

18-05 | August 29, 2018

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 U.S. Department of the 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: August 28, 2018

Abstract

Over-the-counter (OTC) markets for financial assets are dominated by a relatively small number of core intermediaries and a large number of peripheral customers. In this paper, we develop a model of trade in a core-periphery network and estimate its key structural parameters using proprietary credit default swap data from the Depository Trust & Clearing Corporation (DTCC). Using our calibrated model, we provide quantitative estimates of: (1) the effect of network frictions on the level of OTC derivatives prices; (2) the key determinants of cross sectional dispersion in bilateral prices; and (3) how prices and risk-sharing change in response to the failure of a dealer.

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

*The views expressed in this paper are those of the authors and do not reflect the position of the 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 Darrell Duffie (discussant), Maryam Farboodi (discussant), Francisco Gomes (discussant), Felipe Iachan (discussant), Pete Kyle (discussant) as well as participants at the Laboratory for Aggregate Economics and Finance OTC Conference, Laboratory for Aggregate Economics and Finance Women in Macroeconomics Conference, 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 conference, the Macro Financial Modeling Annual Meeting, Rome Junior Finance Conference, Lubrafin, the Maryland Junior Finance Conference, and the NBER 2018 Summer Asset Pricing Meetings. We also thank seminar participants at UCLA Macro Finance lunch, UCLA Finance Brown Bag, Pontifical Catholic University of Rio de Janeiro (PUC-Rio), and University of Washington 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 over-the-counter (OTC), including virtually all corporate bonds, sovereign fixed-income instruments, and swaps (e.g. interest rate, currency, and credit). Trade in OTC markets occurs bilaterally between pairs of counterparties who are connected through a trading network. These networks are usually incomplete in the sense that trade does not occur between every pair of counterparties.¹ Indeed, OTC networks are typically characterized by a core-periphery structure, in which a core set of interconnected dealers trade with each other and with a peripheral set of clients (Li and Schürhoff (2018); Di Maggio, Kermani, and Song (2017)). In contrast, the clients in the periphery of the network are largely unconnected to each other and essentially trade exclusively with core dealers. The fact that trade occurs bilaterally also means that traders carefully consider how to size their positions with their counterparties. For example, it is common risk management practice to impose line limits that prevent large exposures to any single counterparty.

How does the core-periphery network, and the institutional features that characterize trade within it, impact equilibrium prices and allocations? In this paper, we build and calibrate a model of OTC trading in credit default swaps (CDS) to answer these questions. We model an incomplete trading network—with the core-periphery as a special case—and add two other practical features of OTC market participants. Investors in our model are risk averse over future payoffs and are endowed with different exposures to aggregate default risk, meaning that agents are heterogeneous in their effective willingness to bear default risk. Endowing traders with varying exposure to fundamental default risk is our way of modeling, say, a cash position in risky bonds or loans. Agents can trade away from this initial exposure by entering into bilateral credit default swap (CDS) contracts with their connected trading partners. However, we also assume that traders have an aversion to concentrated bilateral exposures. Thus, in equilibrium, each trader must balance its desire to hedge aggregate default risk against exposing itself too much to any single counterparty.²

We estimate the structural parameters of our model by matching observed prices, quantities, and the empirical network structure in the CDS market. This allows us to quantitatively assess how much the trading network structure, along with aversion to both asset payoff risk and concentrated bilateral trades, matters for equilibrium outcomes. Our estimation approach relies on proprietary data covering reported CDS trades in the United States from 2010 to 2013. The data is provided to the Office of Financial Research (OFR) from the Depository Trust & Clearing Corporation (DTCC).

Our quantitative inference is driven by three key facts about the CDS market. First, we confirm that, like in many OTC markets, the CDS market displays a core-periphery structure with a small number of dealers serving the entire market. Second, during our sample, we find that, as a

¹OTC trading networks are naturally incomplete because, at least in the short run, establishing bilateral trade agreements is costly.

²The idea that traders manage counterparty risk through quantities is broadly consistent with the evidence in Du, Gadgil, Brody, and Vega (2017).

group, the core set of dealers is on average a net seller of credit protection to periphery customers. Moreover, the net provision of credit insurance is highly concentrated in a small subset of dealers within the core. Finally, we document price dispersion, meaning that the same underlying CDS contract can trade at different prices. In particular, we find that customer-dealer transaction prices are on average higher than dealer-dealer transaction prices.

These three facts allow us to infer the structural parameters of our model, namely investors' risk-bearing capacity and aversion to counterparty concentration. With these parameters in hand, we provide a quantitative assessment of how much the network structure and frictions matter for both the average level of credit spreads in the economy and the amount of price dispersion in the market. In addition, we explore how the trading network interacts with dealer failure, a key concern of OTC market regulators.

We start by benchmarking our economy against one in which all agents can trade with each other. Surprisingly, it turns out that the average level of credit spreads in customer-dealer trades is *lower* than it would be if the network were complete. Quantitatively, we estimate that credit spreads in customer-dealer trades are 5 percent lower than they would be if all agents in the CDS market could trade with each other. The reason is that the core-periphery structure implies that customers have to trade with dealers and are averse to concentrating their trades with a small number of them. Consequently, equilibrium prices must be attractive for customers to incentivize them to purchase the observed amount of protection. Consistent with this logic, if we leave the network structure untouched but instead shut off aversion to concentrated bilateral exposures, the resulting average level of spreads still converges to the complete-network benchmark. In the absence of aversion to large counterparty exposure, risk sharing is no longer inhibited by the network because agents will trade as much as they want with each other.

Our model also provides an explanation for why we observe price dispersion in the data, and in particular, why dealer-dealer trades occur at lower prices than customer-dealer trades. In the data, dealers are net sellers of protection to customers, implying that, in the model, dealers have lower pre-trade exposure to aggregate credit risk than customers. In our setting, aversion to concentrated exposures prevents perfect risk sharing because agents trade off the costs of concentrated bilateral exposures against the benefits of default insurance. As a result, the post-trade exposure of dealers to aggregate credit risk is still less than that of customers. In turn, dealers pay lower spreads when purchasing credit protection from other dealers in equilibrium because, as a group, they prefer to take credit risk, not hedge it. In a sense, dealers benefit from being a part of the core because they have access to more counterparties and hence better risk sharing opportunities. By our estimates, credit spreads in dealer-dealer trades are nearly 6 percent lower than those in customer-dealer trades and almost 11 percent lower than the average level of spreads that would prevail in a fully-connected economy.

Importantly, our quantitative network model also allows us to measure the systemic importance of a dealer in the CDS market. We argue that it is natural to define a systemically important dealer as one whose failure would have a large impact on equilibrium prices. A major advantage of our

structural approach is that we can easily explore this counterfactual in our model. Specifically, we remove a dealer from the economy by exogenously breaking all of its connections with other agents. Then, we simply re-solve the model for equilibrium prices. We find that removal of a dealer has very different effects on prices depending on the specific dealer’s bearing capacity. The failure of a dealer with median risk bearing capacity would have almost no effect on average prices, both in the dealer-dealer market and the customer-dealer market. In contrast, the failure of a dealer who sells a lot of credit protection would generate large price movements. Our model indicates that when the largest net seller among the dealers fails, the average credit spread in dealer-dealer trades increases roughly 46 basis points, a nearly 40 percent increase over the average observed spread of 124 basis points in the data. Moreover, we find that the effect on spreads can even double or triple if dealer failure is accompanied by a simultaneous increase in either fundamental risk aversion or aversion to concentrated counterparty exposure, both of which seem likely to occur during an actual systemic event like the failure of Lehman Brothers in 2008.

Network frictions play a central role in determining how dealer failure impacts the market prices in equilibrium. Using our estimated parameters, we show that in the counterfactual world where the network is complete, there is almost *no effect* on prices when the dealer that is the largest net seller fails. The effect of a dealer failure is significantly diluted in the complete network case because everyone can trade with each other directly, meaning risk is reallocated relatively easily across all agents.

Our findings highlight why connectivity alone provides an incomplete view of which dealers are systemically important—one must also consider the net exposure when thinking about systemic importance. Put differently, the distribution of risk-bearing capacity matters when measuring the systemic importance of OTC intermediaries. Dealers’ accumulated positions provide information about their equilibrium risk-bearing capacities, as well as their role in reallocating risk through the core-periphery network. Systemically important dealers are both highly connected and provide a large share of credit insurance.

Next, we discuss the related literature. In Section 2, we described the data used and we present key stylized facts about the CDS market. In Section 3, we present our theoretical model. In Section 4, we conduct the dealer removal counterfactual exercise at calibrated parameters, and we conclude in Section 5.

1.1 Related literature

In this paper, we provide a tractable theoretical framework that admits closed-formed solutions, and the simplicity of our model allows us to directly estimate moments from the data statistics. Our setting is related to the framework in Denbee, Julliard, Li, and Yuan (2014), but with significant differences. Our model features endogenous bilateral quantities and endogenous market-clearing prices. Also, while Denbee, Julliard, Li, and Yuan (2014) focus on estimating the degree of strategic complementarity in the U.K.’s interbank borrowing market, we focus on systemically important dealers in the CDS market. Although our setting is not a search-based framework, some of our

findings are consistent with Atkeson, Eisfeldt, and Weill (2015) in terms of generating both price dispersion and intermediation in equilibrium. We, however, develop a network-based model for OTC markets and quantify the effects of dealer removal on equilibrium prices.

More generally, this paper contributes to a fast-growing literature that studies OTC markets through the lens of network frictions.^{3,4} An important contribution of our paper is to quantify the effects of network frictions on equilibrium prices. Also, we identify systemically important dealers as being both very interconnected and large providers of insurance. This is related to findings by Gofman (2015a, 2015b), who studies financial stability when there are financial institutions that are too interconnected to fail. Our contribution is to argue that the risk bearing capacity of dealers plays a key role in determining whether a dealer is systemically important. Our paper also relates to the work by Babus and Hu (2017) who study endogenous intermediation in OTC markets with incomplete information. In contrast, our framework features complete information and endogenous risk reallocation.

A key part of the economic mechanism in our model relies on the assumption that counterparties are averse to concentrated bilateral positions, presumably to avoid some form of counterparty risk. Farboodi (2014) studies counterparty risk in an endogenous network formation model of interbank borrowing and lending, showing that the resulting equilibrium can be constrained inefficient. The economic mechanism in our setting is slightly different because the network is taken as exogenous; instead, agents carefully consider how to size their positions with each other.

We also contribute to the literature on CDS markets by documenting some new stylized facts about the CDS network.⁵ In particular, we show that the average spread between dealers is lower than the average spread between dealers and customers. In addition, we confirm that the CDS trading network is a core-periphery one, which is consistent with the findings by Hollifield, Neklyudov, and Spatt (2017) and Li and Schürhoff (2018).⁶

Finally, our paper has many similar themes to recent research that studies the impact of intermediary risk-bearing capacity on asset prices (He and Krishnamurthy (2013), Siriwardane (2018)). In our model, dealer failure is meaningful because risk sharing is limited by the network and counterparties are averse to large bilateral exposures. Because dealers are situated at the core of the

³There are several recent works on networks and finance, including Herskovic, Kelly, Lustig, and Van Nieuwerburgh (2017), Malamud and Rostek (2014), Herskovic (2017), Babus (2013), Babus and Kondor (2013), Farboodi (2014), Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Salehi (2012), Acemoglu, Ozdaglar, and Tahbaz-Salehi (2013), Aymanns, Georg, Bundesbank, and Golub (2017), Di Maggio, Kermani, and Song (2017), Elliott and Hazell (2016), Elliott, Golub, and Jackson (2014), Erol and Vohra (2017), Wang (2016), Kawakami and Zhong (2016), Hendershott, Li, Livdan, and Schurhoff (2016), Peltonen, Scheicher, and Vuillemeys (2014), Shen, Wei, and Yan (2016), Erol and Vohra (2017), Babus and Hu (2017), Babus (2013), and Babus and Kondor (2013). See Allen and Babus (2009) for a review.

⁴See also Glode and Opp (2017) for a comparison of centralized and decentralized markets.

⁵Recent contributions include Duffie, Li, and Lubke (2011), Duffie, Scheicher, and Vuillemeys (2014), Green, Hollifield, and Schürhoff (2007a, 2007b), Capponi, Cheng, Giglio, and Haynes (2017), Bolton and Oehmke (2013), Chen, Fleming, Jackson, Li, and Sarka (2011), Oehmke and Zawadowski (2013), Schachar (2012), Siriwardane (2018), Atkeson, Eisfeldt, and Weill (2013), Du, Gadgil, Brody, and Vega (2017), Collin-Dufresne, Junge, and Trolle (2016), Gunduz (2018), Arora, Gandhi, and Longstaff (2012), and Longstaff, Mithal, and Neis (2005).

⁶There is a recent set of studies that endogenize network formation and find that a core-periphery network structure in equilibrium, when agents choose their connections unilaterally. See Bala and Goyal (2000), Galeotti and Goyal (2010), and Herskovic and Ramos (2015).

network, these two forces mean that, like in the intermediary asset pricing literature, the risk-bearing capacity of dealers plays a central role in determining asset prices. However, unlike in Walrasian settings with limited participation or intermediary asset pricing models with a representative agent, participation in our core-periphery network model is limited at both the extensive and intensive margins. As a result, overall market risk-bearing capacity depends on individual investor risk-bearing capacities and network limitations to risk sharing – the *distribution* of risk-bearing capacity is critical. Indeed, if the provision of credit risk is shared equally among dealers, we find that the exit of even a fully-connected dealer will have an economically insignificant effect on credit spreads.

2 Data and Supporting Facts

In this section, we start by describing the data that we use to estimate key structural parameters of our model. We then establish three key stylized facts about the CDS market: (1) The network we observe in the DTCC data displays a core-periphery structure, in which a small set of dealers is highly connected to each other, while customers’ only connections are to the dealer sector. (2) On average, dealers are net sellers of credit protection to customers. In addition, net provision of credit risk is highly concentrated in a small number of dealers within the inter-dealer core. (3) Finally, customer-dealer transactions occur at a higher price relative to inter-dealer transactions. Because our model provides closed form expressions for prices and quantities, we then use these three facts to make quantitative inference of the model’s key parameters.

2.1 Data Description

Our primary data on CDS transactions and positions come from 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 from DTCC’s Trade Information Warehouse (TIW) and include CDS transactions and positions reported to DTCC. Transactions represent flows in CDS, and positions represent stocks. The 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 (2013), Siriwardane (2018) and Du, Gadgil, Brody, and Vega (2017). We note that a key difference between our DTCC data and those supplied to the U.S. Federal Reserve System is that ours includes entities that are not Fed-regulated, such as hedge funds. This facet of the data is important for our purpose of measuring the price differences between dealer and customers.

For both transactions and positions, we observe complete information on the counterparties in the trade, pricing terms, size, etc. The 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, the 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.⁷

The mapping between the model and the data requires us to estimate the CDS spread when dealers transact with other dealers, relative to when dealers transact with customers. We focus on data from 2010 through 2013, when central clearing of single names was rare.⁸ 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 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, we estimate the price spread between dealers and customers using single-name transactions on U.S. firms from 2010 to 2013, a time period which pre-dates central clearing of single name contracts but not index contracts.

More generally, the introduction of central clearing does not disrupt the key economic forces in our model. Our model applies whenever agents have an aversion to concentrated bilateral exposures. Such an aversion is likely to persist even with central clearing, since diversifying across trading relationships can reduce the spread of information and can also reduce the risk of hold-up problems, or costly execution delay in the case of counterparty exit. Thus, it is plausible that our model’s main implications are relevant even for markets with central counterparties. However, we leave a full analysis of the impact of central clearing to future work. We also note that, in practice, a large portion of swaps markets (e.g. interest rate, FX, single name and many index CDS) are still cleared bilaterally, and that our model can be easily modified to study these other contracts.

2.2 Fact 1: The CDS Network is Core-Periphery

As in all OTC markets, trades in the CDS market happen bilaterally between pairs of counterparties. A natural way to represent trading relationships in the CDS market is through a matrix G , where element $G_{i,j}$ equals one if counterparties i and j trade with each other, and is zero otherwise. Our model will take this connection matrix, matched to the data, as given. To establish the G matrix in the data, we define two counterparties to be connected, and the entry $G_{i,j}$ to be one, if i and j have any outstanding CDS positions open with each other over our sample.

Our structural network model is general, however in the special case of a core-periphery network we obtain closed form solutions for all equilibrium prices and allocations. Section 3.3 contains the details describing the theoretical core-periphery network. Here, we provide a definition, and establish empirically that the network of trade in credit default swaps is well-represented by a core-periphery network.

Definition 1. *A trading network is core-periphery if there are two groups of agents, i.e., a core*

⁷We 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.

