_d = \frac{1}{n_d} \sum_{i=1}^{n_d} \omega_i$ and $\alpha = \frac{n_d \alpha \sigma^2}{n_d \alpha \sigma^2 + 2}$. Equation (11) is derived from Equation (8) applied to dealers and taking an average across all dealers. The average post-trade exposures of dealers are a convex combination of their own average pre-trade exposure ($\bar{\omega}_d$) and the average pre-trade exposure in the economy ($\bar{\omega}$).

Since $\alpha \in (0, 1)$, dealers are net sellers of protection on average, i.e., $\bar{z}_d > 0$, if, and only if, dealers are less exposed to aggregate default risk, i.e., $\bar{\omega}_d < \bar{\omega}$. Moreover, the average price in the dealer market, i.e., \bar{R}_d , is given by:

$$\bar{R}_d \equiv \frac{1}{n_d^2} \sum_{i=1}^{n_d} \sum_{j=1}^{n_d} R_{ij} = \mu + \alpha \sigma^2 \bar{\omega} \psi (1 - \alpha) \alpha \sigma^2 (\bar{\omega} - \bar{\omega}_d). \quad (12)$$

If dealers are, on average, less exposed to the underlying default risk ($\bar{\omega}_d < \bar{\omega}$), then prices in the dealer market are lower than the complete network benchmark as derived in Equation (9).

Similarly, in the customer market, the average post-trade exposures of customers are given by:

$$\bar{z}_c + \bar{\omega}_c = \alpha \bar{\omega} \psi (1 - \alpha) \bar{\omega}_c + \alpha (1 - \alpha) (\bar{\omega} - \bar{\omega}_d), \quad (13)$$

where $\alpha = \frac{n_d \alpha \sigma^2}{n_d \alpha \sigma^2 + 2}$. We can write the average price in the customer market as follows:

$$\bar{R}_c = \mu + \alpha \sigma^2 \bar{\omega} \psi \frac{1}{2} \alpha \sigma^2 (\bar{\omega} - \bar{\omega}_d) \left[(1 + \alpha) (1 - \alpha) \frac{n_d}{n} (1 - \alpha) \right], \quad (14)$$

$$= \bar{R}_d + \frac{1}{2} \alpha \sigma^2 (1 - \alpha) \left[1 + (1 - \alpha) \frac{n}{n_d} \frac{n_d}{n} \frac{n_d}{n} (\bar{\omega} - \bar{\omega}_d) \right]. \quad (15)$$

The next proposition compares the average price in the dealer market, the average price in the customer market, and the average price in the complete network benchmark.

Proposition 1. *In the core-periphery model with $\frac{n_d}{n} < \frac{1}{2}$, the average pre-trade exposure of dealers is lower than the average exposure in the economy, i.e., $\bar{\omega}_d > \bar{\omega}$, if, and only if,*

$$\mu + \alpha \sigma^2 \bar{\omega}_d < \bar{R}_d < \bar{R}_c < \bar{R}_{\text{Complete Network}},$$

where $\bar{R}_{\text{Complete Network}} = \mu + \alpha \sigma^2 \bar{\omega}$ as in Equation (9). Alternatively, $\bar{\omega}_d < \bar{\omega}$ if, and only if, $\mu + \alpha \sigma^2 \bar{\omega}_d > \bar{R}_d > \bar{R}_c > \bar{R}_{\text{Complete Network}}$.

The proof is in Appendix A.2. Proposition 1 presents two interesting results. First, it shows if dealers have a lower-than-average initial exposure to default risk—or equivalently, if dealers are net sellers of protection—then spreads in the interdealer market will be lower than those in the dealer-customer market. This dispersion occurs because dealers start with a lower initial risk exposure, hence why they want to sell CDS protection to customers. However, trading costs prevents the two groups from perfectly sharing risks, so dealers still have lower total risk exposure even after trading in CDS. In other words, dealers will have a lower shadow cost of risk bearing than customers in equilibrium. Equation 7 states that the equilibrium price between equals their average shadow cost of bearing default risk. Because dealers have a lower shadow cost, Equation 7 therefore implies that interdealer spreads will be lower on average than customer-dealer spreads. Importantly, this dispersion occurs despite the fact that competition is perfect in our setting.

The second result builds on the same logic but is a bit more subtle. This result shows that if the number of dealers is sufficiently small, the average price in *both* the dealer and the customer markets is below the complete network benchmark average price.¹³ The intuition reflects two considerations. First, due to imperfect risk sharing, dealers have less post-trade risk exposure than compared to customers. Second, there are a small number of dealers in the economy. From Equation (5), the average prices in customer-dealer trades will be equal to the average between two components: (i) dealers average shadow cost of insurance and (ii) customers average shadow cost of insurance. Hence, if average shadow cost of insurance among dealers is sufficiently low, then average dealer-customer prices are lower as well. This effect is stronger when there are fewer dealers. As a consequence, the average between (i) and (ii) is below the spread that would obtain in a complete network.

2.5 Mapping the Core-Periphery Model to the Data

Our goal is to map the model’s two structural preference parameters, α and λ , to the data. This turns out to be straightforward in the case of the core-periphery network. For example, by rearranging equations (11), (12), and (14), we can express α as a function of observables:

$$\alpha = \frac{1}{2\bar{\omega}_d} \left[\bar{R}_d - \mu + (\bar{R}_c - \bar{R}_d) \left(1 - \frac{n_d}{n} \right) \right]. \quad (16)$$

¹³See Gavazza (2016) for a related result.

This equation says that aversion to aggregate credit risk α is primarily identified by the level of equilibrium spreads in the dealer-dealer (\bar{R}_d) and customer-dealer market (\bar{R}_c).

Next, to calibrate the parameter measuring aversion to bilateral concentration, β , we can combine β from Equation (11) with Equations (12) and (14) to obtain:

$$\beta = (n - n_d) \left(\frac{\bar{R}_c - \bar{R}_d}{\bar{z}_d} \right). \quad (17)$$

This equation says that the inference of β depends on the number of customers, $n - n_d$, the customer-dealer spread, $\bar{R}_c - \bar{R}_d$, and the average exposure of dealers, \bar{z}_d . Higher price dispersion ($\bar{R}_c - \bar{R}_d$) coincides with larger values of β , as agents are less willing to build up concentrated bilateral exposures. Conversely, because dealers are net sellers of protection in the data ($\bar{z}_d > 0$), when they sell more protection it reflects better risk sharing in equilibrium, and hence a lower β . The comparative statics in Proposition A3 in the Appendix provides further intuition for the mapping between β and α and observable data. To summarize, in the case of the core-periphery network, α is mainly identified from the level of credit spreads, while β is mainly identified by the dispersion in spreads.

Furthermore, in Appendix A.2, we show that we can use the net position of each dealer to infer their pre-trade exposures to aggregate default. The average net position of dealers also allows us to infer the *average* net position of customers and then compute the implied *average* pre-trade exposures for customers as well.¹⁴

3 Empirical Analysis

In the first part of this section, we provide reduced-form evidence that traders do indeed prefer to smooth their trades across counterparties in the CDS market. This evidence supports our modeling of preferences in the preceding section ($\beta > 0$). We then measure two key properties of the CDS market that are necessary to calibrate the core-periphery network model. Specifically, when the network is core-periphery, we showed in Section 2.5 that we can infer the model's structural parameters based on: (i) the net CDS position of dealers and (ii) the difference between CDS spreads in dealer-dealer transactions and dealer-customer transactions. As we show below, dealers are on average net sellers of credit protection in the CDS market and dealer-dealer transactions typically occur at lower CDS spreads than dealer-customer transactions.

3.1 Data Description

Before proceeding to our main analysis, we start with a basic description of the data used throughout the remainder of the paper. Our primary data on CDS transactions and positions come from Deposit Trust and Clearing Corporation (DTCC), which provides the data to the U.S. Treasury Department's Office of Financial Research (OFR) under a license agreement. The data are derived

¹⁴See Equations (A8), (A10), (A11), and (A16) in Appendix A.2.

from DTCC’s Trade Information Warehouse and include CDS transactions and positions reported to DTCC. Transactions represent flows in CDS, and positions represent stocks. DTCC converts transactions to open positions before delivering both to the OFR. Positions data are updated at the end of each week. DTCC data have been used previously by Oehmke and Zawadowski (2016), Siriwardane (2018) and Du, Gadgil, Brody, and Vega (2017).

For both transactions and positions, we observe complete information on the identities of the counterparties in the trade, pricing terms, size, and all contract details. When working with transactions, we follow industry practice and infer CDS spreads based on the the International Swaps and Derivatives Association standard pricing model. DTCC provides the OFR with data on transactions or positions that meet at least one of two conditions: (i) the underlying firm covered by the swap is U.S. based or (ii) at least one of the counterparties in the swap is U.S. registered. In addition, DTCC CDS data include all North American index swap transactions and positions (i.e. the index family is “CDX.NA.”). The data therefore capture most of the CDS market for U.S. firms.¹⁵ This is the key difference between DTCC data from the OFR vs. from the Federal Reserve Banks; the network coverage from the OFR is more comprehensive in the sense that it covers more than just entities regulated by the Fed. This is important as large hedge funds are important customers in the CDS market.

We focus on data from January 1, 2010 through December 31, 2013, when central clearing of single-name contracts was rare.¹⁶ In our data, we do not observe the ultimate counterparty for contracts that are centrally cleared. For example, a centrally cleared trade between Hedge Fund *A* and Dealer *B* will appear in our data as a trade between Hedge Fund *A* and the central clearing party, plus another (unlinked) trade between the central clearing party and Dealer *B*. This feature of the data matters only when we estimate the difference between inter-dealer prices and customer-dealer prices, as we must observe the ultimate counterparty type. For this reason, when analyzing CDS spreads, we use only single-name transactions on U.S. firms from 2010 to 2013, a time period that pre-dates central clearing of single-name but not index contracts. By contrast, we use all trades when we estimate total net bilateral exposures and credit positions.

It is important to keep in mind that our model applies whenever building bilateral positions is costly. While central clearing should mitigate costs related to counterparty risk, it should have less of an impact on costs associated with information leakage or other hold-up problems. To the extent that these other costs are large, the main insights of our model should therefore remain applicable in a world with central clearing, though this is certainly an interesting question for future work. Moreover, a large portion of the interest rate and FX swaps markets are not centrally cleared, and our model can be easily modified to study these other settings.

Our subsequent analysis also requires us to designate which of the members of the network are dealers and which are customers. In Appendix C, we use a minimum-distance algorithm for

¹⁵We refer to the underlying company whose default is covered by a CDS contract as the “firm” or “underlying firm”. The underlying firm is also often referred to as the reference entity or “name” in the swap.

¹⁶Central clearing of single-name contracts was not prevalent until 2014. See https://www.theice.com/article/cds-growth?utm_source=Insights&utm_medium=tile.

dealer classification. This algorithm generates a counterparty network with 14 dealers. In many applications, we also use DTCC’s classification of dealers as a robustness check.¹⁷

3.2 Do agents prefer to spread out their trades across counterparties?

A key assumption in our model from Section 2 is that institutions prefer to spread their trades out across multiple counterparties (i.e., $\alpha > 0$). There are several reasons why traders may wish to spread trades across counterparties, including risk-management requirements (e.g., position limits) or the desire to conceal private information. Regardless of the precise micro-foundation, assuming that institutions face some form of bilateral trading costs is critical for generating limited risk sharing and price dispersion when the trading network is incomplete, as it is in practice. Using bilateral positions and individual transactions, we now test whether the assumption of $\alpha > 0$ is supported empirically. In the Internet Appendix, we also formally prove that our testing approach is valid inside of our model.

3.2.1 Measuring Bilateral Exposures

We begin by constructing a measure of bilateral CDS exposure between two counterparties, i and j , at time t . At any given point in time, i and j may have several outstanding CDS contracts with each other that cover different firms and maturities. We therefore aggregate i and j ’s bilateral exposures by representing each individual CDS contract in terms of its exposure to aggregate credit risk. This approach to aggregation also accords with our model, where default risk is defined as exposure to a single aggregate risk factor. We provide a detailed account of our methodology in the Internet Appendix, though we briefly describe the key elements now. Specifically, we proceed in three steps. First, we define the aggregate credit risk factor. Second, we compute each bilateral position’s exposure to this aggregate credit risk factor. Finally, we aggregate positions simply by summing the computed exposures from Step 2.

Step 1: Defining an Aggregate Credit Risk Factor On each date t , we define our aggregate credit risk factor as the cross-sectional equal-weighted average of all five-year CDS spreads for U.S. firms in the Markit Ltd. database. We have separately confirmed that this index is a close approximation to the first principal component of credit spreads across all maturities (i.e. a level factor). However, our simpler index is much better at dealing with missing data, which can be an issue for firms with lower volumes of CDS trading. Figure 2 shows that our aggregate credit risk factor evolves as one might expect, peaking at nearly 1000 basis points during the 2007-09 financial crisis. Our factor is over 90 percent correlated with both the 5-year CDX investment grade and high yield indices. The average of our index is a little over 200 basis points, so it is between the investment grade and high yield indices in terms of average credit spread.

¹⁷DTCC also classifies traders based on its list of registered dealer members. In single-name transaction data, DTCC’s set of dealers is responsible for nearly 86 percent of gross volume. The 14 counterparties who we label as a dealer are responsible for about 83 percent.

Step 2: Measuring the Exposure of a Single Contract to Aggregate Credit Risk Next, we compute the exposure of an arbitrary CDS position p to our aggregate credit risk factor. On date t , suppose that the position is written on firm f and has m remaining years till maturity. We first assign each position to a “maturity bucket” b based on its maturity m as follows: $b = 1$ for maturities of less than 2 years, $b = 3$ for maturities between 2 and 4 years, $b = 5$ for maturities between 4 and 6 years, and $b = 7$ for all maturities over 6 years.

Next, for each position p , we match it to the Markit CDS spread database based on the underlying firm f and maturity bucket b . Markit provides constant maturity CDS spreads for maturities ranging from 6 months all the way to 10 years. We match each position’s maturity bucket b to the closest constant maturity spread in Markit. For instance, if we observe a position on Ford Motor Co. that has a maturity bucket $b = 3$, we obtain Ford’s history of three-year CDS spreads up to date t from Markit.¹⁸ Next, we compute the position’s underlying beta with respect to changes in our aggregate credit risk factor via the following rolling regression:

$$\Delta CDS_{f,b,s} = \alpha_{p,t} \times \Delta \text{CDS Index}_s + \varepsilon_{p,b,s}, \quad s \in [t - 2 \text{ years}, t]$$

where CDS Index_s is our aggregate credit risk factor on date s . The regression is run using weekly data over a rolling window of two years. The position’s beta $\alpha_{p,t}$ gives us a gauge of how sensitive the underlying CDS spread of the position is to movements in this index.

We compute $\alpha_{p,t}$ for every position contained in our database sourced from DTCC. Importantly, we account for both index and single name CDS positions. Selling protection on an index is equivalent to selling protection on the individual firms that comprise the index. This distinction is particularly important in the CDS market because index positions are nearly half of the net notional outstanding for the entire CDS market during our sample (Siriwardane (2018)). To account for this fact, we follow Siriwardane (2018) and disaggregate CDS indices into their individual constituents and then combine these “disaggregated” positions with any pure single name positions. We then estimate $\alpha_{p,t}$ for every position and date in this disaggregated data. The Internet Appendix provides more details on how we compute $\alpha_{p,t}$ ’s.

Step 3: Aggregation To compute a position’s overall sensitivity to aggregate credit risk, we simply multiply the estimated $\alpha_{p,t}$ by the notional value $\text{Notional}_{p,t}$ of the position. Formally, we define $E_{p,t} \equiv \alpha_{p,t} \times \text{Notional}_{p,t}$ as the position’s exposure to aggregate credit risk. In some applications, we instead define $E_{p,t}$ by multiplying $\alpha_{p,t}$ by what is referred to in practice as a position’s “DV01”. DV01 refers to how much the market value of the position changes in response to a one-basis point move in the credit spread of firm f – it is similar to the concept of “delta” in options markets. Thus, scaling $\alpha_{p,t}$ by DV01 measures how much the position will change in market value for a one-basis point change in the aggregate credit risk factor.¹⁹

Having defined a position’s exposure to aggregate credit risk, $E_{p,t}$, aggregation is straightforward.

¹⁸We also match positions to Markit using the documentation clause and underlying currency of the position.

¹⁹We describe the computation of and calculation of sensitivity to DV01 in the Internet Appendix.

ward. For example, define $S_{i,j,t}$ as the set of positions where i is a seller to j , and $B_{i,j,t}$ as the set of positions where i is a buyer from j , both as of time t . We define the net and gross notional exposures to aggregate credit risk between counterparty i and j as follows:

$$\begin{aligned}\text{Net}_{i,j,t} &\equiv \sum_{p \in S_{i,j,t}} E_{p,t} - \sum_{p \in B_{i,j,t}} E_{p,t} \\ \text{Gross}_{i,j,t} &\equiv \sum_{p \in S_{i,j,t}} E_{p,t} + \sum_{p \in B_{i,j,t}} E_{p,t}.\end{aligned}$$

By construction, positive values of $\text{Net}_{i,j,t}$ mean that i is a net seller of CDS protection on aggregate credit risk to j .

3.2.2 A Measure of Bilateral Concentration

From i 's perspective, j is a concentrated counterparty if most of i 's gross-notional exposure occurs with j . Thus, a natural measure of bilateral concentration is:

$$\kappa_{i,j,t} \equiv \frac{\text{Gross}_{i,j,t}}{\sum_k \text{Gross}_{i,k,t}}.$$

Intuitively, $\kappa_{i,j,t} \in [0, 1]$ and $\sum_j \kappa_{i,j,t} = 1$. In the extreme case, if all of i 's positions are with j , then $\kappa_{i,j,t} = 1$.

Table 1 provides some basic summary statistics of $\kappa_{i,j,t}$, broken out by whether i is a dealer or not. The median customer spreads its trades somewhat equally across trading partners. Based on the actual degree distribution, an equally-weighted benchmark would predict a κ of 7% for the median customer, while we observe a κ of about 9%. There are of course some customers that have more concentrated exposures on average, with the 75th percentile customer trading about 16% of its gross exposure with a single counterparty. Through the lens of the model, this skewness could arise due to the underlying distribution of ω 's and the fact that some customers have less connections than others. Because dealers are inherently more connected than customers, they naturally face more counterparties and so their average bilateral concentration is somewhat smaller at 2%.

3.2.3 A Measure of Bilateral Price Concessions

Armed with a measure of concentration, we now work through a stylized example of how we measure the price concession that counterparty i gives to j in week t . Suppose that the CDS market consists of contracts for only one firm and one maturity. Define $\underline{S}_{i,t}$ as the minimum CDS spread paid by i in transactions where i bought protection in week t . Similarly, $\bar{S}_{i,t}$ is defined as the maximum CDS spread received by i in transactions where i sold protection in week t . Define $S_{k,i,j,t}$ as the CDS spread in trade k between i and j in week t . A natural way to define how much i concedes in

price to j on trade k is then:

$$PriceConcession_{k,i,j,t} = \begin{cases} Spread_{k,i,j,t} \frac{S_{i,t}}{S_{j,t}} & \text{if } i \text{ buys from } j \\ \bar{S}_{i,t} \frac{Spread_{k,i,j,t}}{S_{j,t}} & \text{if } i \text{ sells to } j \end{cases}$$

Naturally, if i purchases protection from j at a high spread relative to other trades where i bought protection, then i is conceding more in price to j than to its other counterparties. And, if i sells protection to j at a low spread relative to other trades where i sold protection, then i is also conceding in price. We aggregate our measure of price concession across all of the trades between i and j in week t in two different ways. First, we take an equal-weighted average of $PriceConcession_{k,i,j,t}$ across the trades k between the two counterparties. We call this metric $PriceConcession_{i,j,t}^{EW}$. Second, we take a notional-weighted average across their trades, where the weights are defined by the notional in each trade. We call this metric $PriceConcession_{i,j,t}^{NW}$.

While the preceding example focused on defining our measure of price concession when there are CDS traded on only one firm and one maturity, the logic of our approach generalizes in a straightforward way to account for trades on different firms and of different maturities by averaging over these dimensions as well. Table 1 provides some basic summary statistics of price concession ($PriceConcession$), again broken out by whether i is a dealer or not. Here, it is readily apparent that there is substantial variation in how much traders concede in price to their counterparties. Our model assumes that part of this variation is driven by differences in bilateral concentration.

3.2.4 Results

We explore the relationship between price concessions and concentration using a panel regression:

$$PriceConcession_{i,j,t} = \frac{S_j}{S_{i,t}} + \alpha \times \kappa_{i,j,t} + 1 + \varepsilon_{i,j,t} \quad (18)$$

where $\frac{S_j}{S_{i,t}}$ is a fixed effect based on whether i is a net seller or net buyer to j in week t . Thus, α in the regression is identified from variation within the set of counterparties to whom i either sells, or buys, protection.

