

Federal Reserve Bank of Dallas
Globalization and Monetary Policy Institute
Working Paper No. 313

<https://www.dallasfed.org/~media/documents/institute/wpapers/2017/0313.pdf>

The Double-Edged Sword of Global Integration: Robustness, Fragility & Contagion in the International Firm Network*

Everett Grant
Federal Reserve Bank of Dallas

Julieta Yung
Federal Reserve Bank of Dallas

May 2017

Abstract

We estimate global inter-firm networks across all major industries from 1981 through 2016 and provide the first empirical tests for both robust (beneficial) and fragile (harmful) network behavior, relating firms' health with global integration. More connected firms are less likely to be in distress and have higher profit growth and equity returns, but are also more exposed to direct contagion from distressed neighboring firms and network level crises. Our analysis reveals the centrality of finance in the international firm network and increased globalization, with greater potential for crises to spread globally when they do occur.

JEL codes: C3, F36, F61, G15

* Everett Grant, Research Department, Federal Reserve Bank of Dallas, 2200 N. Pearl Street, Dallas, TX 75201. 214-922-5622. egrant.econ@gmail.com. Julieta Yung, Research Department, Federal Reserve Bank of Dallas, 2200 N. Pearl Street, Dallas, TX 75201. 214-922-5443. julieta.yung@dal.frb.org. We thank Alex Chudik, Galina Hale, Karen Lewis, Eric van Wincoop, Kei-Mu Yi, and Eric Young for helpful comments and suggestions. We also thank participants of the 2016 Federal Reserve Bank of Dallas Seminar Series, 2017 Federal Reserve System Committee on International Economic Analysis, and Spring 2017 University of Houston Economic Workshop. The views in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System.

1 Introduction

After a series of events during the past decade with swift and severe consequences for firms across the globe — including the U.S. Sub-Prime Mortgage and Eurozone Debt Crises — researchers and policymakers have reexamined our understanding of systemic risk and contagion across financial firms given their perceived centrality to the transmission of these episodes. To this end, the Bank of England’s then Executive Director for Financial Stability, Andrew G. Haldane, drew parallels between financial systems and complex systems in other fields and postulated that financial networks may be “robust-yet-fragile” (RyF).¹ This term describes the duality wherein greater network connectedness acts to cushion and smooth the effects of small shocks; however, these connections may also spread crises by passing on large shocks and propagating contagion.

What this research often misses, however, is the striking extent and breadth of these events, not only across the global banking system, but across all business sectors. Reflecting this interconnectedness across non-finance firms, Ford CEO Alan R. Mulally appeared before the Senate Banking Committee in November 2008, lobbying for bailouts for General Motors and Chrysler in the midst of the ongoing financial turmoil, though not Ford itself. He worried that because of the major automobile manufacturers’ significant overlap in suppliers, dealers, and other business partners that “the collapse of one of our competitors would have a severe impact on Ford... because the domestic auto industry is highly interdependent. It would also have devastating ripple effects across the entire U.S. economy.”² The synchronicity across firms worldwide during recent crisis events emphasizes the importance of the network of connections between them. A better understanding of the features of firm networks could have broad implications, as systemic importance has been used to support policies ranging from these government bailouts, to import tariffs, to the sweeping Dodd-Frank reforms aimed at preventing future crises.

We estimate global inter-firm networks in real time and test the robustness and fragility of these systems, particularly as they relate to crisis susceptibility and firm health. Our work expands the analysis of inter-firm networks beyond financial companies with a set of global firms from 1981-2016 across up to 49 countries and in all major industry categories. Both firms’ industry and locality influence their network position, with financial firms typically at the center of the global network. We find evidence for the existence of robustness through more connected firms having better performance and resistance to network level crises, and of fragility through direct contagion, vulnerability to system-wide shocks, and these latter two mechanisms reinforcing one another. Our results suggest that a one standard deviation increase in a firm’s total connections in from other

¹See Haldane (2009) for the full speech text.

²See Mulally (2008) for details.

firms reduces the probability of firm distress by 3.6% and is associated with improved firm health — e.g., increased growth of 1.2% for monthly equity prices, 4.0% for quarterly profits, and 1.7% for both quarterly revenue and annual return on equity. At the same time, if a neighboring firm is in distress then its contagious effect is nine times greater than its diversification benefit. The robust network features and initially low level of connectedness at the beginning of our sample suggest why representative agent and closed economy models have had such success in economics; however, they are not appropriate to study events such as the Global Financial and Eurozone Debt Crises owing to increased global integration, and the network’s exhibited firm heterogeneity and fragility.

It is an open question in economics whether inter-firm networks are a steadying or destabilizing force for the firms within them.³ Pioneering research by Allen and Gale (2000) and Freixas et al. (2000) into the nature of systemic risk in bank networks found that more connected firm networks are robust by mitigating risk across counterparties. Conflicting conclusions were reached by later research, which found that bank network connectedness creates fragility by acting as a conduit for contagion.⁴ Subsequent theoretical work attempted to reconcile these divergent results, but the RyF nature of the global inter-firm network has not yet been explicitly tested nor empirically measured.⁵ Further, while most of the attention in the firm network literature has been given to the financial sector, recent work by Acemoglu et al. (2012) and Baqaee (2016) demonstrated the importance of non-financial firm networks. We therefore estimate firm networks and use a regression framework to test their RyF nature along several dimensions and run simulations to evaluate these properties.

As a first step, we build from the work of Demirer et al. (2015), Diebold and Yilmaz (2009, 2014, 2016), and Scida (2015) by using vector-autoregressive (VAR) techniques to estimate economic networks.⁶ Firms can be connected in myriad ways: shared product markets; sourcing from the same labor markets; obtaining credit from the same financial center; location; having similar business risks; etc. Collecting the data on the multitude of firm interactions would be near impossible, and it

³The evolution of global financial networks was the subject of research by Billio et al. (2012) that analyzed monthly returns of hedge funds, banks, broker-dealers and insurance companies, and that of McGuire and Tarashev (2008) and Minoiu and Reyes (2013) into cross-border banking flows. See Glasserman and Young (2016) for an excellent overview of the development of research into financial networks.

⁴For example, Blume et al. (2011), Blume et al. (2013) and Vivier-Lirimont (2006) find evidence of financial network fragility.

⁵See for example the work on the robustness and fragility of bank networks by Acemoglu et al. (2015b), Gai and Kapadia (2010), Gai et al. (2011), and Elliott et al. (2014).

⁶These papers offer noteworthy insights into a number of areas of economics by using tools from computer science and physics to characterize and visualize the relationships between the network members. For example, the Diebold-Yilmaz series demonstrates the evolution towards greater global integration over time, especially during crisis periods.

would be difficult to find a function to properly aggregate them into a single network.⁷ The efficient markets hypothesis postulates that firms' equity prices reflect all available information about them, including proximity across these many dimensions. If two firms produce in the same labor market, borrow from the same banks, have similar risk profiles, hold similar assets, etc... these should be reflected in linked equity returns for the two. If this were not the case then profit opportunities would exist for investors savvy enough to trade on the omitted information. The equity returns act as both an easily available proxy for collecting data on the different channels through which firms are connected — one that can be updated in real-time — and as an aggregator function to combine them.

To infer the firm networks from daily equity prices, we do VAR estimation using adaptive elastic-net shrinkage methods to avoid the dimensionality problem.⁸ We then use these estimates to calculate generalized impulse response functions (GIRFs) between every pair of firms. In these networks each node represents a firm, and the edges connecting them are the GIRFs from the source firm to the terminal one. These edges capture the strength and direction of interactions — which flow in both directions — making them what is referred to as weighted, directed networks. Using GIRFs, the network edges capture how much other target firms' equity prices are expected to move conditional on a change in the price of the source node, rather than the effect of an orthogonalized shock emanating from the source firm. Therefore, the network edges do not necessarily reflect cause and effect but rather capture co-movements regardless of the nature of the shock.

We start our sample with the union of the top 1% of global firms by equity market capitalization for each year from 1992-2016, representing more than 55% of the global market capitalization of all exchange listed firms. We then examine three subsets of the data: global firms that are continuously traded throughout 1991-2016; U.S. listed firms that are continuously traded throughout 1981-2016; and 5-year rolling windows of global firms that are continuously traded throughout each sub-period. We do the latter in order to deal with any potential issues relating to survivorship bias or new firms entering the market. Once we estimate the networks, we analyze the full time period network connections, as well as rolling windows to study the evolution of the network structure. We identify several empirical features that characterize the global inter-firm network: both industry and locality play important roles in connectedness between firms; firms have become more connected over time; and finance is at the center of the global firm network. These results generalize and refine the

⁷To overcome these data collection challenges, Hoberg and Phillips (2016) use text-based machine learning analysis of firm 10-K product descriptions to identify a new network of industries based on product classifications rather than production processes to study how industries and their competitors change.

⁸Several of the papers in the Diebold-Yilmaz series, as well as Kitwittanachai (2015) and Scida (2015), also use financial market prices, volatility and credit market data to estimate firm networks.

conclusions of Demirer et al. (2015), who found that when looking at the top 150 global banks it is bank location, not bank assets, that matters for network proximity.

In order to test for robustness and fragility in the firm networks, we run a series of regressions where the dependent variables are different measures of firm health, including equity returns, revenue, CDS spreads, return on equity, and profit. The explanatory variables are a firm's aggregate connectedness in the network estimated over the previous five years, these weights multiplied by indicators of neighboring firms' health, and several aggregate variables such as the TED spread, VIX and the share of firms currently in a distressed state. We define as distressed, firms that are in the bottom 10% of the performance distribution. From this analysis, we find evidence of robustness with well connected firms less likely to be in distress, especially during crisis periods. Moreover, beyond the common definition of robustness as reduced susceptibility to firm distress events, greater connectedness has a positive relationship with firm health, including higher profit and revenue growth. We also find fragility along three dimensions: direct contagion between neighboring firms; firm vulnerability to network level shocks; and reinforced contagion wherein the direct contagion effect becomes greater during crisis periods.

We conclude the analysis by running a series of contagion simulations based on these estimated regressions, which yield three important findings about the global firm network. First, the model simulations provide further evidence of network crisis reinforced contagion. Second, the level of global firm connectedness — and with it the potential for global contagion — grew over our sample period. The evolution of the firm network illuminated by these simulations helps explain why research developed around 2000 found robust networks and later work suggested fragility. Finally, we discover a new mechanism for global contagion that has been activated by increased firm network connectedness, the crisis echo effect. When there is extensive distress in an economy that plays an important part in the global firm network, distress is prone to be transmitted to firms abroad. It is likely there will then be a large number of distressed firms abroad as contagion to them occurs, leading to an echo effect wherein spillovers then return from these firms to spread to other domestic ones. This echo effect should be an important consideration for policymakers in large economies when facing potential crises, since these disturbances can be reflected back by the global network to exacerbate domestic conditions.

The existence of the echo effect indicates that it is misleading to study crises with countries in isolation, with the real world implication that, going forward, large economies like the U.S., China and the Eurozone cannot ignore their impact on developments abroad. Our results provide evidence that while increased network connections may make the global economy more resilient and benefit firm growth, they also may act as a conduit for contagion and elevated systemic risk.

2 Robustness and Fragility in Firm Networks

There is a growing body of literature with theoretical models that predict that inter-bank networks have a RyF structure.⁹ Allen and Gale (2000) and Freixas et al. (2000) focused on direct firm claims on one another as the edges in the financial network connecting banks. In both cases, their results suggest that more connected firm networks are robust by mitigating risk when one of them defaults. One can think of this as diversifying counterparty risk, reducing overall systemic risk.