⁸Central clearing of single-name contracts was not prevalent until 2014. See https://www.theice.com/article/cds-growth?utm_source=Insights&utm_medium=tile.

and a periphery, in which all agents are connected to members of the core but not connected to those in the periphery.

Specifically, let n_d be the number of members in the core, and $n_c = n - n_d$ the number of agents in the periphery. We use subscripts d and c to denote dealers and customers as is common in empirical applications. Without loss of generality we set agents $i = 1, \dots, n_d$ to be core agents. Hence, a core-periphery trading network is defined as:

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

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.

Figure 1 presents an illustration of the G matrix of counterparty connections in the CDS market. The upper left corner of the matrix is densely populated, indicating a set of counterparties who are all connected to each other. From the picture, we see that this same set of counterparties (i.e. the core or the dealer sector) is also connected to the other counterparties in the network (i.e. the periphery or the customer sector). Importantly, most of the empirical G matrix is sparse, indicating that the vast majority of counterparties in the data are not connected to each other. The plot provides visual confirmation that the network in the data is closely approximated by a core-periphery network. To make this more precise in the data, we have to take a stand on what we label as a dealer and what we label as a customer. To aid with this task, we use a minimum distance algorithm that is based on the fact that in a pure core-periphery network all dealers should be connected to each other and to every customer, while customers should be connected to all dealers and no one else. Our algorithm proceeds 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 and 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. Such 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.

The algorithm generates a counterparty network with 14 dealers. Our selection algorithm consistently identifies the same 14 dealers even if we focus on subsamples of our data. To verify the consistency of our selection algorithm, we perform the following exercise. We start with a full network matrix that includes all the existing counterparties, and compute who is a dealer based on the algorithm. In a second step, we sort all counterparties based on degree and then transaction volume. We then iteratively remove one counterparty at a time, based on the previous degree-volume sort. Every time we remove a counterparty, we rerun the algorithm for the remaining counterparties. In Panel A of Figure 7, we plot the minimized function against the number of remaining agents in this iterative procedure. In Panel B, we plot the number of dealers as well. The same 14 dealers survive this strict selection procedure for every network with more than 200 counterparties.⁹ For the remainder of the paper, we refer to this set of 14 counterparties as the dealers (or core) in our sample.

2.3 Fact 2: Dealers Provide Credit Protection on Average, and, A Few Dealers Provide Most Net Credit Insurance

In this subsection, we document that dealers are on average net sellers of credit protection to customers over our sample period. In addition, within the dealer sector, the net provision of credit insurance is highly concentrated within a few counterparties.

As in our model, we define net selling for a given counterparty in terms of exposure to a single aggregate credit risk factor. On each day, we define our factor as the equal-weighted average of all five-year CDS spreads for U.S. firms in the Markit Ltd. database. We show in the Appendix that this equally-weighted index is a close approximation to the level factor (first principal component) of credit spreads across all maturities. 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 3 shows that our aggregate credit risk factor evolves as one might expect, peaking at nearly 1000 basis points during the 2007-09 financial crisis. The average of the index is a little over 200 basis points,

⁹The DTCC data also provide a classification for whom DTCC considers a dealer or a customer. This classification is based on DTCC's list of registered dealer members. In our sample of single name transactions, the DTCC's set of dealers is responsibly for nearly 86 percent of gross volume. The 14 counterparties who we label as a dealer are responsible for about 83 percent, which gives us some comfort in the success of our algorithm. For robustness, we also report our results using the DTCC's dealers in the appendix.

so it sits somewhere in between the popular CDX investment grade and high yield indices in terms of average credit risk. In fact, our factor is over 90 percent correlated with both the 5-year CDX investment grade and high yield indices.

With our aggregate credit risk factor in hand, our next task is to define each counterparty’s net overall CDS exposure to our factor. To start, consider an open position p as of date t written on a firm f with m remaining years till maturity. We determine each position’s “maturity bucket” b based on its maturity m , with maturity buckets defined as:

$$b = \begin{cases} 1 & \text{if } m \in [0, 2) \\ 3 & \text{if } m \in [2, 4) \\ 5 & \text{if } m \in [4, 6) \\ 7 & \text{if } m \geq 6 \end{cases}$$

Then 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 previous five year history of three-year CDS spreads up to date t from Markit.¹⁰ Next, we compute the position’s underlying beta with respect to changes in our aggregate credit risk factor via the following regression:

$$\Delta CDS_{f,b,t} = \alpha + \beta_{p,t} \times \Delta \text{CDS Index}_t + \varepsilon_{f,b,t}$$

where CDS Index_t is our aggregate credit risk factor. The position’s beta $\beta_{p,t}$ gives us a gauge of how sensitive the underlying CDS spread of the position is to movements in this index.

We compute $\beta_{p,t}$ for every position contained in our database sourced from DTCC. At this juncture, it is critical to carefully 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 $\beta_{p,t}$ for every position and date in this disaggregated data.

For each dealer i and date t , we compute its beta-weighted net notional sold, denoted $BNS_{i,t}$, as follows:

$$BNS_{i,t} \equiv \sum_{p \in S_{i,t}} \beta_{p,t} \times \text{Notional}_{p,t} - \sum_{p \in B_{i,t}} \beta_{p,t} \times \text{Notional}_{p,t}$$

¹⁰In practice, we also match positions to Markit using the documentation clause and underlying currency of the position.

where $Notional_p$ is the notional amount for position p , $S_{i,t}$ is the set of positions where i is a seller, and $B_{i,t}$ is the set of positions where i is a buyer. Conceptually, BNS is similar to computing the net amount of protection sold by counterparty i , but weights positions based on their exposure to the index (via the $\beta_{p,t}$). The $\beta_{p,t}$'s are therefore useful in this context because they put all positions in terms of exposure to the index, which alleviates any issues that arise from netting positions across different firms. The logic underlying our construction of BNS is similar to computing an equity portfolio's dollar-beta with respect to an aggregate asset return index.

Finally, we scale $BNS_{i,t}$ by dealer i 's market value of equity $E_{i,t}$ to arrive at a normalized measure of dealer exposure:

$$z_{i,t} \equiv \frac{BNS_{i,t}}{E_{i,t}} \quad (1)$$

To focus on lower-frequency movement in equity, $E_{i,t}$ is computed by taking a quarterly moving average of end-of-week market capitalizations. The scaling that we use when constructing $z_{i,t}$ is consistent with the model's use of mean-variance preferences over future returns. Thus, when mapping the model to the data, we interpret z_i as agent i 's net exposure to the underlying default risk in the CDS market relative to i 's portfolio size.

A key quantity for our mapping between the model and the data is the cross-sectional average exposure of dealers, which is easily computed as follows:

$$\bar{z}_{d,t} \equiv \frac{1}{n_d} \times \sum_{i \in \mathcal{D}} z_{i,t}$$

where $n_d = 14$ is the number of dealers and \mathcal{D} is the set of dealers. Table 3 indicates that the average $\bar{z}_{d,t}$ is about 0.04 across all dates in our sample. One way to interpret this number is as follows: on average, dealers' CDS positions represent \$0.04 of notional exposure per dollar of their equity to the overall CDS index. This number is a key input when we map prices back to the underlying structural parameters of our model.

While conceptually simple, the procedure to compute $z_{i,t}$ for each date and each dealer is computationally challenging. It requires us to compute rolling betas for each of the over one billion positions in our sample. We provide some additional details on this process in the Appendix. We also explore other ways of defining dealer credit exposure and draw the same conclusion – dealers are on average net sellers of protection to customers during our sample. For example, we compute the sensitivity of the market value of each dealer's CDS portfolio to movements in our aggregate credit risk factor, which in practice is sometimes referred to as a portfolio DV01. By this metric, dealers lose money on average if the factor increases, indicating that they are net sellers of protection on the index to their customers.

In addition to the provision of credit protection being concentrated within a small core of dealers, even within that core net selling is very concentrated. Figure 4 plots the distribution of z_i within the dealer sector. The highly concentrated nature of net credit provision is clear from the plot, as most of the mass is concentrated around zero (intermediaries tend to have zero net

exposure, as in (Atkeson, Eisfeldt, and Weill 2015)), however, the leftmost part of the plot shows that a few dealers have substantial negative net positions. The fact that a few key dealers provide most of the credit insurance is the reason that a single dealer’s failure can have a large effect.

2.4 Fact 3: Inter-Dealer Prices are Lower than Dealer-Customer Prices

The third key fact we establish is that dealer-dealer (DD) transactions occur on average at lower CDS spreads than dealer-customer (DC) transactions. To estimate this average spread differential, we use all single name transactions for U.S. single name reference identities between 1/1/2010 and 12/31/2013. In addition, we merge each transaction in our sample with an associated 5-year spread in the Markit database.¹¹ For our main set of results, we winsorize both the transaction spreads and Markit spreads at their 5 percent tails. We also confirm that our main conclusions are robust to alternative filtering methods.

Table 1 displays some simple summary statistics on our transaction panel. To construct the reported statistics, we first group transactions into buckets based on the underlying firm and week (f, w) of the transaction. Within each firm-week bucket, we calculate, for instance, the average transaction spread. The reported statistic in the table is simply the average across all (f, w) buckets. In some cases, we also compute a liquidity-weighted average across (f, w) buckets, with the weights determined by the number of trades in each bucket. In addition, we also compute statistics for a subset of our sample where there is a minimum number of trades in an (f, w) bucket.

Table 1 shows that the average transaction spread for our sample is roughly 124 basis points, though this increases slightly for the most actively traded names. We also merge our transaction data with Moody’s expected five-year default frequency (EDF) data, and the average EDF for our sample of firms is 84 basis points.¹² The average EDF will be a key input to our calibration of our model in Section 4. The average maturity of our transactions is around four years for our sample, consistent with the common perception that five-year CDS contracts are the most liquid. The second-to-last row of the table indicates that the CDS market is heavily intermediated—on average, over 70 percent of transactions are between dealers, both in terms of the frequency and the notional amount of trading.

Our ultimate goal is to use this transaction information to measure the difference between DD and DC trades. Ideally, we would do so as follows: on a given date t , compare the average prices in DD transactions to DC transactions on the same firm f . In the data, trades do not occur frequently enough for us to execute this ideal approach without throwing away a large chunk of our

¹¹We compute fair-value spreads from DTCC transactions using the International Swaps and Derivatives Association standard pricing model. To merge with Markit spreads, suppose we observe a transaction on Firm A on date t . We find Firm A ’s quoted five-year CDS spread in Markit on date t , also accounting for the underlying documentation clause and seniority.

¹²For 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.

data. However, we can approximate it fairly well via the following regression:

$$\begin{aligned}
FVS_{k,f,t} = & FE(f) + FE(Rating \times Week_{k,t}) + FE(MaturityBucket_{k,t}) \\
& + \theta_1 \times MarkitSpread_{f,t} + \theta_2 Notional_{k,t} + \sum_{j=1}^2 \gamma_j Maturity_{k,t}^j \\
& + \Phi 1_{k,t}(\text{Dealer-Dealer}) + \varepsilon_{k,f,t}
\end{aligned} \tag{2}$$

where we use k to index each transaction. To keep the notional as compact as possible, we roll transaction characteristics (e.g. seller, buyer, etc.) into k . As before, f denotes the underlying firm in the CDS transaction and t denotes the date of the trade. FVS is the fair-value spread in the trade and $Notional$ indicates the notional amount covered by the trade. We include it in this regression to account for the potential that larger trades have a differential impact on the price that counterparties pay. $Maturity$ is the maturity of the transaction and $MaturityBucket$ is defined as in Section 2.3. These variables enter the regression directly or through fixed effects in order to account for standard maturity effects on the level of spreads.

The key variable in regression (26) is $1_{k,t}(\text{Dealer-Dealer})$, which is just a dummy variable for whether the transaction is between two dealers. Φ in the regression therefore provides an estimate of $-(\bar{R}_c - \bar{R}_d)$. As previously discussed, if we had enough daily transactions per firm, the fixed effect (FE) in regression (26) would be a firm-by-date fixed effect. Instead, we use a rating-by-week fixed effect to avoid wasting too much of the variation in the data. The ratings that we use to define the rating-by-week fixed effect come from Markit. This choice of fixed effect still leaves open the potential that intraweek variation in firm f 's spreads is driven by fundamentals, as opposed to DD or DC trades. To account for this possibility, we include the five-year CDS spread for firm f on date t from Markit, denoted by $MarkitSpread_{f,t}$. Importantly, this control also accounts for differences in credit risk across firms. To this end, we also include a firm fixed effect in all of our regressions.

Table 2 presents the results from running variants of regression (26). Across all specifications, the coefficient on $MarkitSpread$ is slightly less than, but close to one, suggesting that Markit spreads on average are quite close to transaction spreads. It is not surprising that the coefficient on Markit Spreads is slightly less than one because the average maturity of our transactions is roughly four years, whereas Markit spreads are for five-year CDS contracts. The term-structure of credit spreads is on average upward sloping, which means the five-year Markit spread slightly overstates the actual credit spread in the transaction.

The most relevant specification for our purposes is found in column (3), which runs the regression (26) after winsorizing both the transaction spread and Markit data at their 5 percent tails. The estimated Φ of -7.69 basis points indicates that spreads in the dealer-dealer market are lower than spreads in the customer-dealer market. The standard error of the estimate indicates that this coefficient is also statistically significant at conventional levels.¹³

¹³We double cluster our standard errors by firm and year.

To ensure that outliers are not driving our results, we also provide in column (2) results in which we trim (eliminate) the data based on the 5 percent tails of the transaction spread and Markit data. In this case, the point estimate on the coefficient dealer-dealer dummy drops slightly but is relatively close to our baseline estimate when winsorizing.

The remaining columns in the table display the results when using the DTCC’s definition of dealers, as opposed to our custom definition based on the algorithm in Section 2.2. 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. This analysis reveals that our choice of dealers is largely unimportant in terms of estimating the difference in DD versus DC pricing. For the remainder of the paper, we therefore use $\bar{R}_c - \bar{R}_d = 7.69$ basis points when we calibrate our model.

In sum, we have established the following three facts: 1) The CDS network has a core-periphery structure; 2) dealers are net sellers of credit protection, and a few dealers provide a large share of all credit insurance; and 3) trades occur at dispersed prices, with inter-dealer trades occurring at lower prices than dealer-customer trades.¹⁴

3 Model

In this section, we present our OTC network trading model. In Subsection 3.1, we discuss the baseline model. In Subsection 3.2, we consider a three-agent example through which key features of the model are highlighted. Finally, in Subsection 3.3, we focus on the special case of a core-periphery trading network.

3.1 Setup

There are n agents in the economy and one asset with random payoff given by $(1 - D)$, where D is a default component with mean μ and variance σ^2 . An agent i is initially endowed with an exposure, ω_i , to the underlying asset. There is an insurance market where agents trade CDS contracts before the aggregate default is realized. 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 . The price of the contract, namely R_{ij} , is determined in equilibrium.

Let γ_{ij} be the number of contracts agent i sells to agent j . The amount γ_{ij} can be positive or negative depending on whether agent i is selling to or buying insurance from agent j . A positive γ_{ij} means that agent i sells insurance and takes additional exposure to aggregate default risk in the underlying asset.

In addition to CDS payment structure, we assume a counterparty-specific risk to each contract. We assume that the actual payment of a CDS contract between i and j is given by $D + \varepsilon_{ij}$. The term ε_{ij} is independent and identically distributed with mean zero and variance $\varphi > 0$. The counterparty-specific risk can be a benefit (positive ε_{ij}) or a cost (negative ε_{ij}) to the insured counterparty due to non-contractible payoffs such as relationship and information sharing.

¹⁴We provide further evidence of price dispersion in the CDS market in Appendix A.

Therefore, the total payment received by agent i is given by:

$$\pi_i = \omega_i(1 - D) + \sum_{j=1}^n \gamma_{ij}(R_{ij} - D - \varepsilon_{ij}), \quad (3)$$

where the first term represents the pre-trade exposure and the second one represents payments from long and short bilateral positions in the CDS market.

There is a network of trade connections specifying which agents can trade with whom. The network is exogenous and given by an n by n matrix G of zeros and ones, in which each entry (i, j) is given by $g_{ij} \in \{0, 1\}$. If agents i and j can trade, then $g_{ij} = 1$, and, if they cannot trade, then $g_{ij} = 0$. If two agents are not allowed to trade, then it must be that they hold zero bilateral position with each other, that is, we have $\gamma_{ij} = 0$ whenever $g_{ij} = 0$. Furthermore, the network is symmetric, i.e.,

$$g_{ij} = g_{ji} \quad \forall i, j,$$

which means that if agent i can trade with agent j , then j can also trade with i . We assume that $g_{ii} = 1$ for every i , without loss of generality.¹⁵ The exogenous trade structure is consistent with our empirical finding that these connections do not vary over time. Although the trading network is exogenous, the quantities traded (γ 's) are endogenously determined in equilibrium.