In the Internet Appendix, we show within the model that $\alpha > 0$ if and only if $\kappa < 0$ in regression (18). The intuition for this result is simple: If i already has a highly concentrated position with j coming into week t , then i will concede less in price in any trades that occur with j during the week. Another advantage of our approach to testing if $\alpha > 0$ is that it does not depend on any assumptions on the shape of the network. However, as we also show in the Internet Appendix, a limitation of this test is that it only identifies the sign of α , not its magnitude. We leverage the core-periphery structure of the network to pin down the magnitude of α in subsequent sections.

Table 2 contains the results from running several variants of regression (18). In column (1), we use $PriceConcession_{i,j,t}^{EW}$ as our measure of price concession. The estimated α in this case is -5.39 and we can reject the null hypothesis that it is equal to zero with 95% confidence. The negative coefficient implies that traders concede less in price when facing counterparties with whom

they have large pre-existing exposures. The magnitude of the point estimate indicates that a one-standard deviation increase in bilateral concentration (κ) is associated with about 0.8 basis point less in price concessions. As a point of reference, the average of $PriceConcession_{i,j,t}^{EW}$ is 3.4 basis points. Note that if traders with a relationship have a high $\kappa_{i,j,t-1}$ on average, any price concessions to favored counterparties would cut against finding a negative coefficient.

In column (2), we run the regression for the subset of the data where i is a dealer, which sheds some light on whether dealers are also averse to concentrated bilateral concentration. The estimated coefficient is once again negative and measured with statistical precision, implying that dealers charge their counterparties more when they have concentrated pre-existing positions with them. In terms of magnitudes, the point estimates in column (2) are comparable to column (1) because, as shown in Table 1, $\kappa_{i,j,t}$ is less volatile for dealers than non-dealers.

Columns (3) and (4) contain the results using $PriceConcession_{i,j,t}^{NW}$, which puts more weight on price concessions that occur on large trades. We again observe a similar pattern in that larger bilateral concentration is associated with lower price concession. In columns (5)-(8), we repeat the analysis using a different approach to constructing $\kappa_{i,j,t-1}$. We redefine a position p 's exposure to aggregate credit risk based on its DV01, which as previously discussed provides a sense of how much the position will change in market value if our aggregate credit risk factor increases by 1 basis point. We then aggregate positions in the same way as before to arrive at $\kappa_{i,j,t-1}$. Intuitively, when $\kappa_{i,j,t-1}$ is defined in this manner it measures how much of i 's potential margin payments go to or come from j (see both Section 3.2.1 and the Internet Appendix for more details). As columns (5)-(8) show, the negative relationship between bilateral price concessions and concentration is robust to this alternative approach to construction of κ .

A natural objection to this analysis is that dealers may have market power over some counterparties, particularly those with whom they have large exposures. While this mechanism could explain why dealers concede less in price to more concentrated counterparties, it would also predict that on average dealers concede less in price than non-dealers. However, this view is not supported by the simple summary statistics in Table 1. Overall then, the evidence in Table 2 is consistent with the idea that CDS traders prefer to spread their trades out across multiple counterparties. In the model, this is equivalent to assuming that $\beta > 0$.

3.3 Dealer Exposures and Dealer-Dealer versus Dealer-Customer Spreads

With support for our model's key assumptions in hand, we now turn to a more careful calibration. As demonstrated in Section 2.5, mapping the model to the data is straightforward when the CDS network is core-periphery, as it is in practice. In this case, there is a closed-form relationship between the model's structural parameters and two easily observable objects in the data: (i) the net CDS position of dealers and (ii) the average difference between CDS spreads in dealer-dealer transactions and dealer-customer transactions. Our goal in this subsection is to estimate both.

3.3.1 The Net CDS Position of Dealers

In Section 3.2.1, we denoted the net amount of CDS protection that i has sold to j on date t by $\text{Net}_{i,j,t}$. Recall that $\text{Net}_{i,j,t}$ is positive if i is a net seller to j of CDS protection on aggregate credit risk. Summing this metric across all j yields a measure of i 's net overall exposure to aggregate credit risk:

$$\text{Net}_{i,t} \equiv \sum_j \text{Net}_{i,j,t}. \quad (19)$$

To define the net exposure of dealers, we first compute $\text{Net}_{i,t}$ for each of the 14 dealers in our sample. In the model, net exposure is defined as fraction of total wealth, so we then scale each dealer's net CDS exposure by its market capitalization:

$$z_{i,t} \equiv \frac{\text{Net}_{i,t}}{\text{MktCap}_{i,t}}. \quad (20)$$

We compute $\text{MktCap}_{i,t}$ using the end-of-week market capitalizations of each dealer, though we obtain similar results if use within-week averages or moving averages.²⁰ Here, we use $z_{i,t}$ to denote this scaled-exposure measure to match the notation used in the model.

A key quantity for the mapping between the model and the data is the cross-sectional average exposure of dealers, $\bar{z}_{d,t} = n_d^{-1} \sum_{i \in \mathcal{D}} z_{i,t}$, where $n_d = 14$ is the number of dealers and \mathcal{D} is the set of dealers. Table 5 indicates that the average $\bar{z}_{d,t}$ is about 0.045 across all dates in our sample. One way to interpret this number is as follows: on average, dealers sell \$0.045 of net notional CDS protection on the aggregate credit risk factor for each dollar of their equity.

In the Internet Appendix, we explore an alternative way of defining dealer exposure by computing how the market value of each dealer's CDS portfolio responds to movements in our aggregate credit risk factor. As previously mentioned, this is often referred to as a portfolio DV01 in practice. By this metric, U.S. dealers lose money on their CDS positions if the factor increases, indicating that they are net sellers of protection on the index to their customers. These patterns are also consistent with dealer behavior in U.K. markets (Morrison, Vasios, Wilson, and Zikes 2018).

Within the core set of dealers, net selling of CDS protection is itself very concentrated. To visualize this fact, Figure 3 plots the distribution of z_i within the dealer sector for our sample period. The first thing to notice is that much of the mass is concentrated around zero, meaning many intermediaries have roughly zero net exposure (Atkeson, Eisfeldt, and Weill 2015). Still, the highly concentrated nature of net credit provision is clear from the plot, as the right tail shows that a handful of dealers have substantial positive net positions, indicating that they are large net sellers. The data in Figure 3 are a key input into our quantitative analysis, and to our knowledge the empirical distribution of dealers' net exposures is new to the literature. We explore the implications

²⁰ Market capitalization data are not available for most buy-side customers, such as pension funds. This is not an issue in our baseline model because its calibration only requires us to measure the CDS exposure of *dealers* (who are all listed) relative to their equity.

of this extreme skewness for systemic risk measurement in Section 4.

3.3.2 CDS Spreads: Dealer-Dealer versus Dealer-Customer Trades

We estimate the average spread difference between customer-dealer and dealer-dealer trades using all single name transactions on U.S. firms from 2010 through 2013, excluding those with CDS spreads greater than 1000 bps. We further focus on transactions where the tier of debt in the contract is senior unsecured, and the so-called “doc-clause” that determines CDS trigger events is XR (no restructuring), the most commonly used clause. In addition, we match each transaction in our sample to an associated 5-year spread in the Markit database and a 5-year expected default frequency from Moody’s, both based on the date and underlying firm of the contract.²¹ We winsorize both the transaction spreads and Markit spreads at their five percent tails, but have confirmed that our main conclusions are robust to alternative methods for dealing with outliers. Table 3 displays some simple summary statistics on our transaction panel. The average default probability and loss-given-default are 0.65% and 60.6%, respectively. These statistics will be inputs to our calibration in Section 3.4.

We then run the following panel regression on this sample of transactions:

$$Spread_{k,f,t} = \text{FEs} + \theta_{CD} \times 1_{k,t}(\text{Customer-Dealer}) + \beta_1 \times \text{MarkitSpread}_{f,t} + \beta_2 \times \log(\text{Notional}_{k,t}) + \varepsilon_{k,f,t} \quad (21)$$

where FEs denote a series of fixed effects that we describe below. We use the subscript k to index each transaction. As before, f denotes the underlying firm in transaction k and t denotes the date of the trade. $1_{k,t}(\text{Customer-Dealer})$ is an indicator variable for whether transaction k is between a customer and a dealer. The coefficient θ_{CD} therefore estimates the average difference between spreads in the customer-dealer market versus the dealer-dealer market, which we denote by $\bar{R}_c - \bar{R}_d$.

With a large enough sample, we could estimate θ_{CD} based on spread variation within trades on the same firm, date, and maturity. Unfortunately, our data does not afford us enough power to implement this strategy. Instead, we estimate θ_{CD} based on variation within a ratings class, maturity bucket, and week. Within a week, there will certainly be variation in spreads that is not fully captured by ratings or maturity. We account for this source of variation by controlling for each firm’s daily 5-year CDS spread as reported by Markit. Similarly, the firm fixed effect in the regression accounts for time-invariant differences in *Spread* across firms, which would bias our estimation if customer-dealer trades are somehow concentrated in a given name. Finally, we include the notional amount of trade k to account for any size effects on spreads.

Table 4 presents the results from running regression (21) on various subsamples of our data. Columns (1)-(3) of the table all use our definition of dealers based on the algorithm in Appendix C. Column (1) of the table runs the regression for our full sample and provides our baseline estimate of $\theta_{CD} = 6.14$ bps. The point estimate is measured precisely in a statistical sense and implies that, on

²¹Markit reports composite CDS spreads based on quotes from dealers and actual transactions. For firms that do not have a match in Moody’s, we use the average EDF for the set of firms with the same rating during that week.

average, customer-dealer trades occur at 6.14 bps higher than dealer-dealer trades. In column (2), we find that $\theta_{\psi} = 6.40$ bps in the subsample of trades with a maturity between four and six years, which tend to be the most liquid segment of the CDS market. The similarities between the estimates in columns (1) and (2) also suggest that spread differences in the interdealer and dealer-customer markets are not driven by any maturity effects.²² In terms of economic magnitude, 6.14 bps is a large amount of price dispersion considering that the average level of spreads in our transaction is 133 bps. We provide additional context for the magnitude of θ_{ψ} in the Internet Appendix and when we study dealer removal in Section 4.

Market Power A central motivation behind our model in Section 2 is to understand how network frictions – the combination of a sparse network and aversion to bilateral concentration – interact with price dispersion in equilibrium. However, in reality there are certainly other mechanisms that could generate a gap between CDS spreads in dealer-dealer and dealer-customer trades. One natural candidate is market power: if dealers have market power over their customers, then they might sell CDS at higher spreads to customers relative to other dealers.

In column (3) of Table 4, we attempt to strip out the potential impact of market power on our estimate of θ_{ψ} by running regression (21) on a subset of transactions in which prices are more plausibly competitive. To do so, we filter our transaction panel to include only dealer-dealer trades and trades between dealers and large customers, where large customers are defined as those whose share of gross volume is in the top 10% for our sample period. Further, we require that the Market depth for the firms in the panel is at least eight. Market depth is a data field provided by Markit and gives an indication of how many dealers provide quotes on a given firm and date. The cutoff of eight was chosen because it is the top 10% of Market depth for our sample of trades. The fact that many dealers are providing quotes to large customers supports the idea that dealers are unlikely to possess monopolistic pricing power in this subset of trades. The estimate of $\theta_{\psi} = 5.12$ bps that emerges from this regression is not markedly lower than our estimate from column (1), indicating that a sizable portion of our baseline estimate of $\theta_{\psi} = 6.14$ bps is due to forces that are distinct from strategic pricing. Indeed, a crude decomposition suggests that about 17% ($= 5.12/6.14$) is attributable to market power. Still, to ensure that strategic pricing does not influence our results, we take a conservative approach and use $\theta_{\psi} = 5.12$ when we calibrate the model in Section 3.4.

Additional Robustness The remaining columns in Table 4 display the results when using DTCC’s definition of dealers, as opposed to our definition based on the model-implied minimum-distance algorithm discussed in Appendix C. In this case, the point estimates on the dealer-dealer dummy are quite consistent with the ones we obtain when using our custom dealer definition. Thus, our choice of dealers is largely unimportant when estimating differences in dealer-dealer versus dealer-customer pricing.

²²For example, our estimate of θ would be biased if dealer-customer activity correlates with movements in a firm’s term structure that are orthogonal to: (i) common term structure movements within a ratings class; and (ii) movements in that firm’s own 5-year Markit CDS spread. This is not likely and is also not the case empirically.

3.4 Calibration

Armed with estimates for the net position of dealers and the spread between dealer-dealer and dealer-customer spreads, we are now in a position to infer the model’s two structural risk aversion parameters using Equations (16) and (17). Table 5 presents the data estimates we use for our calibration, along with the resulting parameter values. We set $\mu = 0.39\%$ and $\sigma^2 = 4.87\%$ based on our estimates of loss-given-default and the probability of default from Section 3.3.2.²³ Our estimates of $\bar{R}_d = 133$ bps and $\bar{R}_c - \bar{R}_d = 5.12$ basis points also come from the analysis in that section. To be conservative about the amount of price dispersion which results from network frictions alone, we use the estimate which is based only on dealer-dealer trades and trades between dealers and large customers, where large customers are defined as those whose share of gross volume is in the top 10% for our sample period (see Table 4). We set the number of dealers $n_d = 14$ and the total number of counterparties in the CDS market to $n = 723$, the latter of which is based on our sample of transactions. The only variable on the right-hand side of Equation (16) that we do not observe is $\bar{\omega}$, which is the economy-wide exposure to the underlying default risk. We assume that, on average, agents have one unit of exposure to default risk by normalizing $\bar{\omega}$ to one. This leads to an estimated aversion to aggregate default risk of $\alpha\sigma^2\bar{\omega}^2\mu = 4.37$.²⁴

Next, to calibrate the parameter measuring aversion to bilateral concentration, β , we combine our estimates of price dispersion $\bar{R}_c - \bar{R}_d$ with the net position of dealers, $\bar{z}_d = 0.045$. Table 5 shows that this suggests $\beta = 7.98$. In the Internet Appendix, we show that in the calibrated equilibrium of our baseline model this implies average trading costs of 0.04bps, which is about one twentieth of the size of observed average bid-ask spreads of 50-100bps in normal times post-crisis.²⁵

3.5 How Much Do Network Frictions Impact Prices?

Using our calibrated model, we now quantify the impact that the incomplete network has on equilibrium prices. A natural way to do so is to compute equilibrium prices if the network were instead complete, meaning all agents are connected. Rearranging Equation (16) yields the complete network benchmark price as a function of observables:

$$\bar{R}_{\text{Complete Network}} = \alpha\sigma^2\bar{\omega}^2\mu + \bar{R}_d + (\bar{R}_c - \bar{R}_d)2\left(1 - \frac{n_d}{n}\right). \quad (22)$$

Importantly, Equation (22) allows us to infer what the complete network benchmark price would be based solely on observed prices and the number of counterparties in the network. Based on the parameters in Table 5, we infer that the average CDS spread would be $\bar{R}_{\text{Complete Network}} = 143$ bps if the CDS trading network were complete.

We can then use the implied $\bar{R}_{\text{Complete Network}}$ to quantify the size of network frictions in the market. From Table 5, we see that in the core-periphery network customer-dealer trades occur

²³For a probability of default of $p = 0.0065$ and a loss-given-default of $L = 0.6060$, $\mu = Lp = 0.0039$ is the expected loss rate and $\sigma^2 = L^2p(1 - p) = 0.0024$ is its variance.

²⁴We show that our main results hold for other values of $\bar{\omega}$ in the Internet Appendix.

²⁵See Adrian, Fleming, Shachar, and Vogt (2017).

at spread of about $\bar{R}_c = 138.12$ bps. Thus, credit spreads in the customer-dealer market are about 3.4% lower than they would be if the network were complete. Intuitively, spreads in trades between dealers and customers reflect the average post-trade exposures of dealers and customers, but are tilted towards dealers' marginal cost of risk bearing because of their central role as OTC intermediaries. Since, empirically, dealers are net sellers of credit protection (indicating a higher capacity to bear credit risk than that of customers), spreads are depressed downward from the perfect risk-sharing benchmark. Moreover, in equilibrium, these lower prices reflect the fact that customers must be incentivized to purchase protection from a concentrated set of dealers.

Similarly, we estimate average dealer-dealer spreads in the core-periphery network to be $\bar{R}_d = 133$ bps, nearly 7% lower than they would be if the network were complete. Again, in our model, the empirical fact that dealers are net sellers of credit protection in equilibrium implies that they start with less ex-ante exposure to aggregate credit risk than customers. However, aversion to concentrated exposures prevents dealers from selling as much protection as they would like because all agents trade off the costs of concentrated bilateral exposures against the benefits of default insurance (and exploiting price dispersion). As a result, the model suggests that the post-trade exposure of dealers to aggregate credit risk remains less than that of customers, i.e. that there is imperfect risk sharing. In turn, dealers pay lower spreads when purchasing credit protection from other dealers in equilibrium because, as a group, their lower ex-ante exposure results in a higher risk-bearing capacity and a lower shadow cost of risk bearing.

More broadly, the fact that CDS spreads in both the interdealer and customer-dealer markets are lower than the complete network benchmark reflects *worse* risk sharing in equilibrium, and a tilt toward dealers' (who are net sellers of CDS) lower marginal costs of risk bearing. The lower premiums can be interpreted as compensation for concentrated bilateral exposures – absent these costs, agents would compete away any price dispersion.

These results provide further evidence that small trading costs, on the order of 5% of average bid-ask spreads, combined with the core-periphery network, can have large effects. In addition to driving price dispersion $\bar{R}_c - \bar{R}_d$ of 5.12 bps, network frictions lead to customer prices that are 4.88 bps lower, and dealer prices that are 10bps lower than in the complete network benchmark. In the next section, we build on the results in Proposition 1 and use our calibrated model to show that the direction of the price distortion depends crucially on the shadow cost of risk bearing in the dealer core relative to the customer periphery. If the dealer sector's ability to bear aggregate default risk falls below that of the customer sector, perhaps due to dealer exit, then prices are distorted upward.

4 Dealer Removal

Core-periphery networks are often thought to be susceptible to systemic risk. For example, in many accounts of the 2008 financial crisis, the failure of Lehman Brothers caused dislocations in several OTC markets as traders were forced to replace Lehman as a counterparty. Indeed, CDS spreads

uniformly rose immediately after Lehman filed for bankruptcy on September 15, 2008.²⁶ Motivated by the general concern about concentrated dealer markets, as well as the dramatic events in OTC markets in 2008, we now use the calibrated model to assess how the CDS market would respond to the failure of a central dealer in the network.

Specifically, we remove one dealer from the core holding fixed all other model parameters, including each agent’s pre-trade exposure to credit risk. We further assume that the removed dealer’s initial endowment of credit risk exposure (ω) is not reallocated to the remaining dealers after it exits. This implies that the post-dealer-removal risk-bearing capacity of the dealer sector, and overall market, is determined by the remaining traders’ risk-bearing capacities, as determined by their initial exposures to credit risk. We then allow all remaining agents in the model to re-trade in the CDS market and study the resulting new equilibrium. Our stress test thus assesses how markets would respond if an exiting dealer’s CDS positions needed to be absorbed by the remaining market participants.

It is easy to imagine how such a situation could arise in practice. For example, suppose that a dealer is short credit risk, for example from sales of structured credit products, and that it offsets that short exposure with a long credit exposure from selling CDS. Now imagine that this dealer fails, perhaps due to a shock that starts outside of corporate credit markets (e.g. in residential mortgage markets). In a frictionless world, another dealer could potentially absorb all of the failed dealer’s cash and derivatives positions without taking on substantial additional risk. In reality, reallocation in markets for standardized derivatives contracts may occur much faster than markets for more heterogenous cash assets such as structured products, loans, or bonds. Consequently, the failed dealer’s initial short credit position (i.e., their ω and hence their risk-bearing capacity) would be temporarily removed from the CDS market as it re-equilibrates. This medium-run scenario is exactly what our stress test is designed to simulate.

4.1 Dealer Failure in the Baseline Model

A salient feature of Figure 3 is that the sale of CDS protection is highly concentrated. A few key institutions within the dealer sector contribute substantially to the total credit insurance provided by the dealer sector. Building off of this fact, we first use our model to simulate the removal of the largest net-selling dealer and summarize the resulting equilibrium in Table 6. In the table, under each scenario, we report the average dealer-dealer price \bar{R}_d , dealer-customer price \bar{R}_c , and the net CDS position of the remaining dealers \bar{z}_d . In addition, we report the resulting spread under each removal scenario if the network were complete, as opposed to core-periphery. As a benchmark, column (1) of the table shows the same statistics in the baseline model.