On the other hand, later work — especially that in the wake of the Global Financial Crisis — tends to support the idea that bank network connectedness acts as a destabilizing force producing financial fragility.¹⁰ Blume et al. (2011) and Blume et al. (2013) modeled the spread of cascading failures through networks that emerge in many domains, with one focus being the contagious failures that spread among financial institutions during a financial crisis. They found evidence of financial network fragility. By modeling strategic financial network formation, they proposed that stable network systems tend to be ones where agents “over-link,” with contagious risks that have dire consequences for the welfare of the participants. Likewise, Vivier-Lirimont (2006) found that the higher is network connectedness, the larger is the number of banks involved in the contagion process, and the quicker is the contagion phenomenon.

These conflicting views on the role of network connectedness with regard to robustness and fragility were reconciled in subsequent work. Elliott et al. (2014), Acemoglu et al. (2015b), Gai and Kapadia (2010), and Gai et al. (2011) developed models that have “hump shaped” effects of bank network connectedness.¹¹ The former paper breaks connectedness into two measures: integration (greater counterparty exposure) and diversification (more counterparties per organization). They

⁹For example, Acemoglu et al. (2015b), Acemoglu et al. (2015a), Gai and Kapadia (2010), Gai et al. (2011) and the references therein.

¹⁰Hale (2012) is the first study of bank-level networks on a global scale — using syndicated loan data to construct a global banking network from 1980 to 2009 of 7938 banking institutions across 141 countries — finding that bank linkages are less likely to form if a country is experiencing a recession or a banking crisis.

¹¹Caccioli and Farmer (2012) and Caccioli et al. (2015) augment the model of Gai and Kapadia (2010). These “humps” arise because very low levels of connectedness with other firms are not enough to pass on contagion, and very high levels provide enough diversification over many firms to buffer against contagion. It is only for moderate levels of connectedness that the probability and potential breadth of contagion are large. The former paper adds firms that are heterogeneous in their degree distributions, balance sheet size and degree correlations. They find that networks with heterogeneous degree distributions are more resilient to contagion occurrences, so long as the source is not a highly connected bank. Their results recommend a targeted policy aimed at reinforcing the stability of the biggest banks to improve financial market stability. Caccioli et al. (2015) includes both direct counterparty failure risk and contagion via overlapping portfolio exposures, finding that neither channel is significant on its own, but that when both channels are active at once, bankruptcies are much more common and have large systemic effects.

found that these have different, non-monotonic effects on the extent of contagion. Diversification initially is a negative, as the small connections transmit crises; however, as it increases further organizations are better insured against one another’s failures. Integration also entails a trade-off between dependence on other banks versus less sensitivity to a firm’s own idiosyncratic shocks. Acemoglu et al. (2015b) found that the extent of financial contagion is sensitive to a phase transition. For small shocks there is robustness from a more densely connected financial network (corresponding to more diversified counterparties). Beyond a certain threshold, greater firm connectedness dominates in its role as a shock transmission mechanism leading to fragility. In this way, the effects of firm connectedness can be twofold, able to contribute to resilience under certain conditions or act as a source of systemic risk under others, with network integration and concentration being important for which force dominates. The latter two papers — on which we will base an example model to illustrate RyF dynamics in Section 2.1 — conclude that, while with more connectedness the probability of contagion may be lower, the exposure can be extremely widespread when distress events do occur.

The firm network literature has been mostly focused on the financial sector; however, Acemoglu et al. (2012) and Baqaee (2016) showed that it is important to consider the network of all firms in an economy.¹² Despite a long running assumption in macroeconomic modeling that idiosyncratic firm shocks will average out and so can be ignored, Acemoglu et al. (2012) demonstrated how these shocks can generate cascade effects whereby firm problems propagate not only to immediate customers, but also to the aggregate macroeconomy.¹³ This propagation occurs if there is significant asymmetry in the scopes of sectors as intermediate input suppliers to others, with the sparseness of the input-output matrix being largely immaterial to the nature of aggregate fluctuations.¹⁴ Baqaee (2016) likewise found that the interaction of input-output networks with industrial structure affects the propagation and amplification of shocks and is an important consideration for business cycle fluctuations.

¹²Other papers exploring connectedness and transmission of shocks across sectors in the macroeconomy are Bak et al. (1993), Carvalho and Gabaix (2013) and Gabaix (2011).

¹³Acemoglu et al. (2015a) developed a unified framework for studying network propagation and amplification of microeconomic shocks dependent on whether the system has (log) linear agent interaction and network utility functions. This work clarifies that, with linear firm interactions, the desire for reduced output volatility lies at the heart of the results in Acemoglu et al. (2012). Likewise, for non-linear interaction functions — such as with bank defaults in Acemoglu et al. (2015b) — they characterized similar phase transitions transforming the role of network interconnections in shaping aggregate performance.

¹⁴Asymmetry in the scope of actors in a network refers to the relative weights of their edges out, with some nodes being far more influential than others. Sparseness measures how many of the potential edges in a network do not exist.

2.1 Modeling Firm Networks

To help explain what is meant by a RyF network, and to motivate why these dynamics are of consequence, in this section we present a simple one period model of connected firms that exhibits these attributes and is a special case of the global firm network model we develop later. The model is based on that of Gai and Kapadia (2010), however, we expand the framework to include both financial and production (i.e., non-finance) firms.

Financial firms interact by making investments in one another, and production firms interact by using each others' goods as intermediate inputs. In both cases, these are denoted as \mathcal{I}_{ij} . For financial firms, this term represents the investment proceeds owed to firm j by firm i from the former's previous investment in the latter that is to be repaid today. In the case of a production firm, \mathcal{I}_{ij} is the expenditure on firm j 's goods as intermediate inputs in firm i 's production process. Alternately, for a more comprehensive firm structure, the \mathcal{I}_{ij} could represent the gross of both loans and intermediate input payment flows, so that the network simultaneously includes both production and financial firms, or firms taking part in both markets. In any case, \mathcal{I}_{ij} is the cash flow from firm i to firm j , assuming that firm i does not go bankrupt.

These cash flows are the connections between firms and are not netted, creating a network that is both weighted and directed. For now, we abstract from the formation of the network, as well as the firm asset pricing, goods pricing and production problems by focusing solely on these cash flows. In order to remain solvent a firm must have a positive liquidation value. Firm i 's solvency condition derived from its balance sheet is:

$$\underbrace{\mathcal{A}_i + \mathcal{D}_i + \sum_{j=1}^J \mathcal{I}_{ji}}_{\text{Income}} - \underbrace{\sum_{j=1}^J \mathcal{I}_{ij}}_{\text{Outflows}} \geq 0. \quad (1)$$

\mathcal{A}_i is the liquidation value of firm i 's assets. In the case of a financial firm these are investments other than those in other firms. In the case of a production firm, this is the value of the firm's capital. \mathcal{D}_i is firm i 's net change in bank deposits for financial firms, and consumer demand for production firms. All quantities are known at the start of the period except for these \mathcal{D}_i shocks, which are exogenously determined. If a great enough negative shock is drawn for firm i and its balance sheet goes negative, then it goes bankrupt and does not pay its neighboring firms, and $\mathcal{I}_{ij} = 0 \forall j$. This can then cause contagion in the adjoining firms, potentially cascading throughout the network.

We run a series of numerical simulations to demonstrate the RyF nature of the network, and how these two properties vary with the average level of firm connectedness. We assume that all of the firms have an initial asset value of 100, with the outside asset (\mathcal{A}_i) and receivables from other

firms ($\sum_{j=1}^J \mathcal{I}_{ji}$) equal to 80 and 20, respectively.¹⁵ We assume that each firm i has receivables from \mathcal{R}_i other firms, and that they are equally distributed among those debtor firms. This could be interpreted, for example, as financial firms choosing investments in other firms to achieve the maximum level of diversification possible. Since the level of each cash flow is determined using the rule above by the firm that is owed the money, and each inter-firm asset is another's liability, this condition determines the magnitude of all the \mathcal{I}_{ij} once which edges exist in the network is known. Finally, we assume that the unperturbed \mathcal{D}_i values are such that each firm has a 4% cash buffer (i.e., the left hand side of Equation (1) sums to 4).

Before running any simulations, we can already discern the RyF properties of the network. From the model structure, a firm will become insolvent if and only if:

$$\frac{\mathcal{B}_i}{\mathcal{R}_i} > \frac{\mathcal{A}_i + \mathcal{D}_i + \sum_{j=1}^J \mathcal{I}_{ji} - \sum_{j=1}^J \mathcal{I}_{ij}}{\sum_{j=1}^J \mathcal{I}_{ji}} = \frac{\text{Cash Buffer}}{\text{Receivables}} = \frac{1}{5}, \quad (2)$$

where \mathcal{B}_i is the number of bankruptcies among those owing money to firm i . With fewer network connections and a lower \mathcal{R}_i , a firm is less likely to have an insolvent neighbor in the network; however, with few connections, having such a neighbor guarantees that it will go bankrupt itself. Alternately, with more connections and a higher \mathcal{R}_i , a firm is more likely to have an insolvent neighbor, but it will take many of them going bankrupt concurrently to cause insolvency due to the diversification benefit of being more connected. At the aggregate network level, widespread contagion is unlikely at the extremes where there are either very few connections — as the contagion is likely to be quarantined by the lack of connectedness — or with very many connections — as it is unlikely that enough firms will become insolvent to trigger bankruptcies across the network.

For the model simulations, we assume a network of 1000 firms. We examine firms across a range of average \mathcal{I}_{ij} per firm ranging from zero to ten to examine how connectedness affects contagion.¹⁶ To achieve the desired level of connectedness, \mathcal{I}_{ij} edges are randomly added to the network until the average connectedness reaches the target level. This is repeated to get 1000 networks for each target level of connectedness.

To model contagion in each of these networks, one firm is chosen at random to default and shocked with a poor \mathcal{D}_i draw, and all of its neighbors that are left insolvent also default. Next, all of the remaining solvent firms are reconsidered to see if these new defaults will make them insolvent. This is continued until the spread of the contagion through the network ceases and there are no newly insolvent firms, allowing all cascading failures to play out. A default outbreak is said

¹⁵These values, and the cash buffers, are chosen to match those of Gai and Kapadia (2010) and are based on averages across developed country banking systems.

¹⁶This will be the same whether the \mathcal{I}_{ij} are averaged over i or j .

to occur if over five percent of the firms in the network become insolvent to focus on network level effects.

The simulation results can be seen in the two panels of Figure 1 plotting the extent of contagion conditional on an outbreak occurring and the probability of an outbreak. Along the x-axis of each panel is the average number of connections per firm in the simulated networks, ranging from zero to ten, and the dots mark the average value across the 1000 simulations at each given level of connectedness. The top panel plots the extent of contagion when it occurs; that is, conditional on an initial firm bankruptcy that causes over five percent of the firms to become insolvent, what is the ultimate share of firms that go bankrupt. The bottom panel plots the probability of an outbreak where over five percent of the firms become insolvent. Both of these plots demonstrate RyF network behavior, producing “hump” shaped figures, a hallmark of the presence of these dynamics. In particular, there are phase transitions in the model generating a contagion window, highlighted with gray shading, where there is a non-trivial probability of a default outbreak, and if one does occur for it to be widespread. On the other hand, in line with our analysis of Equation (2), for either very low or very high levels of connectivity it is unlikely that a default outbreak occurs. While our model is quite simple, this result is a recurring one in the bank and firm network literatures.

These results illustrate the duality of firm networks, with them having the potential to be either robust or fragile. With our empirical analysis below, we seek to determine where the international firm network is relative to the contagion window to better understand the risks it poses, or possibly the benefits firm connectedness provides for weathering shocks and preventing outbreaks. This knowledge is a first step towards helping contend with — or even prevent — possible outbreaks in the future.