We assume that agents feature mean-variance preferences given by:

$$\begin{aligned} U_i(\gamma_{i1}, \dots, \gamma_{in}) &= \mathbb{E}[\pi_i] - \frac{\alpha}{2} V(\pi_i) \\ &= w_i(1 - \mu) + \sum_{j=1}^n \gamma_{ij}(R_{ij} - \mu) - \frac{\alpha}{2} (w_i + z_i)^2 \sigma^2 - \frac{\phi}{2} \sum_{j=1}^n \gamma_{ij}^2, \end{aligned} \quad (4)$$

where $z_i = \sum_{j=1}^n \gamma_{ij}$ is agent i 's net position in the CDS market, $\alpha > 0$ is a risk aversion parameter and $\phi = \alpha\varphi$ is a measure of aversion to counterparty-specific risk.

The parameters α and ϕ play distinct roles in the model. Risk aversion represented by the parameter α measures aversion to total post-trade exposure to the underlining asset, namely $\omega_i + z_i$, while the parameter ϕ measures aversion of concentrating trades with few counterparties. By assuming that ε_{ij} is independent and identically distributed random variable, we assume that counterparties cannot control non-contractible components of the CDS bilateral contracts. These bilateral non-contractible payments become an additional source of risk and aversion to such risk is measured by ϕ . Ultimately, the parameter ϕ captures the cost of trade concentration net of the any potential benefit. On the one hand, in addition to managing counterparty risk, agents may want to spread trades to prevent others from acquiring information about the trading strategies and exposures, or to minimize price impact. On the other hand, trade concentration can be beneficial as trading parties build relationships.

¹⁵The equilibrium allocation is identical whether we set $g_{ii} = 0$ or $g_{ii} = 1$.

Agent i 's optimization problem is given by:

$$\max_{\{\gamma_{ij}\}_{j=1}^n, z_i} w_i(1 - \mu) + \sum_{j=1}^n \gamma_{ij}(R_{ij} - \mu) - \frac{\alpha}{2}(w_i + z_i)^2 \sigma^2 - \frac{\phi}{2} \sum_{j=1}^n \gamma_{ij}^2 \quad (5)$$

subject to

$$\gamma_{ij} = 0 \text{ if } g_{ij} = 0, \quad (6)$$

and

$$z_i = \sum_{j=1}^n \gamma_{ij}. \quad (7)$$

The first restriction guarantees that agent i can trade with agent j if they are connected, and the second one is the definition of agent i 's net position in the CDS market.

In our framework, agents are identical except for their initial pre-trade exposures and their trading connections, and 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 our model features bilateral clearing conditions for any two counterparties:

$$\gamma_{ij} + \gamma_{ji} = 0 \quad \forall i, j = 1, \dots, n. \quad (8)$$

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. Formally, prices satisfy the following condition:

$$R_{ij} = R_{ji} \quad \forall i, j = 1, \dots, n. \quad (9)$$

We solve this model for a competitive equilibrium, in which agents optimize taking prices as given, and all markets clear. Formally, we use the equilibrium concept below.¹⁶

Definition. *A economy consists of a finite number of agents n , a trading network G , preferences described in Equation (4), and pre-trade exposures given by $\{\omega_i\}_{i=1}^n$. A competitive equilibrium with no transaction costs consists of spot market prices, i.e., $\{R_{ij}\}_{i,j=1,\dots,n}$, and traded quantities, i.e., $\{\gamma_{ij}\}_{i,j=1,\dots,n}$, such that (i) agents optimize, taking the network of trading connections and prices as given (Equation 5), (ii) markets clear (Equation 8), and (iii) there is no transaction costs (Equation 9).*

3.1.1 Equilibrium

In this subsection, we fully characterize the equilibrium of the model. Details of the model's solutions are in Appendix C.1. To characterize the equilibrium, we first solve agents' optimization problem- taking price as given. If agent i can trade with agent j , i.e., $g_{ij} = 1$, then agent i 's

¹⁶In Appendix C.4, we present a version of our theoretical model allowing agents to take into account their impact on equilibrium prices. We show that our analysis holds in an environment with price impact as well.

first-order condition with respect to γ_{ij} is:

$$\underbrace{R_{ij} - \mu}_{\text{MB of selling}} = \underbrace{\phi \gamma_{ij}}_{\text{Counterparty-specific MC}} + \underbrace{\hat{z}_i}_{\text{Shadow price of insurance}}, \quad (10)$$

where

$$\hat{z}_i = (w_i + z_i)\alpha\sigma^2. \quad (11)$$

Equation (10) specifies agent i 's optimal exposure to aggregate default risk as a function 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 (7). Hence, \hat{z}_i is agent i 's willingness to pay to insure against one additional unit of exposure to aggregate default.

Agent i 's first-order condition equalizes the marginal benefit of selling insurance to its own shadow price of insurance combined with the marginal cost associated with the counterparty-specific risk. The risk aversion parameter, α , determines how much agent i values net positions through the shadow price of insurance, while counterparty-specific risk aversion parameter, ϕ , determines how much agent i values bilateral contracts individually. In other words, α drives total net positions, while ϕ defines how much agents sell to and buy from each counterparty.

By combining the first-order condition in Equation (10) with the counterparty clearing conditions, in Equation (8), and the no-transaction cost assumption, in Equation (9), we can write equilibrium prices as a linear combination of counterparties' shadow prices of insurance:

$$\underbrace{R_{ij} - \mu}_{\text{contract premium}} = \frac{\hat{z}_i + \hat{z}_j}{2}, \quad (12)$$

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. Therefore, a contract premium depends on post-trade exposures to default risk, along with the variance of aggregate default and preferences parameters. As a result, whenever there are differences in post-trade exposures, there is price dispersion in the cross section of agents in equilibrium, even if agents have identical preferences. This result is consistent with Atkeson, Eisfeldt, and Weill (2015), however our model features additional sources of price dispersion as well. Specifically, our model generates price dispersion from the structure of the trading network itself.

We model the CDS market payment structure as in Atkeson, Eisfeldt, and Weill (2015). The main difference between these two settings is that we impose a network of trading connections and we assume counterparty-specific risk, instead of a fixed risk bearing capacity. Furthermore, in our framework, agents are price takers and we solve for a competitive equilibrium with market clearing prices instead of Nash bargaining. Interestingly, based on Equation (12), equilibrium prices in our framework are isomorphic to Nash bargaining prices with equal bargaining weights. In our model,

the parameter ϕ drives the marginal cost of trading with a counterparty. We can also interpret this marginal trading cost as a bargaining cost and therefore equilibrium prices are an average of agents' willing to pay for insurance against aggregate default risk. However, in our setting, the shadow prices of insurance are determined in equilibrium.

Our modeling approach is also related to the work by Denbee, Julliard, Li, and Yuan (2014). We share similar mean-variance preferences, however, in their model, banks choose liquidity levels taking as given other banks' liquidity. In our model, agents instead choose how much to trade, with each other, taking prices as given. A key distinction is that, in our framework, the network weights, i.e., the specific traded quantities between counterparties (γ 's) are endogenous. Another important difference is that we have equilibrium prices while they focus on the quantity of bank-liquidity buffer. Therefore, our setting features price dispersion and endogenous bilateral exposures in equilibrium.

The price in Equation (12) is a function of the shadow prices of risk, which are determined in equilibrium. We can use Equations (7), (10), (11), and (12) to solve for equilibrium net positions as a linear combination of initial exposures and the net positions of other agents:

$$z_i + \omega_i = (1 - \lambda_i)\omega_i + \lambda_i \sum_{j=1}^n \tilde{g}_{ij} (z_j + \omega_j) \quad \forall i = 1, \dots, n \quad (13)$$

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

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 λ_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 counterparty-specific risk aversion, ϕ . Risk aversion increases agents' willingness to diversify risk away and makes agent i 's post-trade exposures closer to her neighbors, by increasing λ_i . Counterparty-specific risk aversion, however, makes agents less willing to concentrate trade with their counterparties at the expense of lower risk sharing. As a result, ϕ decreases λ_i , which leads to *lower* risk diversification.

The expression for post-trade exposures in Equation (13) has implications for agents not directly connected to each other as well. Since $\lambda_i \in (0, 1)$, the pre-trade exposure of an agent in the network has less and less influence on the post-trade exposure of other agents the farther away from each other trading partners are in the trading network. This decaying influence of pre-trade exposures on post-trade exposures becomes clear if we write post-trade exposures as a function of pre-trade

exposures:

$$z_i + w_i = (1 - \lambda_i)w_i + \lambda_i \left[\sum_{j=1}^n \tilde{g}_{ij}(1 - \lambda_j)w_j + \sum_{j=1}^n \sum_{s=1}^n \lambda_j \tilde{g}_{ij} \tilde{g}_{js} (1 - \lambda_s)w_s + \sum_{j=1}^n \sum_{s=1}^n \sum_{k=1}^n \lambda_j \lambda_s \tilde{g}_{ij} \tilde{g}_{js} \tilde{g}_{sk} (1 - \lambda_k)w_k + \dots \right].$$

The equilibrium condition from equation (13) as well as the equation above shows that the equilibrium net position of agent i depends on the post-trade exposure of its neighbors. However, agent i 's neighbors' post-trade exposures also depend on their own neighbors' post-trade exposures, and so on. These equilibrium conditions imply a system of equations that we can solve for equilibrium exposures. In Appendix C.1, we characterize equilibrium quantities by rewriting Equation (13) in matrix notation and solving for z 's.

3.1.2 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 to be the highest risk-sharing benchmark because the trading network itself does not impose any additional trading friction. In Appendix C.1, we show that, under the complete network benchmark, the post-trade exposure of agent i in equilibrium is given by

$$z_i + w_i = (1 - \lambda)w_i + \lambda \left(\frac{1}{n} \sum_{j=1}^n w_j \right),$$

where $\lambda = \frac{n\alpha\sigma^2}{n\alpha\sigma^2 + 2\phi}$.

The equilibrium post-trade exposures are a convex combination of the agents' pre-trade exposure and the perfect risk-sharing allocation. The coefficient λ measures how far the equilibrium allocation is from perfect risk sharing due to counterparty-specific risk aversion. As ϕ goes to zero, we have that λ goes to one, and perfect risk sharing is achieved. Alternatively, as ϕ goes to infinity, we have that λ goes to zero, and autarky is achieved in equilibrium. These two limiting cases are discussed in Appendix C.1 for a more general network structure.

Under the complete network benchmark, the average prices in equilibrium would be:

$$\bar{R}_{\text{Complete Network}} \equiv \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n R_{ij} = \sigma^2 \alpha \bar{w} + \mu, \quad (14)$$

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

In Appendix C.1, we show that when $\phi = 0$, there is perfect risk sharing and all equilibrium spreads are equal to $\sigma^2 \alpha \bar{w} + \mu$, under a general network structure as long agents are directly or

indirectly connected to each other.¹⁷ 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 $\phi \neq 0$ and pre-trade exposures are heterogeneous.

3.2 Three-agent example

In this subsection, we consider an example with three agents to highlight key features of our framework. First, the example generates price dispersion and intermediation in equilibrium. Second, it has bid-ask spreads with asymmetric prices. Lastly, the example can generate a counterintuitive trading pattern, in which an agent with higher pre-trade exposure sells protection to someone with lower pre-trade exposure to the underlying asset. The derivations of the three-agent example are in Appendix C.2

We assume that there are three agents in the economy in the example. Agents 1 and 2 can trade with each other, and agents 1 and 3 can also trade with one another. However, agents 2 and 3 cannot trade with each other. Formally, the trading network is given by:

$$G = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}. \quad (15)$$

The trading network is also represented in Figure 6. In this economy, agents 1, 2, and 3 have pre-trade exposures given by ω_1 , ω_2 , and ω_3 , respectively. To keep the example more tractable, we set

$$\omega_1 = 0.$$

In equilibrium, based on Equation (13), agent 1's net position is:

$$z_1 = \frac{\alpha\sigma^2}{3\alpha\sigma^2 + 2\phi} (\omega_2 + \omega_3). \quad (16)$$

Agent 1's net position, z_1 , is a combination of the pre-trade exposures of agents 2 and 3. If agents 2 and 3 have pre-trade exposures greater than agent 1, i.e. $\omega_2 + \omega_3 > 0$, then agent 1 endogenously becomes a net seller of insurance with $z_1 > 0$ in equilibrium.

Using Equation (13) and agent 1's net position, agents 2 and 3 net positions are given by:

$$z_2 = \left(\frac{\alpha\sigma^2}{\alpha\sigma^2 + 2\phi} \right) \left(\frac{\alpha\sigma^2}{3\alpha\sigma^2 + 2\phi} \right) \left[\omega_3 - \omega_2 \left(\frac{2\alpha\sigma^2 + 2\phi}{\alpha\sigma^2} \right) \right],$$

and

$$z_3 = \left(\frac{\alpha\sigma^2}{\alpha\sigma^2 + 2\phi} \right) \left(\frac{\alpha\sigma^2}{3\alpha\sigma^2 + 2\phi} \right) \left[\omega_2 - \omega_3 \left(\frac{2\alpha\sigma^2 + 2\phi}{\alpha\sigma^2} \right) \right].$$

¹⁷See Corollary 1.

Next, we highlight the three aforementioned features of this example. First, notice that if $\omega_3 > 0$ and $\omega_2 = -\omega_3$, then $z_1 = 0$ in equilibrium from Equation (33). Also, in equilibrium, we would have $z_2 > 0$ and $z_3 < 0$. This example generates intermediation in equilibrium as agent 1 buys insurance from agent 2 and sells it to agent 3.

Second, this example generates bid-ask spreads. The difference between the price at which agent 1 sells to agent 2 and the price at which agent 1 buys from agent 3, is positive and given by:

$$R_{13} - R_{12} = \phi \left(\frac{\alpha\sigma^2}{\alpha\sigma^2 + 2\phi} \right) (\omega_3 - \omega_2).$$

If $\omega_3 > 0$ and $\omega_2 < 0$, then such price difference is positive, i.e., $R_{13} - R_{12} > 0$.

Furthermore, prices are tilted towards larger pre-trade exposures, generating asymmetric bid-ask spreads. To show such asymmetry is generated, let R_{11} be the equilibrium price for agent 1 if it would trade with itself. Specifically, let us define R_{11} as

$$R_{11} - \mu = \alpha\sigma^2(z_1 + \omega_1) = \alpha\sigma^2 z_1$$

Hence, we can show that:

$$R_{13} - R_{11} > R_{11} - R_{12} \Leftrightarrow \omega_3 > -\omega_2,$$

which means that the spread between agents 1 and 3 is greater than the spread between agent 1 and 2 if, and only if, agent 3's pre-trade exposure is sufficiently high. In this case, agent 3 has too much exposure relative to other market participants and pays a higher price in equilibrium to buy protection against the underlying default risk.

The third feature of this example is a counterintuitive trading pattern, in which an agent with higher pre-trade exposure sells protection to someone who with lower pre-trade exposure to the underlying asset. Specifically, we have that

$$z_2 > 0 \Leftrightarrow \omega_3 > \frac{2\alpha\sigma^2 + 2\phi}{\alpha\sigma^2} \omega_2.$$

This means that agent 2 sells insurance to agent 1, even if agent 2 is more exposed than agent 1 before trade, i.e., $\omega_2 > \omega_1 = 0$. This is true in equilibrium because agent 3 is significantly more exposed to the underlying default risk. In equilibrium, agent 3 demands more insurance from agent 1, who in order to supply such insurance, has to buy additional protection from agent 2. As a result, agent 1 buys insurance from agent 2 and sells to agent 3 in equilibrium.

3.3 Core-periphery networks

In this subsection, we consider a core-periphery trading network and solve the model for equilibrium exposures and prices in closed form. Definition 1 characterizes a core-periphery network in our setting and Appendix C.3 contains the detailed derivation of the model with a core-periphery

network. In the core-periphery economy, there are two interconnected markets: a dealer-to-dealer market and a customer-to-dealer market. For simplicity, we call them dealer and customer markets, respectively.

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

$$\bar{z}_d + \bar{\omega}_d \equiv \frac{1}{n_d} \sum_{i=1}^{n_d} (z_i + \omega_i) = (1 - \lambda_d)\bar{\omega}_d + \lambda_d\bar{\omega}. \quad (17)$$

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

$$\lambda_d = \frac{n\alpha\sigma^2}{n\alpha\sigma^2 + 2\phi}. \quad (18)$$

Equation (17) is derived from Equation (13) 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, i.e., $\bar{\omega}_d$, and the average pre-trade exposure in the economy, i.e., $\bar{\omega}$. Notice that since $\lambda_d \in (0, 1)$, 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:

$$\bar{R}_d \equiv \frac{1}{n_d^2} \sum_{i=1}^{n_d} \sum_{j=1}^{n_d} R_{ij} = \mu + \alpha\sigma^2\bar{\omega} - (1 - \lambda_d)\alpha\sigma^2(\bar{\omega} - \bar{\omega}_d). \quad (19)$$

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 (14).