In our baseline calibration, removing the largest net-selling dealer would cause spreads in the dealer-dealer market to rise approximately 24% from 133 to 164.31 basis points. The increase in

²⁶While some of those price movements were undoubtedly driven by news about aggregate default risk, the fact that the CDS-Bond basis also rose suggests that Lehman’s failure had an impact on CDS spreads that was distinct from pure default risk considerations.

spreads occurs despite the fact that there is no realized default, no contagion, and agents can trade to reallocate risk after the dealer fails. To better understand the mechanism driving this result, we can approximate the price-impact in the dealer-dealer market by combining Equations (11) and (12) as follows:

$$\Delta \bar{R}_d \approx - \frac{(1 - \alpha) \sigma^2 \psi}{d} \Delta \bar{z}_d, \quad (23)$$

where the approximation comes from the fact that we are ignoring the small effect that reducing the number of dealers has on \bar{z}_d and $\bar{\omega}$. Equation (23) says that the impact of dealer-removal on the inter-dealer market is negatively and linearly related to the change in the net position of the dealer sector. When the largest net seller exits, so to does its capacity to absorb credit risk via CDS (its low ω_i and resulting high risk-bearing capacity). Consequently, without the largest net seller, the overall dealer sector reduces the amount of CDS protection it sells (i.e. $\Delta \bar{z}_d < 0$), thereby causing dealer-dealer spreads to rise. In fact, Figure 3 indicates that most dealers have a low capacity to absorb credit risk via CDS and are therefore net buyers of protection. This means that removing the largest net seller strongly impacts \bar{z}_d . Column (2) of Table 6 shows that the impact is large enough that the dealer sector overall switches from being a net seller of CDS protection to being a net buyer of credit insurance.

Equation (23) also highlights how the size of trading cost α affects the quantitative impact of dealer removal in the our baseline model. To see why, recall from Section 2.4 that α approaches one as β approaches zero. In this limiting case, removing a dealer would have a smaller impact on the level of spreads, regardless of its impact on dealer-sector-level risk-bearing capacity, because risk sharing between dealers and customers improves as β declines. As β declines customers play a more important role in risk sharing and the dealer sector becomes relatively less important for pricing credit risk. Conversely, if β is large and the network is core-periphery, the level of spreads in the economy is more sensitive to the risk-bearing capacity of the core because the core plays a larger role in market-wide risk-bearing capacity.

Heterogeneity in risk-bearing capacity ω_i across dealers plays a very important role in our model. Even without frictions, spreads could in principle react to the failure of a key dealer because of the resulting change in aggregate risk bearing capacity. We measure the strength of this channel by simulating the failure of the largest net-selling dealer when the network is complete. The last row of column (2) shows that spreads are minimally impacted in this case. The reason is that the failed dealer contributes very little to aggregate risk-bearing capacity, which in the complete network (or no trading cost) case is the relevant state variable for prices. The key difference between the complete network/no trading cost model and our model is that trading frictions between the dealer and customer sectors, and lack of connectivity among customers, implies that the risk-bearing capacity of the core, vs. the aggregate, drives the pricing of risk. And, individual dealers can meaningfully contribute to the dealer sector's ability to bear risk. Thus, dealer heterogeneity plays a crucial role in our OTC setting due to trading frictions between dealers and customers, combined

with the incomplete network. We elaborate on this below when we extend our model to consider a frictionless inter-dealer market.

We further characterize the role of dealer heterogeneity by removing dealers who are drawn from various points in the empirical distribution of net CDS positions (see Figure 3). The resulting spread impact is reported in columns (3)-(5) of Table 6. Column (3) presents results for removing a dealer at the 90th percentile of the z_d distribution. Removing this dealer leads to an increase in spreads of about 5 bps in the dealer-dealer sector. The effect is modest because this dealer does not sell very much CDS protection and is therefore a small portion of the core’s total CDS supply. Column (4) indicates that a similar outcome occurs when removing a dealer who is essentially net-neutral in its CDS positions and thus functions more like a true intermediary. Next, in column (5), we remove the dealer who is the largest net buyer of protection. Removing this dealer leads to lower average spreads in both dealer and customer markets. As can be seen from Equation (23), this occurs because removing a large buyer of protection from the dealer sector increases the available risk-bearing capacity of the remaining dealers.

Dealer removal also has an impact on price dispersion, $\bar{R}_d - \bar{R}_c$, in our model. Column (2) of Table 6 shows that when we remove the largest net seller from the core, dealer-customer spreads increase by 11%, compared to 24% for dealer-dealer spreads. As a result, dealer-customer spreads are now lower than those in the inter-dealer market. This reversal reflects the fact that after removing the largest net seller, the remaining dealers become net buyers of protection from customers. In other words, the remaining dealers have less ability to bear credit risk than customers and in equilibrium will pay a higher premium for protection. The effect can be seen directly by combining Equations (11) and (14):

$$\bar{R}_c = \bar{R}_d + \frac{1}{2}\alpha\sigma^2\left(1 - \frac{n_d}{n}\right)\left[1 + \left(1 - \frac{n_d}{n}\right)\frac{n_d}{n} - \frac{n_d}{n}\frac{\bar{z}_d}{d}\right]. \quad (24)$$

As shown in Proposition 1, when the customer becomes a net insurance provider to the dealer sector ($\bar{z}_d < 0$) – as it does when we remove the largest net seller – spreads in the inter-dealer market exceed those in the dealer-customer market.

In summary, our stress tests show that the removal of an OTC intermediary can have a sizable impact on equilibrium outcomes, even when there is no contagion or other knock-on effects from dealer exit. The reason is simple: when customer connectivity is limited to dealers and building bilateral positions is costly, prices depend heavily on the risk-bearing capacity of core dealers. This aspect of our model accords with research that emphasizes the impact of the aggregate dealer sector on asset prices (He and Krishnamurthy (2013)).²⁷ Relative to that literature, however, our analysis suggests that the financial soundness and net positions of dealers *as a sector* is not sufficient for understanding their potential impact on markets. Due to the widely varying roles of dealers, as net demanders or suppliers of credit insurance, individual dealers’ presence or absence can push prices

²⁷There are other channels, such as information revelation (Babus and Kondor (2018), Glode and Opp (2016)), or price impact (Malamud and Rostek (2017)), through which dealers could impact pricing. We have abstracted away from these mechanisms in order to focus on how and when dealer risk-bearing capacity impacts market outcomes.

in different directions.

To see why, recall the thought experiment discussed at the outset of the paper: should regulators monitor and collect data on the activity of individual OTC dealers? Our model suggests there is value in doing so because heterogeneity within the core is critical for predicting the impact of dealer failure on markets. The argument for regulators to monitor intermediaries as a group would be stronger if each was net neutral in CDS, though this is not the case empirically.

Our findings also imply that the risk-bearing capacity of dealers may not be fully captured by their net worth, as is usually assumed in the intermediary-based asset pricing literature (Adrian, Etula, and Muir 2014). Dealers in our model sell protection if they start with low exposures to aggregate credit risk. Differences across dealers in initial risk exposures could be driven by factors unrelated to size or net worth, such as heterogeneity in their cash asset (non-derivative) business lines. In particular, dealers that sell structured products inherit a short credit exposure from that business, while dealers with large bond portfolios enter the derivatives market with a long credit position.²⁸ This observation highlights why data on dealer-level derivative and non-derivative positions can be useful for measuring the systemic importance of dealers in OTC markets. In addition, we show below that heterogeneity in risk aversion can also lead to heterogeneity in dealer risk-bearing capacity per unit of net worth. The fact that business-induced ex-ante exposures and/or heterogeneity in risk aversion or risk-management capabilities may play a part in determining dealer risk-bearing capacity suggests that refinements to the simple equity-weighted measures of intermediary net worth used in asset pricing may improve their empirical performance.

Sensitivity to estimates In our model, the impact of dealer exit on markets depends heavily on the magnitude of trading costs, τ . As discussed in Section 3.4, the estimate of τ that we use in our baseline dealer-removal analysis is based on the amount of observed price dispersion ($\bar{R}_d - \bar{R}_c$) that is observed between dealers and customers in the top 10% of the trading volume distribution. We focused on this subset of transactions as a way to minimize the impact of imperfect competition on our estimates. For robustness, in Figure 4, we plot the impact of removing the largest dealer against the amount of price dispersion attributable solely to network frictions. In the plot, lower values of price dispersion implicitly correspond to lower values of τ in our calibration and hence a smaller impact of removing the largest net seller. The figure reveals that even small network frictions τ can still amplify shocks to the risk-bearing capacity of the core. For example, even if only half (roughly 2.5 basis points) of the observed 5.12 bps of price dispersion in the data is attributable to network frictions, removing the largest net-seller would still increase the level of spreads by about 15 basis points, or over 10%.

²⁸Cross-sectional variation in z_i is largely unexplained by dealer leverage (book assets-to-market equity) during our sample. This can be seen by running a panel regression of z_i on a time fixed effect and dealer leverage, which yields a low within- R^2 and a statistically insignificant slope coefficient.

4.2 Preference Heterogeneity and Dealer Failure

The baseline model used thus far has two main sources of heterogeneity: position in the network (i.e. core vs periphery) and initial exposure to aggregate default risk. This formulation of the model is useful because it can be solved in closed-form, which allows us to transparently fit key elements of the data and characterize when and why the failure of a dealer can impact prices. Thus, this baseline provides a workhorse for quantitative applications in OTC markets. In this subsection, we consider a few extensions of the baseline model that allow for heterogeneity in two other key parameters, risk aversion and bilateral trading costs. We then study how these extensions impact the pricing effects of dealer exit.

4.2.1 Risk Aversion versus Initial Endowments

In Appendix B.1, we extend the baseline model by allowing both risk aversion α_i and initial endowments ω_i to vary across agents. That is, each agent now solves the following portfolio optimization problem:

$$\max_{\{x_{ij}\}_{j=1}^n, z_i} \omega_i(1 - \mu) + \sum_{j=1}^n x_{ij}(R_{ij} - \mu) - \frac{\alpha_i}{2} (\omega_i + z_i)^2 - \frac{1}{2} \sum_{j=1}^n x_{ij}^2. \quad (25)$$

Within this general setup, we show an equivalence between a model with heterogeneous α_i and one with heterogeneous ω_i . Specifically, for a fixed initial endowment $\omega_i = \omega$, there is a distribution of α_i that can match any observed set of dealer-dealer prices, dealer-customer prices, and net CDS positions z_i .²⁹ Conversely, for any fixed $\alpha_i = \alpha$, there is a distribution of ω_i that can do the same.³⁰ We leverage this latter property when calibrating our baseline model. More generally, the equivalence result implies that the model with fully heterogeneous α_i and ω_i is not fully identified by our data on CDS pricing, positions, and transactions. Intuitively, this is because we only observe CDS positions, and a trader who purchases large amounts of CDS protection could do so because of hedging motives, risk aversion, or both.

The fact that α_i and ω_i are not separably identifiable with our data does not mean the distinction between risk aversion and initial default exposure is not economically meaningful. To see why, imagine that a regulator is trying to determine whether a set of dealers is likely to fail or to become impaired in the event of default. Further suppose that agents differ in risk aversion but have the same initial exposure to default risk. In this case, dealers who sell a large amount of protection are

²⁹ In order to fully calibrate the heterogeneous risk aversion model, we must also make an assumption about the distribution of net worth in the customer sector. This is because z_i in the model is expressed as i 's net notional CDS position per dollar of net worth, which is difficult to measure for customers. We therefore make the simplifying assumption that each customer's net worth is proportional to the size of its net notional position. For instance, if one customer sells twice as much protection as another, we assume it has twice the net worth. In principle, a regulator with enough data could simply assess the size of each customer's balance sheet to handle this measurement issue more carefully.

³⁰ Formally, Proposition A4 in Appendix B.1 shows that we can use either risk aversion or pre-trade exposure heterogeneity to match net positions of agents. To match prices, we calibrate the average risk aversion and the trading cost parameters.

fairly vulnerable because they are taking on net default risk by selling CDS protection to agents with high risk aversion. On the other hand, if agents differ only in their initial exposure to default risk, then dealers who sell large amounts of protection may not be so vulnerable because they are using the CDS market for hedging purposes – that is, they are short corporate credit risk in other parts of their portfolio and use the CDS market to achieve a more neutral position.

Importantly, the empirical distinction between $\alpha\psi$ and ω_i is less relevant for us because we are not trying to predict which dealers are vulnerable to realized default. Instead, we ask a related, yet conceptually different question: how would OTC markets react if a dealer were to fail? Our analysis suggests the answer depends mainly on whether the exiting dealer contributes positively or negatively to the core’s overall risk-bearing capacity, not whether it does so because of its risk aversion or initial risk exposure. We illustrate this point more formally in Appendix B.1, where we show that the model with heterogeneous risk aversion (α) delivers dealer removal effects that are quantitatively similar to our baseline model.

4.2.2 Bilateral Trading Costs

Next, we consider an extension extension of the baseline model in which trading costs are lower in the inter-dealer market than they are when dealers trade with customers. This type of cost differential could arise if adverse selection concerns due to information asymmetry (Kyle (1985)) are higher in dealer-customer transactions, for example because hedge funds trade based on speculative motives while dealers trade to manage risk. The limit as inter-dealer trading costs tend to zero is a natural benchmark to consider because it implies a frictionless inter-dealer market, which is a standard assumption in search-based models of OTC markets following the classic work of Du e, Gârleanu, and Pedersen (2005). This extension also allows us to highlight the different roles of dealer-dealer vs. dealer-customer trading frictions in OTC markets.

We formally model heterogeneity in trading costs by adjusting each agent’s optimization problem as follows:

$$\max_{\{i_j\}_{j=1}^n, z_i} \omega_i(1 - \mu) + \sum_{j=1}^n i_j(R_{ij} - \mu) - \frac{\alpha\psi}{2}(\omega_i + z_i)^2 - \sum_{j=1}^n \frac{i_j}{2} \frac{z_j}{i_j} \quad (26)$$

where $i_j = d$ if both i and j are dealers and $i_j = c$ otherwise. To understand the quantitative implications of this extension, we focus on the $d \approx 0$ case, which corresponds to a frictionless inter-dealer market. For any value of d , we can back out the values of c and $\alpha\psi$ that are necessary to match the observed level and dispersion of spreads, just as we did in Section 3.4. We discuss the details of the model and calibration at length in Appendix B.2.

Interestingly, as we vary d , parameters c and $\alpha\psi$ do not need to adjust in order to match the level and dispersion in spreads. In fact, in the appendix, we show that the calibration of $\alpha\psi$ and c as outlined above delivers exactly the same parameter values as in our benchmark calibration.³¹

³¹See Proposition A5 in Appendix B.2.

Also, the average pre-trade exposure of dealers, $\bar{\omega}_d$, is identical to the one in our baseline model at calibrated parameters.

The reason for this result is that trading costs between dealers do not impact their total CDS position as a group in equilibrium, but do alter the distribution of net CDS positions within the core. By contrast, risk aversion, trading costs between customers and dealers, along with the relative risk-bearing capacity of the two groups, are a central determinant of the average level of spreads, how much protection is sold by dealers, and how spreads differ on average in dealer-dealer and dealer-customer trades. Thus, α , $\bar{\omega}_c$, $\bar{\omega}_d$, and $\bar{\omega}_c$ do not change because we continue to match the average net position of dealers and average spreads in the dealer-dealer and dealer-customer markets.

The preceding equilibrium logic also highlights why the calibrated distribution of ω_i across dealers varies with λ_d . Intuitively, when trading frictions are higher within the dealer sector, a wider distribution of ω_i s is required to match the observed distribution of CDS purchases in sales. For instance, for a dealer to purchase a lot of protection, its desire to hedge credit risk must be high enough that it is willing to pay the large utility cost λ_d of trading. Similarly, a dealer who sells lots of protection must start with a relatively low exposure to default risk, ω . On the other hand, when the cost of trading between dealers is low, the distribution of ω_i s does not need to be as dispersed in order to match the empirical distribution of net CDS positions.

Table 7 summarizes the effects of removing dealers from different points in the distribution of net positions in the case that λ_d is negligible, and equal to 10^{-8} . In this calibration dealers have numerically indistinguishable post-trade net positions – that is, they share risk perfectly. We simulate dealer failure in this setting based on the same procedure used in Section 4.1. In our baseline model, dealer-dealer spreads and dealer-customer spreads rise by 31 bps (24%) and 16 bps (11%), respectively, when the largest net-selling dealer is removed. In the case of a frictionless inter-dealer market, dealer-dealer and dealer-customer spreads rise by 10 bps (8%) and 5 bps (4%), respectively. While the impact of failure on spreads in both markets is diminished by a factor of about 3, it is not zero as in the case of a complete network or no trading costs. The reason is twofold. First, even with perfect interdealer risk-sharing, the observed large amount of protection sold by the largest net-selling dealer implies enough dealer heterogeneity that losing this dealer’s risk-bearing capacity still causes spreads to increase by a material amount. Second, trading costs between dealers and customers prevent customers from replacing the (still substantial) lost capacity from this dealer.

The dampened effect from the loss of the largest net-selling dealer is driven entirely by the fact that the dispersion in dealers’ pre-trade exposures is lower when λ_d is low. When λ_d is small, the exiting dealer’s ω_i s are less extreme compared to the rest of the dealer sector. In turn, there is a smaller change in the sector’s risk bearing capacity when this dealer exits. When λ_d is equal to λ_c as in our baseline model, removing the largest net selling dealer switches the dealer sector overall from being a net seller of credit insurance to demanding insurance from the customer sector. In contrast, when $\lambda_d \approx 0$, the change in the net position of dealers is still sizable, however it is not

enough to change the role of the dealer sector as a net provider of credit insurance to the market overall.

The impact of removing the largest net-*buying* dealer is also dampened in the case of a frictionless inter-dealer market. The lower dispersion in initial exposures implied by observed net CDS trades given the lower level of inter-dealer trading costs makes both the net sellers and net buyers less extreme in terms of their positions. Removing the largest net demander of credit protection within the dealer sector still leads to a decrease in spreads as the dealer sector gains in overall risk-bearing capacity through the loss of this dealer, however the effect is more modest than in the baseline model.

Overall, the extended model with heterogenous trading costs reveals two unique insights relative to the prior literature on OTC markets and intermediary asset pricing. First, frictions between dealers and customers are enough to lead to large effects from the loss of a single dealer, even if dealers can perfectly share risk with one another. Second, dealer heterogeneity is crucial for generating these effects. The data on net positions reveals substantial heterogeneity in different dealers' roles in CDS markets. There is no effect from removing *any* dealer when the network is complete or trading costs are zero. Moreover, when the network is core-periphery and trading costs differ across agents, the effect of dealer removal depends critically on the identity of the dealer that exits the market.

4.3 Dealer Failure and Increases in Fundamental Risk

Building off of the preceding analysis, we now use the calibrated model with heterogeneous λ to study how the exit of a dealer would impact CDS markets if it were also accompanied by a change in preferences. We start from the calibrations in Tables 6 and 7. These scenarios correspond to our baseline model in which $\lambda_d = \lambda_c$ and our extended model in which $\lambda_d \approx 0$. We then remove the dealer who sells the most protection and simultaneously increase the cost of customer-dealer trades λ_c and risk aversion α . An increase in λ_c could be driven by heightened concerns about counterparty risk exposure or adverse selection. An increase in α could occur if the shock that led the dealer to exit also impaired the balance sheets of remaining traders. Both types of preference shifts seem plausible in a full-blown economic crisis like the fall of 2008 or March 2020.

Table 8 reports the resulting impact on dealer-dealer spreads and the top and bottom panels correspond to the baseline $\lambda_d = \lambda_c$ model and the extended model in which $\lambda_d \approx 0$, respectively. The first thing to notice is that even small changes in risk aversion α can dramatically amplify the impact of dealer failure. In our baseline model where $\lambda_d = 7.78$, spreads would rise 48 bps upon dealer exit if risk aversion α were to increase 14% from 4.37 to 5. For low values of λ_d (Panel B), the same change in risk aversion would cause spreads to rise by 25 bps, compared to a 10 bp change without any associated shift in risk aversion.

Interestingly, Table 8 shows that for a given level of α , a simultaneous rise in λ_c has a much smaller amplifying effect. The reason network frictions (λ_c) have a small additive impact on prices is because there is less risk to be shared between customers and dealers after the largest net-selling

dealer fails. To see why, consider the case where \bar{d} is low in Panel B. In this case, prior to failure, the average pre-trade exposure in the dealer sector is below the economy-wide average, i.e. $\bar{w}_d = 0.86 < 1$. This is largely driven by the one dealer – the one whose failure we simulate – and is also why dealers sell protection on average. After this dealer exits, the remaining dealers have an average exposure of close to 1. In other words, their pre-trade exposure to default risk is very close to that of customers. Consequently, there is not much risk to be shared via the CDS market, and network frictions τ_c do not have much scope to impact pricing.

5 Conclusion

We develop a model of pricing in OTC markets. We emphasize two key trading frictions, namely, network incompleteness (i.e. a core-periphery structure), and bilateral trading costs. We show that the model’s main assumptions are supported in the data using detailed transaction and position-level data from DTCC on credit default swaps. We then use the key pricing and quantity equations from the closed-form solution to our model in equilibrium to calibrate our model to this data. Our calibrated model then allows us to answer two important questions in OTC asset pricing quantitatively.

First, we show how network frictions distort OTC prices away from their Walrasian, or complete network, benchmark. As long as the network is incomplete, and aversion to concentrated bilateral positions limits risk sharing, there is always a distortion. However, the sign depends on the relative risk bearing capacity of OTC intermediaries, or the dealer sector of our core-periphery network. If the dealer has a higher risk-bearing capacity than the customer sector, then network frictions distort prices downward, toward dealers’ shadow price of risk. By contrast, if the risk-bearing capacity of the OTC intermediary sector becomes impaired, for example through the loss of a key dealer, then prices are distorted upward. As a result, we argue that network frictions can lead to changes in credit spreads that do not rely on changes in parameters governing risk or risk aversion.