3 Data

Our sample is the set of the top 1% of global firms by equity market capitalization from Bloomberg over the past two and a half decades. We queried this list on December 31st every year from 1992-2015 with all market values converted to U.S. dollars at market exchange rates, removed exchange traded funds and took the remaining 1,182 equity securities in the union of these annual sets as our sample.¹⁷ Our focus is on the largest firms in order to ensure that they all have actively traded, liquid equity securities that are highly researched and followed, providing them with accurate price

¹⁷If there were multiple tickers for a single firm then we first filtered by keeping only the securities marked as primary, and if there was still more than one ticker we took the one with the best data coverage. The ticker search begins in 1992 because of Bloomberg data limitations on extending it further back in time.

discovery. Further, we limit the number of firms so that they do not overwhelm our network estimation methodology and feel that this set is sufficient given that the 919 of these firms traded on December 31st 2015 represented just over 55% of the global market capitalization of all exchange listed firms according to data from Bloomberg and the World Federation of Exchanges.¹⁸

For each firm we collect a set of daily, quarterly and annual data. One of the benefits of our methodology is that it can be implemented using easily obtainable data. First, we include the daily equity closing price for each firm from January 1st, 1981 through September 30th, 2016. We also gathered daily 5-year CDS spreads from October 2013 through November 2015 from Markit. From firm quarterly statement data on Bloomberg we collected EBITDA (profit), revenue (Sales/Revenue/Turnover), and total firm equity. We also include the annual return on equity (RoE) over each calendar year. From these we calculate log changes as well as distress/crisis state indicator variables for the worst 10% performance within our sample by each measure. All prices and market values are in local currency, so their log changes do not conflate currency movements with individual firm performance. For a summary of the data see Online Appendix Table A1. The headquarters country, currency of equity issuance, and industry firm meta-data is also included.

In addition to the firm data, we collect daily data to gauge the overall health of the global economy from Global Financial Data: the 3-month TED spread; S&P 500 Total Return Index; and the CBOE S&P 500 Volatility Index (VIX).¹⁹ The log changes of each variable and indicators for the worst 10% of outcomes over monthly, quarterly and annual tenures are included in the dataset.

4 Global Network Estimation

In this section, we first describe how to estimate inter-firm networks for the sample of global firms reviewed in Section 3, explore the features of the network from 1991 through 2016, and finally examine how it evolves this sample period. As mentioned in the introduction with the example of Ford’s CEO lobbying the U.S. Congress for bailouts of GM and Chrysler, firms are highly connected and dependent on one-another in nuanced ways. Ford and the automotive industry provided further reason to consider the full context of the inter-firm network structure when such an eventuality actually played out in Australia. In May 2013, Ford Australia announced it would stop manufacturing cars in the country and did cease doing so in October 2016. Two months after the Ford announcement, Toyota followed suit by stating that it would terminate its manufacturing,

¹⁸Supporting restricting the sample to only the largest firms is work by Gabaix (2011), which found that the idiosyncratic movements of the 100 largest U.S. firms explained about one-third of all variation in output growth, implying that many macroeconomic questions can be resolved by looking at only the behavior of large firms.

¹⁹The TED spread is the difference between 3-month USD LIBOR and the 3-month U.S. Treasury bill yield.

too, citing a lack of volume from those in the component supply chain to keep its Australian production operations economically viable.²⁰ Within a few months, GM Australia announced that it would also have to shutdown, along with many of the parts suppliers, effectively ending the Australian automotive manufacturing industry.

There are many other examples demonstrating inter-firm network effects. Some, as in the case of the August 2016 bankruptcy of Hanjin Shipping, can have dramatic and sweeping echoes through the global economy. The disruptions emanating from the Hanjin bankruptcy stranded \$14 billion of goods at sea around the world.²¹ In others, the reliance of firms on each other are more local but no less real, even when they are not in the same industry nor supply chain. For example, work by Benmelech et al. (2014) demonstrated how local businesses rely upon one another to attract customers when they showed that the closure of a retail chain weakens the economies of agglomeration in the local area. They further showed that companies with greater geographic exposure to the closed retailer are more likely to close themselves, with the externality being higher for smaller stores. These results corroborate recent worries about what the challenges confronting large department stores in the face of increased online competition and over development of retail might mean for shopping malls that rely upon them as anchor stores, and validates the usage of that term, “anchor,” when describing them.²²

On the labor side, Goins and Gruca (2008) used firm stock price as an aggregate measure of firm well-being to find that when a firm announces a significant permanent layoff there is a spillover beyond the announcing firm to others in the industry. They studied a longitudinal sample of layoff announcements in the U.S. oil and gas industry from 1989 to 1996 and found that if the equity price of the announcing firm responds negatively, there will be simultaneous price declines for non-announcing firms; however, close rivals see a countervailing competitive effect dampening this contagion. There are many other examples we could go into, but at this point it should not be difficult to appreciate the importance and variety of inter-firm connections.

4.1 Network Estimation Methodology

We investigate connections between firms; however, doing so is challenging as firms are connected via numerous different networks, many of which cannot be observed in real time or are difficult

²⁰<http://www.smh.com.au/business/the-economy/australias-car-industry-one-year-from-closing-its-doors-20151012-gk7ip0.html>

²¹<http://www.reuters.com/article/us-hanjin-shipping-debt-usa-bankruptcy-idUSKCN11Q2TA>

²²See for example <http://www.businessinsider.com/macys-store-closures-and-what-it-means-for-malls-2015-9>, <https://www.washingtonpost.com/news/digger/wp/2016/08/30/macys-is-closing-100-stores-does-yours-stand-a-chance/>, and http://www.nytimes.com/2015/01/04/business/the-economics-and-nostalgia-of-dead-malls.html?_r=1.

to measure.²³ Further complicating matters is that there has not been a single clear technique proposed to combine the diverse dimensions of connectedness into one inter-firm network.

To overcome these issues we utilize the efficient markets hypothesis, which states that firms' equity prices reflect all available information about them, for if that were not the case then there would be profit opportunities for shrewd investors trading on the omitted information. If two firms are in close proximity along any of the dimensions of the network multiplex connecting them — similar asset holdings, geography, labor skill requirements, distribution networks, etc. — then it should be revealed in their equity returns. This includes not only direct connections between firms, but also exposure to similar common shocks. An aggregate shock can only affect two firms if there is an underlying connection that is a conduit for the shock to affect both, and we intend for our estimated networks to include these latent exposure channels. For example, greater globalization would increase firms' responses to shocks abroad, which we capture in greater estimated network connections. We use equity returns as both a proxy for collecting data on the different channels through which firms are connected and as an aggregator function to combine them, much as Goins and Gruca (2008) exploit equity returns to measure changes in overall firm well-being.

We estimate inter-firm networks that represent co-movements in the firms' equity prices, reflecting similarities in the firms themselves, regardless of the source of any underlying shock(s). Our first step in doing so is to estimate VAR models of the daily equity log returns on lags of themselves and of the returns of other firms. The primary challenge facing this approach is the curse of dimensionality from having numerous firms and potentially multiple lags of their equity returns in the VAR and is likely the obstacle which has prevented similar global network research in the past. As we show, this impediment is worth overcoming as the international network connections between firms are often important conduits for contagion and propagation of shocks.

To contend with this issue, we follow the work of Demirer et al. (2015) in using shrinkage methods borrowed from machine learning to obtain our VAR estimates. The particular shrinkage method that we use is the adaptive elastic-net (AEN) estimator from Zou and Zhang (2009). Their method uses an elastic-net estimation combining the L_1 and L_2 penalties of the LASSO and ridge methods, with the adaptive label in the name referring to the manner in which weights are selected to further penalize coefficients that are smaller in magnitude to aid in the shrinkage. The AEN

²³Collecting off-balance sheet items is challenging in its own right, as described by Demirer et al. (2015); however, tracking all of the inter-firm connections through the supply chain and financing of numerous companies, countries and sectors is virtually impossible.

estimation procedure solves the following problem for dependent firms $j \in \{1, 2, \dots, I\}$:

$$\hat{\beta}_j = \arg \min_{\{\beta_{il}^j\}} \sum_{t=L+1}^T \left(x_{j,t} - \beta_{00}^j - \sum_{i=1}^I \sum_{l=1}^L \beta_{il}^j x_{i,t-l} \right)^2 + \rho_j \sum_{i=1}^I \sum_{l=1}^L w_{il}^j (\alpha_j |\beta_{il}^j| + (1 - \alpha_j) \beta_{il}^j{}^2), \quad (3)$$

with $\tilde{\beta}_{il}^j$ being the standard OLS coefficient estimates, $w_{il}^j \equiv \frac{1}{|\tilde{\beta}_{il}^j|}$, $\{\beta_{00}^j\}$ the constant terms, $\{\beta_{il}^j\}$ the set of elements of the coefficient matrix β_j , $\{x_{it}\}$ the standardized daily log returns, and L the maximum lag considered. The $\{\rho_j\}$ and α_j penalty weights are set to fit the data.²⁴ Each firm’s daily log return series is standardized independently over the full sample, so that they all have mean zero and a standard deviation of one to make the GIRF values comparable across firms and over time. We do not remove the average market returns because these include important information about broader shocks, sensitivity to which we incorporate into our inter-firm networks. With this approach we utilize the AEN for both the shrinkage and selection of the VAR lag order.

When $\alpha_j = 1$ only the LASSO penalty is included, and when $\alpha_j = 0$ only the ridge penalty is applied, otherwise the two are mixed in standard elastic-net. To benefit from both techniques, we set α_j such that each penalty term is on average equal. In particular, by including both penalty arguments strongly-correlated predictors are shrunk in or out of the model together unlike LASSO, where having only the absolute value term makes the optimization more likely to hit an multidimensional vertex where only the “best” explanatory variable will be kept among highly correlated variables, with the others driven towards zero. We run 10-fold cross-validation to determine the ρ_j value for each firm.

The AEN estimator has several drawbacks, including that there are not well-defined values to use for the penalization parameters, and that there is not asymptotic econometric theory under serial and cross-sectional dependence of regressors, and cross-sectionally correlated errors. There are three leading estimation procedures that we could have also implemented to estimate our VAR system. The first is a large Bayesian VAR; however, it is not clear what priors and distributions should be assumed.²⁵ The second option is a Global VAR in following Pesaran et al. (2004); however, the manner by which this approach would reduce the dimensionality of the VAR would be by estimating regressions of each firm’s equity return series on synthetic stock indexes of the other firm’s equity returns with ex-ante assumed weights that would enter the regressions, instead of each individual firm’s return series. This again requires making prior assumptions about the connections between the firms, as we would have to decide upon the weights for these synthetic stock indexes. Finally, another alternative to our estimation is the OCMT procedure of Chudik et al. (2016) for model selection, but this method was not developed for a large-scale VAR system.

²⁴See Online Appendix Section A1.1 for further details of the estimation procedure.

²⁵For more on large Bayesian VARs see Banbura et al. (2010) and Koop (2013).

Like AEN, each of these other leading estimation approaches to contend with the curse of dimensionality has its own drawbacks. The benefits of AEN — requiring few assumptions while dealing with the model shrinkage and selection — leave the AEN method as the best option for our purposes from what is currently available.²⁶

We deviate from the Demirer et al. (2015) procedure in how we derive our inter-firm network from the estimated VAR model. While they use the generalized forecast error variance decompositions based on their estimated VARs to measure the weights of the edges connecting firms, we instead use generalized impulse response functions (GIRFs).²⁷ There are two main reasons for using the GIRFs rather than the variance decompositions: the GIRFs do not necessarily add to near 100% at any forecast horizon as the forecast error variance decompositions do, which may artificially increase firms’ estimated connectedness; and the GIRFs have an economic interpretation that more closely quantifies the underlying relationships between the different firm returns that we are interested in. The precise definition of the GIRF is the effect of a shock to the return of firm i in the VAR at time t (ϵ_{it}) on the vector of firm returns at a horizon h (X_{t+h}), which is given by:

$$GIRF(h, i, \Omega_{t-1}) = \mathbb{E}(X_{t+h} | \epsilon_{it} = \sigma_i, \Omega_{t-1}) - \mathbb{E}(X_{t+h} | \Omega_{t-1}), \quad (4)$$

where Ω_{t-1} is the non-decreasing information set known at time $t - 1$, and σ_i is the standard deviation of the error term ϵ_{it} .²⁸ Once we have the estimated VAR coefficients and residuals it is straightforward to calculate the GIRFs between all of the firms. The $I \times I$ matrix of the GIRFs between all of the sample firms’ returns created by appending the $GIRF(h, i, \Omega_{t-1})$ vectors is the basis of our estimate of the firm network adjacency matrix. Because we care about the magnitude of the degree of connectivity — rather than the positive or negative signs of the relationships — we take the matrix of the absolute values of all of the paired GIRFs as the adjacency matrix representing the inter-firm network, A , with the value in each a_{ji} cell being the network weight from firm i to firm j .²⁹ Each firm, or node in our network, is listed along the rows and columns, with the edge weights between them being the entries in the matrix, generating a network that is both weighted and directed.