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

$$\bar{z}_c + \bar{\omega}_c = \lambda_c\bar{\omega} + (1 - \lambda_c)\bar{\omega}_c - \lambda_c(1 - \lambda_d)(\bar{\omega} - \bar{\omega}_d), \quad (20)$$

where $\lambda_c = \frac{n_d\alpha\sigma^2}{n_d\alpha\sigma^2 + 2\phi}$.

The average price in the customer market, i.e., \bar{R}_c , is given by:

$$\begin{aligned} \bar{R}_c &\equiv \mu + \frac{1}{n_d(n - n_d)} \sum_{j=1}^{n_d} \sum_{i=n_d+1}^n R_{ij} \\ &= \mu + \alpha\sigma^2\bar{\omega} - \frac{1}{2}\alpha\sigma^2(\bar{\omega} - \bar{\omega}_d) \left[(1 + \lambda_c)(1 - \lambda_d) - \frac{n_d}{n - n_d}(1 - \lambda_c) \right], \end{aligned} \quad (21)$$

where the last term in brackets is positive if, and only if, $\frac{n_d}{n} < \frac{1}{2}$.

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

$$\bar{R}_c = \bar{R}_d + \frac{1}{2}\alpha\sigma^2(1 - \lambda_c) \left[1 + (1 - \lambda_d)\frac{n - n_d}{n_d} \right] \frac{n_d}{n - n_d}(\bar{\omega} - \bar{\omega}_d) \quad (22)$$

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., $\bar{\omega} > \bar{\omega}_d$, 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 (14).

Alternatively, $\bar{\omega} < \bar{\omega}_d$ if, and only if,

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

Proof. This is a direct implication of Equations (19), (21) and (22). \square

Proposition 1 presents two interesting results. First, it shows that if dealers are less exposed to the underlying asset, the average price in the dealer market is lower than in the customer market. The intuition is that dealers are less exposed to the underlying risk and in equilibrium, they are still less exposed post-trade. As a result, the dealers trade at a lower price among themselves.

The second result is more subtle. It 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. The intuition reflects two considerations. First, dealers are less exposed post-trade when compared to customers. Second, there is a smaller number of dealers in the economy. The core-periphery structure implies that customers have to trade with dealers but are increasingly averse to trading with a small number of them because of agents' aversion to counterparty risk. As a result, equilibrium prices ought to be attractive for customers to incentivize trade in the customer market and for market clearing.

3.3.1 Comparative statics

In this subsection, we analyze how equilibrium prices depend on risk aversion as well as on aversion to counterparty risk. We assume dealers to be less exposed to the underlying default risk, i.e., $\bar{\omega} > \bar{\omega}_d$. This assumption implies that dealers are net sellers in equilibrium, i.e., $\bar{z}_d > 0$, which is consistent with the data. The following proposition shows how equilibrium prices and spreads depend on the risk aversion parameter α and on the aversion to counterparty risk parameter ϕ .

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

- (i) $\frac{\partial}{\partial \alpha} \bar{R}_d > 0$ and $\frac{\partial}{\partial \phi} \bar{R}_d < 0$, i.e., the average price in the dealer market is increasing in α but decreasing in ϕ ;
- (ii) $\frac{\partial}{\partial \alpha} \bar{R}_c > 0$ and $\frac{\partial}{\partial \phi} \bar{R}_c < 0$, i.e., the average price in the customer market is increasing in α but decreasing in ϕ ;

- (iii) $\frac{\partial}{\partial \alpha} \bar{R}_{\text{Complete Network}} > 0$ and $\frac{\partial}{\partial \phi} \bar{R}_{\text{Complete Network}} = 0$, i.e., the complete network benchmark price is increasing in α but does not depend on ϕ .
- (iv) $\frac{\partial}{\partial \alpha} (\bar{R}_{\text{Complete Network}} - \bar{R}_d) > 0$ and $\frac{\partial}{\partial \phi} (\bar{R}_{\text{Complete Network}} - \bar{R}_d) > 0$, i.e., the spread between the complete network benchmark price and the average price in the dealer market is increasing in both α and ϕ ;
- (v) $\frac{\partial}{\partial \alpha} (\bar{R}_{\text{Complete Network}} - \bar{R}_c) > 0$ and $\frac{\partial}{\partial \phi} (\bar{R}_{\text{Complete Network}} - \bar{R}_c) > 0$, i.e., the spread between the complete network benchmark price and the average price in the customer market is increasing in both α and ϕ ;
- (vi) $\frac{\partial}{\partial \alpha} (\bar{R}_c - \bar{R}_d) > 0$ and $\frac{\partial}{\partial \phi} (\bar{R}_c - \bar{R}_d) > 0$, i.e., the spread between the the average price in the customer and dealer markets is increasing in both α and ϕ ;

The proof consists of taking these derivatives using Equations (14), (19), (21), and (22).

Aversion to counterparty risk 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 (items *i*, *ii*, and *iii*). As ϕ increases, agents are more averse to trading too much with one counterparty. Hence, there is less risk sharing in the equilibrium with a higher ϕ , 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 ϕ increases.

The deterioration in risk sharing caused by an increase in ϕ changes equilibrium prices. For average customer prices, the decrease in dealer exposures and increase in customer exposures as ϕ increases generates two offsetting effects. This is because the average price in the customer market reflects a weighted average of the resulting lower mean post-trade exposure of dealers and the resulting higher mean post-trade exposure of customers. However, because the population of dealers is small relative to the population of customers in the economy, market clearing implies that the average post-trade exposure of dealers decreases by more than the average post-trade exposure of customers increases. 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 counterparty risk costs. These costs are disproportionately borne by customers due to their small number of connections and associated limited ability to spread trades across counterparties. Note that, in the dealer market, both lower post-trade exposures, and the burden of higher counterparty risk costs, 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 spread increases between average dealer and customer market prices.

Risk aversion increases the average price in both dealer and customer markets, as well as the complete benchmark price (items *i*, *ii*, and *iii*). As agents become more risk-averse, protection against aggregate default risk becomes more expensive. Similar to the effect of an increase in ϕ , 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, dealers' post-trade exposures increase on average, while customers' post-trade exposures decrease on average. Dealers will trade at a higher price because of their higher post-trade exposures on average. 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, and the fact that dealers are less exposed to the underlying asset than customers, means that their shadow price of insurance 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. 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 2, item *vi*, 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.

4 Calibration and Dealer Removal

In this section, we calibrate our model to evaluate the effect that removing a dealer has on equilibrium prices. In Section 4.1, we detail our calibration procedure. In Section 4.2, we discuss how to implement a dealer removal in our setting, and finally in Section 4.3 we quantify the effects.

4.1 Calibration

As discussed in Section 2.3, dealers are net sellers of protection and we can estimate the average exposure of dealers to the underlying asset, namely \bar{z}_d . In Section 2.4, we discuss how to estimate the average price in the dealer market, \bar{R}_d , and the average customer-dealer price spread, $\bar{R}_c - \bar{R}_d$. Furthermore, we can use estimates of loss-given-default along with probabilities of default to estimate μ and σ^2 .¹⁸ In this section, we use the model derivations discussed in the previous section to write key parameters as a function of data observables. Table 3 reports the calibrated

¹⁸Given a probability of default given by p and a loss-given-default given by L , we have $\mu = Lp$ as the unconditional expected default and $\sigma^2 = L^2p(1 - p)$ as the unconditional variance of aggregate default.

parameters.

By rearranging Equations (17), (19), and (21), we can write the complete network benchmark price as a function of observables:

$$\bar{R}_{\text{Complete Network}} = \alpha\sigma^2\bar{\omega} + \mu = \bar{R}_d + (\bar{R}_c - \bar{R}_d)2\left(1 - \frac{n_d}{n}\right). \quad (23)$$

Equation (23) allows us to infer from the data what the complete network benchmark price would be, irrespective of the risk aversion parameter. The complete network average price is the average price in the dealer market adjusted for the dealer-customer spread. This expression allows us to quantify the network friction by measuring the how the network *per se* affects equilibrium prices. The complete network benchmark price implied by the model, $\bar{R}_{\text{Complete Network}}$, equals to 139 basis points, which is 15 basis points above the average spread in the dealer market.

Our calibration depends on the two risk aversion parameters, namely ϕ and α . To calibrate the risk aversion parameter, α , we rearrange Equation (23) as follows:

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

Risk aversion is mainly driven by equilibrium spreads in the dealer market in excess of expected default, i.e., $\bar{R}_d - \mu$. A higher average spread indicates that agents are willing to pay a higher price for insurance against aggregate default risk. In Equation (24), the only right-hand-side variable we do not observe is $\bar{\omega}$, which is the economy-wide exposure to the underlying default risk. We normalize $\bar{\omega}$ to one. On average, agents have one unit of exposure to default risk—which is equivalent to having agents collectively owning all the aggregate default risk. As a robustness check, we show that our results are not sensitive to changes to this normalization.

To calibrate the parameter driving aversion to counterparty risk, ϕ , we can combine Equations (18), (19), and (21) and solve for ϕ :

$$\phi = (n - n_d) \left(\frac{\bar{R}_c - \bar{R}_d}{\bar{z}_d} \right). \quad (25)$$

Counterparty risk aversion 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 . A higher aggregate customer-dealer spread indicates that agents face more price heterogeneity in equilibrium, which is consistent with higher aversion to counterparty risk as agents are unwilling to take larger bilateral positions. Since dealers are net sellers of protection, higher average net position of dealers, i.e., higher \bar{z}_d , means more risk sharing in equilibrium, which is consistent with lower ϕ .

4.2 Dealer Removal - Details

Core-periphery networks are often thought to be susceptible to systemic risk. To quantitatively evaluate systemic risk in our network, we use the core-periphery model at estimated parameters to provide counterfactual prices and risk allocations when a dealer fails. Specifically, we conduct

the following exercise: We remove one dealer from the core, without changing any other model parameter or agents' pre-trade exposures. We then analyze what the model implies for prices and risk reallocation in the equilibrium after the loss of a dealer. We use the empirical distribution of net CDS positions to calibrate the distribution of dealer characteristics. One result we show is that, to measure systemic risk, net positions are as important or more important than number of connections. In particular, all core dealers are connected to all other dealers, and to all customers, in our model. That is, they all have the same number of connections. However, removing a dealer that is a large net seller has substantial effects: CDS spreads increase by about 46 basis points. On the other hand, removing a central dealer with a neutral net position has only a minor effect. Finally, removing a dealer who is well-connected but is a net buyer of protection actually lowers CDS spreads.

We use the following steps to quantify the effects of removing a dealer. First, we follow the procedure described in Subsection 4.1 to compute model parameters as well as model-implied dealers' pre-trade exposures in excess of the average economy-wide exposure. Table 3 displays the estimated and calibrated parameters. In the second step, we compute the average exposure of dealers in excess of the average economy-wide exposure, excluding the failed dealer. We also recalculate λ_c since the number of dealers is reduced by one. Finally, we compute the new equilibrium prices and allocations by using the updated values of $\bar{\omega} - \bar{\omega}_d$, λ_c , and λ_d .¹⁹

A key question in this exercise is which dealer to remove. Removing a dealer has the same impact on λ_d , regardless of which dealer is removed. However, removing the dealer with largest net selling position, i.e., the dealer with the largest z_i , will increase equilibrium prices the most. The reason is that by removing the largest net-seller dealer, we are effectively removing the dealer with the lowest pre-trade exposure. Therefore, the removal of the largest net-seller dealer would lead to the largest increase in $\bar{\omega} - \bar{\omega}_d$, which would lead to the largest increase in the average price in the dealer and customer markets, \bar{R}_d and \bar{R}_c from Equations (19) and (21), respectively. Removal of the largest net-seller would increase prices by the most. Similarly, removal of the largest net-buyer dealer would lower prices by the most. The total price effect of dealer removal depends on the cross-sectional distribution of dealers' net positions. If the cross-sectional distribution of net positions is positively skewed with a fat right tail, which is true in the data, then the largest net-seller dealer is a few standard deviations away from the average. In this case, the removal of the largest net-seller dealer would lead to significant changes in the average net position of dealers and equilibrium prices.

4.3 Dealer Removal - Results

Table 4 displays the results for the baseline calibration in Column (1), and three counterfactuals in Columns (2)-(4). In the first counterfactual, we remove the dealer that is the largest net seller, which is reported in Column (2). Then, in Column (3), we remove a dealer with the median net

¹⁹Specifically, we use Equations (19) and (21) for prices, and Equations (35) and (36) from the Appendix for dealer's net positions.

position in the benchmark case. Finally, in Column (4), we alternatively remove the dealer that is the largest net buyer in the baseline model.

In our benchmark calibration, we match the average net position of dealers and average spread in both dealer and customer markets. If we remove the largest net-seller dealer from the economy (Column 2), the average spread in the dealer market increases 46 basis points, from 124 to 170 basis points. Out of all 723 counterparties in the economy, the removal of this one agent is quite disruptive to the CDS market, leading to an economically significant change in average spreads. For some additional context, note that a 46 basis point increase rivals what was observed for the CDX IG index after the collapse of Lehman Brothers in 2008.

The removal of the largest net-seller dealer has a large impact on risk sharing in the market. After the removal, dealers become net-buyers of protection of average with an average negative net position. If instead we remove the the largest net-buyer dealer (Column 4), we have opposite effects with lower average spreads in both dealer and customer markets. These results confirm that the net position of the removed dealer has significant effects on the average spreads. Our results do not rely on the normalized of $\bar{\omega}$ to one. In Table 5, we report the same set of results assuming $\bar{\omega} = 0.5$ and $\bar{\omega} = 3$ and our findings are unchanged.²⁰

When a dealer fails and is no longer able to trade with other counterparties, it is likely that the economy faces some sort of financial stress. It is reasonable to assume a dealer failure can be accompanied by an increase in aversion to aggregate default risk and counterparty risk. In our setting, risk aversion parameters, namely α and ϕ , further amplify the dealer removal effects. In Panel A of Figure 8, we plot the average dealer market spreads as a function of both ϕ and α , without removing any dealer. In Panel B, we make a similar plot, but remove the largest net-seller dealer. Aversion to aggregate default risk, α , increases equilibrium spreads with and without the dealer removal. However, the effects of dealer removal on average dealer market spreads are further amplified by higher risk aversion. For example, a 50 percent increase in the risk aversion α increases average spreads from 124 basis points to about 162 basis points, without removing any dealer. After the dealer removal, spreads further increase to 226 basis points. Thus, the impact of dealer failure is amplified substantially if it is accompanied by an increase in risk aversion, resulting in credit spreads rising by a total of 102 basis points.

The effects of changing the counterparty risk aversion parameter, ϕ , on equilibrium spreads before and after the dealer removal are more subtle. Based on Panel A, the parameter ϕ lower the average spreads in the dealer market. As discussed in detail in Section 3,²¹ as ϕ increases, agents become increasingly more averse to trading too much with one counterparty. Thus, a higher ϕ lowers risk sharing and equilibrium prices will be more representative of agents' pre-trade exposures. Since dealers are net sellers of protection, $\bar{z}_d > 0$, they are less exposed to aggregate default risk before

²⁰As another robustness exercise, in Table 6, we report the same set of results but use the DTCC dealer definition, which includes a total of 26 dealers. We also use the DTCC dealers' cross-sectional distribution of net positions in this robustness exercise. Under the DTCC dealer definition, there are more dealers in the economy and risk is therefore more efficiently allocated. As a result, the effect of dealer removal on the average dealer market spreads is about 20 bps, which is lower but still economically significant.

²¹See discussion about Proposition 2 results.

trade, that is, $\bar{\omega}_d < \bar{\omega}$. As a result, ϕ decreases the average spread in the dealer market, which makes it closer to the lower bound identified in Proposition 1. Formally, ϕ lowers λ_d , which, from Equation (19), lowers the average spreads in the dealer market, as observed in Panel A. In Panel B, however, the parameter ϕ increases the average spreads in the dealer market. After removing the largest net-seller dealer from the economy, dealers become, on average, net buyers of protection. The same argument holds, but now dealers are more exposed to aggregate default risk, on average. A higher ϕ increases the average spreads in the dealer market, which is what we observe in Panel B.

In Panel C of Figure 8, we plot the average spreads under the complete network benchmark as a function of both ϕ and α , without removing any dealer. In Panel D, we make a similar plot, but remove the largest net-seller dealer. Panels C and D are almost identical, which means that the dealer removal would have minimal effects on average spreads if the network of trading connections were complete. Removing the largest net-seller dealer has significant effects on average spreads under a core-periphery trading network.

5 Conclusion


This paper develops a model of OTC trading that emphasizes the role of core-periphery networks in these markets. We find that network frictions play an important role in distorting risk sharing and thus equilibrium pricing in OTC markets.