Second, we use our calibrated model as a laboratory for regulators to evaluate the systemic importance of key OTC intermediaries. We measure the impact of dealer removal on prices and risk sharing in the OTC network. We show that the loss of a systemically important dealer causes a large increase in credit spreads. Crucially, this systemically important dealer is not only identified by its centrality as defined by measures of connectedness, as all dealers are fully connected to all other agents in our model. The key features of the most systemically important dealer in our model are that they are both a dealer (fully connected) *and* a large net provider of credit insurance. We find empirically that different dealers play very different roles within the intermediary sector. A small number of dealers are responsible for the fact that the dealer sector as a whole is a net provider of credit insurance, while many dealers maintain net neutral positions or are net buyers of protection. The effect of dealer removal on prices depends critically on their net position, and removing a net demander of insurance can actually lead to lower prices. Importantly, we show that the effects of dealer removal on prices are entirely due to network frictions. If the network is

complete, or if there are no bilateral trading costs, then there is no effect of removing any dealer on prices or risk sharing.

References

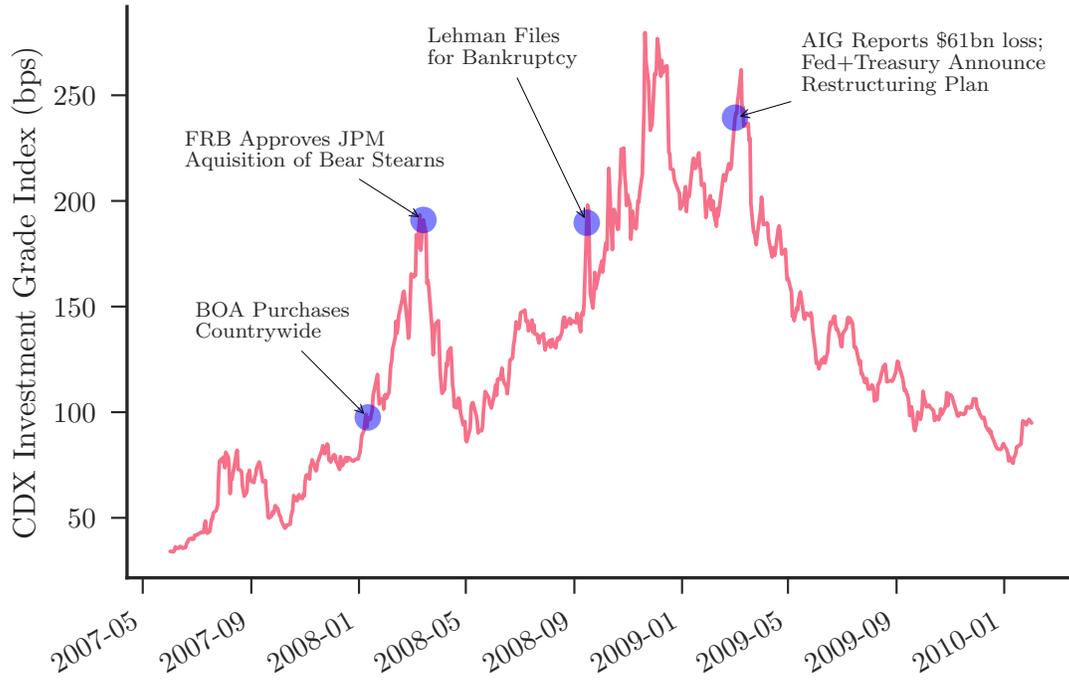
- ACEMOGLU, D., A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2015): “Systemic risk and stability in financial networks,” *American Economic Review*, 105(2), 564–608.
- ADRIAN, T., E. ETULA, AND T. MUIR (2014): “Financial Intermediaries and the Cross-Section of Asset Returns,” *The Journal of Finance*, 69(6), 2557–2596.
- ADRIAN, T., M. FLEMING, O. SHACHAR, AND E. VOGT (2017): “Market liquidity after the financial crisis,” *Annual Review of Financial Economics*, 9, 43–83.
- AKERLOF, R., AND R. HOLDEN (2016): “Movers and shakers,” *The Quarterly Journal of Economics*, 131(4), 1849–1874.
- ALLEN, F., AND D. GALE (2000): “Financial contagion,” *Journal of political economy*, 108(1), 1–33.
- ATKESON, A. G., A. L. EISFELDT, AND P.-O. WEILL (2015): “Entry and Exit in OTC Derivatives Markets,” *Econometrica*, 83(6), 2231–2292.
- BABUS, A., AND P. KONDOR (2018): “Trading and information diffusion in over-the-counter markets,” *Econometrica*, 86(5), 1727–1769.
- BALASUBRAMANIAM, S., A. GOMES, AND S. LEE (2019): “Asset Reallocation with Intermediaries Under Selling Pressure,” *Working Paper*.
- COLLIARD, J.-E., T. FOUCAULT, AND P. HOFFMANN (2018): “Inventory Management, Dealers’ Connections, and Prices in OTC Markets,” *SSRN*: <https://ssrn.com/abstract=3211285>.
- CONT, R., AND A. MINCA (2016): “Credit default swaps and systemic risk,” *Annals of Operations Research*, 247(2), 523–547.
- DENBEE, E., C. JULLIARD, Y. LI, AND K. YUAN (2014): “Network risk and key players: a structural analysis of interbank liquidity,” *Working Paper*.
- DI MAGGIO, M., A. KERMANI, AND Z. SONG (2017): “The value of trading relations in turbulent times,” *Journal of Financial Economics*, 124(2), 266–284.
- DU, W., S. GADGIL, M. BRODY, AND C. VEGA (2017): “Counterparty Risk and Counterparty Choice in the Credit Default Swap Market,” *Working Paper*.
- DUFFIE, D., N. GÂRLEANU, AND L. H. PEDERSEN (2005): “Over-the-counter markets,” *Econometrica*, 73(6), 1815–1847.
- EISENBERG, L., AND T. H. NOE (2001): “Systemic risk in financial systems,” *Management Science*, 47(2), 236–249.

- FARBOODI, M. (2017): “Intermediation and Voluntary Exposure to Counterparty Risk,” Working Paper.
- GAVAZZA, A. (2016): “An empirical equilibrium model of a decentralized asset market,” *Econometrica*, 84(5), 1755–1798.
- GLODE, V., AND C. OPP (2016): “Asymmetric information and intermediation chains,” *American Economic Review*, 106(9), 2699–2721.
- GOLUB, B., AND Y. LIVNE (2010): “Strategic Random Networks and Tipping Points in Social Institutions,” Working Paper.
- HADDAD, V., AND T. MUIR (2018): “Do intermediaries matter for aggregate asset prices?,” *NBER Working Paper*.
- HAU, H., P. HOFFMANN, S. LANGFIELD, AND Y. TIMMER (2018): “Discriminatory Pricing of Over-the-Counter Derivatives,” *SSRN: <https://ssrn.com/abstract=3095575>*.
- HE, Z., AND A. KRISHNAMURTHY (2013): “Intermediary Asset Pricing,” *American Economic Review*, 103(2), 732–70.
- HENDERSHOTT, T., D. LI, D. LIVDAN, AND N. SCHÜRHOFF (2017): “Relationship trading in OTC markets,” *SSRN: <https://ssrn.com/abstract=3051859>*.
- HERSKOVIC, B. (2018): “Networks in production: Asset pricing implications,” *The Journal of Finance*, 73(4), 1785–1818.
- HERSKOVIC, B., B. T. KELLY, H. LUSTIG, AND S. VAN NIEUWERBURGH (2017): “Firm Volatility in Granular Networks,” Working Paper.
- HOLLIFIELD, B., A. NEKLYUDOV, AND C. SPATT (2017): “Bid-ask spreads, trading networks, and the pricing of securitizations,” *The Review of Financial Studies*, 30(9), 3048–3085.
- KYLE, A. S. (1985): “Continuous auctions and insider trading,” *Econometrica: Journal of the Econometric Society*, pp. 1315–1335.
- LEWIS, K. F., F. A. LONGSTAFF, AND L. PETRASEK (2017): “Asset mispricing,” Discussion paper, National Bureau of Economic Research.
- LI, D., AND N. SCHÜRHOFF (2018): “Dealer networks,” *Forthcoming, The Journal of Finance*.
- MALAMUD, S., AND M. ROSTEK (2017): “Decentralized Exchange,” *American Economic Review*, 107(11), 3320–62.
- MORRISON, A., M. VASIOS, M. WILSON, AND F. ZIKES (2018): “Identifying Contagion in a Banking Network,” *SSRN: <https://ssrn.com/abstract=2904966>*.

- OEHMKE, M., AND A. ZAWADOWSKI (2016): “The anatomy of the CDS market,” *The Review of Financial Studies*, 30(1), 80–119.
- PELTONEN, T. A., M. SCHEICHER, AND G. VUILLEMEY (2014): “The network structure of the CDS market and its determinants,” *Journal of Financial Stability*, 13, 118–133.
- SAITA, F. (2007): *Value at Risk and Bank Capital Management*. Academic Press.
- SIRIWARDANE, E. N. (2018): “Limited Investment Capital and Credit Spreads,” *Forthcoming, The Journal of Finance*.

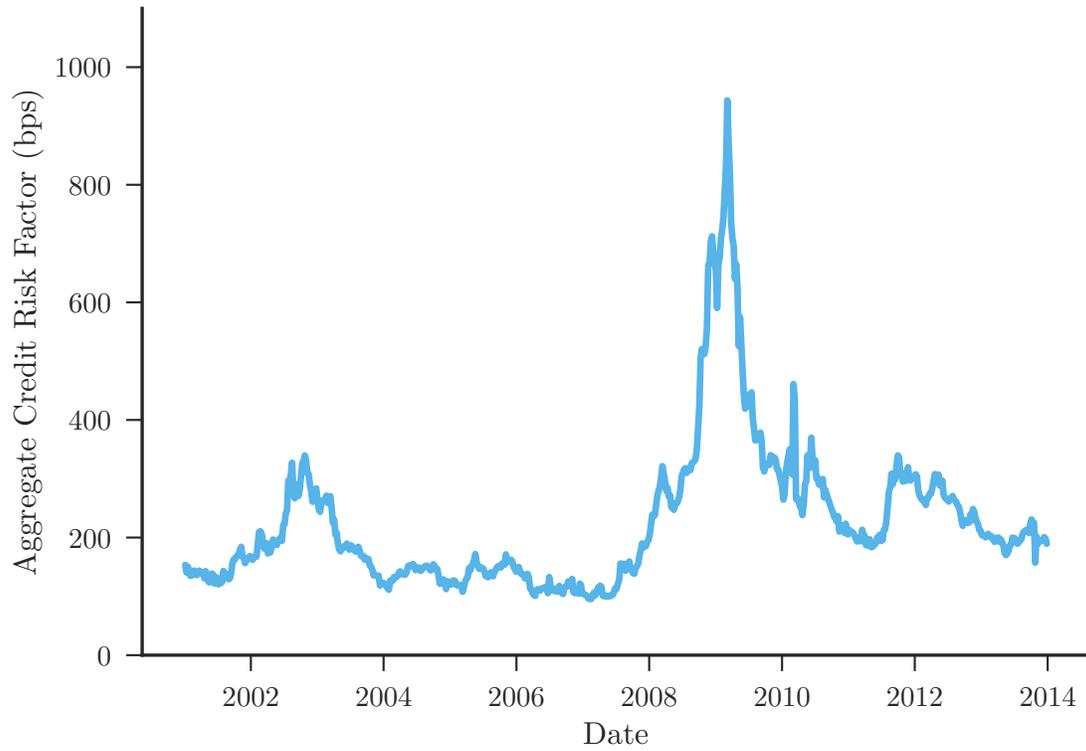
Figures

Figure 1: CDS Markets During the 2008 Global Financial Crisis



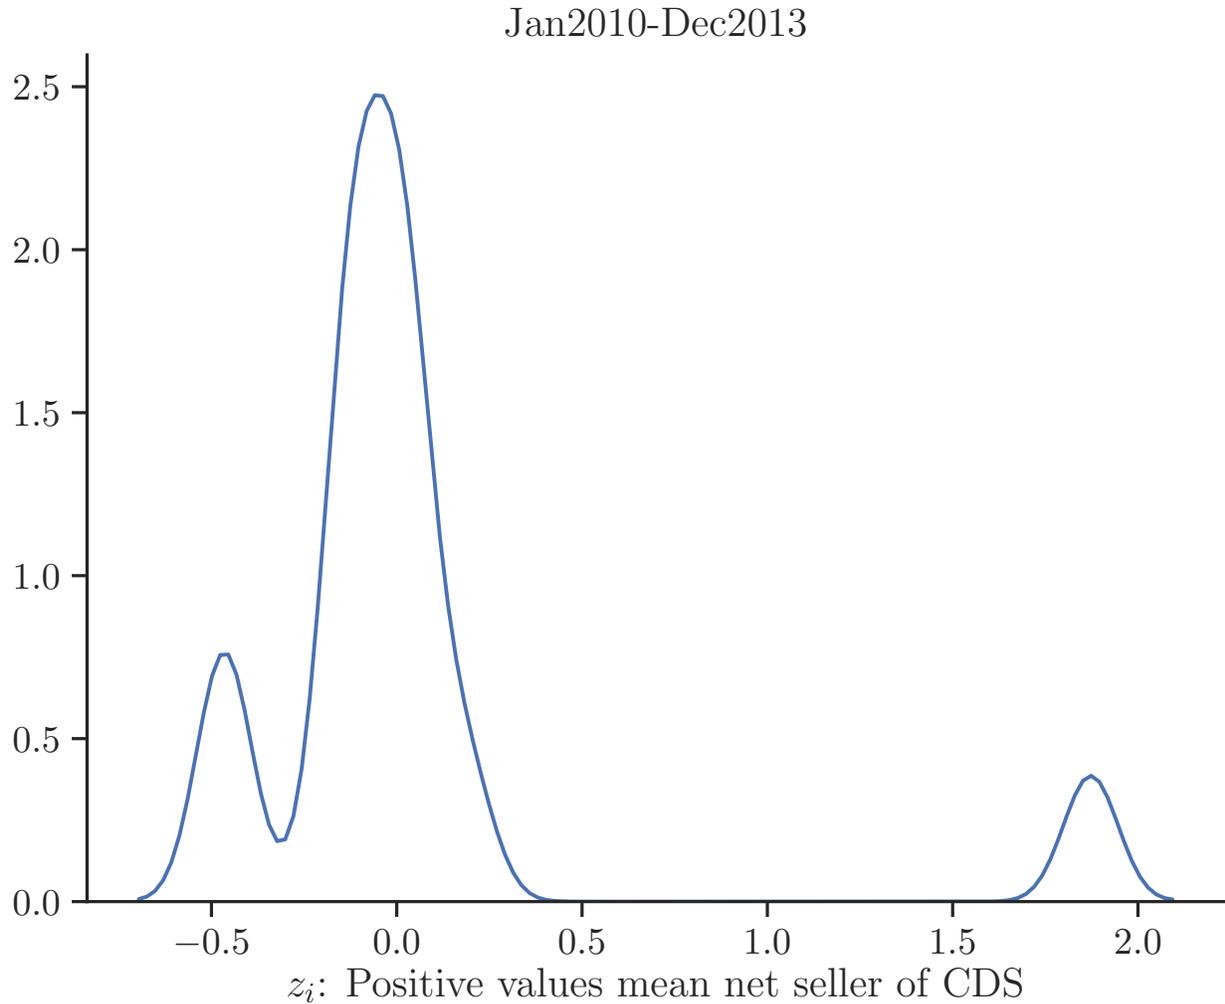
Notes: This figure plots the 5-year CDX Investment Grade Index from June 1, 2006 through January 31, 2010. Source: Authors' analysis, which uses data provided to the OFR by The Depository Trust & Clearing Corporation.

Figure 2: Aggregate Credit Risk Factor



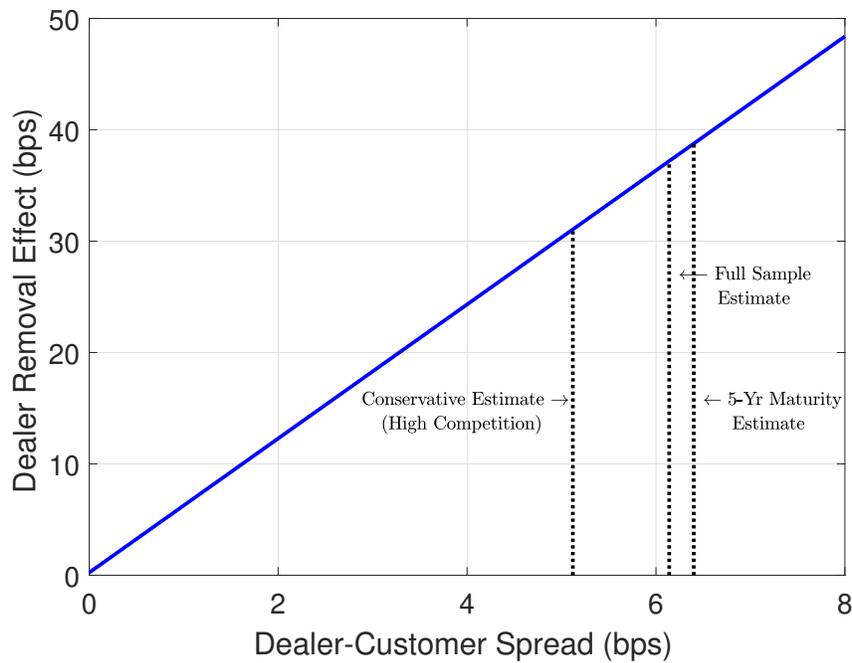
Notes: This figure plots our aggregate credit risk factor from 2002 through 2013. We construct the factor on each date by taking a cross-sectional average of 5-year CDS spreads for all U.S. firms. CDS spreads are obtained from Markit Ltd.

Figure 3: The Distribution of Post-Trade Net Exposures, z_i for Dealers



Notes: This figure plots the distribution of dealer exposure to aggregate credit risk, scaled by their equity, for the period 2010-01-04 through 2013-12-31. We compute each dealer's notional exposure to aggregate credit risk across all of its bilateral positions using the methodology described in Section 3. The aggregate credit risk factor we use is the cross-sectional average of 5-year CDS spreads for all U.S. firms from the Markit database. We then scale each dealer's notional exposure to aggregate credit risk by the market value of their equity to obtain z_i on each date. The plot shows the distribution of the average z_i across dealers for our sample. Positive values of z_i indicate that dealer i is a net seller of credit protection on the aggregate credit risk factor. For readability, we've scaled the y -axis so that the probability distribution integrates to 1000, not 1. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Figure 4: Dealer Removal Effect for Different Dealer-Customer Spread Estimates



Notes: This figure plots the change in the average dealer market spreads upon the removal of the largest net-seller dealer. On the x-axis, we vary the customer-dealer spread, namely $\bar{R}_d - \bar{R}_c$, which in turn change the model-implied values for α and β . Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

TABLES

Table 1: Summary Statistics of Bilateral Concentration and Price Concessions

		p25	p50	p75	Mean	StdDev
$\kappa_{i,j,t}$	Non-Dealers	0.03	0.09	0.18	0.16	0.21
	Dealers	0.00	0.01	0.03	0.02	0.04
	All	0.00	0.02	0.07	0.07	0.14
$PriceConcession^{EW}$	Non-Dealers	0.00	0.00	2.48	2.93	8.53
	Dealers	0.00	0.74	3.21	3.61	9.99
	All	0.00	0.51	3.01	3.41	9.58
$PriceConcession^{NW}$	Non-Dealers	0.00	0.00	2.32	2.90	8.58
	Dealers	0.00	0.70	3.13	3.62	10.12
	All	0.00	0.49	2.96	3.40	9.69

Notes: This table presents summary statistics of bilateral concentration and weekly price concessions between counterparties. $\kappa_{i,j,t}$ measures the fraction of i 's total gross exposure to aggregate credit risk is with j as of the end of week t . $PriceConcession_{i,j,t}$ is the average amount of basis points that agent i conceded in price to agent j over their trades in week t . The superscript on $PriceConcession$ indicates whether it is computed using an equal-weighted (EW) or notional-weighted (NW) average over the trades between i and j in week t . We winsorize both measures at their 99% tails. See Section 3 of the paper for more detail on how we construct $PriceConcession$, κ , and the aggregate credit risk factor. Dealers are defined according to the algorithm in Appendix C. The sample size of the panel with both concentration and price concession measures is 45,074. The sample runs from 2010-01-04 to 2013-12-31. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Table 2: Are Traders Averse to Concentrated Bilateral Exposures?