²⁶Also, as described below, we use the bilateral equity return correlations between each firm pair to create a model-free network matrix as a robustness check and find similar results.

²⁷In Online Appendix Figure A19 we provide a comparison of the network spring plots by industry and headquarters country using both our method and the scaled generalized forecast error decomposition based networks of Demirer et al. (2015). Both methods utilize the results of our baseline reduced VAR estimates and have qualitatively similar results as far as the importance of both industry and locality in network organization.

²⁸Specifically, we use the scaled generalized impulse response functions of Pesaran and Shin (1998).

²⁹We apply shrinkage on the approximating VAR model, but this does not mean we are shrinking the network itself.

We use generalized responses — rather than orthogonalized responses such as those found using Cholesky decompositions to capture shocks from individual firms — so their values are not estimating the cause and effect of a shock emanating from one firm to another.³⁰ The proper way to interpret the edges in our estimated networks is how many standard deviations one would expect other target firms’ daily equity returns to move conditional on a one standard deviation innovation for the source node’s equity return. In this way, the network edges flexibly capture co-movements and firm connectedness regardless of the source of the shocks, whether they be idiosyncratic firm shocks or common shocks.

When investigating the RyF nature of the inter-firm networks in Sections 5 to 8, we examine three sub-samples of the data: all firms that are publicly traded throughout 1991-2016; U.S. dollar listed firms that are publicly traded throughout 1981-2016; and 5-year rolling samples of firms that are publicly traded throughout each sub-period. The rolling samples help us deal with survivorship bias from the former two samples. Additionally, these rolling samples allow us to capture new firms entering the network — like Facebook and Google — that would not be in the longer networks, also potentially skewing our results. We estimate a series of rolling 5-year networks for each of these samples in order to incorporate how they might evolve over time into our test regressions. In the remainder of this section we assess the characteristics of the inter-firm networks, and what they can tell us about global firm connectedness before analyzing their robustness and fragility.

4.2 Features of the Global Firm Network

Our analysis of the global inter-firm network focuses on the 382 firms in our sample that were actively traded from 1991 through 2016 to provide continuity when analyzing the network over time. Our primary network is from one period ahead GIRFs derived from the AEN estimation procedure in the previous section run on the firms’ daily log equity returns with one lag. To help picture the global firm network, we borrow network visualization methods used in various other disciplines (e.g., biology, physics, computer science). Specifically, we generate a series of spring plots using the ForceAtlas2 method from Bastian et al. (2014). ForceAtlas2 is a force-directed layout algorithm to display network spatialization, transforming a network into a map where nodes with greater connectedness are closer together. At a high level, all of the nodes are repulsed from one another like charged particles, while edges attract their nodes, like springs — yielding the name for

³⁰Orthogonalization of the shocks is not appropriate for our purposes, since we are not trying to recover the specific structural shocks; rather, we are estimating how the system behaves for shocks similar to historical ones, taking into account the correlations among the innovations. Additionally, an often overlooked drawback of Cholesky decomposition based impulse response functions is that they do not necessarily identify the proper shocks, even if they do orthogonalize them and allocate them to particular variables.

this class of algorithms — with greater edge weights producing greater attraction. The final node positions provide a balanced state, helping to interpret the data without having to incorporate any other attributes of the nodes.

The maps of our estimated global firm networks are quite revealing about the nature of the interactions between the firms. To begin, Figure 2 shows the results of applying the ForceAtlas2 algorithm to the global firm network from 1991-2016, with the nodes colored based on four sets of firm characteristics. In the top left panel, the firms are colored based on their industry classification. From the quasi pie-chart appearance of the main cluster it is evident that industry plays a key role in connecting firms in the network. The firms on the edges of the network are generally near those in the same industry in the central cluster, with ICT firms at the top right, base materials and energy firms at the bottom left, and industrial diversified and consumer cyclical firms on the left. Further, the black center of the figure identifying a cluster of financial firms at the heart of the network indicates that finance has played a central role within it.

The next two panels show firms colored based on locality related attributes, namely the country of a firm’s headquarters and the currency in which the firm’s equity was issued. There are again clear, but very different, color patterns to the figures. Both of these plots have U.S. based firms at the center, with other advanced western markets closely integrated with them. There are a number of South African, Canadian, and U.K. firms on the fringes of that central cluster, with Australian, Japanese and other Asian firms stationed in orbits farther outside, on the periphery of the panels.³¹ Since all of the equity prices are in local currency, currency movements are not captured in the equity returns we use so they cannot be the cause of the regional organization of the plots. Further, in related work, Grant and Yung (2017) finds that firm size by several measures does not play an important part in our global inter-firm network structure nor in each firm’s degree of network connectedness.

The fourth panel shows the firms grouped using a spectral clustering algorithm on their daily equity return series. The spectral clustering algorithm produced four clusters, which appear to be associated with locality, perhaps reflecting local equity market betas. The black region appears to cover the U.S. and other closely related advanced western firms, the orange encompasses those firms that are on the periphery of that cluster, and the red ones are the Australian, Japanese and other Asian companies. There is one cluster of two firms in green near the middle of the diagram. These two firms are the U.S. agencies Fannie Mae and Freddie Mac. While they are close to several other U.S. and financial firms, their tight connections with one another and unique return profiles

³¹That Asian and Australian equities are on the periphery with generally low out weights but higher in ones matches findings from Diebold and Yilmaz (2015) studying the network between country level equity indices.

during the Global Financial Crisis and subsequent government intervention placed them outside the main knot of firms.

Figure 3 has the same global network spring plot with the nodes colored by four different categories based on the estimated network weight matrix. The top two plots show the firms colored based on quartiles of each node’s sum of in and out edge weights. Unsurprisingly, those nodes with the greatest weight out sums are at the center, smoothly transitioning to the lower weight out sums on the periphery. On the other hand, the firms with the lowest edge weight in sums are at the center of the second panel, but there is not as smooth an evolution as firms become farther from the center, suggesting that leading firms are in the main cluster while firms on the edges are there because they do the least to drive other firms. The bottom two plots color firms based on spectral and k-means clustering of the network adjacency matrix. Both a firm’s industry and its locality are important for which clusters it is in.³²

Exploring the estimated network further, Table 1 provides the top 25 firms in the global network by their edge weight out sums. There are a few aspects of this list that stand out. The first is that all of the firms are dollar based.³³ The second is the preponderance of financial firms at the top of the list, with over half of the top 25 being in that industry. In fact, even though the second firm, General Electric, is not a financial firm, up until it significantly changed its businesses in June 2016, its financial dealings were great enough that the Financial Stability Oversight Council designated it a nonbank systemically important financial institution. Taking all of these results into account, it is notable how the relatively straightforward estimation procedure we use based off of easily available equity price data seems to reflect so many attributes of the firms.

Given that firms’ importance appears to be closely connected with their currency and industry, Table 2 groups the firms at that level and has a similar ranking of the top 25 by their edge weight out sums. The prominent role of the U.S. markets is again displayed, with the top six entries — and eight of the top ten — all being based in dollars. Within the currencies, the finance sector is the most important sector for the USD, EUR, CAD, GBP, and CHF, reinforcing its central role in global connectedness. In fact, there is not a single utility, consumer non-cyclical, ICT or Energy sector present. This confirms the results of Grant (2016) showing finance as a primary international transmission channel. The Online Appendix contains similar analysis for a network of only USD issued firms, and in that network other sectors are better represented in the list of most influential

³²This refines the results of the long-standing literature in international finance studying asset prices, which finds that country factors are more important drivers of co-movement and volatility than industry factors, as established by the seminal work of Heston and Rouwenhorst (1994).

³³The most connected firms are actually Asian and Australian ones, but their connectedness comes from having large in degrees. Online Appendix Tables A7 and A8 provide the rankings by in degree.

firms, indicating the particularly global nature of finance.³⁴

The centrality of the U.S. and the financial sector over this time leads to the question of whether the network estimation is driven by the Global Financial Crisis. In order to answer this, and to understand the evolution of the global firm network more generally, Figures 4 and 5 examine the global firm network estimated in non-overlapping 5-year sub-periods during this timeframe. The former figure has the firms colored by industry, and the latter by currency. In both figures, the first panel has the longer term network spring plot we have been reviewing for reference, followed by the sequential 5-year plots. There is a distinct pattern of consolidation over the first four plots, suggesting increasing globalization throughout the sample period. The 1992-1996 panel shows a loose cloud of firms that consolidates over time, hitting a maximal level of concentration over the 2007-2011 crisis period, before slightly expanding outwards again in the following five years. There is evidence of a particularly large increase in agglomeration between 1997-2001 and 2002-2006, notably with European firms moving into the center, possibly reflecting the adoption of the Euro and the associated integration. The 2007-2011 crisis period particularly brought together North American and European firms, but many Asian firms were far out on the network periphery.

The firm positioning with regards to industry and location in these 5-year sub-periods appears generally similar to the longer term plots, and the top firms for these sub-periods — provided in the Online Appendix — also exhibit similar characteristics to those for the full sample network.³⁵ An exception is that over the 5-year period from 2007-2011 covering the Global Financial Crisis, the firms with the largest weight out sums were not financial companies. Rather, over this period the top firms were in the Energy and Base Materials industries, as well as supporting firms such as Caterpillar, Inc. That firms in these industries were the most central in these years is likely due to the tremendous run-up in commodity prices over this period. Combined with the fact that in the Online Appendix section on networks of only USD issued firms finance does not dominate the list of the most central firms, these results suggest that there is something uniquely global about financial industry connectedness and network centrality, and this is not solely an artifact of the Global Financial Crisis. Additionally, the overall characteristics we observed in the longer term network are not driven by the Global Financial Crisis.

The Online Appendix includes a number of figures that delve into the industry level dynamics of rolling 5-year sub-periods ending in 1995 through 2016 for these firms. The total industry in and out edge weight sums, as well as the net of the two, can be seen in Figures A9 and A10. These show the generally increased level of connectedness in the late 2000s, and the rise in Energy as a source

³⁴See Online Appendix Section A1.3 for the results of our analysis of other networks, including only the USD issued firms.

³⁵The list of top 25 firms by 5-year period can be found in Online Appendix Tables A9-A13.

of edge weights and Finance as a sink for in edge weights. The average network connectedness by year varies substantially, so we control for this by plotting the average ratio of the firms in each industry’s sum of edge weights in and out relative to means of those sums in that year in Figure A11. To help see the evolution of industry importance as edge sources and sinks, Figures A12 and A13 show the industry shares of in and out edge weights as a fraction of total edge weights in and out for each year. The edge weight out shares are nearly static over time, but there is more variation across the in weights — particularly for Finance around 1999 and 2010. Finally, Figure A14 plots all bilateral pairs of average edge weight out sums from industries in the columns to those listed in the rows, relative to the mean edge weights for each year. This figure reveals several interesting patterns, such as the increase in Energy weights to Finance during the Global Financial Crisis — matching the observations mentioned above — a peak of ICT to ICT edge weights around the 2000 Dot-Com Boom, and that the Energy sector consistently receives edge weights in from other industries that are roughly only half the annual average.