We use detailed transaction and position-level data from DTCC on credit default swaps to calibrate the structural parameters of our model. Our estimation relies on the fact that during our sample, dealers are on average net sellers of credit protection to customers, with most of the net selling done by a small subset of dealers. Dealers transact at lower credit spreads with each other compared to dealer-customer trades. These two facts allow us to infer how averse agents are to concentrating trades in our model, and their relative risk aversion.

The calibrated model provides a natural framework for regulators to evaluate whether a dealer is systemically important to an OTC market: if a given dealer fails, what will be the impact on market prices? We use our estimated model to answer this question, finding that the inter-dealer market spread can rise by nearly 40 percent if the largest net-seller dealer fails. The takeaway from this analysis is that the most systemically important dealers – in the sense that their failure distorts market prices – are highly connected, and provide a large amount of credit insurance to the market. It is the combination of having many interconnections and large positions in the CDS market that makes a dealer systemically important. We show that in a complete network, the failure of an agent with large net position generates almost no effect on equilibrium prices. This observation suggests that the regulatory focus should take into account total net exposures and not just institutions' interconnectedness.

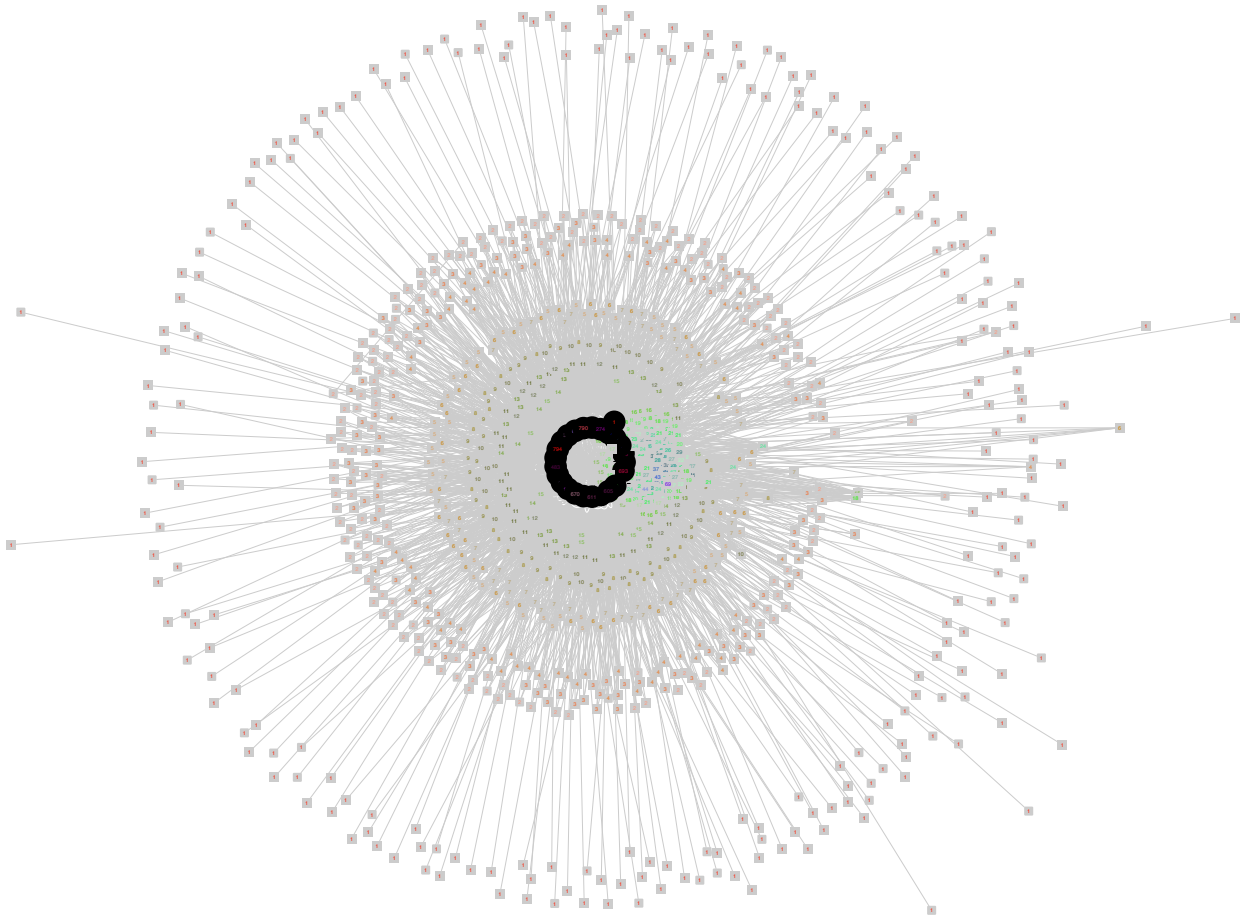
Figures

Figure 1: The Empirical G Matrix




Notes: This figure plots the matrix G where element $G_{i,j}$ equals one if i and j have an open position with each other in our sample, for all counterparties with an open position in the investment grade index. If i and j do not have an open position, $G_{i,j}$ equals zero. Counterparties are ordered by their total number of connections, highest to lowest. Theoretically, a core-periphery network has a structure as in Definition 1, with ones along the diagonal, a core of dealers each represented by a column and row of ones, and zeros elsewhere. This plot shows the close approximation in the data to the theoretical core-periphery structure. Dealers are represented by the left-most columns, and top-most rows, and customers are connected to these dealers, but not each other. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Figure 2: Numbers of Connections in the CDS Trading Network




Notes: This figure shows the numbers of connections in the trading network in the U.S. CDS market on 12/30/2011. The picture is based on the matrix G where element $G_{i,j}$ equals one if i and j have an open position with each other on 12/30/2011. If i and j do not have an open position, $G_{i,j}$ equals zero. We include all CDS positions, both single name and index, when defining G . Counterparties are located on the plot based on their total number of connections, with the most connected counterparties in the middle. The core-periphery structure can be seen by the fact that core dealers have hundreds of connections, while customers are only connected to dealers. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Figure 3: Aggregate Credit Risk Factor



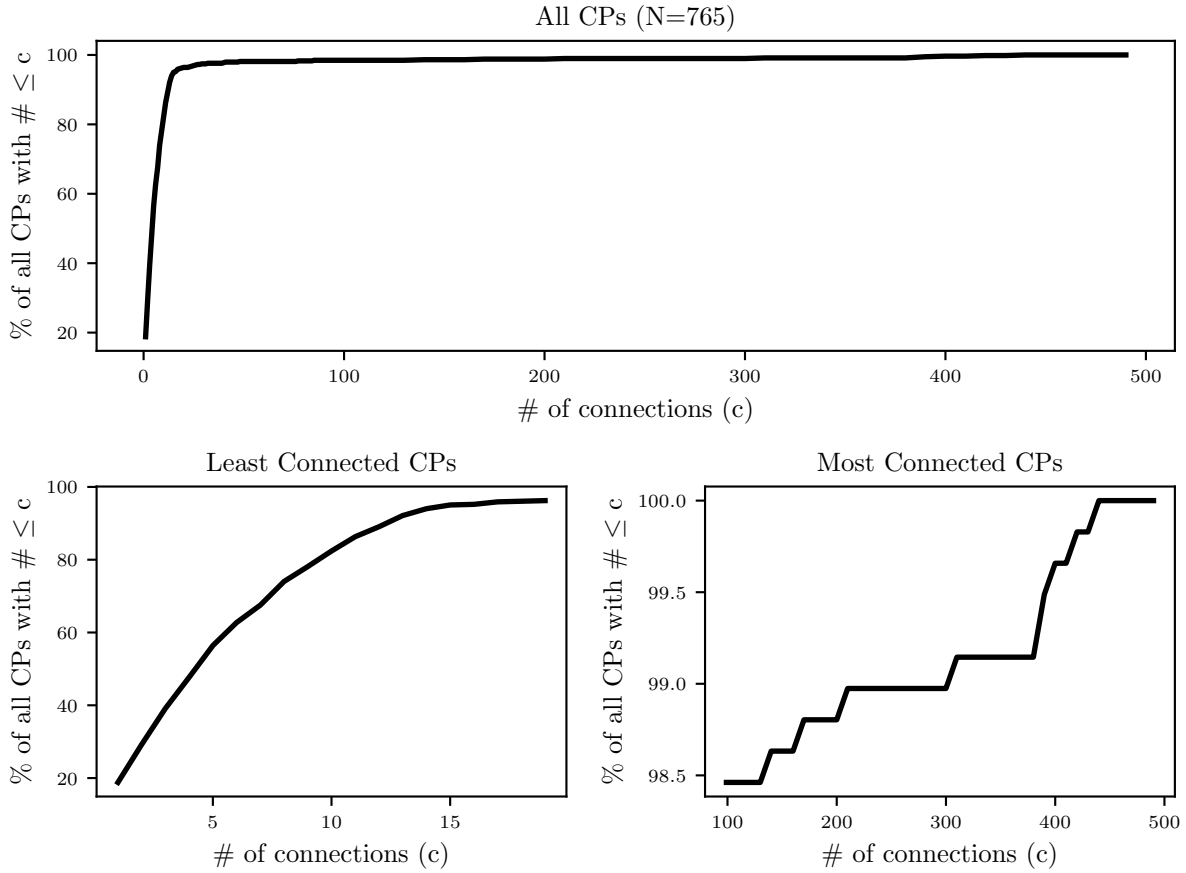
Notes: This figure plots our aggregate credit risk factor from 2002 through 2013. We construct the factor on each date by taking a cross-sectional average of 5-year CDS spreads for all U.S. firms. CDS spreads are obtained from Markit Ltd.

Figure 4: The Distribution of Post-Trade Net Exposures, z_i for Dealers



Notes: This figure plots the matrix G where element $G_{i,j}$ equals one if i and j have an open position with each other in our sample, for all counterparties with an open position in the investment grade index. If i and j do not have an open position, $G_{i,j}$ equals zero. Counterparties are ordered by their total number of connections, highest to lowest. Theoretically, a core-periphery network has a structure as in Definition 1, with ones along the diagonal, a core of dealers each represented by a column and row of ones, and zeros elsewhere. This plot shows the close approximation in the data to the theoretical core-periphery structure. Dealers are represented by the left-most columns, and top-most rows, and customers are connected to these dealers, but not each other. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Figure 5: Degree Distribution in the CDS Network



Notes: This figure shows the degree distribution across counterparties in the U.S. CDS market from January 2010 through December 2013. The picture is based on the matrix G_t where element $G_{i,j,t}$ equals one if i and j have an open position with each other on date t . If i and j do not have an open position, $G_{i,j,t}$ equals zero. The picture is then constructed by averaging G_t over all dates to arrive at what we call \bar{G} . We include all reported CDS positions, both single name and index, when defining G_t on each date. We define the degree, or number of connections, for counterparty i as $D_i \equiv \sum_j \bar{G}_{i,j}$. The top subplot shows the percent of counterparties who have D_i less than or equal to various thresholds. The bottom left subplot shows the same data as the top subplot, but zooms in on those counterparties who have $D_i \leq 10$. The bottom right subplot shows the same data as the top subplot, but zooms in on those counterparties who have $D_i \geq 100$. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Figure 6: Three-agent example

This figure represents an economy with three agents, in which agents 2 and 4 are connected to agent 1. The network matrix in this example is given by Equation (32)






Figure 7: Dealer Selection




Notes: In this figure, we report our selection algorithm outcome for different subsamples. We start with full network matrix that includes all the existing counterparties, compute who is a dealer based on the algorithm. In a second step, we sort all counterparties based on degree and then transaction volume. We then interactively remove one counterparty at a time, based on the previous degree-volume sort. Every time we remove a counterparty, we rerun the algorithm for the remaining counterparties. In Panel A, we plot the minimized function against the number of remaining agents in this iterative procedure. In Panel B, we plot the number of dealers against the number of agents. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Figure 8: Dealer Removal


Panel A: \bar{R}_d without dealer removal




Panel B: \bar{R}_d with dealer removal



Panel C: $\bar{R}_{\text{Complete Network}}$ without dealer removal



Panel D: $\bar{R}_{\text{Complete Network}}$ with dealer removal



Notes: This figure plots the average spreads, levels of risk aversion, α , and aversion to counterparty risk, ϕ . The lowest ϕ and α plotted in both graphs are the benchmark calibrated parameters, also available in Table 3. Panel A plots the average spread in the dealer market, \bar{R}_d , while Panel B plots \bar{R}_d after the removal of the highest-net-seller dealer. Panel C plots the model-implied complete network average spread, $\bar{R}_{\text{Complete Network}}$, while Panel D plots $\bar{R}_{\text{Complete Network}}$ after the removal of the highest-net-seller dealer. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

TABLES

Table 1: 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	7	25	15	30	23	37	37	50
Avg. Spread (bps)	124	146	147	153	157	157	165	160
Avg. EDF (bps)	84	90	89	92	94	94	98	95
Avg. Recovery (%)	38.7	39.1	39.1	39.3	39.3	39.4	39.5	39.5
Avg. Maturity	3.67	4.02	4.07	4.13	4.13	4.16	4.14	4.18
Avg. Notional (mm)	6.32	5.90	6.15	5.77	5.83	5.57	5.44	5.30
σ (Notional)	4.25	5.09	5.05	5.36	5.34	5.52	5.54	5.69
% Trades Dealer-to-Dealer	76	73	72	72	73	72	73	72
% Notional Dealer-to-Dealer	77	75	74	74	75	75	76	75
# of (f, w) groups	80,852	80,852	29,909	29,909	14,993	14,993	5,880	5,880

Notes: This table presents summary statistics of spreads and trading activity across firm-week buckets (f, w) pairs. Within each (f, w) group, we compute each statistic (e.g. average spread, etc.). We then take an equal-weighted (EW) average of these statistics across (f, w) groups. 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 Section 2.2. Notional values are reported in \$ millions and 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, drop contracts between nondealers and nondealers (only 0.31% of total), and pool contracts of different tiers and doc-clauses. We winsorize the 5% tails of outliers based on Markit and fair-value transaction spreads. 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: Dealer and Customer Prices

Dependent Variable:	Fair-Value CDS Transaction Spread					
	Robust Dealer Selection			DTCC Dealers		
	(1)	(2)	(3)	(4)	(5)	(6)
Market Spread (bps)	0.42** (0.03)	0.82** (0.01)	0.80** (0.02)	0.42** (0.03)	0.82** (0.01)	0.80** (0.02)
1 (Dealer-Dealer)	-9.00** (1.60)	-5.95** (1.55)	-7.69** (1.39)	-9.81** (1.76)	-5.97** (1.47)	-8.28** (1.52)
Filtering Method	None	Trim	Winsor	None	Trim	Winsor
Overall R^2	0.87	0.92	0.90	0.87	0.92	0.90
ARMSE (bps)	53.64	35.83	46.63	53.65	35.86	46.65
N	431,614	353,297	431,614	432,882	354,237	432,882

Notes: This table presents regressions of the following form:

$$FVS_{k,f,t} = FE(\text{Firm}) + FE(\text{Maturity Bucket}_{k,t}) + FE(\text{Rating}_{k,t} \times \text{Week}_t) + \theta_1 \text{MarketSpread}_{f,t} + \theta_2 \text{Notional}_{k,t} + \sum_{j=1}^2 \gamma_j \text{Maturity}_{k,t}^j + \Phi 1_{k,t}(\text{Dealer-Dealer}) + \varepsilon_{k,f,t}$$

$FVS_{k,f,t}$ is the fair-value spread for transaction k , written on firm f , and executed on date t . $\text{MarketSpread}_{f,t}$ is the 5-year CDS spread from Markit's singlename database that is associated with firm f on date t . $\text{Notional}_{k,t}$ is the notional amount in the transaction and $\text{Maturity}_{k,t}$ is the maturity of the transaction. $1_{k,t}$ (Dealer-Dealer) is a dummy variable that equals 1 if the transaction is between two dealers and is zero otherwise. We define dealers precisely in Section 2.2. 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, drop contracts between nondealers and nondealers (only 0.31% of total), and pool contracts of different tiers and doc-clauses. When we trim or winsorize spreads, we do so based on the 5% tails of both $FVS_{k,f,t}$ and $\text{MarketSpread}_{f,t}$. Each transaction is grouped into one of the following four buckets based on its maturity: (i) 0-2 years; (ii) 2-4 years; (iii) 4-6 years; and (iv) 7+ years. The fixed effects for maturity bucket are based on this grouping. The credit ratings used in the Rating-by-Week fixed effect come from Markit. The sample runs from 2010-01-04 to 2013-12-31. 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%. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Table 3: Calibration