Dependent Variable:	$PriceConcession_{i,j,t}$ (bps)							
	Notional-Based $\kappa\psi$				DV01-Based $\kappa\psi$			
	EW	NW	NW	NW	EW	NW	NW	NW
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\kappa_{i,j,t}$	-5.39**	-10.25**	-5.11**	-9.95**	-5.94**	-11.01**	-5.73**	-10.94**
	(1.73)	(2.17)	(1.80)	(2.20)	(1.77)	(2.34)	(1.81)	(2.25)
Sample	Full	Dealer	Full	Dealer	Full	Dealer	Full	Dealer
Overall R^2	0.21	0.20	0.20	0.19	0.21	0.20	0.20	0.19
N	44,785	39,318	44,785	39,318	44,785	39,318	44,785	39,318

Notes: This table presents regressions of the following form:

$$GiveUp_{i,j,t} = \frac{S_j}{i,t} + \kappa_{i,j,t} + \epsilon_{ij,t}$$

where $PriceConcession_{i,j,t}$ is the average (equal or notional-weighted) amount of basis points that agent i conceded in price to agent j over their trades in week t . $\kappa_{i,j,t}$ measures the fraction of i 's total gross exposure to aggregate credit risk is with j as of the end of week t . See Section 3 of the paper for more detail on how we construct $PriceConcession$, κ , and the aggregate credit risk factor. $\frac{S_j}{i,t}$ is an $i \times t$ fixed effect interacted with a dummy for whether i sells protection to j . Columns (1)-(4) run the regression when $\kappa\psi$ s defined using notional values. Columns (5)-(8) run it when $\kappa\psi$ s defined using DV01s with respect to aggregate credit risk, which measure the amount of margin payments that i would make or receive if our aggregate credit risk index were to go up by one basis point. See Section 3 of the paper for more detail on how we construct $GiveUp$, κ , and the aggregate credit risk factor. To be included in the regression, i must have at least five trades in week t . We also winsorize the price concession variables at their 1% tails. In all regressions, we report standard errors that are clustered by i below point estimates. * indicates a p -value of less than 10% and ** indicates a p -value of less than 5%. The full sample runs from 2010-01-04 to 2013-12-31. The dealer subsample spans the same period, but runs the regression only for the subset where i is a dealer. Dealers are defined according to the algorithm in Appendix C. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Table 3: Summary Statistics of Spreads by Firm-Week Buckets

	All		5 per (f, w)		10 per (f, w)		20 per (f, w)	
	EW	LW	EW	LW	EW	LW	EW	LW
# Trades	8	23	14	27	22	33	35	45
Avg. Spread (bps)	133	153	150	159	161	164	173	170
Avg. EDF (bps)	65	77	68	78	75	82	90	90
Avg. Loss-Given-Default (%)	60.6	60.3	60.4	60.2	60.2	60.0	59.9	59.8
Avg. Maturity (Years)	3.7	4.0	4.1	4.1	4.1	4.1	4.1	4.1
Avg. Notional (mm)	6.3	5.7	5.9	5.6	5.6	5.4	5.2	5.1
% Trades Dealer-to-Dealer	76	76	74	76	75	76	78	79
% Notional Dealer-to-Dealer	77	77	75	77	76	78	79	80
# of (f, w) groups	56,409	56,409	25,053	25,053	12,733	12,733	4,910	4,910

Notes: This table presents summary statistics of spreads and trading activity across firm-week pairs (f, w). Within each (f, w) group, we compute each statistic (e.g., average spread). We then average these statistics across (f, w) groups using equal weights. We also liquidity-weight (LW) across groups, where a groups' liquidity weight is determined by the number of trades in that (f, w) group. For the % of dealer-dealer trades, we define use our definition of dealers from Appendix C. Notional values are reported in \$ millions and CDS spreads are reported in basis points. Our sample contains only single name transactions on firms that are domiciled in the United States. In addition, we consider trades that are in denominated in USD, have documentation clause XR (no restructuring), and are for senior unsecured debt (tier = SNFOR). We also drop contracts between nondealers and nondealers (only 0.31% of total) and those whose fair-value spread is above 1000 basis points. We then winsorize fair-value transaction CDS spreads at their the 5% tails. The average EDF (bps) row corresponds to the 5-year expected default frequency from Moody's. For transactions where the underlying firm does not have a match in Moody's EDF database, we fill in the missing value with the average EDF for firms in the same rating during the week of the trade. The loss-given-default data comes from Markit. The sample runs from 2010-01-04 to 2013-12-31. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Table 4: Dealer and Customer Prices

Dependent Variable:	Robust Dealer Selection			Fair-Value CDS Transaction Spread (bps)			DTCC Dealers		
	Full Sample (1)	5-yr Maturity (2)	High Competition (3)	Full Sample (4)	5-yr Maturity (5)	High Competition (6)	Full Sample (4)	5-yr Maturity (5)	High Competition (6)
Market Spread (bps)	0.82** (0.01)	0.90** (0.04)	0.84** (0.02)	0.82** (0.01)	0.90** (0.04)	0.84** (0.02)	0.82** (0.01)	0.90** (0.04)	0.84** (0.02)
1 (Customer-Dealer)	6.14** (1.12)	6.40** (1.28)	5.12** (0.90)	7.38** (0.90)	6.88** (1.20)	6.28** (0.92)	7.38** (0.90)	6.88** (1.20)	6.28** (0.92)
Overall R^2	0.92	0.97	0.93	0.92	0.97	0.93	0.92	0.97	0.93
N	424,485	229,086	79,095	424,485	229,086	79,095	424,485	229,086	79,095

Notes: This table presents regressions of the following form:

$$Spread_{k,f,t} = FE(\text{Firm}) + FE(\text{IG} \times \text{Mat Bucket}_k \times \text{Week}) + \theta_1 \times \text{MarketSpread}_{f,t} + \theta_2 \times \log(\text{Notional}_{k,t}) + \times 1_{k,t}(\text{Customer-Dealer}) + \varepsilon_{k,f,t}$$

$Spread_{k,f,t}$ is the fair-value spread (in basis points) for transaction k , written on firm f , and executed on date t . $MarketSpread_{f,t}$ is the 5-year CDS spread from Markit's singlename database that is associated with firm f on date t . $Notional_{k,t}$ is the notional amount in the transaction. $1_{k,t}(\text{Dealer-Dealer})$ is a dummy variable that equals 1 if the transaction is between two dealers and is zero otherwise. In columns (1)-(4), dealers are defined according to the algorithm in Appendix C and in columns (5)-(8) we use the DTCC's labeling of dealers. $FE(\text{Firm})$ is a fixed effect based on the underlying firm f in the transaction. $\text{Firm}(\text{IG} \times \text{MatBucket}_k \times \text{Week})$ is a fixed effect based on the interaction of whether the firm in the trade is rated as investment grade, the maturity bucket of the trade, and the week of the trade. To define a maturity bucket, we first group transactions into one of the following four buckets based on their maturity: (i) 0-2 years; (ii) 2-4 years; (iii) 4-6 years; and (iv) 7+ years. Column (1) runs the regression for the full sample of transactions. Column (2) uses transactions whose maturity is between four and six years. Column (3) uses transactions where Markit's liquidity factor (the number of dealers providing quotes) for the firm on date t is at least eight (the 10% of tail of Markit depth for the full sample). In addition, column (3) focuses on dealer-dealer trades and trades between dealers and large customers, where large customers are defined as those in the top 10% of gross transaction volume for our sample period. Columns (4)-(6) mirror the filters used in columns (1)-(3), except these regressions use the DTCC's definition of dealer. Our main sample contains only single name transactions on firms that are domiciled in the United States. In addition, we consider trades that are in denominated in USD, have documentation clause XR (no restructuring), and are for senior unsecured debt (tier = SNFOR). We also drop contracts between nondealers and nondealers (only 0.31% of total) and those where the fair-value spread is over 1000 basis points. We then winsorize fair-value transaction CDS spreads and 5-year CDS Market spreads at their 5% tails. The credit ratings used to determine whether a firm is investment-grade fixed effect come from Markit. Standard errors, which are double clustered by year and firm, are listed below point estimates. * indicates a p -value of less than 10% and ** indicates a p -value of less than 5%. The sample runs from 2010-01-04 to 2013-12-31. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Table 5: Calibration

Parameter	Value	Source
\bar{z}_d	0.045	DTCC Data 2010-2013
\bar{R}_c \bar{R}_d (bps)	5.12	DTCC Data 2010-2013
\bar{R}_d (bps)	133.00	DTCC Data 2010-2013
n	723	DTCC Data 2010-2013
n_d	14	DTCC Data 2010-2013
$L = \text{Loss-Given-Default}$	60.60%	Markit
$p = \text{Probability of Default}$	0.65%	Moody's
d	0.32	Model Implied
$\alpha\sigma^2\bar{w}_i\psi + \mu$	143.04	Model Implied
$\alpha\psi$	4.37	Model Implied
	7.98	Model Implied

Notes: This table shows parameters used to calibrate the model. \bar{z}_d is the time-series average of dealer exposure. For each week, we compute the average dealer \bar{z}_d across dealers, then report the time-series average for the full sample in the table. Section 3.3.1 contains a full description of this procedure. Dealers are those identified by the algorithm described in Appendix C. \bar{R}_c \bar{R}_d is the estimate that comes out of a regression of transaction spreads on a dummy variable for if the transaction is a customer-dealer trade (see Table 4 for complete details). \bar{R}_d is the average transaction spread in the CDS market from Table 3. n is the total number of counterparties in the network. n_d is the number of dealers. L and p are the physical loss-given-default and probability of default for the firms that are included in our estimation of \bar{R}_c \bar{R}_d . See Table 4 for more details on this set of firms. The remaining parameters in the table are implied by our structural model. Specifically, the values of d and $\alpha\psi$ are computed according to Equations (16) and (17), where \bar{w}_i is normalized to one. The expected loss rate is given by $\mu = Lp$ and its variance given by $\sigma^2 = L^2p(1-p)$. The average complete network benchmark price, $\bar{R}_{\text{Complete Network}} = \alpha\sigma^2\bar{w}_i\psi + \mu$ is in basis points. From Equation (11), we have $d = \frac{n\alpha\sigma^2}{n\alpha\sigma^2+2}$. See Appendix A for detailed derivations of the model. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Table 6: Dealer Removal

	Benchmark	Top	90 th prc.	Median	Bottom
	(1)	(2)	(3)	(4)	(5)
Number of dealers	14	13	13	13	13
Complete network \bar{R} (bps)	143.04	143.88	143.21	143.02	142.80
\bar{R}_d (bps):	133.00	164.31	138.61	131.35	123.11
\bar{R}_c (bps):	138.12	153.91	140.95	137.29	133.14
\bar{z}_d	0.045	0.092	0.021	0.053	0.089

Notes: This table reports the number of dealers, the average spreads under the complete network, the average spreads in the dealer market, the average spreads in the customer market, and the average net position of dealers. We define dealers precisely in Section 1.1.1.2. Column (1) reports our benchmark calibration. In Column (2) reports the results after removing the largest net-seller. Column (3) reports results after removing one dealer at the 90th percentile. Column (4) reports results after removing the dealer with the median net position, and Column (5) reports results after removing the dealer that is the largest net buyer in the baseline model. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Table 7: Dealer removal effects and heterogeneous trading costs

	Benchmark	Top	90 th prc.	Median	Bottom
	(1)	(2)	(3)	(4)	(5)
Number of dealers	14	13	13	13	13
Complete network \bar{R} (bps)	143.04	143.32	143.11	143.05	142.98
\bar{R}_d (bps):	133.00	143.00	134.79	132.49	129.84
\bar{R}_c (bps):	138.12	143.17	139.02	137.87	136.53
\bar{z}_d	0.045	0.001	0.038	0.048	0.059

Notes: This table reports the number of dealers, the average spreads under the complete network, the average spreads in the dealer market, the average spreads in the customer market, and the average net position of dealers for the model extension featuring heterogeneous trading costs discussed in Section 4.2.2. See the heterogeneous trading cost model derivations in Appendix B.2. Parameter α_d is set to 10^{-8} . Parameters α_c and α_d are calibrated so that the model-implied dealer and customer average prices prior to any dealer removal match the value observed in the data of 138.12 and 133 basis points, respectively. We define dealers precisely in Section 1.1.1.2. Column (1) reports our calibration without dealer removal. In Column (2) reports the results after removing the largest net-seller. Column (3) reports results after removing one dealer at the 90th percentile. Column (4) reports results after removing the dealer with the median net position, and Column (5) reports results after removing the dealer that is the largest net buyer in the baseline model. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Table 8: Dealer removal effect varying risk aversion and trading costs

Trading Costs	Risk aversion				
	$\alpha_d = 4.37$	$\alpha_d = 4.50$	$\alpha_d = 5.00$	$\alpha_d = 5.50$	$\alpha_d = 6.00$
Panel A: effect on dealer prices (bps) under $\alpha_c = 7.98$					
$\alpha_c = 7.98$	31.31	34.81	48.27	61.63	74.90
$\alpha_c = 15$	34.90	38.56	52.63	66.62	80.52
$\alpha_c = 20$	36.16	39.88	54.20	68.45	82.61
$\alpha_c = 30$	37.56	41.35	55.97	70.52	85.01
Panel B: effect on dealer prices (bps) under $\alpha_c = 10^{-8}$					
$\alpha_c = 7.98$	10.00	13.07	24.94	36.81	48.68
$\alpha_c = 15$	9.95	13.01	24.87	36.73	48.59
$\alpha_c = 20$	9.93	12.99	24.85	36.70	48.56
$\alpha_c = 30$	9.91	12.97	24.82	36.67	48.52

Notes: This table reports dealer removal effects on prices when combined with increases in risk aversion and trading costs. The model specification considered here features heterogeneous trading costs discussed in Section 4.2.2. See the heterogeneous trading cost model derivations in Appendix B.2. Parameter α_d varies across Panels, from its benchmark value in Panel A to 10^{-8} in Panel B. In each panel, the value reported in the upper left corner is the dealer removal effect on dealer prices at calibrated parameters from Table 7—that is, parameters α_c and α_d are calibrated so that the model-implied dealer and customer average prices before the dealer removal match the value observed in the data of 138.12 and 133 basis points, respectively. As we move across columns (rows), we report the dealer removal effect on prices when combined with an increases in α_c (α_c). Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Appendix

A Model Derivations

A.1 Solving the Model

Agent i 's optimization problem is given by Equation (1):

$$\max_{\{i_j\}_{j=1}^n, z_i} w_i(1 - \mu) + \sum_{j=1}^n i_j(R_{ij} - \mu) - \frac{\alpha w_i}{2}(w_i + z_i)^2 - \frac{\alpha}{2} \sum_{j=1}^n i_j^2$$

subject to

$$i_j = 0 \text{ if } g_{ij} = 0,$$

and

$$z_i = \sum_{j=1}^n i_j.$$

Agent i 's first-order conditions give us:

$$i_j = \begin{cases} \frac{1}{K_i}(R_{ij} - \mu) - \frac{1}{2}\hat{z}_i & \text{if } g_{ij} = 1 \\ 0 & \text{if } g_{ij} = 0, \end{cases} \quad (\text{A1})$$

where

$$\hat{z}_i = (w_i + z_i)\alpha\sigma^2, \quad (\text{A2})$$

$$z_i = \sum_{j=1}^n i_j = \frac{\frac{1}{K_i} \sum_{j=1}^n g_{ij}(R_{ij} - \mu) - \alpha w_i}{\frac{1}{K_i} + \alpha\sigma^2}, \quad (\text{A3})$$

and

$$K_i = \sum_{j=1}^n g_{ij}. \quad (\text{A4})$$

We can derive Equation (8) by combining Equations (2), (6), and (7). Furthermore, to fully characterize the equilibrium, we solve for equilibrium quantities by rewriting Equation (8) in matrix notation as follows:

$$z + \omega\psi = (\mathbf{I} - \Lambda)\omega\psi + \Lambda\tilde{G}(z + \omega),$$

where $z = [z_1, \dots, z_n]'$ and $w = [w_1, \dots, w_n]'$ are column vectors of net positions and pre-trade exposures, Λ is a diagonal matrix with the i th element given by $\frac{\alpha w_i}{K_i + \alpha\sigma^2}$, and \tilde{G} is a $n \times n$ matrix with the ij th entry given by $\frac{g_{ij}}{K_i}$.

We can solve the system of equations for the equilibrium net positions and post-trade exposures:

$$z + w = (\mathbf{I} - \Lambda\tilde{G})^{-1}(\mathbf{I} - \Lambda)w, \quad (\text{A5})$$

which fully characterize the solution of the model. Equation (A5) defines the map between agents' pre-trade exposures to the underlying asset on the right-hand side and their post-trade exposures on left-hand side. The right-hand side of the equation depends only on exogenous parameters of the model.

A.1.1 Complete Network

Under the complete network benchmark, we have $g_{ij} = 1$ for every i and j . In this case, $K_i = n$ for every i , and $\tilde{g}_{ij} = \frac{1}{n}$ for every i and j . Also, $\frac{\alpha w_i}{K_i + \alpha\sigma^2} \equiv \frac{\alpha w_i}{n + \alpha\sigma^2} \equiv \bar{\alpha} w_i$ for every i , and the matrix \tilde{G} is becomes idempotent,

i.e., $\tilde{G}^2 = \tilde{G}$. Therefore, the vector of net positions becomes:

$$\begin{aligned} z + \omega &= (\mathbf{I} - \Lambda \tilde{G})^{-1} (\mathbf{I} - \Lambda) w \\ &= (\mathbf{I} - \tilde{G})^{-1} \omega \\ &= (\mathbf{I} - \tilde{G}) \left(\mathbf{I} + \frac{\tilde{G}}{1} \right) \omega \\ &= (\mathbf{I} - \tilde{G}) \omega. \end{aligned}$$

Specifically, the post-trade exposure of agent i is given by:

$$z_i + w_i = (\mathbf{I} - \tilde{G}) w_i + \left(\frac{1}{n} \sum_{j=1}^n w_j \right).$$

The average prices in equilibrium becomes:

$$\begin{aligned} \bar{R}_{\text{Complete Network}} &\equiv \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n R_{ij} \\ &= \frac{\alpha \sigma^2 \psi}{n^2} \sum_{i=1}^n \sum_{j=1}^n \frac{1}{2} (\hat{z}_i + \hat{z}_j) + \mu \\ &= \frac{\alpha \sigma^2 \psi}{n^2} \sum_{i=1}^n \sum_{j=1}^n \frac{1}{2} (z_i + w_i + z_j + w_j) + \mu \\ &= \frac{2}{n} \alpha \bar{\omega} + \mu, \end{aligned}$$

where $\bar{\omega} = \frac{1}{n} \sum_{i=1}^n \omega_i$.

A.1.2 Equilibrium properties

Although the model features closed-formed solutions, the equilibrium variables still depend on the entire trading network. In this subsection, we exploit some limiting cases of the model. First, we define what it means for two agents to be path-connected.

Definition A1. *Two players i and j are path-connected if there is a sequence of agents $\{s_1, s_2, \dots, s_k\}$ such that:*

$$g_{is_1} = g_{s_1 s_2} = \dots = g_{s_{k-1} s_k} = g_{s_k j} = 1.$$

The following proposition shows that when there is no counterparty-specific risk aversion, i.e., $\alpha = 0$, then there is perfect risk sharing among path-connected agents. The corollary following the proposition shows that if all agents are path-connected, then perfect risk sharing among all agents is achieved in equilibrium.

Proposition A1. *If $\alpha = 0$ for every $i = 1, \dots, n$, then any two path-connected agents have the same post-trade exposure:*

$$(z_i + w_i) = (z_j + w_j)$$

for any i and j who are path connected.

Proof. Suppose players i and j are path-connected, but

$$(z_i + w_i) \neq (z_j + w_j).$$

Then, there are two agents, say s and l , that are directly connected with each other (i.e., $g_{sl} = 1$) and have different post-trade exposure (i.e., $z_s + w_s \neq z_l + w_l$). If both agents are maximizing and their first-order conditions hold with equality, then we have that:

$$R_{sl} - \mu = \alpha(z_s + w_s)^2 = \alpha(z_l + w_l)^2 \implies z_i + w_i = z_j + w_j$$

□

Corollary A1. *If $\psi = 0$ for every $i = 1, \dots, n$, and all agents are path connected, then there is perfect risk-sharing in equilibrium, i.e.,*

$$z_i + w_i = \frac{1}{n} \sum_j w_j,$$

and equilibrium prices are given by:

$$R_{ij} - \mu = \alpha \bar{\omega} \psi \quad \forall i, j,$$

where $\bar{\omega} = \frac{1}{n} \sum_{i=1}^n \omega_i$.

Proof. We know that:

$$z_i + w_i = z_j + w_j = \bar{z},$$

where \bar{z} is a constant. We also know that

$$\sum_j z_j = 0,$$

from the clearing conditions. □

Finally, the next proposition shows that when counterparty-specific risk aversion goes to infinity, then the equilibrium features autarky, regardless of the trading network in place.