The results in this section are robust to several variations of the estimation procedure. First, Figures A1-A4 repeat the analysis for a forecast horizon of zero, producing similar spring plots. Figure A5 shows the spring plots for horizons from zero to five, showing that the connections dissipate as the horizon is increased, but that the same basic patterns emerge. Finally, Figure A6 shows the network re-estimated allowing up to five lags of the explanatory variables, the results of which are nearly indistinguishable from our main ones using a single lag.³⁶

5 Robust-yet-Fragile Behavior of Firm Networks

Many theoretical models have predicted that inter-firm networks may have both positive and negative implications for their members: greater network connectedness insures against and diversifies away the effects of smaller shocks, but these connections may also propagate contagion and make firms more susceptible to larger shocks. In this section, we identify and measure these two counter-acting dynamics — referred to as robustness and fragility, respectively — in the inter-firm networks that we estimated in Section 4, with the goal of better understanding the nature of firm networks and systemic risk.

³⁶We also experimented with lag orders and forecast horizons from zero to ten days each and found our results were robust to these choices.

5.1 Identifying Robust-yet-Fragile Behavior in Firm Networks

To study the structure of the network and how shocks are transmitted across firms, we relate a firm’s health to its network connectedness, the health of its neighboring firms, and of the system as a whole.³⁷ Specifically, we use a probit regression framework to measure how these other quantities predict the incidence of firm distress:

$$Pr(D_{it}) = \Phi \left(\underbrace{\beta + \hat{\phi} \sum_{j \neq i} w_{jit}}_{\text{Robustness: Diversification}} + \underbrace{\hat{\gamma} \sum_{j \neq i} w_{jit} D_{jt}}_{\text{Fragility: Direct Contagion}} + \underbrace{\hat{\lambda} N_t}_{\text{Fragility: Network Vulnerability}} + \underbrace{\hat{\omega} N_t \sum_{j \neq i} w_{jit}}_{\text{Robustness: Network Resistance}} + \underbrace{\hat{\theta} N_t \sum_{j \neq i} w_{jit} D_{jt}}_{\text{Fragility: Network Crisis Reinforced Contagion}} + \epsilon_{it} \right), \quad (5)$$

where D_{it} is a distressed state indicator for firm i (e.g., monthly equity log return in the bottom 10% of the overall sample distribution), w_{jit} is the network weight from firm j to i estimated over the preceding 5-year window (e.g., for January 2010 the network weights estimated from the daily equity log returns from January 2005-December 2009 are used), and N_t is an indicator of the aggregate network state (e.g., TED spread or VIX in the top 10% of the sample distribution).³⁸ We use rolling lagged 5-year network weights in order to capture how the inter-firm networks evolve over time, and to avoid simultaneity bias in the estimation. Additionally, we interpret this regression as a descriptive measurement equation to empirically decompose the relationships between the various firm and network attributes, not to explicitly determine causality, which could be subject to issues of endogeneity.

The model in Equation (5) can be used to answer several questions about the nature of the estimated inter-firm networks. There are five aspects of RyF behavior that we investigate, broadly divided into two categories depending on whether they are expected to be beneficial or detrimental to firms. The labels for each have been included under the corresponding terms in the equation. Throughout this discussion we focus on the dependent variable in the model being a distressed state indicator where positive values reflect poor firm health; however, when we instead examine firm

³⁷The model presented in Section 2.1 is a subset of this probit latent linear equation based model. To place the previous model within this framework the constant should be 16, the diversification term coefficient should be -1, and the direct contagion term coefficient should be 1 for firms that are owed money by others in the network, with a constant of -4 added if not a creditor of any other firms. All remaining coefficients should be zero, and the network weights should be calculated based on the size of borrowing as described in the text.

³⁸The variables with the hats are the coefficient estimates from the latent equation. When presenting the regression results, we instead focus on the more appropriate marginal effects for the probit regressions, which will be listed without the hats. In the case of standard linear regressions, the actual coefficients will be what is referred to by the non-hat variables.

returns, revenue, return on equity and profits, the expected signs under the null hypotheses should accordingly be reversed, as in those cases higher levels for the dependent variables are a positive rather than a negative outcome. Additionally, we assume above that N_t is a distress indicator, but in alternate formulations we include other measures of network level economic conditions where the signs on the related coefficients should be negated.

The first form of RyF behavior that we investigate is the diversification channel of Elliott et al. (2014).³⁹ This channel in their model is initially expected to be a negative, as small connections transmit crises; however, as connectedness increases further it provides better insurance and diversification benefits that make this channel a net positive. We capture this in the first term of the regression, answering the question of whether greater firm connectedness is associated with better firm outcomes. While the theory does not provide a definitive prediction for the sign of the coefficient on this term in all cases, we conjecture that the inter-firm network is integrated enough to provide the insurance and diversification benefits and that $\phi < 0$.

The second term in the regression captures the next facet of RyF behavior that we are interested in, the degree of pass-through to a firm from its neighbors, which we label direct contagion. This is what is commonly measured in both the banking and cross-country crisis contagion literatures.⁴⁰ We expect that — in following with the consensuses of these literatures — $\gamma > 0$ indicating direct contagion.

The final three channels we look at incorporate network wide disturbances, N_t , rather than only the direct neighbor relationships in the previous two terms. We use several gauges of network level disturbances, including the TED spread, VIX and within-period averages of the D_{jt} measures across all firms, giving the share of firms currently in a distressed state. The degree to which firm outcomes are related with these is captured by λ , and we expect this coefficient to be positive, with firms exhibiting vulnerability to the overall health of the network. On the other hand, we expect that greater connectedness may help to buffer this effect, as in the right hand portions of the plots in Figure 1. We call this dynamic network resistance and predict that $\omega < 0$, since we expect the inter-firm networks to be highly connected.

The fifth and final aspect of RyF network behavior that we examine is whether network wide disturbances and direct contagion from a firm’s neighbors act to amplify one another when they coincide. We refer to this situation as network crisis reinforced contagion and expect that each of these would make a firm more sensitive to the other. This would, for example, be similar to the

³⁹Diversification here refers to more network connectedness diversifying away individual shocks, not the distribution of the network edges themselves.

⁴⁰For examples of direct bank contagion see Elliott et al. (2014), Acemoglu et al. (2015b), Gai and Kapadia (2010), and Gai et al. (2011), and for cross-country contagion see Glick and Rose (1999) and Hernandez and Valdes (2001).

network dynamics of Caccioli et al. (2015) investigating both direct counterparty failure risk and contagion via overlapping portfolio exposures in a banking network. This paper found that neither channel is significant on its own, but that when both channels are active at once bankruptcies are much more common and have large systemic effects. If reinforced contagion is an extant phenomenon in the global firm network, then it would be reflected in a positive value of θ .

5.2 Regression Results

In this section, we run various permutations of Equation (5) to delve into the RyF nature of the firm networks. In particular, we measure firm health using equity returns, 5-year CDS spreads, revenue, return on equity, and profits. Distressed firm states are defined as being in the worst 10% tail of the overall sample distribution for the corresponding underlying health measures. Also, we run standard linear regressions with the continuous versions of these measures as the dependent variables in order to see if — without specifically isolating distressed states — there is evidence that network structure is related to firm performance.

Our baseline estimates of Equation (5) testing the RyF network structure of our global firm networks using monthly log equity return distress firm health indicators can be found in Table 3. The table has two panels: the top panel has the estimated marginal effects of each explanatory variable (not estimates of coefficients in the latent linear model); and the bottom panel has the estimated changes in the probability of a distressed state from a one standard deviation increase in each explanatory variable. To help interpret the values, the right hand columns provide the signs one would expect if the associated type of RyF behavior exists. The explanatory variables are introduced across the regressions. The first two regressions have only the average sum of weights into a firm and direct contagion variables, respectively. If one were to examine these individually then it would seem that greater network connectedness was a bad thing; however, as column (3) shows, more connectedness actually has a robust, diversifying effect, and it is the direct contagion that occurs through these same connections that has a positive impact on the probability of a firm being in distress.⁴¹ Both of these results are in line with there being RyF behavior. The next column then adds the network crisis state variable — measured as having a monthly average TED spread in the top 10% of its sample distribution — and the final column adds interactions of the network crisis state with the former two variables. All of the estimated effects are in the hypothesized directions

⁴¹Hale et al. (2016) provide evidence of similar direct contagion effects. The authors construct a yearly global network of interbank exposures from 1997 to 2012 of more than 6,000 banks using long-term interbank loan data to study the transmission of financial sector shocks across borders, finding that direct exposures to crisis countries squeeze banks' profit margins, thereby reducing their returns.

and, with the exception of the network crisis reinforced contagion in column (5), are statistically significant.⁴² Additionally, there is substantial agreement between the estimated effects in the latter three regressions.

Since it is not necessarily clear what a one unit increase in the explanatory variables means, we focus on the standardized effects in the bottom panel. For the network crisis state the effect is of an increase from zero to one — indicating the difference between being in a crisis versus not — and the interactions of this increase multiplied by the standard deviation of the interacted sum term. These results suggest that a one standard deviation increase in network connectedness reduces the probability of a firm being in distress by about 3.7%, while an increase in direct contagion of one standard deviation increases the probability by 5.9%. From the estimates in columns (3)-(5), it is clear that the direct contagion effect is much larger than the diversification one. These coefficients indicate that if a neighbor is in crisis, then it would take other non-crisis neighbors with in edge weights roughly nine times as great to counteract the impact. Being in a network crisis state increases the probability of being in distress by 8.1%; however, when controlling for network resistance one can see that greater connectedness can significantly offset this effect. These effects are quite large given that the unconditional probability of being in distress is 10%, with both direct contagion and network vulnerability being over half of this value.⁴³

Table 4 presents the estimates of similar linear regressions where rather than a distressed indicator dependent variable, the dependent variable is the continuous monthly log equity return. The two panels are similar to the previous table, but now the bottom panel has standardized coefficients, which are more appropriate to understanding continuous linear models. The estimation produces statistically significant results in line with what one would expect given the RyF theory we outlined above in all cases, including the network crisis reinforced contagion.

Focusing on the values in the lower panel, a one standard deviation increase in network connectedness corresponds with increased expected monthly returns of about 12% of a standard deviation, or roughly 1.2%. This result is interesting because it suggests that being connected in the network is positively related with increased firm equity returns, going beyond just reducing the probability of a firm distress event. The direct contagion effect is nearly three times as large as the diversification one, with a standardized coefficient of -31%, correlating with a 3.1% equity price decline. Finally, there again is a negative expected effect of a network crisis state, with evidence that network

⁴²All standard errors are clustered at the firm level. Additionally, we did not find evidence of significant serial correlation in the estimation residuals and rerunning the regressions allowing for AR(1)-AR(10) correlation in the error terms does not alter our findings.

⁴³Plots of the average sums of weights and direct contagion variables used as the first two dependent variables by various firm characteristics and overall can be found in Online Appendix Figures A7-A8.

connectedness can aid, and direct contagion can hinder, a firm’s response to such a state.

It could be that firm, industry or location effects might in part be responsible for these results. To control for these possible effects, Online Appendix Table A2 presents the results of running this estimation in first differences, effectively controlling for various fixed effects, and the results are not materially different.⁴⁴ Additionally, we estimate both the discrete and continuous versions of the model with industry interactions on all of the coefficients and do not find any substantive differences across industries. Also, the coefficients on the industry interactions are not stationary over time, further challenging the idea that there are underlying structural differences across industries.