Parameter	Value	Source
\bar{z}_d - Robust Dealers	0.045	DTCC Data 2010-2013
\bar{z}_d - DTCC Dealers	0.087	DTCC Data 2010-2013
$\bar{R}_c - \bar{R}_d$	7.69	DTCC Data 2010-2013
\bar{R}_d	124	DTCC Data 2010-2013
n	723	DTCC Data 2010-2013
n_d	14	DTCC Data 2010-2013
$L = \text{Loss-Given-Default}$	61.3%	Moody's
$p = \text{Probability of Default}$	0.84%	Moody's
λ_d	0.21	Model Implied
$\alpha\sigma^2\bar{\omega} + \mu$	139	Model Implied
α	2.80	Model Implied
ϕ	12.02	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 2.3 contains a full description of this procedure. DTCC Dealers are the set of firms identified in our data from DTCC as being dealers, and Robust Dealers are those identified by the procedure described in Section 2.2. $\bar{R}_c - \bar{R}_d$ 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 Section 2.4 for complete details). \bar{R}_d is the fitted value for dealer-dealer trades that comes out of the same regression. 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 investment grade firms from Moody's. The remaining parameters in the table are implied by our structural model. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Table 4: Dealer Removal

	Benchmark	Top	Mediam	Bottom
	(1)	(2)	(3)	(4)
Number of dealers	14	13	13	13
Complete network \bar{R} (bps)	139.08	140.16	139.05	138.78
\bar{R}_d (bps):	124.00	170.46	121.60	109.41
\bar{R}_c (bps):	131.69	155.03	130.49	124.36
\bar{z}_d	0.05	-0.09	0.05	0.09

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 2.2. Column (1) reports our benchmark calibration. In Column (2) reports the results after removing the largest net-seller. Column (3) reports results after removing the dealer with the median net position, and Column (4) 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 5: Dealer Removal Robustness: $\bar{\omega}$

	Benchmark	Top	Mediam	Bottom
	(1)	(2)	(3)	(4)
Panel A: Assuming $\bar{\omega} = 0.5$				
Number of dealers	14	13	13	13
Complete network \bar{R} (bps)	139.08	140.39	139.05	138.71
\bar{R}_d (bps):	124.00	170.69	121.59	109.34
\bar{R}_c (bps):	131.69	155.26	130.48	124.30
\bar{z}_d	0.05	-0.09	0.05	0.09
Panel B: Assuming $\bar{\omega} = 3$				
Number of dealers	14	13	13	13
Complete network \bar{R} (bps)	139.08	140.01	139.06	138.82
\bar{R}_d (bps):	124.00	170.31	121.61	109.46
\bar{R}_c (bps):	131.69	154.88	130.49	124.41
\bar{z}_d	0.05	-0.09	0.05	0.09

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 2.2. In Column (1), we report our benchmark calibration. Column (2), we report the results after removing the largest net-seller. In Column (3), we report results after removing the dealer with the median net position, and, in Column (4), we report results after removing the dealer which is the largest net buyer in the baseline model. In Panel A, we report results assuming $\bar{\omega} = 0.5$, while in Panel B we report the results assuming $\bar{\omega} = 3$. In each Panel, given the assumed value for $\bar{\omega}$, we recalibrate the model following the procedure described in Section 4.1. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Table 6: Dealer Removal Robustness: DTCC Dealers

	Benchmark	Top	Mediam	Bottom
	(1)	(2)	(3)	(4)
Number of dealers	26	25	25	25
Complete network \bar{R} (bps)	139.96	140.89	139.96	139.62
\bar{R}_d (bps):	124.00	143.69	123.16	115.79
\bar{R}_c (bps):	132.28	142.24	131.86	128.13
\bar{z}_d	0.06	-0.01	0.06	0.09

Notes: This table uses the DTCC definition of dealers and recalibrates the model following the procedure described in Section 4.1. We report 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. In Column (1), we report our benchmark calibration. In Column (2), we report the results after removing the largest net-seller. In Column (3), we report results after removing the dealer with the median net position, and, in Column (4), we report results after removing the dealer which 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.

Appendix

A Ancillary Facts

Degree Distribution in the CDS Network Figure A1 plots the distribution of D_i across counterparties based on the average G -matrix from January 2010 through December 2013. For all of the subplots in the figure, the x -axis contains the number of connections, c . The y -axis plots $P(c)$, defined as the percent of all counterparties with $D_i \leq c$. It is immediately clear from the figure that most counterparties in the CDS market are only connected to a handful of other counterparties: over 80 percent have less than 10 connections and nearly 60 percent have less than 3. However, there are also a very concentrated subset of counterparties who have lots of connections. For example, the top 3 percent of connected counterparties have well over 100 connections each.

These highly-connected counterparties are central in the sense that they are also connected to each other. To visualize this fact, Figure 2 plots G on 12/31/2011. The figure displays all counterparties in the market and their connections to other counterparties. The most connected counterparties are arranged in the core of the plot, and the less connected counterparties are on the periphery. It is easy to see from the figure that the core traders are all connected to each other, as well as the periphery. This type of core-periphery trading structure is common to most OTC markets (e.g. Maggio, Kermani, and Song (2017)).

The CDS Network is Static Our model is static. However, it is interesting to ask how persistent trading relationships in CDS markets are in the data. In practice, trading relationships are governed by what is called an International Swaps and Derivatives Association (ISDA) master agreement, which sets standard trading protocols (e.g. collateral requirements, netting agreements, etc.) between each counterparty pair. ISDA agreement takes quite a bit of time and effort to put in place, and consequently, there are a large fixed costs to forming new trading connections in this market. For this reason, the CDS network is sticky in the sense that new trading relationships are not formed very often, especially for the most active traders. As a simple way to demonstrate this in the data, we first compute the following statistic for each counterparty:

$$\text{Gross Turnover}_{i,t} = \frac{\# \text{ New Connections}_{i,t} + \# \text{ Lost Connections}_{i,t}}{\# \text{ Connections}_{i,t-1}}$$

In words, gross turnover is just number of new or lost connections made by i at time t , scaled by the number of i 's connections in the previous period. For example, if counterparty i has 100 connections and makes a new one this week, then gross turnover will be 1 percent. Large values of gross turnover indicate that there are lots of connections being formed or lost at each point in time, whereas small numbers indicate a relatively stable network dynamic.

We compute gross turnover for every counterparty-date in our sample. At each date, we compute cross-sectional statistics (e.g. percentiles, mean) across counterparties. Table A1 presents

these statistics, averaged over all time periods. Because dealers and customers might exhibit different behavior in terms of forming connections, we also repeat the exercise for each subset of counterparties. The first thing to notice from the table is that gross turnover is zero for most counterparty-date pairs in our sample. For example, the 75th percentile of gross turnover across all counterparties is zero for every week in our sample. Even for dealers, the 75th percentile of gross turnover is incredibly low at around 2.6 percent. Considering that many dealers have well over 100 connections, this means that in an average week most dealers add less than 2 connections. For customers, the 95th percentile of gross turnover is relatively higher at 13.7 percent, but this is to be expected because a counterparty with a single connection who adds another will have a gross turnover of 100 percent.

Overall, this analysis indicates that the CDS network is quite stable in terms the trading network – connections are rarely added or deleted. This fact is helpful when mapping the data to the model for two reasons: (i) in order to complete the mapping, we need to pick a G matrix for a given week. Because the network is static, averaging the G matrix over all weeks or picking a random week will not impact our overall results; and (ii) ultimately, we want to quantitatively assess how prices respond when a dealer in the market fails. Given that G is relatively static in normal times, it seems reasonable for us to assume that new connections could not be formed after a dealer failure, at least in the short run. More broadly, this feature of the data justifies our decision to use a static network model to represent the CDS market.

Price Dispersion A key implication of our model is that there will be price dispersion in the CDS market. By price dispersion we specifically mean that, on a given date, trades on the same firm and for the same maturity may occur at different prices. In our model, this occurs in equilibrium due to the fact that there is limited risk sharing and trading costs (ϕ). Still, one may be skeptical that there is actually meaningful price dispersion in the data.

It is important to note that our estimate of $\bar{R}_c - \bar{R}_d$ is only possible because there is price dispersion in the data. Recall that we estimate $\bar{R}_c - \bar{R}_d$ via the following regression:

$$\begin{aligned}
FVS_{k,f,t} = & FE(f) + FE(Rating \times Week_{k,t}) + FE(MaturityBucket_{k,t}) \\
& + \theta_1 \times MarketSpread_{f,t} + \theta_2 Notional_{k,t} + \sum_{j=1}^2 \gamma_j Maturity_{k,t}^j \\
& + \Phi 1_{k,t}(\text{Dealer-Dealer}) + \varepsilon_{k,f,t}
\end{aligned} \tag{26}$$

The regression works by comparing firms in the same rating and same week, then checking whether within this group dealer-dealer trades occur at different prices from dealer-customer trades (via Φ). A natural concern here is whether our estimate of $\bar{R}_c - \bar{R}_d$ is biased by the fact that the rating-week fixed effect does not control for variation across firms of the same rating within a week or even within-week spread variation for the same firm. In order for this to be problematic for us, dealer-dealer trades would have to be systematically correlated with some temporal pattern in

spreads that occurs within a week and across firms (e.g. spreads happen to be higher on average later in the week and this just happens to be when dealers trade with each other). While this seems unlikely on its face, we control for it directly by including the 5-year Markit spread for firm f on each date t in the regression. Even in this specification, we get a statistically significant estimate of $\bar{R}_c - \bar{R}_d = 7.69$ bps. Compared to an average spread of roughly 150 bps from Table 1, this means dealers receive approximately a 5 percent discount compared to customers. This itself indicates a fair amount of price dispersion in the market.

To add further support for this argument, we can group transactions at more granular level. For example, suppose we observe N trades on Ford on date t for a 5-year maturity. The existence of price dispersion means that we would observe different prices across these N trades. To quantify the size of this effect in the data we first group transactions by firm, date, and maturity bucket, (f, t, m) . Keep in mind that this limits the set of transactions we are able to study because it requires multiple trades in the same (f, t, m) bucket.²² We then compute measures of price dispersion like standard deviation or range within each (f, t, m) bucket. Because we are grouping at the level of transaction date t , any observed dispersion within the (f, t, m) bucket cannot be driven by within-week variation in prices for firm f of maturity m . This grouping is therefore sharper than the firm-week buckets we use to construct our summary statistics in the main text. Finally, we average across the buckets to get a broad sense of price dispersion in the data.

Table A3 contains the results of this analysis. Regardless of how we measure it, the table indicates there is a meaningful amount of price dispersion in the market. For example, the average interquartile range within each group is about 10 basis points. Put differently, the ratio of the interquartile range to the average spread in a bucket is roughly 15 percent. We interpret these statistics as follows: on a given day, consider all CDS prices on the same firm and maturity. For this set of trades, the majority of prices range from a 7.5 percent discount to a 7.5 percent premium, relative to the average price. We view this as a fairly large band and thus conclude that there is a reasonable amount of price dispersion in the CDS market, as predicted by our model.

B The Net Position of Dealers

In this section, we provide additional details and robustness checks on how we construct our exposure measure $z_{i,t}$ for each dealer i and date t . As a reminder, our preferred measure is defined as:

$$z_{i,t} \equiv \frac{BNS_{i,t}}{E_{i,t}}$$

$$BNS_{i,t} \equiv \sum_{p \in S_{i,t}} \beta_{p,t} \times \text{Notional}_p - \sum_{p \in B_{i,t}} \beta_{p,t} \times \text{Notional}_p$$

²²As discussed in the main text, it is for this reason that we do not estimate $\bar{R}_c - \bar{R}_d$ using this level of granularity.

where p indexes the positions of dealer i on date t . $S_{i,t}$ is the set of positions where i is a net seller of protection and $B_{p,t}$ is the set of positions where i is a buyer. $E_{i,t}$ is the trailing quarterly average market capitalization of dealer i , computed using weekly data. $\beta_{p,t}$ is the position’s beta with respect to the on-the-run 5-year CDX IG index; we now provide more details on how we compute these β ’s.

B.1 Estimating Betas

To keep this appendix self-contained, we repeat some details of our methodology that are presented in the main text.

To start, consider an open position p as of date t written on a firm f with m remaining years till maturity. We determine each position’s “maturity bucket” b based on its maturity m , with maturity buckets defined as:

$$b = \begin{cases} 1 & \text{if } m \in [0, 2) \\ 3 & \text{if } m \in [2, 4) \\ 5 & \text{if } m \in [4, 6) \\ 7 & \text{if } m \geq 6 \end{cases}$$

Then 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 previous four year history of 3-year CDS spreads up to date t from Markit. In addition, we match position p to Markit based on a plethora of characteristics. These characteristics include Markit RED id (i.e. the underlying the firm), currency, capital structure tier, and documentation clause relating to the CDS default trigger. For instance, holding all other characteristics equal, Ford CDS quoted in USD and EUR would be matched to two different records in Markit. Similarly, Ford CDS on senior and junior debt, holding all other characteristics equal, would be matched to two different records in Markit.

Next, we compute the position’s underlying beta with respect to changes in the on-the-run CDX IG index via the following regression:

$$\Delta CDS_{f,b,t} = \alpha + \beta_{p,t} \times \Delta \text{CDS-IG}_t + \varepsilon_{f,b,t}$$

The position’s beta $\beta_{p,t}$ gives us a gauge of how sensitive the underlying CDS spread of the position is to movements in the CDX IG index. We compute $\beta_{p,t}$ using four-year rolling regressions for every position in the DTCC database, including direct single name transactions and indirect single name transactions that occur through the index (see Siriwardane (2018)). We obtain the 5-year on-the-run CDX IG index from Bloomberg Finance L.P.

B.2 Alternative Methodologies

We now turn to alternative ways of computing $z_{i,t}$ for each dealer i and date t . All of our alternative approaches take the following form:

$$z_{i,t} \equiv \frac{X_{i,t}}{E_{i,t}}$$

where $X_{i,t}$ is one of several ways of defining net credit exposure, and $E_{i,t}$ is the smooth market-capitalization that we described at the beginning of the section.

B.2.1 Index Positions Only

An easy way to define $X_{i,t}$ is simply by computing the net notional amount of CDS protection sold directly on the CDX IG index. To operationalize this in the data, we compute the net notional sold by i on any CDX index product written on the IG index with a maturity between four and six years. For this measure of exposure, we use positive values of CDS-IG $_t$ to represent net selling and negative values to represent net buying.

B.2.2 Disaggregated Positions

Defining $X_{i,t}$ using only index positions is potentially misleading because counterparties can have exposure to the index via single name trades. The reason is that there is an equivalence to selling protection on an index and selling protection on the individual firms that comprise the index. As discussed in Siriwardane (2018), this distinction is particularly important in the CDS market because single name positions represent well over half of the net notional outstanding for the entire CDS market during our sample. To account for this fact, we therefore disaggregate CDS indices into their individual constituents and then combine these “disaggregated” positions with any pure single name positions. Using the disaggregated DTCC data, we then compute $X_{i,t}$ by summing the net amount sold on firms that are in the CDX IG index, conditional on the maturity of the position being between four and six years.

B.2.3 Beta-Weighted DV01

In addition, we adopt a more general approach that accounts for each counterparty’s entire portfolio (e.g. positions of all maturities and all firms, even those not in the IG index) and market values. As with our preferred measure of $z_{i,t}$, we start with the β ’s of each position (see Section B.1). After matching all DTCC positions to a β , we then compute each position’s “DV01”. Analogous to an option delta, DV01 is the standard way that industry professionals quantify the dollar change of a position with respect to a move in the positions’s underlying credit spread. For example, suppose that a fictitious position on Xerox Corp. has a notional value of \$1. The DV01 tells how many dollars the seller in the swap would gain/lose if Xerox’s credit spread falls by 1 basis point.²³

²³Following with industry standard, we consider a one basis point decrease in the entire term structure of Xerox’s CDS spread.

We then use $DV01_{p,t}^f$ to denote the position’s DV01 as of date t .²⁴ The superscript f denotes that this DV01 is computed for a one basis point move *in firm f ’s CDS spread*. See Section B.3 for details on how we compute $DV01_{p,t}^f$. In all cases, we define $DV01_p^f$ from the perspective of the protection seller, meaning that it is always positive for sellers and is negative for buyers (e.g. a decrease in CDS spreads always helps the seller).

Once we compute $DV01_{p,t}^f$, it is easy to ask how much the seller would lose if there is a fall in the IG index by a basis point:

$$DV01_{p,t}^{IG} = DV01_{p,t}^f \times \beta_{p,t}$$

$DV01_{p,t}^{IG}$ is useful because we can sum it across positions – its units are dollars per *one basis point fall in the IG index*. Once we compute $DV01_{p,t}^{IG}$ for all positions, we aggregate it within each counterparty i to determine i ’s portfolio DV01, denoted by $DV01_{i,t}^{IG}$. We use $DV01_{i,t}^{IG}$ as one of our measures of $X_{i,t}$. Keep in mind that $DV01_{i,t}^{IG}$ is positive for dealers who are net sellers of CDX IG credit risk.