Proposition A2. *If $\psi \rightarrow \infty$ for every $i = 1, \dots, n$, then there is no trade in equilibrium, regardless of the network structure.*

Proof. From the first-order conditions, we get that $z_{ij} = 0$ for any two agents i and j . □

A.2 Model with Core-Periphery Network

A.2.1 Dealer market

Using the clearing conditions to get $\sum_{s=1}^n z_s = 0$, along with having $g_{ij} = 1$ for every j if i is a dealer, we can use Equation (8) to get following expression for post-trade exposures of dealers:

$$z_i + \omega_i = (1 - \alpha_d) \omega_i + \alpha_d \bar{\omega} \psi \quad \forall i = 1, \dots, n_d, \quad (\text{A6})$$

where

$$\alpha_d = \frac{n \alpha \psi^2}{n \alpha \psi^2 + 2}$$

and $\bar{\omega} = \frac{1}{n} \sum_{j=1}^n \omega_j$. Hence, dealers' post-trade exposures are a convex combination of their own pre-trade exposure, i.e., ω_i , and the average pre-trade exposure in the economy, i.e., $\bar{\omega}$.

The average post-trade exposure in the dealer market is given by:

$$\bar{z}_d + \bar{\omega}_d = \frac{1}{n_d} \sum_{i=1}^{n_d} (z_i + \omega_i) = (1 - \alpha_d) \bar{\omega}_d + \alpha_d \bar{\omega}.$$

The equilibrium price of a contract between dealers i and j is given by:

$$R_{ij} - \mu = \alpha \sigma^2 \psi \left[\alpha_d \bar{\omega} \psi + (1 - \alpha_d) \frac{\omega_i + \omega_j}{2} \right],$$

and the average price in the dealer market, i.e., $\bar{R}_d = \frac{1}{n_d^2} \sum_{i=1}^{n_d} \sum_{j=1}^{n_d} R_{ij}$, is given by:

$$\bar{R}_d - \mu = \alpha \sigma^2 \bar{\omega} \psi + (1 - \alpha_d) \alpha \sigma^2 (\bar{\omega} - \bar{\omega}_d),$$

where $\bar{\omega}_d = \frac{1}{n_d} \sum_{i=1}^{n_d} \omega_i$.

A.2.2 Customer market

Applying Equation (8) to customers gives us the following expression for their post-trade exposures:

$$z_i + \omega_i = (1 - c)\omega_i + \frac{1}{n_d + 1} \left[\sum_{j=1}^{n_d} (z_j + \omega_j) + z_i + \omega_i \right] \quad \forall i = n_d + 1, \dots, n,$$

where $\bar{c} = \frac{(n_d+1)\alpha\sigma^2}{(n_d+1)\alpha\sigma^2+2}$. We can use Equation (11) to write the post-trade exposures as follows:

$$z_i + \omega_i = \bar{c}\bar{\omega}_i + (1 - \bar{c})\omega_i - \bar{c}(1 - d)(\bar{\omega} - \bar{\omega}_d) \quad \forall i = n_d + 1, \dots, n, \quad (\text{A7})$$

where $\bar{c} = \frac{n_d\alpha\sigma^2}{n_d\alpha\sigma^2+2}$.

The equilibrium price of contract between a customer $i \in \{n_d + 1, \dots, n\}$ and a dealer $j \in \{1, \dots, n_d\}$ is given by:

$$\begin{aligned} R_{ij} - \mu &= \alpha\sigma^2\psi \left(\frac{z_i + \omega_i + z_j + \omega_j}{2} \right) \\ &= \alpha\sigma^2\bar{\omega} - \frac{\alpha\sigma^2\psi}{2} (1 - d)(\bar{\omega} - \bar{\omega}_d) + \frac{\alpha\sigma^2\psi}{2} [(1 - c)(\omega_i - \bar{\omega}) + (1 - d)(\omega_j - \bar{\omega})], \end{aligned}$$

where we used Equations (A6) and (A7) to derive the last expression.

Hence, the average price in the customer market, i.e.,

$$\bar{R}_c = \frac{1}{n_d(n - n_d)} \sum_{j=1}^{n_d} \sum_{i=n_d+1}^n R_{ij},$$

is given by:

$$\bar{R}_c - \mu = \alpha\sigma^2\bar{\omega} - \frac{1}{2}\alpha\sigma^2\psi(\bar{\omega} - \bar{\omega}_d) \underbrace{\left[(1 + c)(1 - d) \frac{n_d}{n} \frac{(1 - c)}{n_d} \right]}_{>0 \text{ iff } \frac{n_d}{n} < \frac{1}{2}}.$$

We can also write the average price in the customer market as a function of the average price in the dealer market as follows:

$$\bar{R}_c = \bar{R}_d + \frac{1}{2}\alpha\sigma^2\psi(1 - c) \left[1 + (1 - d) \frac{n - n_d}{n_d} \frac{n_d}{n} \frac{(1 - c)}{n_d} (\bar{\omega} - \bar{\omega}_d) \right]$$

Proposition 1 is a direct implication of Equations (12), (14) and (15).

A.2.3 Calibration

From Equation (11), we can compute $(\bar{\omega} - \bar{\omega}_d)$ as a function of \bar{z}_d and d :

$$\bar{\omega} - \bar{\omega}_d = \frac{\bar{z}_d}{d}. \quad (\text{A8})$$

Furthermore, we can write d as follows:

$$d = \frac{\alpha\sigma^2\bar{z}_d}{\alpha\sigma^2\bar{z}_d + \bar{R}_c - \bar{R}_d - 2 \frac{1 - c}{n}}, \quad (\text{A9})$$

by taking the difference between Equations (12) and (14) and solving for d .

We can use market clearing conditions and the definition of $\bar{\omega}_i$ to write \bar{z}_c and $\bar{\omega}_c$ as a function of \bar{z}_d , $\bar{\omega}_d$,

and $\bar{\omega}$:

$$\bar{z}_c = \frac{n_d}{n} \bar{z}_d \quad (\text{A10})$$

$$\bar{\omega}_c = \frac{n \bar{\omega} \psi}{n} \frac{n_d \bar{\omega}_d}{n_d} \quad (\text{A11})$$

The complete network benchmark price can be written as:

$$\alpha \sigma \psi \bar{\omega} \bar{z}_c = \bar{R}_d \quad \mu + \frac{1}{d} \bar{z}_d \alpha \sigma \psi = \bar{R}_d \quad \mu + (\bar{R}_c - \bar{R}_d) 2 \left(1 - \frac{n_d}{n} \right),$$

where the first equality is derived by combining Equations (12) and (A8), and the second equality is obtained by substituting in Equation (A9).

We can rearrange Equation (22) to compute $\alpha \psi$ as follows:

$$\alpha \psi = \frac{1}{2 \bar{\omega} \psi} \left[\bar{R}_d - \mu + (\bar{R}_c - \bar{R}_d) 2 \left(1 - \frac{n_d}{n} \right) \right], \quad (\text{A12})$$

and we can rearrange \bar{z}_d defined in Equation (11) to compute $\alpha \psi$ as follows:

$$= \frac{1}{2} n \alpha \psi^2 \left(\frac{1}{d} - \frac{d}{n} \right) \quad (\text{A13})$$

$$= \frac{1}{2} n \alpha \psi^2 \left[\frac{\bar{R}_c - \bar{R}_d 2 \left(1 - \frac{n_d}{n} \right)}{\alpha \sigma \psi \bar{z}_d} \right] \quad (\text{A14})$$

$$= (n - n_d) \left(\frac{\bar{R}_c - \bar{R}_d}{\bar{z}_d} \right) \quad (\text{A15})$$

where \bar{z}_d is computed from Equation (A9).

To compute the model-implied dealers' pre-trade exposures, we can rearrange Equation (A6) as follows:

$$\bar{\omega} \quad \omega_i = \frac{z_i}{d} \quad \forall i = 1, \dots, n_d, \quad (\text{A16})$$

and to compute customers' pre-trade exposure we can rearrange Equation (A7) as well:

$$\bar{\omega} \quad \omega_i = \frac{z_i}{c} + (1 - \frac{d}{n}) (\bar{\omega} - \bar{\omega}_d) \quad \forall i = n_d + 1, \dots, n. \quad (\text{A17})$$

These are useful objects on their own. These are agents' model-implied pre-trade exposures, and they allow us to *measure* which market participant is more or less risky ex-ante.

A.2.4 Comparative statics

In this subsection, we analyze how equilibrium prices depend on risk aversion and aversion to holding concentrated position risk. The following proposition shows how equilibrium prices and spreads depend on α and ψ .

Proposition A3. *If $\frac{d}{n} < \frac{1}{2}$ and $\bar{\omega} \psi > \bar{\omega}_d > 0$, then the following comparative statics hold:*

- (i) $\frac{\partial}{\partial \psi} \bar{R}_d < 0$, $\frac{\partial}{\partial \psi} \bar{R}_c < 0$, $\frac{\partial}{\partial \psi} \bar{R}_{\text{Complete Network}} = 0$, $\frac{\partial}{\partial \psi} \bar{R}_{\text{Complete Network}} - \bar{R}_d > 0$,
 $\frac{\partial}{\partial \psi} \bar{R}_{\text{Complete Network}} - \bar{R}_c > 0$, and $\frac{\partial}{\partial \psi} \bar{R}_c - \bar{R}_d > 0$,
- (ii) $\frac{\partial}{\partial \alpha} \bar{R}_d > 0$, $\frac{\partial}{\partial \alpha} \bar{R}_c > 0$, $\frac{\partial}{\partial \alpha} \bar{R}_{\text{Complete Network}} > 0$, $\frac{\partial}{\partial \alpha} \bar{R}_{\text{Complete Network}} - \bar{R}_d > 0$,
 $\frac{\partial}{\partial \alpha} \bar{R}_{\text{Complete Network}} - \bar{R}_c > 0$, and $\frac{\partial}{\partial \alpha} \bar{R}_c - \bar{R}_d > 0$.

The proof consists of taking these derivatives using Equations (9), (12), (14), and (15). In this proposition, we assume dealers to be less exposed to the underlying default risk, i.e., $\bar{\omega}_c > \bar{\omega}_d$, which implies dealers as net sellers in equilibrium, i.e., $\bar{z}_d > 0$. This is consistent with the evidence presented in Section 3.3.

Aversion to bilateral concentration has no effect on the complete network benchmark average price and has a negative effect on the average price in the dealer and customer markets (item i). As α increases, agents are more averse to trading too much with one counterparty. Hence, there is less risk sharing in the equilibrium with a higher α , which means that both customer and dealer post-trade exposures are closer to their pre-trade exposures. When dealers are net sellers of protection, this implies lower post-trade exposures for dealers and higher post-trade exposures for customers when α increases.

The deterioration in risk sharing caused by an increase in α changes equilibrium prices. For average customer prices, α has two offsetting effects. On the one hand, it increases customer exposures which increases customer prices. On the other hand, because dealers absorb less risk from the customer sector, they have lower post-trade exposures to aggregate default risk, which pushes customer prices down. Given that dealers are small in number relative to customers, market clearing implies that the average post-trade exposure of dealers decreases by more than the increase in average post-trade exposure of customers. As a result, equilibrium prices in both the dealer and customer markets then decline. Intuitively, equilibrium prices are lower in the customer market to offset the higher costs of holding concentrated positions. Note that, in the dealer market, both lower post-trade exposures, and the burden of higher marginal costs of bilateral concentration, drive prices down. Since the two effects work in the same direction, average prices in the dealer market decrease by more than in the customer market. As a result, the gap between CDS premiums in customer-dealer and dealer-dealer trades widens.

Risk aversion increases the average price in both dealer and customer markets, as well as the complete benchmark price (items ii). As agents become more risk-averse, protection against aggregate default risk becomes more expensive. Similar to the effect of an increase in α , risk aversion also increases the spread between the average price in the dealer and customer markets. However, the economic mechanism behind the comparative statics for α is entirely different.

If risk aversion goes up, agents with high exposures have a higher demand for aggregate default risk protection and as a result there will be *more* risk-sharing in equilibrium. Given the improved risk reallocation, the dealer sector absorbs more risk from the customer sector. Dealers will trade at a higher price because of their higher average post-trade exposures. Customers will also trade at a higher price not only because risk aversion *per se* is higher but also because the average post-trade exposure of dealers increases by more than the decline in customers' average post-trade exposure, due to their smaller number. This means that contracts between dealers and customers will be executed at higher prices on average.

Risk aversion increases both dealer and customer markets' average prices. However, it increases the average price more in the customer market than in the dealer market. An increase in aversion to aggregate default risk also increases the spread in prices across dealer versus customer trades. There are two distinct offsetting effects driving this result. First, there is the direct effect of an increase in risk aversion on shadow prices of risk. More risk-averse agents have a higher shadow price of risk bearing for a given net exposure. Comparing the effect on dealers versus customers, the fact that dealers are less exposed to the underlying asset than customers means that their shadow cost of risk bearing is less sensitive to changes in risk aversion. The effect of higher risk aversion on shadow prices of risk increases average prices in the customer market by more than in the dealer market, due to imperfect risk sharing. The second effect is more subtle and is dominated by the first one. The higher demand for risk sharing resulting from higher risk aversion implies that market participants become more similar in their post-trade exposures. Less dispersion in post-trade exposures implies less dispersion in the average prices observed in dealer versus customer markets. However, Proposition A3 shows that the first effect dominates the second one. Thus, as risk aversion increases, the spread between the average price in dealer and customer markets widens.

B Model Extensions

B.1 Risk Aversion Heterogeneity

In this section, we present a variation of our benchmark model where agents are heterogeneous in both risk their aversion and pre-trade exposures. Formally, agent i 's optimization problem is given by:

$$\begin{aligned}
& \max_{z_i, \{i_j\}_{j=1}^n} \omega_i(1 - \mu) + \sum_{j=1}^n i_j(R_{ij} - \mu) - \frac{\alpha_i}{2}(\omega_i + z_i)^2 - \frac{1}{2} \sum_{j=1}^n i_j^2 \\
& \text{s.t.} \quad i_j = 0 \text{ if } g_{ij} = 0 \\
& \quad \quad z_i - \sum_{j=1}^n i_j = 0,
\end{aligned}$$

From the first-order conditions, we can solve agent i 's optimization problem for agents' bilateral and net positions:

$$i_j = \begin{cases} \frac{1}{2} [R_{ij} - \mu - \alpha_i(\omega_i + z_i)] & \text{if } g_{ij} = 1 \\ 0 & \text{if } g_{ij} = 0 \end{cases} \quad (\text{A18})$$

where

$$z_i = \sum_{j=1}^n i_j = \frac{\frac{1}{K_i} \sum_{j=1}^n g_{ij}(R_{ij} - \mu) - \alpha_i \omega_i}{\bar{K}_i + \alpha_i}, \quad (\text{A19})$$

and $K_i = \sum_{j=1}^n g_{ij}$.

In this variant of the model, the market clearing conditions are the same as in the benchmark model and are given by:

$$i_j + j_i = 0 \quad \forall i, j = 1, \dots, n. \quad (\text{A20})$$

As in the benchmark model, we assume no transaction costs between any two counterparties:

$$R_{ij} = R_{ji}.$$

If agents i and j can trade (i.e., $g_{ij} = g_{ji} = 1$), we can use their optimality conditions from Equation (A18) and market clearing condition from Equation (A20) to solve for the equilibrium price of a contract between agents i and j :

$$\underbrace{R_{ij} - \mu}_{\text{contract premium}} = \left[\frac{\frac{1}{2}\alpha_i(\omega_i + z_i) + \frac{1}{2}\alpha_j(\omega_j + z_j)}{\frac{1}{2} + \frac{1}{2}} \right] - \frac{1}{2} = \left[\frac{\alpha_i(\omega_i + z_i) + \alpha_j(\omega_j + z_j)}{2} - \frac{1}{2} \right]. \quad (\text{A21})$$

If we substitute the contract premium (Equation A21) into the agents' first-order condition (Equation A18), then whenever $g_{ij} = 1$ we have:

$$i_j = \frac{1}{2} [R_{ij} - \mu - \alpha_i(\omega_i + z_i)] = \frac{1}{2} [\alpha_j(\omega_j + z_j) - \alpha_i(\omega_i + z_i)]$$

which means that agent i sells insurance to agent j (i.e., $i_j > 0$) if, and only if, agent j is more exposed to default than agent i after trade and correcting for differences in risk aversion, that is, $\alpha_j(\omega_j + z_j) > \alpha_i(\omega_i + z_i)$. The trade volume between two counterparties depends on the difference of their post-trade exposures scaled by risk aversion and on the trading cost parameter, μ .

The net positions, $\{z_i\}_i$, are determined in equilibrium. We can use Equations (A19) and (A21), and solve for equilibrium net positions. From Equation (A21), we have:

$$\begin{aligned}
R_{ij} - \mu &= \frac{1}{2} \left[\frac{\alpha_i}{2}(\omega_i + z_i) + \frac{\alpha_j}{2}(\omega_j + z_j) \right] \\
\sum_{j=1}^n g_{ij}(R_{ij} - \mu) &= \frac{1}{2} \sum_{j=1}^n g_{ij} \left[\frac{\alpha_i}{2}(\omega_i + z_i) + \frac{\alpha_j}{2}(\omega_j + z_j) \right]
\end{aligned}$$

Using the same notation as in the benchmark model, let $\tilde{g}_{ij} = \frac{g_{ij}}{K_i}$ and $K_i = \sum_{j=1}^n g_{ij}$, then we have:

$$\sum_{j=1}^n g_{ij}(R_{ij} - \mu) = \frac{1}{2} \alpha_i (z_i + \omega_i) K_i + \frac{1}{2} \sum_{j=1}^n K_i \alpha_j \tilde{g}_{ij} (z_j + \omega_j)$$

We can rewrite Equation (A19) as follows:

$$z_i = \frac{\frac{1}{K_i} \sum_{j=1}^n g_{ij}(R_{ij} - \mu) - \alpha_i \omega_i}{\frac{1}{K_i} + \alpha_i}$$

$$\sum_{j=1}^n g_{ij}(R_{ij} - \mu) = z_i (K_i + \alpha_i K_i) + \alpha_i \omega_i K_i$$

and have:

$$z_i (K_i + \alpha_i K_i) + \alpha_i \omega_i K_i = \frac{1}{2} \alpha_i (z_i + \omega_i) K_i + \frac{1}{2} \sum_{j=1}^n K_i \alpha_j \tilde{g}_{ij} (z_j + \omega_j)$$

$$z_i \left(1 + \frac{\alpha_i}{K_i}\right) + \omega_i = \frac{1}{2} (z_i + \omega_i) + \frac{1}{2} \sum_{j=1}^n \frac{\alpha_j}{\alpha_i} \tilde{g}_{ij} (z_j + \omega_j).$$

The system above then becomes:

$$z_i \left(1 + \frac{\alpha_i}{K_i}\right) = \omega_i + \sum_{j=1}^n \frac{\alpha_j}{\alpha_i} \tilde{g}_{ij} (z_j + \omega_j) \quad \forall i = 1, \dots, n. \quad (\text{A22})$$

The system of equations in (A22) holds in equilibrium and it pins down the equilibrium net positions. In matrix notation, the system of equations becomes:

$$\underbrace{\begin{bmatrix} 1 + \frac{2}{K_1 \alpha_1} & 0 & \dots & 0 \\ 0 & 1 + \frac{2}{K_2 \alpha_2} & \dots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & \dots & & 1 + \frac{2}{K_n \alpha_n} \end{bmatrix}}_{\equiv \Lambda^{-1}} z = \omega \psi \underbrace{\begin{bmatrix} \frac{\tilde{g}_{11} \alpha_1}{\alpha_1} & \frac{\tilde{g}_{12} \alpha_2}{\alpha_2} & \dots & \frac{\tilde{g}_{1n} \alpha_n}{\alpha_n} \\ \frac{\tilde{g}_{21} \alpha_1}{\alpha_2} & \frac{\tilde{g}_{22} \alpha_2}{\alpha_2} & \dots & \frac{\tilde{g}_{2n} \alpha_n}{\alpha_2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\tilde{g}_{n1} \alpha_1}{\alpha_n} & \frac{\tilde{g}_{n2} \alpha_2}{\alpha_n} & \dots & \frac{\tilde{g}_{nn} \alpha_n}{\alpha_n} \end{bmatrix}}_{\equiv G^*} (z + \omega)$$

$$\Lambda^{-1} z = \omega \psi G^* (z + \omega), \quad (\text{A23})$$

where $z = [z_1, \dots, z_n]'$ and $\omega \psi = [\omega_1, \dots, \omega_n]'$ are column vectors of net positions and pre-trade exposures.

Net positions as a function of pre-trades exposures. We can rearrange Equation (A23) and solve the system of equations for the equilibrium net positions and post-trade exposures:

$$z = (\text{I} - \Lambda G^*)^{-1} \Lambda (G^* - \text{I}) \omega \psi \quad (\text{A24})$$

and

$$z + \omega \psi = (\text{I} - \Lambda G^*)^{-1} (\text{I} - \Lambda) \omega. \quad (\text{A25})$$

As in the benchmark model, post-trade exposures are determined jointly in equilibrium. They depend on the network, risk aversion and trading cost parameters, and pre-trade exposures.

B.1.1 Equivalence result and dealer removal

Next we show an equivalence result in which we can use either heterogeneity in risk aversion or pre-trade exposures to match observed net positions. We formalize this result in Proposition A4. The proposition states that for a given distribution of risk aversion parameters, we can find a distribution of pre-trade exposures that will result in a calibrated model that matches the net positions target by the model. In our benchmark model, we effectively use this result by assuming homogeneous risk aversion parameters and we then back out pre-trade exposures by matching the empirically observed net positions.