Beyond the effects on the equity returns, we wish to know whether there is “real” evidence of RyF firm network behavior, so we examine similar continuous regressions with the dependent variables and neighboring firm distress indicators based on the quarterly change in profit, as measured by EBITDA, the 5-year CDS spread, the quarterly change in revenue, and the annual return on equity (RoE). We include the monthly average 5-year CDS spread and RoE in levels and first differences, since these variables likely incorporate firm specific level effects that should be removed. The results of these regressions are presented in Table 5, where it can be seen that even across these other firm health measures and differing time frequencies the RyF behavior still upholds. Not surprisingly, the CDS spread is strongly influenced by direct contagion, with a one standard deviation increase in direct contagion associated with a 46.0% standardized CDS increase. The small diversification and large direct contagion responses are intuitive, as CDS spreads are based on discrete underlying default choices which are most pertinent in very weak economic environments.

Profit and revenue growth have standardized diversification coefficients of just under 10%. In practical terms, a one standard deviation higher network in weight sum corresponds with improved firm health via increased growth of 1.2% for monthly equity prices, 4.0% for quarterly profits, 1.7% for quarterly revenue, and 1.7% for annual return on equity. Interestingly, profit and RoE growth demonstrate significantly larger effects from a network crisis than does revenue growth, -23.0% and -17.5% to -6.6%. For context, the standardized coefficients are reported in the bottom panel of Table 5. These results suggest that profitability is more susceptible to network crises than is income and are consistent with the fact that the two profit measures are significantly more volatile than revenue. All three dependent variables, however, are consistent in showing a positive relationship with connectedness and a negative one with the fragility terms. These continuous regressions suggest not just level improvements for the real profit and revenue quantities, but growth rate increases — especially meaningful effects.

⁴⁴In untabulated results, we also tried controlling for time fixed effects, obtaining similar results with regards to diversification and direct contagion.

One might expect that growth in these real firm health indicators would correspond to greater inter-firm connectedness, with a fast growing firm being a healthy one that drives other firms, especially as it comprises a larger and larger share of its market; however, our regressions are finding a more nuanced relationship in the reverse direction. Our measures of connectedness included in the regressions focus on how firms are influenced by others in the network via the sums of edge weights in from other firms only. The correlation between these sums of edge weights into firms used in the regressions and the corresponding sums of edge weights emanating from them is -0.183. Therefore, these results indicate that firms whose equity prices are more liable to move along with those of other firms are likely to have higher profit, revenue and RoE growth, as well as lower CDS spreads, a novel finding in the literature on the global firm network. This is an important distinction when thinking about the relationships between global integration and firm performance, indicating significant value to being more integrated. Additionally, Grant and Yung (2017) show that the sum of a firm's in weights — as well as the out weight sum — is not correlated with firm size in the estimated networks by numerous measures, including total equity, sales, profits and number of employees, so the beneficial real effects are not simply because the biggest firms necessarily have the largest weights.

Together, these results provide strong evidence of RyF behavior in the global firm network and benefits to global integration. The increased inter-firm network connectedness over time that we find signs of should then be a net positive for global firms — outside of neighboring firm distress events — via the diversification channel. Further, the direct contagion channel appears to overpower the diversification channel, and severe credit tightness substantially weakens firm health. These results suggest that globalization has had positive consequences for firm health, but that policymakers should be wary of greater connectedness allowing for increased contagion. This wariness is particularly critical when there are signs of network wide distress, given the evidence we find for network crisis reinforced contagion in Table 4, with amplification between network crisis states and direct contagion.

We examine two other sets of firm networks in the Online Appendix. The first is the set of USD equities from 1981-2016, with the results provided in Tables A18 and A19. The other is comprised of rolling 5-year windows of the top global firms that are continuously traded within each window. These results are available in Tables A22 and A23. The results from these other samples are markedly similar to our main ones, providing evidence for RyF behavior, and rejecting the possibility that sample selection drove our results.

We then check whether our results are robust to network crisis measures other than our indicator for a credit crisis indicated by a high TED spread. Tables 6 and 7 provide the results when using

the following crisis indicators: the continuous monthly average TED spread level; the log of this level, given the non-linear relationship that it tends to have versus the state of the macroeconomy; a crisis indicator for the VIX being in the top 10% of its monthly distribution; the monthly average VIX; a crisis indicator for the SP500 return being in the bottom 10% of its monthly distribution; the monthly SP500 return; the fraction of firms in a distressed state in a given month; and an indicator for months with the top 10% of the fraction of firms in a distressed state.⁴⁵ These results show that our main findings are robust to the choice of network crisis state measure, with nearly every estimate matching the expected sign associated with RyF behavior and being statistically significant.⁴⁶

We used the GIRFs for a one day horizon throughout the above analysis. To check the robustness of our results with respect to the choice of network, Tables 8 and 9 show the results of running our main regressions for forecast horizons from 0-5 days. The tables also include the results of using the rolling 5-year bilateral equity return correlations between each firm pair as the network edges. All of our main results still obtain, including the proper signs with statistical significance, the direct contagion channel being roughly an order of magnitude larger than the diversification one, and there being network vulnerability.

This is particularly interesting because the bilateral equity return correlation and zero day horizon GIRF networks are positively correlated at a statistically significant level with one another using the quadratic assignment network comparison procedure; both the overall and all of the 5-year sub-period networks. The network correlations are calculated using position-by-position correlations of the entries in the two network adjacency matrices. The quadratic assignment procedure then effectively bootstraps the distribution of correlations one would expect for two similar network structures by repeatedly reordering the rows and columns of one matrix and recalculating the network correlations. This has the benefit of incorporating the likely correlations between networks of similar sparsity, distribution of network edge weights, etc.⁴⁷ Likewise, the 1-5 day horizon GIRF networks are positively correlated at statistically significant levels. However, these two sets of networks have statistically significant and negative correlations with one another. This suggests a great deal of robustness of our results to the measure of the network structure.

Finally, we seek to determine whether these RyF relationships vary over time. Specifically, were there significantly different relationships over the Great Moderation and is fragility a unique

⁴⁵Online Appendix Tables A3-A6 provide results of rerunning the estimation using continuous TED spread measures.

⁴⁶The exceptions are the diversification term for the fraction of firms in crisis in Tables 6 and 7, and the network vulnerability term for the log TED spread in Table 7.

⁴⁷See Online Appendix Table A14 for the network correlation results.

feature of the Global Financial Crisis period, driving our main results? In the Online Appendix we reran the estimation for our main, fixed sample of global firms as well as the USD one, breaking the sample into ten year periods.⁴⁸ The overall results are similar to those obtained in our longer, full samples suggesting that the results are not solely driven by the events of the Global Financial Crisis.

6 Fall 2008 Global Contagion

During the latter half of 2007 and the beginning of 2008 the effects of the U.S. Sub-Prime Mortgage Crisis were starting to be felt in the U.S., with heightened credit spreads, equity market declines, and, most notably, the March 2008 collapse of Bear Stearns, with its ultimately being taken over by J.P. Morgan in a government orchestrated deal. It was during the fall of 2008, however, that the crisis climaxed. Within a week in September 2008, Fannie Mae and Freddie Mac were taken under government conservatorship, Lehman Brothers announced its bankruptcy, and the government pressed a buyout of Merrill Lynch by Bank of America, escalating the U.S. Sub-Prime Mortgage Crisis into a global financial crisis. The fall of 2008 saw equity markets around the world plummet, credit spreads surge, numerous firms require government support, and both the VIX and VSTOXX reach their all-time highs.

Given the importance of this period as the height of the Global Financial Crisis, we take the data for the fall of 2008 to our model to see how well it can account for the facts. As the basis for our model, we take the estimated latent linear model from a quarterly probit regression including our five RyF terms with the average VIX as the network state, as well as the average TED spread as its own separate additional term.⁴⁹ We include the VIX as a measure of the stress or crisis level of the global economy, leaving the TED spread as a measure of global credit spreads. If the value of the latent linear model is positive for a firm, that is taken to indicate a distressed state for it and vice-versa. The sample that we focus on is the global sample of top firms continuously traded throughout the previous five years, 2003-2007.⁵⁰ We use this sample instead of our long continuously traded firm sample to include several now defunct firms which were central to the events of late 2008, and to have the fullest picture possible of the firm network over this period. For the network weights, we use the sub-sample's pre-crisis network connections entering 2008, the

⁴⁸See Tables A15, A16, A20 and A21.

⁴⁹We also ran the model with an interaction term on the TED spread for USD issued equities, and found that the coefficient on that term was small and not statistically significant. Further, the simulation results were little changed when including it.

⁵⁰This sample includes 756 firms across 40 countries, with equities issued in 25 currencies.

estimation of which ran from 2003-2007 and does not include the actual crisis contagion period we are attempting to match.

Using this model, we start with 17 firms that are initially assumed to be in a distressed state, and the VIX and TED spread levels from this period in late 2008.⁵¹ These firms are selected from the worst performing firms over this period — with 2008 equity returns from -49.81% to -99.96% — that also went bankrupt or needed significant government intervention and were recognizably central to the crisis, such as Lehman Brothers, Washington Mutual, and Bear Stearns. For details on the experiences of these firms during this period see Online Appendix Table A24. We use an interacting agent modeling approach to simulate the contagion, reevaluating the latent linear model for every firm to determine whether it is in a distressed state or not, and iterating until a steady state is reached. Note that the only firms we assume to be in a distressed state are the original 17 firms, and it is possible for other firms to go into and out of a distressed state dependent on the network state, their network connections and the status of their neighbors.

Figure 6 contains a series of spring plots based on the global firm network from 2003-2007 that we use, which present features of the actual network and illustrate the modeled contagion. The first two panels show the industry and locality of each firm, which demonstrate distributions that are extremely similar to what we saw earlier for the continuously traded global firm sample. For example, both industry and region are important for firm connectedness, the Asian and Australian firms are on the periphery, and finance is at the center of the network. The next panel shows which firms actually experienced equity return distress over late 2008, with those affected in red. The substantial scope of the contagion across regions and industries in the fall of 2008 is evident in the plot. The final panel in the top row shows the positions of the 17 initially distressed firms in red. These are predominantly financial firms and have USD issued equities, so it is not surprising that they are all located relatively near the center of the network plot. The second row of panels then shows the modeled contagion spread over each iteration through convergence in the fourth one.

While not an exact match for the actual firm distressed states, at an aggregate level the simulation's share of 88.6% of firms in distress is quite similar to the actual level of 88.0%. This is a considerable degree of contagion given that the initially distressed firms represented only 2.2% of all firms in the sample. Delving deeper, 78.2% of firms' actual states matched their final simulated ones, with a higher rate of core firms and a lower rate of periphery firms being in distress in the model than in the data. In this way, greater distance was even less of a buffer in actual contagion

⁵¹By selecting these firms we are not stating that we can identify them as the source of the shocks causing the Global Financial Crisis, but rather exploring how a shock that put them all into distress would be expected to spread across the network.

than in the model simulations. In the end, the contagion swept across virtually all but the most remote firms. Industry does not seem to be very important for contagion beyond what was captured in the network weights, though locality does seem to matter with fewer Asian and Australian firms affected. It is promising that our model, though relatively sparse, is able to match the high level aspects of the crisis.

As the spring plots show, the contagion quickly spread from the initially distressed firms in the first iteration, particularly to other USD and financial firms. The second iteration saw a significant increase in the spread to European firms with the U.S. market having hit a critical mass. At this point, the large number of distressed firms abroad then led to an echo effect, where the contagion spread more widely across U.S. firms that were connected to foreign ones.

The echo effect should be an important consideration for policymakers in large economies when facing potential crises. First, given the evidence we have found for the importance of locality in firm connectedness, such economies are likely to contain most of the firms that have important connections to others in their economy, passing on direct contagion as well as, by definition, leaving a smaller rest of the world to provide connections that offer robustness and buffer against firm distress. Second, a large economy is likely to have a strong influence on other economies, making it more likely to spillover to other firms abroad. Combined, these two influences make it so that when a large economy is in a crisis it is liable to be severe, as there are likely to be returning echoes from contagion sent abroad that the economy is not well buffered against, leading issues to snowball at home and exacerbating the already tenuous economic situation. In this manner, crises that occur within major economies may become more globally contagious going forward, conditional on them becoming widespread internally, with increased global integration creating the circumstances for further situations like the U.S. Sub-Prime Mortgage and Eurozone Debt Crises to develop into having a global impact instead of being smaller, regional issues. How such episodes develop depends on the RyF nature of the global firm network, highlighting the importance of understanding these properties.