B.2.4 Average Dealer Exposure

Table A2 presents some simple time-series averages of \bar{z}_d for each of our construction methodologies. The biggest observation from the table is that all of the \bar{z}_d are positive on average. Thus, regardless of how we measure exposure, dealers are on average sellers of credit protection during our sample. When using the notional-based measures, there is a clear increase in \bar{z}_d when moving from using only CDX IG index positions, then disaggregated positions for CDX IG index constituents, and finally beta-weighted notional. This progression is not surprising because each successive step includes a wider range of positions over which we aggregate credit exposure. The DV01-based metric indicates that a 100 basis point move in the CDX IG index would cause the average dealer to lose 0.99 percent of their equity value.²⁵ Again, the larger point here is that dealers are exposed to the underlying credit risk of the CDX IG index during our sample. This basic fact is important in how we infer the structural parameters of our model based on the prices paid by dealers versus customers.

B.3 Computing Credit Sensitivities (DV01)

We define a position’s credit portfolio sensitivity, $DV01_p^f$, as the sensitivity of the position to a change in the underlying reference entity’s credit spread. We arrive at this measure by applying the ISDA Standard Model for pricing credit derivative contracts (CDS) and the implementation detailed in the Appendix of Paddrik et. al. (2016). A CDS position p written on firm f can be expressed as the difference between premium leg $Prem^{sf}$ and pay leg Pay^{sf} , calibrated from market spread s_f (baseline). From the perspective of the seller, $Prem^{sf}$ is the discounted present value of

²⁴As in the main text, to avoid notational clutter, the position’s characteristics (e.g. the maturity) are rolled in the position index p .

²⁵In the table, we have scaled the DV01-based measure so that it corresponds to a 100 basis point move in the index.

the buyer's incoming payments, while Pay^{sf} is the present value of outgoing payouts contingent on default of f . Both components are functions of the underlying (risk-neutral) default risk of the firm, which is inferred from prevailing credit spreads s_f . (We suppress in our notation other characteristics which uniquely identify the market spread such as term, documentation clause, currency, and date of observation.) The position can be revalued under a differential shock to market spreads, $s'_f = s_f + ds_f$ (shock). Following industry practice, we adopt 1 basis point change. This permits us to express the $DV01_p^f$ from the protection seller's perspective as

$$DV01_p^f = (Prem^{s'_f} - Prem^{s_f}) - (Pay^{s'_f} - Pay^{s_f}) \cdot N_p$$

The $DV01$ expresses the difference between the baseline and a scenario in which credit spreads (e.g. default risk) rise. By this definition, it is therefore always negative from the perspective of the seller.

We rely on multiple data sources to identify contractual inputs for pricing positions. We use the underlying's reference entity's term structure of credit spreads, contract currency, floating risk-free rates, and capital structure of the CDS' underlying reference obligation. We source credit spreads from Markit, contract currency from DTCC, the term structure of risk-free rates for contract currencies from Haver Analytics, and reference entity capital structure from bond information provided by Bloomberg.

C Model Derivations and Extensions

C.1 Solving the Model

Agent i 's optimization problem is given by Equation (5):

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

subject to

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

and

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

Agent i 's first-order conditions give us:

$$\gamma_{ij} = \begin{cases} \frac{1}{\phi}(R_{ij} - \mu) - \frac{1}{\phi}\hat{z}_i & \text{if } g_{ij} = 1 \\ 0 & \text{if } g_{ij} = 0, \end{cases} \quad (27)$$

where

$$\hat{z}_i = (w_i + z_i)\alpha\sigma^2, \quad (28)$$

$$z_i = \sum_{j=1}^n \gamma_{ij} = \frac{\frac{1}{K_i} \sum_{j=1}^n g_{ij}(R_{ij} - \mu) - \alpha w_i \sigma^2}{\frac{\phi}{K_i} + \alpha \sigma^2}, \quad (29)$$

and

$$K_i = \sum_{j=1}^n g_{ij}. \quad (30)$$

We can derive Equation (13) by combining Equations (7), (28), and (12). Furthermore, to fully characterize the equilibrium, we solve for equilibrium quantities by rewriting Equation (13) in matrix notation as follows:

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

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 λ_i , and \tilde{G} is a $n \times n$ matrix with the ij th entry given by \tilde{g}_{ij} .

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

$$z + w = (\mathbf{I} - \Lambda\tilde{G})^{-1}(\mathbf{I} - \Lambda)\omega, \quad (31)$$

which fully characterize the solution of the model. Equation (31) 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.

C.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, $\lambda_i = \frac{n\alpha\sigma^2}{n\alpha\sigma^2+2\phi} \equiv \lambda$ 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:

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

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

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

The average prices in equilibrium becomes:

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

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

C.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 2. *Two players i and j are path-connected if there is a sequence of agents $\{s_1, s_2, \dots, s_k\}$ such that:*

$$g_{is_1} = g_{s_1s_2} = \dots = g_{s_{k-1}s_k} = g_{s_kj} = 1.$$

The following proposition shows that when there is no counterparty-specific risk aversion, i.e., $\phi = 0$, 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 3. *If $\phi = 0$ 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:

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

□

Corollary 1. *If $\phi = 0$ 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:

$$R_{ij} - \mu = \sigma^2 \alpha \bar{w} \quad \forall i, j,$$

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

Proof. We know that:

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

where $\bar{z}w$ 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 4. *If $\phi \rightarrow \infty$ 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 $\gamma_{ij} = 0$ for any two agents i and j . □

C.2 Three-agent example

In the three-agent example, the trading network is given by:

$$G = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}. \quad (32)$$

The trading network is also represented in Figure 6. Agents 1, 2, and 3 have pre-trade exposures given by ω_1 , ω_2 , and ω_3 , respectively. Furthermore, we assume $\omega_1 = 0$.

Let us solve for agent 1's net position using Equation (13):

$$z_1 = \lambda_1 \frac{1}{3} (\omega_2 + \omega_3 + z_1 + z_2 + z_3) = \frac{\alpha\sigma^2}{3\alpha\sigma^2 + 2\phi} (\omega_2 + \omega_3). \quad (33)$$

The derivation above uses the fact that $\omega_1 = 0$, along with the clearing condition given by: $z_1 + z_2 + z_3 = 0$.

Using Equation (13) for agent 2, we have that agent 2's post-trade exposure, $z_2 + \omega_2$, is given by:

$$z_2 + \omega_2 = \frac{\alpha\sigma^2 z_1 + 2\phi\omega_2}{\alpha\sigma^2 + 2\phi}$$

and, using z_1 from Equation (33), agent 2's net position, z_2 , is given by:

$$z_2 = \left(\frac{\alpha\sigma^2}{\alpha\sigma^2 + 2\phi} \right) \left(\frac{\alpha\sigma^2}{3\alpha\sigma^2 + 2\phi} \right) \left[\omega_3 - \omega_2 \left(\frac{2\alpha\sigma^2 + 2\phi}{\alpha\sigma^2} \right) \right]. \quad (34)$$

Similarly, agent 3's net position and post-trade exposure are given by:

$$z_3 + \omega_3 = \frac{\alpha\sigma^2 z_1 + 2\phi\omega_3}{\alpha\sigma^2 + 2\phi},$$

and

$$z_3 = \left(\frac{\alpha\sigma^2}{\alpha\sigma^2 + 2\phi} \right) \left(\frac{\alpha\sigma^2}{3\alpha\sigma^2 + 2\phi} \right) \left[\omega_2 - \omega_3 \left(\frac{2\alpha\sigma^2 + 2\phi}{\alpha\sigma^2} \right) \right].$$

Equilibrium prices are given by:

$$\begin{aligned} R_{12} - \mu &= \frac{1}{2}\alpha\sigma^2 (z_1 + \omega_1 + z_2 + \omega_2) \\ &= \frac{1}{2}\alpha\sigma^2 \times \frac{2\alpha\sigma^2 z_1 + 2\phi\omega_2 + 2\phi z_1}{\alpha\sigma^2 + 2\phi} \end{aligned}$$

and

$$R_{13} - \mu = \frac{1}{2}\alpha\sigma^2 \times \frac{2\alpha\sigma^2 z_1 + 2\phi\omega_3 + 2\phi z_1}{\alpha\sigma^2 + 2\phi}$$

Taking the difference, we have:

$$R_{13} - R_{12} = \frac{1}{2}\alpha\sigma^2 \left(\frac{2\phi}{\alpha\sigma^2 + 2\phi} \right) (\omega_3 - \omega_2) > 0 \quad \Leftrightarrow \omega_3 > \omega_2.$$

The equilibrium price for agent 1 if she would trade with herself is given by:

$$\begin{aligned} R_{11} - \mu &= \alpha\sigma^2(z_1 + \omega_1) \\ &= \alpha\sigma^2 z_1 \\ &= \alpha\sigma^2 \frac{\alpha\sigma^2}{3\alpha\sigma^2 + 2\phi} (\omega_2 + \omega_3). \end{aligned}$$

Hence, we have:

$$\begin{aligned} R_{12} - \mu &= \frac{1}{2}\alpha\sigma^2 \frac{2\alpha\sigma^2 z_1 + 2\phi\omega_2 + 2\phi z_1}{\alpha\sigma^2 + 2\phi}, \\ R_{13} - \mu &= \frac{1}{2}\alpha\sigma^2 \times \frac{2\alpha\sigma^2 z_1 + 2\phi\omega_3 + 2\phi z_1}{\alpha\sigma^2 + 2\phi}, \end{aligned}$$

and

$$\begin{aligned} R_{13} + R_{12} - 2R_{11} &= \alpha\sigma^2 \frac{2\alpha\sigma^2 z_1 + \phi\omega_2 + \phi\omega_3 + 2\phi z_1}{\alpha\sigma^2 + 2\phi} - 2\alpha\sigma^2 z_1 \\ &= 2\alpha\sigma^2 \frac{\alpha\sigma^2 z_1 + \phi\frac{1}{2}(\omega_2 + \omega_3) + \phi z_1 - z_1(\alpha\sigma^2 + 2\phi)}{\alpha\sigma^2 + 2\phi} \\ &= \phi\alpha\sigma^2 \frac{(\omega_2 + \omega_3) - 2z_1}{\alpha\sigma^2 + 2\phi} \\ &= \phi\alpha\sigma^2 \underbrace{\frac{1 - \frac{2\alpha\sigma^2}{3\alpha\sigma^2 + 2\phi}}{\alpha\sigma^2 + 2\phi}}_{>0} (\omega_2 + \omega_3). \end{aligned}$$

Thus:

$$R_{13} - R_{11} > R_{11} - R_{12} \Leftrightarrow \omega_3 > -\omega_2.$$

Based on Equation (34), notice that

$$z_2 > 0 \Leftrightarrow \omega_3 > \omega_2 \frac{2\alpha\sigma^2 + 2\phi}{},$$

which shows the third feature of the three-agent example.

C.3 Model with Core-Periphery Network

C.3.1 Dealer market

Applying Equation (13) to dealers gives us the following expression for the post-trade exposures of dealers:

$$z_i + \omega_i = (1 - \lambda_d)\omega_i + \lambda_d\bar{\omega} \quad \forall i = 1, \dots, n_d, \quad (35)$$

where

$$\lambda_d = \frac{n\alpha\sigma^2}{n\alpha\sigma^2 + 2\phi}$$

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:

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

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

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

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:

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

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

C.3.2 Customer market

Applying Equation (13) to customers gives us the following expression for their post-trade exposures:

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

where $\bar{\lambda}_c = \frac{(n_d+1)\alpha\sigma^2}{(n_d+1)\alpha\sigma^2+2\phi}$. We can use Equation (17) to write the post-trade exposures as follows:

$$z_i + \omega_i = \lambda_c \bar{\omega} + (1 - \lambda_c)\omega_i - \lambda_c(1 - \lambda_d)(\bar{\omega} - \bar{\omega}_d) \quad \forall i = n_d + 1, \dots, n, \quad (36)$$

where $\lambda_c = \frac{n_d\alpha\sigma^2}{n_d\alpha\sigma^2+2\phi}$.

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:

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

where we used Equations (35) and (36) 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:

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

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

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

C.3.3 Calibration

From Equation (17), we can compute $(\bar{\omega} - \bar{\omega}_d)$ as a function of \bar{z}_d and λ_d :

$$\bar{\omega} - \bar{\omega}_d = \frac{\bar{z}_d}{\lambda_d}. \quad (37)$$

Furthermore, we can write λ_d as follows:

$$\lambda_d = \frac{\alpha\sigma^2 \bar{z}_d}{\alpha\sigma^2 \bar{z}_d + (\bar{R}_c - \bar{R}_d) 2 \left(1 - \frac{n_d}{n}\right)}, \quad (38)$$

by taking the difference between Equations (19) and (21) and solving for λ_d .

The complete network benchmark price can be written as:

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

where the first equality is derived by combining Equations (19) and (37), and the second equality is obtained by substituting in Equation (38).

We can rearrange Equation (23) to compute α as follows:

$$\alpha = \frac{1}{\sigma^2\bar{\omega}} \left[\bar{R}_d - \mu + (\bar{R}_c - \bar{R}_d)2 \left(1 - \frac{n_d}{n}\right) \right], \quad (39)$$

and we can rearrange Equation (18) to compute ϕ as follows:

$$\phi = \frac{1}{2} n \alpha \sigma^2 \left(\frac{1 - \lambda_d}{\lambda_d} \right) \quad (40)$$

$$= \frac{1}{2} n \alpha \sigma^2 \left[\frac{(\bar{R}_c - \bar{R}_d) 2 \left(1 - \frac{n_d}{n}\right)}{\alpha \sigma^2 \bar{z}_d} \right] \quad (41)$$

$$= (n - n_d) \left(\frac{\bar{R}_c - \bar{R}_d}{\bar{z}_d} \right) \quad (42)$$

where λ_d is computed from Equation (38).

To compute the model-implied dealers' pre-trade exposures, we can rearrange Equation (35) as follows:

$$\bar{\omega} - \omega_i = \frac{z_i}{\lambda_d} \quad \forall i = 1, \dots, n_d, \quad (43)$$

and to compute customers' pre-trade exposure we can rearrange Equation (36) as well:

$$\bar{\omega} - \omega_i = \frac{z_i}{\lambda_c} + (1 - \lambda_d)(\bar{\omega} - \bar{\omega}_d) \quad \forall i = n_d + 1, \dots, n. \quad (44)$$

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.

C.4 Price Impact

In this subsection, we derive an alternative version of the benchmark model in which agents internalize the effect of their own exposure to the underlying risk on equilibrium prices. In the benchmark model, equilibrium prices are given by Equation (12), which means that when agent i sells insurance to agent j , then she receives R_{ij} as payment. Notice, however, this equilibrium price depends on both agents' post-trade exposures. Notice that agent i optimally chooses the total net exposure to the underlying default risk, i.e., z_i , but takes equilibrium prices as given. In this subsection, we derive equilibrium allocations and prices when agents take into account the effect of their net exposure to the underlying default risk on prices.

To solve this model, we guess and verify that the equilibrium price in a bilateral trade will be a linear combination of the counterparties' post-trade exposures. Specifically, we assume that:

$$R_{ij} - \mu = A + B\alpha\sigma^2 z_i + C\alpha\sigma^2 z_j + D\alpha\sigma^2 \omega_i + E\alpha\sigma^2 \omega_j,$$

where A , B , C , D , and E are coefficients to be determined. The assumption here is similar to a Cournot competition model in which firms take their competitors' quantities as given and equilibrium is pinned by the fixed point of best-responses. In our setting, agent i take j 's exposure and all pre-trade exposures as given but internalize the effect of i 's exposure on equilibrium prices.

Formally, agent i solves the following optimization problem:

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

subject to

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

$$z_i = \sum_{j=1}^n \gamma_{ij},$$

and

$$R_{ij} - \mu = A + B\alpha\sigma^2 (z_i + \omega_i + z_j + \omega_j).$$

Hence, the first-order conditions imply:

$$\begin{aligned} R_{ij} - \mu + \sum_s \gamma_{is} \frac{\partial}{\partial \gamma_{ij}} R_{ij} &= \alpha\sigma^2 (z_i + \omega_i) + \phi\gamma_{ij} \\ \implies R_{ij} - \mu &= \alpha\sigma^2 (z_i + \omega_i - Bz_i) + \phi\gamma_{ij} \end{aligned}$$

Under the no transaction cost assumption, i.e., $R_{ij} = R_{ji}$, along with the bilateral clearing condition, i.e., $\gamma_{ij} + \gamma_{ji} = 0$, we can write equilibrium prices as follows:

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

Applying the method of undetermined coefficients to our initial guess gives

$$\begin{aligned} A &= 0, \\ B = C &= \frac{1}{3}, \end{aligned}$$

and

$$D = E = \frac{1}{2}.$$

Hence, equilibrium prices are given by:

$$R_{ij} - \mu = \alpha\sigma^2 \frac{1}{2} [\omega_i + \omega_j + \tilde{z}_i + \tilde{z}_j], \quad (45)$$

and first-order condition can be written as:

$$R_{ij} - \mu = \alpha\sigma^2 (\omega_i + \tilde{z}_i) + \phi\gamma_{ij}, \quad (46)$$

where $\tilde{z}_i = \frac{2}{3}z_i$.