The second part of the proposition is crucial for the equivalence between pre-trade exposure and risk aversion. It states that for a given distribution of the pre-trade exposures, we can find a distribution of risk aversion parameters that will result in a calibrated model that matches the net positions target by the model. This result, however, requires additional assumptions in order to keep risk aversion parameters positive.

Proposition A4. *Part 1: For a given distribution of risk aversion parameters, a given average average pre-trade exposure ($\bar{\omega}$) and trading cost parameter (ψ), and a given target for net positions to be matched by the model, there is a distribution of pre-trade exposures that result in a calibrated model with model-implied net positions matching the targeted net positions.*

Part 2: For a given distribution of pre-trade exposures, a given average risk aversion ($\bar{\alpha}$) and trading cost parameter (ψ), and a given target for net positions to be matched by the model, if (i) $A_0 - A_2$ is nonsingular, (ii) $[A_0 - A_2]^{-1} A_1$ has only positive entries, and (iii) the sum of all elements in $[A_0 - A_2]^{-1} A_1$ are lower than $n\bar{\alpha}$, where A_0 , A_1 and A_2 are defined in Equation (A30), then there is a distribution of risk aversion parameters that result in a calibrated model with model-implied net positions matching the targeted net positions.

Proof. Part 1: Pre-trades exposures as a function of net positions. Notice that the matrix $(G^* - I)$ is singular because its rank is less than n given that each row in G^* sum to one. Hence, we cannot use directly Equation (A24) to solve for pre-trade positions (ω 's) as a function of net positions (z 's). To overcome this issue, we rearrange Equation (A22) for $i = 1, \dots, n-1$ as follows:

$$\begin{aligned}
 z_i \left(1 + \frac{2}{K_i \alpha_i \psi}\right) &= \omega_i + \sum_{j=1}^n \frac{\tilde{g}_{ij} \alpha_j}{\alpha_i} (z_j + \omega_j) \\
 &= \omega_i + \sum_{j=1}^{n-1} \frac{\tilde{g}_{ij} \alpha_j}{\alpha_i} (z_j + \omega_j) + \frac{\tilde{g}_{in} \alpha_n}{\alpha_i} (z_n + \omega_n) \\
 &= \omega_i + \sum_{j=1}^{n-1} \frac{\tilde{g}_{ij} \alpha_j}{\alpha_i} (z_j + \omega_j) + \frac{\tilde{g}_{in} \alpha_n}{\alpha_i} \left(n\bar{\omega} \psi \sum_{j=1}^{n-1} (z_j + \omega_j) \right) \\
 &= \omega_i + \sum_{j=1}^{n-1} \left[\frac{\tilde{g}_{ij} \alpha_j}{\alpha_i} - \frac{\tilde{g}_{in} \alpha_n}{\alpha_i} \right] (z_j + \omega_j) + \frac{\tilde{g}_{in} \alpha_n}{\alpha_i} n\bar{\omega}, \tag{A26}
 \end{aligned}$$

where the third step uses that fact that $\sum_{j=1}^n \omega_j = n\bar{\omega} \psi$ and that in equilibrium $\sum_{j=1}^n z_j = 0$, which imply:

$$z_n = - \sum_{j=1}^{n-1} z_j \tag{A27}$$

$$\omega_n = n\bar{\omega} \psi - \sum_{j=1}^{n-1} \omega_j. \tag{A28}$$

In matrix notation, Equation (A26) becomes:

$$\begin{aligned}
& \underbrace{\begin{bmatrix} 1 + \frac{1}{K_1 \alpha_1^2} & 0 & \cdots & 0 \\ 0 & 1 + \frac{1}{K_2 \alpha_2^2} & \cdots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & \cdots & & 1 + \frac{1}{K_{n-1} \alpha_{n-1}^2} \end{bmatrix}}_{\equiv \Lambda^{-1}} z = \omega \psi \\
& + \underbrace{\begin{bmatrix} \frac{\tilde{g}_{11} \alpha_1}{\alpha_1} & \frac{\tilde{g}_{1n} \alpha_n}{\alpha_1} & \frac{\tilde{g}_{12} \alpha_2}{\alpha_1} & \frac{\tilde{g}_{1n} \alpha_n}{\alpha_1} & \cdots & \frac{\tilde{g}_{1,n-1} \alpha_{n-1}}{\alpha_1} & \frac{\tilde{g}_{1n} \alpha_n}{\alpha_1} \\ \frac{\tilde{g}_{21} \alpha_1}{\alpha_2} & \frac{\tilde{g}_{2n} \alpha_n}{\alpha_2} & \frac{\tilde{g}_{22} \alpha_2}{\alpha_2} & \frac{\tilde{g}_{2n} \alpha_n}{\alpha_2} & \cdots & \frac{\tilde{g}_{2,n-1} \alpha_{n-1}}{\alpha_2} & \frac{\tilde{g}_{2n} \alpha_n}{\alpha_2} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \frac{\tilde{g}_{n-1,1} \alpha_1}{\alpha_{n-1}} & \frac{\tilde{g}_{n-1,n} \alpha_n}{\alpha_{n-1}} & \frac{\tilde{g}_{n-1,2} \alpha_2}{\alpha_{n-1}} & \frac{\tilde{g}_{n-1,n} \alpha_n}{\alpha_{n-1}} & \cdots & \frac{\tilde{g}_{n-1,n-1} \alpha_{n-1}}{\alpha_{n-1}} & \frac{\tilde{g}_{n-1,n} \alpha_n}{\alpha_{n-1}} \end{bmatrix}}_{\equiv B} (z + \omega \psi) + n \bar{\omega} \psi \underbrace{\begin{bmatrix} \frac{\tilde{g}_{1n} \alpha_n}{\alpha_1} \\ \frac{\tilde{g}_{2n} \alpha_n}{\alpha_2} \\ \vdots \\ \frac{\tilde{g}_{n-1,n} \alpha_n}{\alpha_{n-1}} \end{bmatrix}}_{\equiv C},
\end{aligned}$$

where $z = (z_1, z_2, \dots, z_{n-1})'$ is a $(n-1)$ by 1 column vector of net positions and $\omega \psi = (\omega_1, \omega_2, \dots, \omega_{n-1})'$ is a $(n-1)$ by 1 column vector of pre-trade exposures. We can rearrange the system of $(n-1)$ and solve for $\omega \psi$ as a function of z as follows:

$$\begin{aligned}
\Lambda^{-1} z &= \omega \psi + B(z + \omega \psi) + C \\
\omega \psi &= (I - B)^{-1} [(B - \Lambda^{-1} z) + C].
\end{aligned} \tag{A29}$$

Taking the model parameters as given, Equations (A27), (A28), and (A29) allow us to define pre-trade exposures that match observed net positions in the data.

Part 2: Risk aversion parameters as a function of net positions and pre-trade exposures. Starting from Equation (A22):

$$\begin{aligned}
z_i \left(1 + \frac{2}{K_i \alpha_i^2} \right) &= \omega_i + \sum_{j=1}^n \frac{\alpha_j}{\alpha_i} \tilde{g}_{ij} (z_j + \omega_j) \\
z_i \frac{2}{K_i} &= (z_i + \omega_i) \alpha_i + \sum_{j=1}^n \tilde{g}_{ij} (z_j + \omega_j) \alpha_j \\
(z_i + \omega_i) \alpha_i &= z_i \frac{2}{K_i} + \sum_{j=1}^n \tilde{g}_{ij} (z_j + \omega_j) \alpha_j \\
(z_i + \omega_i) \alpha_i &= z_i \frac{2}{K_i} + \sum_{j=1}^{n-1} \tilde{g}_{ij} (z_j + \omega_j) \alpha_j + \tilde{g}_{i,n} (z_n + \omega_n) \alpha_n \\
(z_i + \omega_i) \alpha_i &= z_i \frac{2}{K_i} + \sum_{j=1}^{n-1} \tilde{g}_{ij} (z_j + \omega_j) \alpha_j + \tilde{g}_{i,n} (z_n + \omega_n) \left[n \bar{\alpha} \sum_{j=1}^{n-1} \alpha_j \right] \\
(z_i + \omega_i) \alpha_i &= \tilde{g}_{i,n} (z_n + \omega_n) n \bar{\alpha} \psi + z_i \frac{2}{K_i} + \sum_{j=1}^{n-1} [\tilde{g}_{ij} (z_j + \omega_j) - \tilde{g}_{i,n} (z_n + \omega_n)] \alpha_j,
\end{aligned}$$

where $\bar{\alpha} \psi = \frac{1}{n} \sum_{j=1}^n \alpha_j$.

In matrix notation, we have:

$$\begin{aligned}
& \underbrace{\begin{bmatrix} z_1 + \omega_1 & 0 & \dots & 0 \\ 0 & z_2 + \omega_2 & \dots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & \dots & & z_{n-1} + \omega_{n-1} \end{bmatrix}}_{A_0} \underbrace{\begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_{n-1} \end{bmatrix}}_{\alpha} = \underbrace{\begin{bmatrix} \tilde{g}_{1,n}(z_n + \omega_n) n \bar{\alpha} \psi z_1 \frac{2}{K_1^2} \\ \tilde{g}_{2,n}(z_n + \omega_n) n \bar{\alpha} \psi z_2 \frac{2}{K_2^2} \\ \vdots \\ \tilde{g}_{n-1,n}(z_n + \omega_n) n \bar{\alpha} \psi z_{n-1} \frac{2}{K_{n-1}^2} \end{bmatrix}}_{A_1} \\
& + \underbrace{\begin{bmatrix} \tilde{g}_{11}(z_1 + \omega_1) & \tilde{g}_{1,n}(z_n + \omega_n) & \dots & \tilde{g}_{1,n-1}(z_{n-1} + \omega_{n-1}) & \tilde{g}_{1,n}(z_n + \omega_n) \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \tilde{g}_{n-1,1}(z_1 + \omega_1) & \tilde{g}_{n-1,n}(z_n + \omega_n) & \dots & \tilde{g}_{n-1,n-1}(z_{n-1} + \omega_{n-1}) & \tilde{g}_{n-1,n}(z_n + \omega_n) \end{bmatrix}}_{A_2} \underbrace{\begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_{n-1} \end{bmatrix}}_{\alpha}
\end{aligned} \tag{A30}$$

If $A_0 - A_2$ is not singular, we have:

$$\begin{aligned}
A_0 \alpha \psi &= A_1 + A_2 \alpha \psi \\
\alpha \psi &= [A_0 - A_2]^{-1} A_1
\end{aligned} \tag{A31}$$

Notice that, in addition of having $A_0 - A_2$ nonsingular, the model requires $\alpha_i > 0$ for every agent i so that agents' objective function is always concave and its first-order conditions are valid. These conditions translate into $[A_0 - A_2]^{-1} A_1$ having only positive entries and its sum being lower than $n\bar{\alpha}$. \square

B.1.2 Calibration

To calibrate the heterogeneous risk aversion model, we choose $\bar{\alpha}$ and ψ to match the spreads \bar{R}_d and $\bar{R}_c - \bar{R}_d$. We set agents' risk aversion according to Equation (A31). However, we need to know the net positions of all agents as well as their pre-trade exposures. We observe dealers' net positions in the data. For costumers, we observe only their net notional position in the CDS market but not their equity value. We assume customers net notional positional is proportional to the equity value, that is, we assume that customers have the same net notional per unit of equity in absolute value. In the model, this translates into having $|z_j| = |z_k|$ for any two customers j and k , and we can infer the absolute values of customers net positions from the fact the all net positions sum to zero.

Finally, we need to define agents' pre-trade exposures. If we assume identical pre-trade exposure for every agents, that is, $\omega_i = \bar{\omega}$ for every agent i , then condition 2 of Part 2 of Proposition A4 is violated when trying to match the data. Specifically, the risk aversion parameter implied by Equation (A31) is negative for the largest net seller dealer. Hence, heterogeneity in risk aversion alone is unable to explain net positions observed in the data.

To overcome the inability of the heterogeneous risk aversion model to match the empirically observed net positions of dealers, we singled out the largest net seller dealer and adjust its pre-trade exposure. Formally, let the largest net seller dealer be agent 1. We reduce its pre-trade exposure in order to satisfy the conditions in Part 2 of Proposition A4. We keep the average pre-trade exposure unchanged by equally adjusting other agents pre-trade exposure. For $j = 2, \dots, n$, we have $\omega_j = (n\bar{\omega} - \omega_1)/(n - 1)$. We solve the model for various values of ω_1 , all which will satisfy proposition conditions. We report results in Table A1.

Across columns, we vary the pre-trade exposure of the largest net seller dealer from $\omega_1 = 2$, which is the highest value of ω_1 that satisfy the propositions conditions, to $\omega_1 = 30$. In each column, we matches both \bar{R}_d and $\bar{R}_c - \bar{R}_d$ to the values observed in the data. Panel A shows calibrated parameters. Trading cost parameter remains unchanged across columns. However, the average risk aversion of all agents declines as pre-trade exposure of dealer 1 decreases. With a lower pre-trade exposure, dealer 1 becomes more willing to sell protection to other market participants. Each calibration also adjusts agents' risk aversion to match their net positions. This implies a lower overall risk aversion to make other agents less willing to buy protection

from dealer 1.

In Panel B, we report the largest net seller dealer removal effects on prices and risk allocation. The effect of dealer prices are nearly identical to the one observed in our benchmark model. They vary from 31.37 to 31.38 basis points, while this effect is 31.31 basis points in our baseline model. The effect on customer prices is also nearly identical, varying from 15.81 to 15.82 basis points. This effect in our baseline model is 15.79 basis points. Additionally, the effects on risk allocation are identical to the baseline model despite having larger variation in dealers' pre-trade exposures. The change in dealers' pre-trade exposure vary from 0.21 to 2.22, whereas it is 0.44 in our benchmark model. However, the change in dealers' average net position is the same as in our benchmark model at 0.138. That is, the dealer being removed changes dealers' average pre-trade exposures significantly but the risk reallocation after the dealer removal is similar to the benchmark model. This confirms that risk aversion acts as a substitute for pre-trade exposures that delivers similar dealer removal effects on prices.

B.2 Trading Costs Heterogeneity

In this section, we present a variation of our benchmark model where trading cost is specific to each counterparty type. Formally, agent i 's optimization problem is given by:

$$\begin{aligned} \max_{z_i, \{i_j\}_{j=1}^n} \quad & \omega_i(1 - \mu) + \sum_{j=1}^n i_j(R_{ij} - \mu) - \frac{\alpha\psi}{2}(\omega_i + z_i)^2 - \frac{1}{2} \sum_{j=1}^n i_j z_j^2 \\ \text{s.t.} \quad & i_j = 0 \text{ if } g_{ij} = 0 \\ & z_i - \sum_{j=1}^n i_j = 0, \end{aligned}$$

where x are the state variables.

From the first-order conditions, we can solve agent i 's optimization problem for agents' bilateral and net positions:

$$i_j = \begin{cases} \frac{1}{i_j} [R_{ij} - \mu - \alpha(\omega_i + z_i)] & \text{if } g_{ij} = 1 \\ 0 & \text{if } g_{ij} = 0 \end{cases} \quad (\text{A32})$$

where

$$z_i = \sum_{j=1}^n i_j = \frac{\frac{1}{\sum_{j=1}^n \frac{g_{ij}}{\phi_{ij}}} \sum_{j=1}^n \frac{g_{ij}}{i_j} (R_{ij} - \mu) - \alpha\omega_i}{\frac{1}{\sum_{j=1}^n \frac{g_{ij}}{\phi_{ij}}} + \alpha\psi} \quad (\text{A33})$$

Market clearing conditions are the same as in the benchmark model and are given by:

$$i_j + j_i = 0 \quad \forall i, j = 1, \dots, n. \quad (\text{A34})$$

As in the benchmark model, we assume no transaction costs between any two counterparties:

$$R_{ij} = R_{ji}.$$

If agents i and j can trade (i.e., $g_{ij} = g_{ji} = 1$), we can use their optimality conditions from Equation (A32) and market clearing condition from Equation (A34) to solve for the equilibrium price of a contract between agents i and j :

$$\underbrace{R_{ij} - \mu}_{\text{contract premium}} = \left[\frac{\frac{1}{i_j}(\omega_i + z_i) + \frac{1}{j_i}(\omega_j + z_j)}{\frac{1}{i_j} + \frac{1}{j_i}} \right] \alpha\psi = \left[\frac{j_i(\omega_i + z_i) + i_j(\omega_j + z_j)}{i_j + j_i} \right] \alpha\psi. \quad (\text{A35})$$

If we substitute the contract premium (Equation A35) into the agents' first-order condition (Equation

(A32), then whenever $g_{ij} = 1$ we get:

$$g_{ij} = \frac{1}{ij} [R_{ij} - \mu - \alpha(\omega_i + z_i)^2] = \frac{\alpha\sigma^2\psi}{ij + ij} [(\omega_j + z_j) - (\omega_i + z_i)]$$

which means that agent i sells insurance to agent j (i.e., $g_{ij} > 0$) whenever agent j is more exposure to default than agent i after trade, that is, $\omega_j + z_j > \omega_i + z_i$. The trade volume between two counterparties depends on the difference of their post-trade exposures scaled by risk aversion and on the sum of their trading cost parameters.

The net positions, $\{z_i\}_i$, are determined in equilibrium. We can use Equations (A33) and (A35) to solve for equilibrium net positions:

$$R_{ij} - \mu = \alpha \left[\frac{ji\alpha\psi}{ij + ji} (\omega_i + z_i) + \frac{ij\alpha\psi}{ij + ji} (\omega_j + z_j) \right]$$

$$\sum_{j=1}^n \frac{g_{ij}}{ij} (R_{ij} - \mu) = \alpha \sum_{j=1}^n \frac{g_{ij}}{ij} \left[\frac{ji\alpha\psi}{ij + ji} (\omega_i + z_i) + \frac{ij\alpha\psi}{ij + ji} (\omega_j + z_j) \right].$$

Let $\bar{g}_{ij} = \frac{g_{ij}}{ij + ji} \frac{1}{\bar{K}_i}$ and $\bar{K}_i = \sum_{j=1}^n \frac{g_{ij}}{ij + ji}$, then:

$$\sum_{j=1}^n \frac{g_{ij}}{ij} (R_{ij} - \mu) = \alpha (z_i + \omega_i) \sum_{j=1}^n \bar{K}_j \alpha \bar{g}_{ji} \frac{ji}{ij} + \alpha \sum_{j=1}^n \bar{K}_i \alpha \bar{g}_{ij} (z_j + \omega_j).$$

Using:

$$\sum_{j=1}^n \frac{g_{ij}}{ij} (R_{ij} - \mu) = z_i \left(1 + \alpha\sigma^2\psi \sum_{j=1}^n \frac{g_{ij}}{ij} \right) + \alpha\sigma^2\psi \omega_i \sum_{j=1}^n \frac{g_{ij}}{ij}$$

from Equation (A33) and using the fact that $g_{ij} = g_{ji}$, we have:

$$z_i \left(1 + \alpha\sigma^2\psi \sum_{j=1}^n \frac{g_{ij}}{ij} \right) + \alpha\sigma^2\psi \omega_i \sum_{j=1}^n \frac{g_{ij}}{ij} = \alpha (z_i + \omega_i) \sum_{j=1}^n \bar{K}_j \alpha \bar{g}_{ji} \frac{ji}{ij} + \alpha \sum_{j=1}^n \bar{K}_i \alpha \bar{g}_{ij} (z_j + \omega_j)$$

$$z_i \left(\sum_{j=1}^n \frac{g_{ij}}{ij} + \frac{1}{2\alpha\psi} \right) + \omega_i \sum_{j=1}^n \frac{g_{ij}}{ij} = (z_i + \omega_i) \sum_{j=1}^n \bar{K}_j \bar{g}_{ji} \frac{ji}{ij} + \sum_{j=1}^n \bar{K}_i \bar{g}_{ij} (z_j + \omega_j)$$

$$z_i \left(\sum_{j=1}^n \frac{g_{ij}}{ij} - \sum_{j=1}^n \bar{K}_j \bar{g}_{ji} \frac{ji}{ij} + \frac{1}{2\alpha} \right) = \omega_i \left(\sum_{j=1}^n \frac{g_{ji}}{ij} - \sum_{j=1}^n \bar{K}_j \bar{g}_{ji} \frac{ji}{ij} \right) + \sum_{j=1}^n \bar{K}_i \bar{g}_{ij} (z_j + \omega_j).$$

Notice that

$$\sum_{j=1}^n \frac{g_{ij}}{ij} - \sum_{j=1}^n \bar{K}_j \bar{g}_{ji} \frac{ji}{ij} = \sum_{j=1}^n \frac{g_{ij}}{ij} \left(1 - \frac{ji}{ij + ji} \right) = \sum_{j=1}^n \frac{g_{ij}}{ij} \frac{ij}{ij + ji} = \sum_{j=1}^n \frac{g_{ij}}{ij + ji} = \bar{K}_i.$$

The system of equation can be simplified to:

$$z_i \left(\bar{K}_i + \frac{1}{\alpha\sigma^2\psi} \right) = \omega_i \bar{K}_i + \sum_{j=1}^n \bar{K}_i \bar{g}_{ij} (z_j + \omega_j)$$

$$z_i \left(1 + \frac{1}{\bar{K}_i \alpha\sigma^2\psi} \right) = \omega_i + \sum_{j=1}^n \bar{g}_{ij} (z_j + \omega_j) \quad \forall i = 1, \dots, n \quad (\text{A36})$$

In matrix notation, we have:

$$\underbrace{\begin{bmatrix} 1 + \frac{1}{K_1 \alpha \sigma^2} & 0 & \cdots & 0 \\ 0 & 1 + \frac{1}{K_2 \alpha \sigma^2} & \cdots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & \cdots & & 1 + \frac{1}{K_n \alpha \sigma^2} \end{bmatrix}}_{\equiv \Lambda^{-1}} z = \omega \psi + \underbrace{\begin{bmatrix} \bar{g}_{11} & \bar{g}_{12} & \cdots & \bar{g}_{1n} \\ \bar{g}_{21} & \bar{g}_{22} & \cdots & \bar{g}_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \bar{g}_{n1} & \bar{g}_{n2} & \cdots & \bar{g}_{nn} \end{bmatrix}}_{\equiv G^*} (z + \omega)$$

$$\Lambda^{-1} z = \omega \psi + G^* (z + \omega), \quad (\text{A37})$$

where $z = [z_1, \dots, z_n]'$ and $\omega \psi = [\omega_1, \dots, \omega_n]'$ are column vectors of net positions and pre-trade exposures. We can rearrange Equation (A37) and solve the system of equations for the equilibrium net positions and post-trade exposures:

$$z = (\mathbf{I} - \Lambda G^*)^{-1} \Lambda (G^* - \mathbf{I}) \omega \psi \quad (\text{A38})$$

and

$$z + \omega \psi = (\mathbf{I} - \Lambda G^*)^{-1} (\mathbf{I} - \Lambda) \omega. \quad (\text{A39})$$

As in the benchmark model, post-trade exposures are determined jointly in equilibrium. They depend on the network, risk aversion and trading cost parameters, and pre-trade exposures.