7 Simulating Sectoral Shocks

Table 2 reveals that there are a handful of currency-industry pairs with outsized network impact. For example, the number one pair, USD-finance, has a weight out sum that is more than double the pair five slots down the list (USD-Energy), which in turn has a weight out sum that is more than double the pair five slots below it (EUR-consumer non-cyclical). The work of Acemoglu et al. (2012) and Acemoglu et al. (2015a) indicates that asymmetries in the scope of network members

can potentially lead to network instability, so in this section we simulate how a shock to each of the currency-industry pairs propagates through the global firm network to study the consequences of having such a high concentration of influence among the top sectors.⁵²

We follow the simulation approach from Section 6, using the estimated probit latent linear model from Column (5) of Online Appendix Table A5 as the foundation for simulating sectoral shocks with distress for all firms in each currency-industry pair. The simulations use a network of our long, continuous global firm sample from 1991-2016, with lag and GIRF orders of one. For the network crisis state, we use the value of the average monthly TED spread in the top percentile of our sample distribution to gauge the worst case reach of contagion from each sectoral shock.

Figure 7 plots the simulation results, with the share of firms in each currency-industry pair along the x-axis, and the ultimate share of firms in distress along the y-axis. The color of the marker indicates the currency region, with small share currencies that are nearby one-another given the same color, and the shape of the marker indicates the industry. For reference, the 45-degree line is also plotted in gray, with the distance above that line representing the degree of contagion from the sector experiencing the shock.

The plot shows that there is a generally positive relationship between the share of firms in a sector and the final share of firms in distress, but the relationship is far from linear. The large gap in the middle of the plot — with contagion to firms outside of a sector being less than 12% or greater than 92% — suggests a critical mass type, all-or-nothing relationship, and indicates that there are only a handful of sectors where a shock would be expected to have broad ramifications. Finance tends to have the greatest contagious effect for each major currency (e.g., for the USD, CAD, GBP, EUR), but that is not the case for Japan where the industrial-diversified and consumer-cyclical industries have more influence. Unsurprisingly, the USD industries are the largest by the share of firms in them, encompassing the top six observations, with many of the most contagious sectors. Also, USD-finance is again at the top of the influence list — in a tie with USD-consumer non-cyclical, USD-ICT, USD-consumer cyclical, and EUR-finance — hinting at the significance of the echo effect for the U.S. and Eurozone.

The EUR-finance sector is an especially interesting case, as a shock to this sector results in far broader contagion than a shock to either the USD-Energy or USD-IndDiv sectors, even though those two sectors have out weight sums that are 1.7 and 2.3 times that of the EUR-finance sector. In separate results examining the pre-Eurozone period, no one currency-finance sector shock for the future member countries is simulated to have had nearly the same contagious effect as seen

⁵²Similar results are found when grouping at the country-industry level, with many more small firm share pairs with little contagion near the origin.

here for the EUR-finance sector. Although further work is needed to assess the impact of European integration on firm-level spillovers, our simulation results suggest that the propagation of shocks was intensified through higher integration during this period.⁵³

8 Network Contagion Simulations

As mentioned above, the degree of integration in the global firm network varied over the past three decades, with a generally increasing trend, especially at the upper tail of the connectivity distribution.⁵⁴ In this section our goal is to provide context for what that increase means and study its implications for potential crisis contagion. Has increased firm connectedness been a stabilizing force producing a more robust global economy, or has the distribution been such that it has engendered greater fragility?

The analysis is similar to that performed in the previous two sections in that we use an estimated probit latent linear model as the basis for running a series of interacting agent model simulations. Specifically, we use the estimates of the latent linear model associated with Column (5) in Table 3, including all five of our RyF terms with the network crisis state measured as having an average monthly TED spread in the top decile of the sample distribution.⁵⁵

To analyze the global firm network over time, we take the same estimated network weights that the 5-year sub-period spring plots in Figures 4 and 5 are based on. We then examine how these 5-year networks respond to crises that originate in one through all 382 firms in the continuously traded global firm sample. Across all of these cases the coefficients from the latent linear model are the same ones estimated over the full sample period, with only the network weights changing. We then randomly select the crisis source firms, with the process repeated 200 times to get simulated contagion distributions for each share of initially distressed firms. This is performed with and without a network level crisis.

The results of these model simulations are presented in Figures 8 and 9. The first column of Figure 8 plots the average share of firms that end up in a distressed state across the 200 iterations

⁵³The classification and acceptance of Greek government debt as top tier collateral by banks in other Eurozone countries has been suggested to have been an especially grievous aspect of financial integration within the single market. See Farhi and Tirole (2016) for a discussion of this issue.

⁵⁴See Online Appendix Figures A7 and A17 for time series of the average global and USD continuously traded firm samples' connectedness showing the increased average network connectedness, and the large increases at the upper tails of the distributions.

⁵⁵When the network crisis state is zero, these estimates are extremely similar to the estimated probit latent linear model when only the first two diversification and direct contagion terms are included. Therefore, the cases with $N_t = 0$ can be viewed as the results from a more concise model without the three network crisis state terms.

for each set of network edge weights and share of firms that are crisis sources. The dashed 45-degree lines denote the share of source firms that are initially distressed, and the height of each average line above that is the share of firms experiencing contagion.⁵⁶ The second column helps put the contagion shares into context, since as the share of firms initially in crisis increases, there are fewer firms remaining to potentially experience contagion. Rather than the simple difference between the simulation averages and the 45-degree lines, the second column shows those differences divided by the share of firms that are not assumed to initially be in crisis. In other words, these plots show the fraction of firms that could potentially experience contagion that actually do. Figure 9 shows more detailed information for each case, with the same simulation averages, the minimum and maximum shares of firms in distress, and the contagion shares. For comparison, these plots also include simulation results for the full 1991-2016 global network.

The top two rows in Figure 9 show the case in which there is not a network level crisis as measured by a high TED spread. Starting at the beginning of the sample period, the 1992-1996 based network experiences almost no contagion. The average, minimum and maximum contagion lines are right on the 45-degree line, with only a handful of non-source firms experiencing distress as the share of source firms approaches one. This result reinforces the idea that at the start of the period there was a significantly lower level of global network integration.

At the other extreme of the contagion plots is the 2007-2011 period covering the Global Financial Crisis. In this case, contagion begins to occur with an initial share of crisis source firms as low as 16%, with 80% of firms affected when the share of crisis source firms is only 34%. At this level the contagion spread fraction is already at 70%, and plateaus around 80% starting at a 54% initial share of crisis source firms. Figure 9 sheds some light on why the simulations produce so much more contagion over this period than others. Examining the maximum and minimum curves for the 2007-2011 period plots, it is clear that they are far wider than for the other periods, especially as the contagion initially begins to occur. The greater variety of simulation outcomes for this period suggests that during this time, more than any other, there was a set of firms with far more outward connectedness that were central to the particular outcomes at lower initial crisis share levels.⁵⁷ At higher shares of crisis source firms more of these central firms are likely to be in crisis, leaving the network to be quickly overcome with contagion.⁵⁸ These results are not surprising given the

⁵⁶See Online Appendix Figure A18 for plots of the contagion shares, the average curves minus the 45-degree line.

⁵⁷Tables A9 through A13 of the top globally connected firms and the sum of their weights out to other firms over each 5-year sub-period support this conclusion.

⁵⁸The work of Acemoglu et al. (2012), Acemoglu et al. (2017), and Gai et al. (2011) studying different subsets of the firm network come to similar conclusions about crisis risk, with greater asymmetry and concentration in specific firms in a network amplifying the network fragility.

perceived instability at this time, and that our estimated networks reflect all manners in which firms were integrated over this period.

The other three sub-periods were relatively more tame than the two discussed above. The late 1990s period covering the Dot-Com Boom and various emerging market crises was the next most contagious period, but in the simulations without a network level crisis there was not decidedly more contagion than in the periods immediately around 2007-2011. Conversely, the late 1990s period also exhibited robustness, with contagion occurring at a higher initial crisis share (28%) than either the 2002-2006 (23%) or 2012-2016 (20%) sub-periods. Additionally, the degrees of contagion in the sub-periods on either side of 2007-2011 were nearly identical, with maximum contagion fractions under 45%, well below the 80% plateau for the 2007-2011 sub-period discussed above.

In our regression estimates we found evidence that being in a network crisis state can have serious consequences, so we now focus on contagion simulations of the same 5-year sub-periods assuming that there is such a state with high credit spreads. These results are shown in the bottom plots of Figures 8 and 9. The most noticeable change is the accelerated contagion for the 2007-2011 network, with 90% of firms on average in distress starting with an initial firm crisis share of only 30%. The contagion fraction then plateaus at 95% when the initial firm crisis share is only 32%. That is a considerable level of contagion, indicating an extremely fragile network.

The 1997-2001 contagion fraction is the second highest, and in this case it is markedly higher than those for the 2002-2006 and 2012-2016 periods. An interesting similarity between the plots with and without a crisis state is the nearly identical share of initial crisis firms at which contagion begins for all of the sub-periods. Additionally, with a network crisis the 1992-1996 period approaches a contagion fraction near the 55-60% for the 2002-2006 and 2012-2016 periods; however, the early 1990s period approaches that level slowly in a shallow convex manner, while the latter two approach it far quicker with concave contagion fraction curves. Altogether, these behaviors suggest significant network crisis reinforced contagion when compared to the upper plots without a network crisis.

Overall, these simulation results reflect increased global firm integration, confirming what we observed looking at the 5-year sub-period spring plots. Generally, the simulations imply greater direct contagion effects, but greater integration also materializes in some cases as more robustness for moderate initial firm crisis shares. These results can guide future macroeconomic modeling choices. Specifically, firm level microeconomic dynamics may be unimportant during normal times due to the robustness of the inter-firm network; however, firm heterogeneity has meaningful aggregate consequences during crises owing to network fragility. Further, our simulations suggest that over time global firm network fragility has increased during crisis periods, underscoring the importance of capturing these effects in a model. Additionally, a network crisis state is found to be a warning

sign, steepening the contagion fraction curves and increasing the maximum contagion shares as the network crisis reinforced contagion effect leads to rapid, non-linear contagion effects.

These results should guide investors and policymakers to be particularly alert for potential contagion when there are signs of system wide stress and high firm connectedness, especially when the firm connectedness distribution’s upper tail fattens. However, this is not the whole story. In particular, Section 5.2 showed more connectedness corresponds with a reduced probability of distress and improved stock returns, CDS spreads, profit, RoE, and revenue growth. Therefore, while contagion can spread extremely rapidly and widely with greater global firm integration — and high levels of connectedness can be a red flag — policymakers should not overreact and cutoff international connections given the benefits of a more integrated network, forming the double-edged sword of greater global integration.

9 Conclusion

In this paper, we estimate inter-firm networks across countries and industries by using readily available data from financial markets and the latest developments in machine learning to resolve the curse of dimensionality associated with large panel VAR models. Most importantly, the global inter-firm network provides essential information on how firms are connected with one another and how these linkages have changed over time. This allows us to study how network connections may transmit or mitigate idiosyncratic — as well as system-wide — shocks, consistent with the idea of integration being a double-edge sword. On the one hand, we find that integration makes the system more robust through diversification and network resistance, increasing companies’ growth prospects and reducing the incidence of firm distress events. On the other hand, more connected firms become more fragile through higher exposure to direct contagion from a distressed neighbor, vulnerability to system-wide shocks, and network crisis reinforced contagion.