To get derive equilibrium allocations, we can combined Equations (45) and (46), along with the fact that $z_i = \sum_{j=1}^n \gamma_{ij}$:

$$\tilde{z}_i + \omega_i = \left(1 - \tilde{\lambda}_i\right) \omega_i + \tilde{\lambda}_i \sum_{j=1}^n \tilde{g}_{ij} (\tilde{z}_j + \omega_j) \quad \forall i = 1, \dots, n \quad (47)$$

where $\tilde{z}_i = \frac{2}{3}z_i$, $\tilde{g}_{ij} = \frac{g_{ij}}{K_i}$, $K_i = \sum_{j=1}^n g_{ij}$, and $\tilde{\lambda}_i = \frac{K_i\alpha\sigma^2}{K_i\alpha\sigma^2 + 3\phi} \in (0, 1)$.

Notice that Equation (47) is extremely similar to Equation (13), except that under price impact we have \tilde{z}_i and $\tilde{\lambda}_i$ instead of z_i and λ_i . As a result, the analyses discussed in the paper hold in a price impact environment as well.

APPENDIX TABLES

Table A1: Connection Turnover

Percentile	Gross Turnover (%)		
	All	Customers	Dealers
p5	0.0	0.0	0.9
p10	0.0	0.0	1.2
p25	0.0	0.0	1.6
p50	0.0	0.0	2.1
p75	0.0	0.0	2.8
p90	1.0	1.0	3.6
p95	13.6	13.7	4.5
Mean	2.8	2.8	2.4

Notes: This table presents statistics related to network connection formation and destruction. For each counterparty i and date t , we compute:

$$\text{Gross Turnover}_{it} = \frac{(\text{Number of New Connections to } i \text{ at } t + \text{Number of Destroyed Connections to } i \text{ at } t)}{(\text{Total Number of } i\text{'s connections at } t - 1)}$$

Next, for each date t , we compute the cross-sectional distribution of $\text{Gross Turnover}_{it}$ across counterparties i . Finally, we take time-series averages of each cross-sectional statistic and report them in the table. Connections are determined based on open positions at the end of each week. The data covers the period January 2010 through December 2013. Dealers and customers are defined according to Section 2.2. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

Table A2: Average Dealer CDS Exposure

Method	\bar{z}_d
Notional, Index Positions Only	0.020
Notional, Disaggregated Positions	0.071
Notional, Beta-Weighted	0.045
DV01, Beta-Weighted (%)	0.99

Notes: This table presents some basic summary statistics about the average credit exposure of dealers to the CDX investment grade (IG) index, denoted by \bar{z}_d . We define exposure to the CDX IG index in four ways: (i) using the net notional sold directly on the CDX IG index; (ii) using the net notional sold on constituents of the IG index, accounting for direct single name positions and indirect positions that come from CDS index exposure; (iii) a beta-weighted average of the net notional sold across all CDS positions, with betas computed with respect to the CDX IG index; and (iv) a beta-weighted average DV01 across all positions, which just measures how much the entire CDS portfolio would lose if there was a one hundred basis point move in the CDX IG index. See Section B for complete details. In all cases, positive values indicates that dealers are on average net sellers. For all metrics, we compute the exposure of dealers in our sample, then scale this exposure by their market capitalization. This is what we call a dealer-specific z_i . \bar{z}_d in each week is the cross-sectional average of each z_i across dealers. The table reports average weekly \bar{z}_d over the period January 2010 through December 2013. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

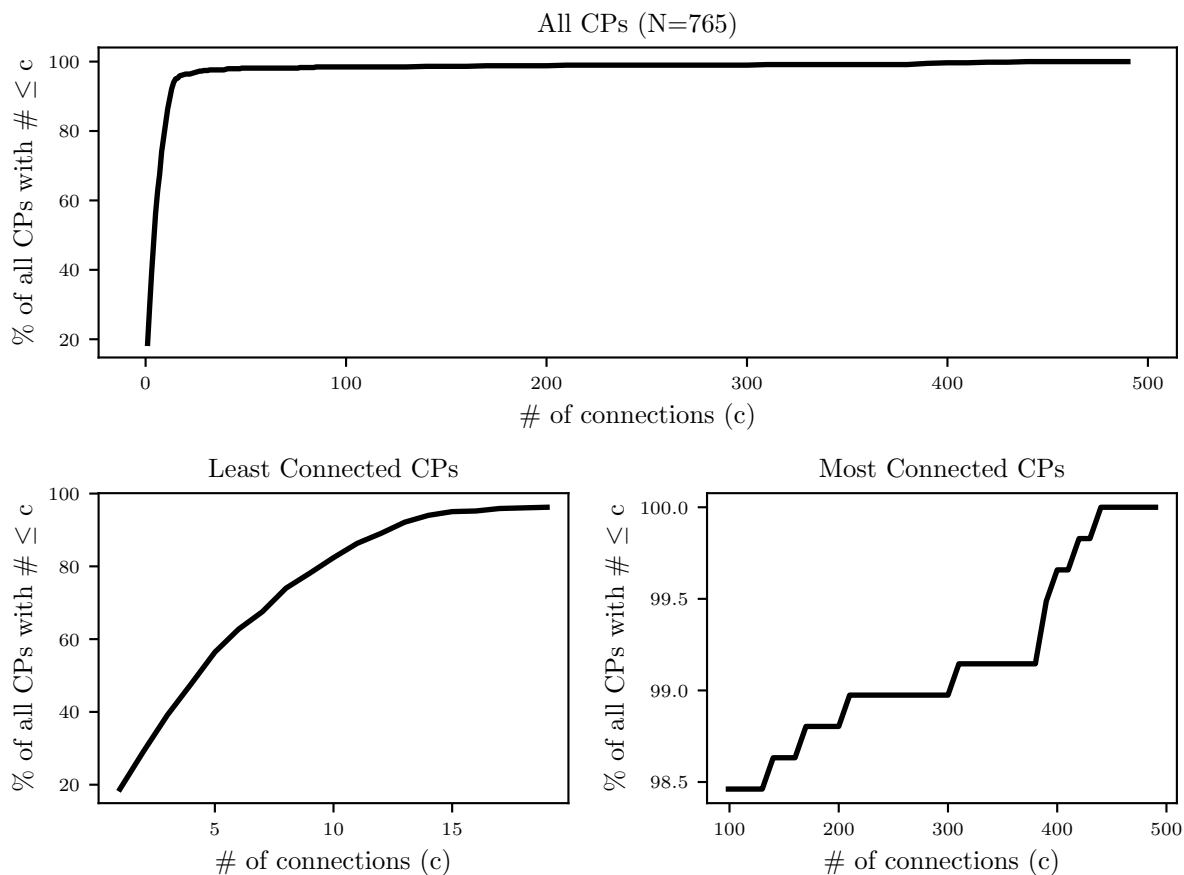
Table A3: Summary Statistics of Spreads by Firm-Date-Maturity Buckets

	≥ 2 per (f, t, m)		≥ 5 per (f, t, m)		≥ 10 per (f, t, m)	
	EW	LW	EW	LW	EW	LW
# Trades	4	9	9	13	16	19
Avg. Spread (bps)	154	160	171	166	157	156
σ (Spread)	9	11	12	12	13	12
σ -to-Avg. Spread (%)	8	9	10	10	11	11
<i>IQR</i> (Spread)	13	15	16	17	18	17
<i>IQR</i> -to-Avg. Spread (%)	12	13	14	15	17	16
<i>Range</i> (Spread)	18	26	32	34	38	38
<i>Range</i> -to-Avg. Spread (%)	16	22	27	29	33	33
# of (f, t, m) groups	78,539	78,539	20,535	20,535	6,611	6,611

Notes: This table presents summary statistics of spreads and trading activity across firm-date-maturity buckets (f, t, m) pairs. Maturity buckets m are defined by grouping transactions into one of the following four buckets based on its maturity: (i) 0-2 years; (ii) 2-4 years; (iii) 4-6 years; and (iv) 7+ years. Within each (f, t, m) group, we compute each statistic (e.g. average spread, σ -to-average spread, etc.). We then take an equal-weighted (EW) average of these statistics across (f, t, m) groups. We also liquidity-weight (LW) across groups, where a groups' liquidity weight is determined by the number of trades in that (f, t, m) group. *IQR* is the interquartile range and *Range* is the maximum-minus-minimum spread. 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, drop contracts between nondealers and nondealers (only 0.31% of total), and pool contracts of different tiers and doc-clauses. We winsorize the 5% tails of outliers based on Markit and fair-value transaction spreads. 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.

APPENDIX FIGURES

Figure A1: Degree Distribution in the CDS Network



Notes: This figure shows the degree distribution across counterparties in the U.S. CDS market from January 2010 through December 2013. The picture is based on the matrix G_t where element $G_{i,j,t}$ equals one if i and j have an open position with each other on date t . If i and j do not have an open position, $G_{i,j,t}$ equals zero. The picture is then constructed by averaging G_t over all dates to arrive at what we call \bar{G} . We include all CDS positions, both single name and index, when defining G_t on each date. We define the degree, or number of connections, for counterparty i as $D_i \equiv \sum_j \bar{G}_{i,j}$. The top subplot shows the percent of counterparties who have D_i less than or equal to various thresholds. The bottom left subplot shows the same data as the top subplot, but zooms in on those counterparties who have $D_i \leq 10$. The bottom right subplot shows the same data as the top subplot, but zooms in on those counterparties who have $D_i \geq 100$. Source: Authors' analysis, which uses data provided to the OFR by the Depository Trust & Clearing Corporation.

References

- ACEMOGLU, D., V. M. CARVALHO, A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2012): “The network origins of aggregate fluctuations,” *Econometrica*, 80(5), 1977–2016.
- ACEMOGLU, D., A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2013): “Systemic risk and stability in financial networks,” Discussion paper, National Bureau of Economic Research.
- ALLEN, F., AND A. BABUS (2009): “Networks in Finance,” in *Network-based Strategies and Competencies*, ed. by P. Kleindorfer, and J. Wind, pp. 367–382.
- ARORA, N., P. GANDHI, AND F. A. LONGSTAFF (2012): “Counterparty credit risk and the credit default swap market,” *Journal of Financial Economics*, 103(2), 280–293.
- ATKESON, A., A. EISFELDT, AND P.-O. WEILL (2013): “The Market for OTC Derivatives,” Working Paper.
- (2015): “Entry and exit in OTC Derivatives Markets,” *Econometrica*, Forthcoming.
- AYMANN, C., C.-P. GEORG, D. BUNDESBANK, AND B. GOLUB (2017): “Illiquidity Spirals in Coupled Over-the-Counter Markets,” .
- BABUS, A. (2013): “The Formation of Financial Networks,” *RAND Journal of Economics*, Forthcoming.
- BABUS, A., AND T.-W. HU (2017): “Endogenous intermediation in over-the-counter markets,” *Journal of Financial Economics*, 125(1), 200–215.
- BABUS, A., AND P. KONDOR (2013): “Trading and information diffusion in over-the-counter markets,” Discussion paper, Working paper.
- BALA, V., AND S. GOYAL (2000): “A noncooperative model of network formation,” *Econometrica*, 68(5), 1181–1229.
- BOLTON, P., AND M. OEHMKE (2013): “Strategic conduct in credit derivative markets,” *International Journal of Industrial Organization*, 31(5), 652–658.
- CAPPONI, A., W. CHENG, S. GIGLIO, AND R. HAYNES (2017): “The Collateral Rule: An Empirical Analysis of the CDS Market,” .
- CHEN, K., M. FLEMING, J. JACKSON, A. LI, AND A. SARKA (2011): “An Analysis of CDS Transactions: Implications for Public Reporting,” *NY Fed Working Paper*, (517).
- COLLIN-DUFRESNE, P., B. JUNGE, AND A. B. TROLLE (2016): “Market structure and transaction costs of index CDSs,” .

- 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., A. LI, AND T. LUBKE (2011): “Policy Perspectives on OTC Derivatives Market Infrastructure,” *Federal Reserve Staff Report*, (424).
- DUFFIE, D., M. SCHEICHER, AND G. VUILLEMEY (2014): “Central Clearing and Collateral Demand,” *Journal of Financial Economics*.
- ELLIOTT, M., B. GOLUB, AND M. O. JACKSON (2014): “Financial networks and contagion,” *American Economic Review*, 104(10), 3115–53.
- ELLIOTT, M., AND J. HAZELL (2016): “Endogenous financial networks: Efficient modularity and why shareholders prevent it,” .
- EROL, S., AND R. VOHRA (2017): “Network formation and systemic risk,” .
- FARBOODI, M. (2014): “Intermediation and Voluntary Exposure to Counterparty Risk,” .
- GALEOTTI, A., AND S. GOYAL (2010): “The law of the few,” *The American Economic Review*, pp. 1468–1492.
- GLODE, V., AND C. OPP (2017): “Can Centralized Markets be More Efficient?,” Working Paper.
- GOFMAN, M. (2015a): “Efficiency and Stability of a Financial Architecture with Too-Interconnected-To-Fail Institutions,” Working Paper.
- (2015b): “A Network-Based Analysis of Over-the-Counter Markets,” .
- GREEN, R. C., B. HOLLIFIELD, AND N. SCHÜRHOFF (2007a): “Dealer intermediation and price behavior in the aftermarket for new bond issues,” *Journal of Financial Economics*, 86(3), 643–682.
- GREEN, R. C., B. HOLLIFIELD, AND N. SCHÜRHOFF (2007b): “Financial Intermediation and the Costs of Trading in an Opaque Market,” *Review of Financial Studies*, 20(2), 275–314.
- GUNDUZ, Y. (2018): “Mitigating Counterparty Risk,” Working Paper.
- HE, Z., AND A. KRISHNAMURTHY (2013): “Intermediary Asset Pricing,” *American Economic Review*, 103(2), 732–70.

- HENDERSHOTT, T., D. LI, D. LIVDAN, AND N. SCHURHOFF (2016): “Relationship Trading in OTC Markets,” *Working Paper*.
- HERSKOVIC, B. (2017): “Networks in Production: Asset Pricing Implications,” *Forthcoming, The Journal of Finance*.
- HERSKOVIC, B., B. T. KELLY, H. LUSTIG, AND S. VAN NIEUWERBURGH (2017): “Firm Volatility in Granular Networks,” *Working Paper*.
- HERSKOVIC, B., AND J. RAMOS (2015): “Information through Peers: Coordination Impacts on Network Formation,” *Working Paper*.
- HOLLIFIELD, B., A. NEKLYUDOV, AND C. SPATT (2017): “Bid-Ask Spreads, Trading Networks and the Pricing of Securitizations,” *Forthcoming, Review of Financial Studies*.
- KAWAKAMI, K., AND Z. ZHONG (2016): “The Risk Sharing Benefit versus the Collateral Cost: The Formation of the Inter-Dealer Network in Over-the-Counter Trading,” *Working Paper*.
- LI, D., AND N. SCHÜRHOFF (2018): “Dealer networks,” *Forthcoming, The Journal of Finance*.
- LONGSTAFF, F. A., S. MITHAL, AND E. NEIS (2005): “Corporate yield spreads: Default risk or liquidity? New evidence from the credit default swap market,” *The Journal of Finance*, 60(5), 2213–2253.
- MAGGIO, M. D., A. KERMANI, AND Z. SONG (2017): “The value of trading relations in turbulent times,” *Journal of Financial Economics*, 124(2), 266 – 284.
- MALAMUD, S., AND M. ROSTEK (2014): “Decentralized Exchange,” *Working Paper*.
- OEHMKE, M., AND A. ZAWADOWSKI (2013): “The Anatomy of the CDS Market,” *SSRN*: <http://ssrn.com/abstract=2023108>.
- PELTONEN, T. A., M. SCHEICHER, AND G. VUILLEMEY (2014): “The network structure of the CDS market and its determinants,” *Journal of Financial Stability*.
- SCHACHAR, O. (2012): “Exposing The Exposed: Intermediation Capacity in the Credit Default Swap Market,” *NY Fed Working Paper*.
- SHEN, J., B. WEI, AND H. YAN (2016): “Financial intermediation chains in an OTC market,” .
- SIRIWARDANE, E. N. (2018): “Limited Investment Capital and Credit Spreads,” *Forthcoming, The Journal of Finance*.
- WANG, C. (2016): “Core-periphery trading networks,” .