Notice that the matrix $(G^* - \mathbf{I})$ is singular because its rank is less than n given that each row in G^* sum to one. Hence, we cannot use directly Equation (A38) to solve for pre-trade positions (ω 's) as a function of net positions (z 's). To overcome this issue, we rearrange Equation (A36) for $i = 1, \dots, n-1$ as follows:

$$\begin{aligned} z_i \left(1 + \frac{1}{K_i \alpha \sigma^2} \right) &= \omega_i + \sum_{j=1}^n \bar{g}_{ij} (z_j + \omega_j) \\ &= \omega_i + \sum_{j=1}^{n-1} \bar{g}_{ij} (z_j + \omega_j) + \bar{g}_{in} (z_n + \omega_n) \\ &= \omega_i + \sum_{j=1}^{n-1} \bar{g}_{ij} (z_j + \omega_j) + \bar{g}_{in} \left(n \bar{\omega} \psi - \sum_{j=1}^{n-1} (z_j + \omega_j) \right) \\ &= \omega_i + \sum_{j=1}^{n-1} [\bar{g}_{ij} - \bar{g}_{in}] (z_j + \omega_j) + \bar{g}_{in} n \bar{\omega} \end{aligned}$$

where the third step uses that fact that $\sum_{j=1}^n \omega_j = n \bar{\omega} \psi$ and that in equilibrium $\sum_{j=1}^n z_j = 0$, which imply:

$$z_n = - \sum_{j=1}^{n-1} z_j \quad (\text{A40})$$

$$\omega_n = n \bar{\omega} \psi - \sum_{j=1}^{n-1} \omega_j \quad (\text{A41})$$

In matrix notation this system of equations becomes:

$$\underbrace{\begin{bmatrix} 1 + \frac{1}{\bar{K}_1 \alpha \sigma^2} & 0 & \cdots & 0 \\ 0 & 1 + \frac{1}{\bar{K}_2 \alpha \sigma^2} & \cdots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & \cdots & & 1 + \frac{1}{\bar{K}_{n-1} \alpha \sigma^2} \end{bmatrix}}_{\equiv \Lambda^{-1}} z = \omega \psi$$

$$+ \underbrace{\begin{bmatrix} \bar{g}_{11} & \bar{g}_{1n} & \bar{g}_{12} & \bar{g}_{1n} & \cdots & \bar{g}_{1,n-1} & \bar{g}_{1n} \\ \bar{g}_{21} & \bar{g}_{2n} & \bar{g}_{22} & \bar{g}_{2n} & \cdots & \bar{g}_{2,n-1} & \bar{g}_{2n} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \bar{g}_{n-1,1} & \bar{g}_{n-1,n} & \bar{g}_{n2} & \bar{g}_{n-1,n} & \cdots & \bar{g}_{n-1,n-1} & \bar{g}_{n-1,n} \end{bmatrix}}_{\equiv B} (z + \omega \psi) + \underbrace{n \bar{\omega} \psi}_{\equiv C} \begin{bmatrix} \bar{g}_{1n} \\ \bar{g}_{2n} \\ \vdots \\ \bar{g}_{n-1,n} \end{bmatrix},$$

where $z = (z_1, z_2, \dots, z_{n-1})'$ is a $n-1$ by 1 column vector of net positions and $\omega \psi = (\omega_1, \omega_2, \dots, \omega_{n-1})'$ is a $n-1$ by 1 column vector of pre-trade exposures. We can rearrange the system of $n-1$ and solve for $\omega \psi$ as a function of z as follows:

$$\Lambda^{-1} z = \omega \psi + B(z + \omega \psi) + C$$

$$\omega \psi = (I - B)^{-1} [(B - \Lambda^{-1} z) + C]. \quad (\text{A42})$$

Taking the model parameters as given, Equations (A40), (A41), and (A42) allow us to define pre-trade exposure to match observed net positions in the data.

B.2.1 Different trading costs for dealers and customer trades

Core-periphery Network Next we derive a closed-form expression for the average net position of dealers, namely \bar{z}_d , under a core-periphery network. Let us assume that i_j depends on whether the counterparties are dealers or not. Specifically, let i_j be specified as follows:

$$i_j = \begin{cases} d & \text{if } i \text{ and } j \text{ are both dealers} \\ c & \text{otherwise} \end{cases}, \quad (\text{A43})$$

which implies

$$\bar{K}_i = \sum_{j=1}^n \frac{g_{ij}}{i_j + j_i} = \begin{cases} \frac{n_d}{2} + \frac{n_c}{2} & \text{for } i = 1, \dots, n_d \text{ (dealers)} \\ \frac{n_d+1}{2} & \text{for } i = n_d+1, \dots, n \text{ (customers)} \end{cases} \quad (\text{A44})$$

$$\bar{g}_{ij} = \frac{g_{ij}}{i_j + j_i} \frac{1}{\bar{K}_i} = \begin{cases} \frac{g_{ij}}{d+d} \frac{1}{\bar{K}_i} & i \text{ and } j \text{ are dealers} \\ \frac{g_{ij}}{c+c} \frac{1}{\bar{K}_i} & i \text{ and } j \text{ are customers} \\ \frac{g_{ij}}{c+d} \frac{1}{\bar{K}_i} & \text{dealer and customer} \end{cases} \quad (\text{A45})$$

For dealer i , Equation (A36) becomes:

$$z_i \left(1 + \frac{1}{\bar{K}_i \alpha \sigma^2} \right) = \omega_i + \sum_{j=1}^n \bar{g}_{ij} (z_j + \omega_j) \quad \forall i = 1, \dots, n_d$$

$$z_i \left(\bar{K}_i + \frac{1}{\alpha \sigma^2} \right) = \bar{K}_i \omega_i + \sum_{j=1}^n \frac{g_{ij}}{i_j + j_i} (z_j + \omega_j)$$

$$z_i \frac{1}{\alpha \sigma^2} = \bar{K}_i (z_i + \omega_i) + \sum_{j=1}^n \frac{g_{ij}}{i_j + j_i} (z_j + \omega_j)$$

$$\begin{aligned}
z_i \frac{1}{\alpha \sigma \psi} &= \left[\frac{n_d}{2} + \frac{n}{2} \frac{n_d}{c} \right] (z_i + \omega_i) + \sum_{j=1}^{n_d} \frac{1}{2} \frac{1}{d} (z_j + \omega_j) + \sum_{j=n_d+1}^n \frac{1}{2} \frac{1}{c} (z_j + \omega_j) \\
z_i \frac{1}{\alpha \sigma \psi} &= \left[\frac{n_d}{2} + \frac{n}{2} \frac{n_d}{c} \right] (z_i + \omega_i) + \frac{n_d}{2} (\bar{z}_d + \bar{\omega}_d) + \frac{n}{2} \frac{n_d}{c} (\bar{z}_c + \bar{\omega}_c), \tag{A46}
\end{aligned}$$

and taking average over dealers, we have:

$$\begin{aligned}
\bar{z}_d \frac{1}{\alpha \sigma \psi} &= \left[\frac{n_d}{2} + \frac{n}{2} \frac{n_d}{c} \right] (\bar{z}_d + \bar{\omega}_d) + \frac{n_d}{2} (\bar{z}_d + \bar{\omega}_d) + \frac{n}{2} \frac{n_d}{c} (\bar{z}_c + \bar{\omega}_c) \\
\bar{z}_d \frac{1}{\alpha \sigma \psi} &= \frac{n}{2} \frac{n_d}{c} (\bar{z}_d + \bar{\omega}_d) + \frac{n}{2} \frac{n_d}{c} (\bar{z}_c + \bar{\omega}_c) \\
\bar{z}_d \frac{1}{\alpha \sigma \psi} &= \frac{n}{2} \frac{n_d}{c} (\bar{z}_d + \bar{\omega}_d) + \frac{n}{2} \frac{n_d}{c} \left(\frac{n_d}{n} \frac{\bar{z}_d}{n_d} + \frac{n}{n} \frac{\bar{\omega}}{n_d} \frac{n_d}{n} \frac{\bar{\omega}_d}{n_d} \right) \\
\bar{z}_d \frac{1}{\alpha \sigma \psi} &= \frac{n}{2} \frac{n_d}{c} (\bar{z}_d + \bar{\omega}_d) + \frac{1}{2} \frac{1}{c} (n_d \bar{z}_d + n \bar{\omega} \psi \frac{n_d}{n} \bar{\omega}_d) \\
\frac{1}{2} \bar{z}_d \frac{1}{\alpha \sigma \psi} &= (n \frac{n_d}{c} (\bar{z}_d + \bar{\omega}_d) + n_d \bar{z}_d + n \bar{\omega} \psi \frac{n_d}{n} \bar{\omega}_d) \\
\frac{1}{2} \bar{z}_d \frac{1}{\alpha \sigma \psi} &= n (\bar{z}_d + \bar{\omega}_d) + n \bar{\omega} \\
\bar{z}_d &= \frac{\bar{\omega} + \bar{\omega}_d}{\frac{2}{n \alpha \sigma^2} + 1}, \tag{A47}
\end{aligned}$$

which means that \bar{z}_d does not depends on d .

To get some intuition behind this result, let us write Equation (A36) as follows:

$$\begin{aligned}
z_i \left(1 + \frac{1}{\bar{K}_i \alpha \sigma \psi} \right) &= \omega_i + \sum_{j=1}^n \bar{g}_{ij} (z_j + \omega_j) \quad \forall i = 1, \dots, n_d \\
z_i \frac{\bar{K}_i \alpha \sigma \psi + 1}{\bar{K}_i \alpha \sigma \psi} &= \omega_i + \sum_{j=1}^n \bar{g}_{ij} (z_j + \omega_j) \\
z_i &= \bar{\omega}_d \omega_i + \bar{\omega}_d \sum_{j=1}^n \bar{g}_{ij} (z_j + \omega_j) \\
z_i + \omega_i &= \frac{1}{\bar{\omega}_d} \omega_i + \bar{\omega}_d \sum_{j=1}^n \bar{g}_{ij} (z_j + \omega_j), \tag{A48}
\end{aligned}$$

where $\bar{\omega}_d = \frac{\bar{K}_d \alpha \sigma^2}{\bar{K}_d \alpha \sigma^2 + 1}$ and $\bar{K}_d = \frac{n_d}{2} + \frac{n}{2} \frac{n_d}{c}$.

When d decreases, then $\bar{\omega}_d$ increases because \bar{K}_d increases. This has different implications for risk sharing from the perspective of dealer i . On the one hand, the post-trade exposure of dealer i depends less on its pre-trade exposure. Intuitively, as dealer i faces lower trading costs, there is more risk-sharing as it becomes cheaper to trade and, as a result, its post-trade exposure moves away from its pre-trade exposure to become more similar to its trading counterparties. On the other hand, the increase in risk sharing is asymmetric between dealer i 's counterparties. The term $\sum_{j=1}^n \bar{g}_{ij} (z_j + \omega_j)$ is a weighted average of dealer i trading counterparties. As d decreases, this average puts more weight on dealers' post-trade exposure and less weight on customers' post-trade exposures. This is because it is cheaper to trade with dealers when d decreases but not with customers. As a result, dealer i will share more risk with other dealers. When we average across dealers, the increase in risk sharing is concentrated among them and d has no effect on their average net position and no effect on how much protection they sell to customers.

Equation (A47) also has implications for the calibration of c and α . Notice that the average prices

between dealers is given by:

$$\bar{R}_d = \mu + \alpha\sigma\psi(\bar{z}_d + \bar{\omega}_d), \quad (\text{A49})$$

while the average prices between customers and dealer is given by:

$$\bar{R}_c = \mu + \alpha\sigma\psi\frac{1}{2}(\bar{z}_d + \bar{\omega}_d + \bar{z}_c + \bar{\omega}_c) = \mu + \alpha\sigma\psi\frac{1}{2}\left(\bar{z}_d + \bar{\omega}_d + \frac{n_d\bar{z}_d}{n} + \frac{\bar{\omega}\psi}{n} + \frac{n_d\bar{\omega}_d}{n_d}\right), \quad (\text{A50})$$

where the second equality uses the fact that $\bar{\omega}_c = \frac{\bar{\omega}}{n} + \frac{n_d\bar{\omega}_d}{n_d}$ and $\bar{z}_c = \frac{n_d\bar{z}_d}{n}$.

Given that we observe \bar{z}_d , \bar{R}_d , and \bar{R}_c , and we normalized $\bar{\omega}\psi$ to one, we can solve Equations (A47), (A49), and (A50) for α , σ , and ψ . Similarly, we can compute dealer removal effects by changing $\bar{\omega}_d$. We still need to infer the model-implied value of ω_i of the dealer being removed. We can use Equation (A46) to write ω_i as a function of net positions and other parameters of the model. Upon the removal of a dealer, we compute $\bar{\omega}_d$ and the new equilibrium average net position of dealer according to Equation (A47). Then we use Equations (A49) and (A50) to assess the effects on prices.

Proposition A5. *In the model with heterogeneous trading costs, $\bar{\omega}_d$ has no effect on the average net positions of dealer (\bar{z}_d), the average dealer prices (\bar{R}_d), and the average customer prices (\bar{R}_c). Furthermore, let α^* and σ^* be the parameters that make the average prices in the heterogeneous trading cost model to match a specific target. These values are identical to the values of α and σ that match the same targeted prices under the benchmark model.*

Proof. Equations (A47), (A49) and (A50) show that $\bar{\omega}_d$ has no effect on the average net positions of dealer (\bar{z}_d), the average dealer prices (\bar{R}_d), and the average customer prices (\bar{R}_c).

For the second part of the proposition, notice that $\bar{\omega}_c = \frac{\bar{\omega}}{n} + \frac{n_d\bar{\omega}_d}{n_d}$, we have that the heterogeneous trading cost model is identical to our benchmark framework. Therefore the calibrated parameters would be identical in this case. Given that $\bar{\omega}_d$ does not affect the equilibrium average prices in both dealer and customer markets, we have that $\bar{\omega}_d$ does not affect the calibrated parameters. \square

Complete Network Average Prices Next, we compute the average price under the assumption of a complete trading networks—that is, $g_{ij} = 1$ for every i and j . In this case, the previous derivations of average net positions and average prices among dealers and customer (Equations (A47), (A49) and (A50)) remain unchanged. These derivations rely on the assumption that dealers can trade with all other agents, that is, $g_{ij} = 1$ for every j whenever i is a dealer. Therefore, in a complete network, the average net position among dealers is given by Equation (A47), and average prices in dealer-dealer and customer-dealer trades are given by Equations (A49) and (A50). For customer-customer trades, the average price is given by:

$$\bar{R}_{cc} = \mu + \alpha\sigma\psi(\bar{z}_c + \bar{\omega}_c) = \mu + \alpha\sigma\psi\left(\frac{n_d\bar{z}_d}{n} + \frac{\bar{\omega}\psi}{n} + \frac{n_d\bar{\omega}_d}{n_d}\right). \quad (\text{A51})$$

Finally, we can infer the average price in the economy from the average prices in dealer-dealer, customer-dealer and customer-customer trades:

$$\bar{R}_{\text{Complete Network}} = \frac{n_d^2\bar{R}_d + 2n_d(n - n_d)\bar{R}_c + (n - n_d)^2\bar{R}_{cc}}{n^2}, \quad (\text{A52})$$

where \bar{R}_d , \bar{R}_c and \bar{R}_{cc} are defined in Equations (A49), (A50) and (A51).

C Algorithm to Classify Market Participants as Dealers

In this section, we describe a minimum-distance algorithm that we use to determine the size of the empirical core in the CDS market. The algorithm uses the fact that a pure core-periphery network requires that all dealers should be connected to each other and to every customer. Moreover, a pure core-periphery network stipulates that customers should be connected to all dealers and no one else. Based on these two observations, we select the number and identity of dealers and assign remaining agents to the periphery as follows:

1. Choose a threshold number of connections, m , above which a counterparty will be classified as a dealer. If the number of connections is below this threshold, we label that agent as a customer. Define $D_{i,t} \equiv \sum_j G_{i,j,t}$ as counterparty i 's degree on date t . In words, $D_{i,t}$ just counts the number of i 's trading partners. For a given threshold m , agent i is a dealer if $D_i \geq m$ and i is a customer otherwise.
2. For each threshold m and its implied definition of dealers and customers, we construct a counterfactual network that is perfectly core-periphery, that is, a network in which everyone is connected to all dealers but not to other customers. Let this counterfactual core-periphery network be $G^{CP(m)} = (g_{ij}^{CP(m)})_{ij}$. Formally, $g_{ii}^{CP(m)} = 1$ for every i , and for $i \neq j$

$$g_{ij}^{CP(m)} = \begin{cases} 1 & \text{if } D_{j,t} \geq m \\ 0 & \text{otherwise} \end{cases}.$$

3. We then compute the number of connections that should exist under a perfect core-periphery network but do not exist in the data, as well as the number of connections that do not exist in the data but should exist under a perfect core-periphery network. This is the number of elements of t that are not consistent with a core-periphery network. We then minimize over choices of m the average number of connections inconsistent with a core-periphery relative to the total number of connections under a perfect core-periphery network. Hence, the minimization problem is given by:

$$\min_m \frac{1}{N} \sum_{i=1}^N \frac{\sum_{j=1}^N G_{i,j,t} - g_{ij}^{CP(m)}}{\sum_{j=1}^N g_{ij}^{CP(m)}},$$

where N is the total number of counterparties.

Empirically, the algorithm generates a counterparty network with 14 dealers. Furthermore, in the Internet Appendix we show for robustness that our selection algorithm consistently identifies the same set of dealers even if we focus on subsamples of our data.

clearpage

Table A1: Dealer Removal Effects with Heterogeneous $\alpha\psi$

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Calibrated Parameters and Equilibrium Summary						
	7.983	7.983	7.983	7.983	7.983	7.983
Avg. $\alpha\psi$	4.55	4.39	4.36	4.33	4.30	4.24
ω_1	2	5	10	15	20	30
Panel B: Dealer Removal Effects						
$\Delta\bar{R}_d$	31.38	31.38	31.37	31.37	31.37	31.37
$\Delta\bar{R}_c$	15.82	15.82	15.82	15.82	15.82	15.81
$\Delta\bar{\omega}_d$	0.21	0.43	0.79	1.14	1.50	2.22
Δz_d	0.138	0.138	0.138	0.138	0.138	0.138

Notes: This table reports different calibration of the model with heterogeneous risk aversion parameters. We assume all agents have the same pre-trade exposure, except for the largest net seller dealer (agent 1). That is, $\omega_j = (n\bar{\omega} - \omega_1)/(n - 1)\forall j = 2, \dots, n$. We solve the model for various values of ω_1 , all which will satisfy the conditions in Proposition A4. Parameters $\bar{\omega}$ and the average risk aversion are calibrated to match the average dealer and customer prices at 133 bps and 138.12 bps, respectively. Across columns, we vary the pre-trade exposure of the largest net seller dealer from 2, which is the highest value of ω_1 that satisfy the conditions in Proposition A4, to 30. Panel A shows calibrated parameters, and Panel B shows dealer removal effects on dealers' average prices, customers' average prices, dealers' average pre-trade exposure and dealers' average net positions. See Section B.1.2 for details. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.