We conclude our study with a series of counterfactual experiments and simulations that reveal non-trivial echo effects that aggravate domestic conditions during a crisis through international connections that are not well buffered. Understanding such RyF qualities of the network is one way that our modeling approach can inform better policy decisions when future crises threaten.

By shedding light on the nature of the firm network, we help identify the classes of models appropriate to account for global firm dynamics, confirming the findings of Acemoglu et al. (2012) that “classic” macroeconomic models that ignore firm heterogeneity, assuming differences will average out, are inadequate to capture real-world developments.⁵⁹ Firm level heterogeneity is especially

⁵⁹See, for example, Acemoglu et al. (2015a) on how modeling assumptions interact with the RyF behavior of

significant during crisis periods and when there is greater globalization and firm network integration, which our work suggests has increased over the past quarter century. Finally, given that finance is at the center of the global inter-firm network, open economy models should be designed to incorporate international connections via this pivotal sector.

Although our work identifies how firms' responses to shocks are related, we do not determine the source or nature of the initial shocks. We believe the identification of such shocks is an important avenue that should be explored as the literature continues to expand to understand systemic risk and contagion through global networks. Moreover, further investigating the connections between inter-firm linkages and the international co-movement of business cycles can enlighten the discussion of how macroeconomic shocks are absorbed through the system, and their repercussions in the global economy. Finally, another extension of our analysis would be a structural model that can build from our empirical findings to formally model inter-firm connections in a general equilibrium context. We believe these are only some of the directions this line of work can take, as we expand our knowledge of the benefits and drawbacks of higher global integration, and its role in the propagation of shocks.

References

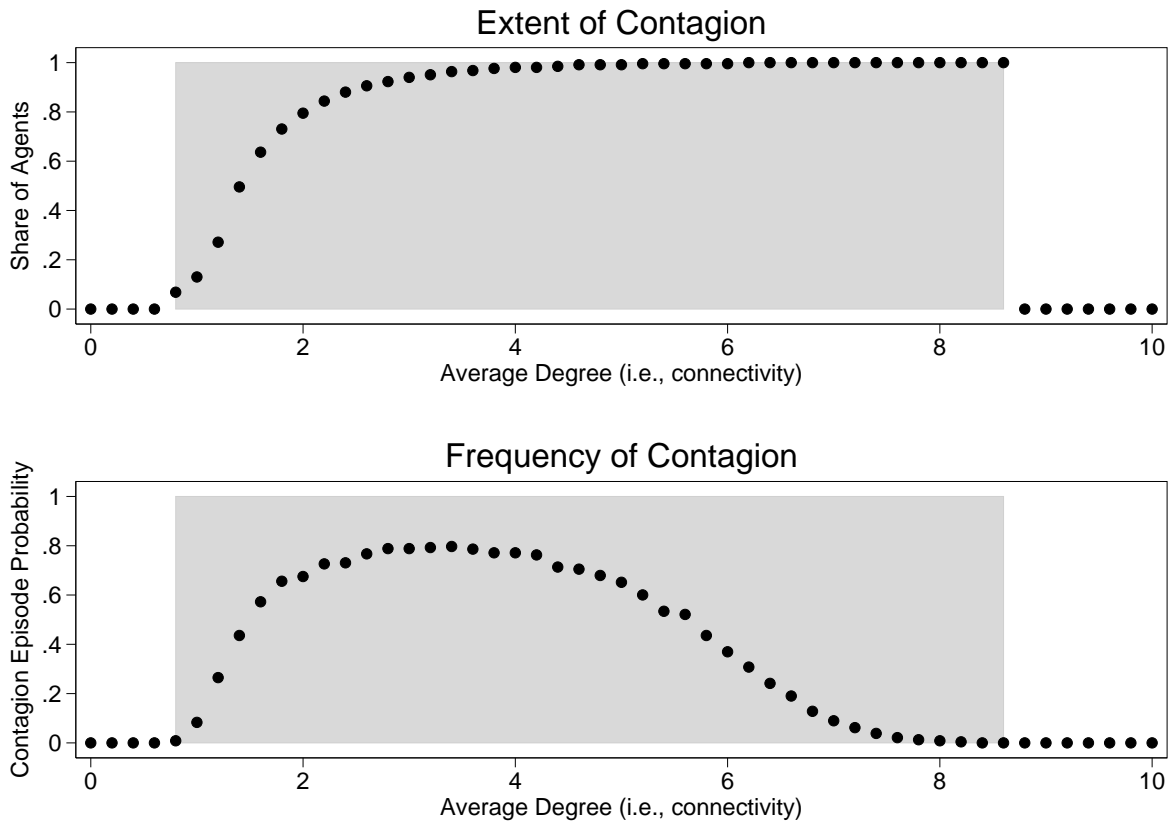
- Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016.
- Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2015a). Networks, shocks, and systemic risk. *NBER WP*, 2015(20931).
- Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2015b). Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608.
- Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2017). Microeconomic origins of macroeconomic tail risks. *American Economic Review*, 107(1):54–108.
- Allen, F. and Gale, D. (2000). Financial contagion. *Journal of Political Economy*, 108(1):1–33.
- Bak, P., Chen, K., Scheinkman, J., and Woodford, M. (1993). Aggregate fluctuations from independent sectoral shocks: self-organized criticality in a model of production and inventory dynamics. *Ricerche Economiche*, 47(1):3 – 30.
- Banbura, M., Giannone, D., and Reichlin, L. (2010). Large bayesian vector auto regressions. *Journal of Applied Econometrics*, 25(1):71–92.
- Baqae, D. R. (2016). Cascading failures in production networks. *Working Paper*.
- Bastian, M., Heymann, S., Jacomy, M., and Venturini, T. (2014). Forceatlas2, a continuous graph layout algorithm for handy network visualization designed for the gephi software. *PLoS ONE*, 9(6).
- Benmelech, E., Bergman, N., Milanez, A., and Mukharlyamov, V. (2014). The agglomeration of bankruptcy. *NBER WP*, 2014(20254).
- Billio, M., Getmansky, M., Lo, A. W., and Pelizzon, L. (2012). Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, 104(3):535 – 559.

modeled networks.

- Blume, L., Easley, D., Kleinberg, J., Kleinberg, R., and Tardos, E. (2011). Which networks are least susceptible to cascading failures? *Proceedings of the 2011 IEEE 52Nd Annual Symposium on Foundations of Computer Science*, pages 393–402.
- Blume, L., Easley, D., Kleinberg, J., Kleinberg, R., and Tardos, E. (2013). Network formation in the presence of contagious risk. *ACM Transactions on Economics and Computation*, 1(2):6:1–6:20.
- Caccioli, F., Farmer, J. D., Foti, N., and Rockmore, D. (2015). Overlapping portfolios, contagion, and financial stability. *Journal of Economic Dynamics and Control*, 51:50 – 63.
- Caccioli, F. and Farmer, T. A. C. J. D. (2012). Heterogeneity, correlations and financial contagion. *Advances in Complex Systems*, 15(2).
- Carvalho, V. and Gabaix, X. (2013). The great diversification and its undoing. *American Economic Review*, 103(5):1697–1727.
- Chudik, A., Kapetanios, G., and Pesaran, M. H. (2016). Big Data Analytics: A New Perspective (OCMT). *Globalization & Monetary Policy Institute Working Paper*. No. 268.
- Demirer, M., Diebold, F. X., Liu, L., and Yilmaz, K. (2015). Estimating global bank network connectedness. *Manuscript, MIT, University of Pennsylvania, and Koc University*.
- Diebold, F. X. and Yilmaz, K. (2009). Measuring financial asset return and volatility spillovers, with application to global equity markets. *The Economic Journal*, 119(534):158–171.
- Diebold, F. X. and Yilmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, 182(1):119 – 134.
- Diebold, F. X. and Yilmaz, K. (2015). *Financial and Macroeconomic Connectedness: A Network Approach to Measurement and Monitoring*. Oxford University Press.
- Diebold, F. X. and Yilmaz, K. (2016). Trans-Atlantic Equity Volatility Connectedness: U.S. and European Financial Institutions, 2004-2014. *Journal of Financial Econometrics*, 14(1):81–127.
- Elliott, M., Golub, B., and Jackson, M. O. (2014). Financial networks and contagion. *American Economic Review*, 104(10):3115–53.
- Farhi, E. and Tirole, J. (2016). Deadly Embrace: Sovereign and Financial Balance Sheets Doom Loops. *NBER WP*, 2016(21843).
- Freixas, X., Parigi, B. M., and Rochet, J.-C. (2000). Systemic risk, interbank relations, and liquidity provision by the central bank. *Journal of Money, Credit and Banking*, 32(3):611–638.
- Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica*, 79(3):733–772.
- Gai, P., Haldane, A., and Kapadia, S. (2011). Complexity, concentration and contagion. *Journal of Monetary Economics*, 58(5):453–470.
- Gai, P. and Kapadia, S. (2010). Contagion in financial networks. *Proceedings of the Royal Society of London A*, 466:2401–2423.
- Glasserman, P. and Young, H. P. (2016). Contagion in financial networks. *Journal of Economic Literature*, 54(3):779–831.
- Glick, R. and Rose, A. K. (1999). Contagion and trade: Why are currency crises regional? *Journal of International Money and Finance*, 18(4):603 – 617.
- Goins, S. and Gruca, T. S. (2008). Understanding competitive and contagion effects of layoff announcements. *Corporate Reputation Review*, 11(1):12–34.
- Grant, E. (2016). Exposure to International Crises: Trade vs. Financial Contagion. *Globalization & Monetary Policy Institute Working Paper*. No. 280.

- Grant, E. and Yung, J. (2017). The Role of Firm Size in Global Network Connectedness. *Working Paper*.
- Haldane, A. G. (2009). Rethinking the financial network. *Speech presented at the Financial Student Association, Amsterdam*.
- Hale, G. (2012). Bank relationships, business cycles, and financial crises. *Journal of International Economics*, 88(2):312 – 325. {NBER} Global.
- Hale, G., Kapan, T., and Minoiu, C. (2016). Crisis transmission through the global banking network. *Federal Reserve Bank of San Francisco Working Paper*, 2016(2016-1).
- Hernandez, L. F. and Valdes, R. O. (2001). What drives contagion: Trade, neighborhood, or financial links? *International Review of Financial Analysis*, 10(3):203 – 218.
- Heston, S. L. and Rouwenhorst, K. (1994). Does industrial structure explain the benefits of international diversification? *Journal of Financial Economics*, 36(1):3 – 27.
- Hoberg, G. and Phillips, G. (2016). Text-based network industries and endogenous product differentiation. *Journal of Political Economy*, 124(5):1423–1465.
- Kitwiwattanachai, C. (2015). Learning Network Structure of Financial Institutions from CDS Data. *Working Paper*.
- Koop, G. M. (2013). Forecasting with medium and large bayesian vars. *Journal of Applied Econometrics*, 28(2):177–203.
- McGuire, P. and Tarashev, N. (2008). Global monitoring with the BIS international banking statistics. *BIS Working Papers*. No. 244.
- Minoiu, C. and Reyes, J. A. (2013). A network analysis of global banking: 1978-2010. *Journal of Financial Stability*, 9(2):168 – 184.
- Mulally, A. R. (2008). Examining the State of the Domestic Automobile Industry. *Hearing of the United States Senate Committee on Banking, Housing, and Urban Affairs*.
- Pesaran, H. and Shin, Y. (1998). Generalized impulse response analysis in linear multivariate models. *Economics Letters*, 58(1):17–29.
- Pesaran, M. H., Schuermann, T., and Weiner, S. M. (2004). Modeling regional interdependencies using a global error-correcting macroeconomic model. *Journal of Business & Economic Statistics*, 22(2):129–162.
- Scida, D. (2015). Structural var and financial networks: A minimum distance approach to spatial modeling. *Job Market Paper*.
- Vivier-Lirimont, S. (2006). Contagion in interbank debt networks. *Working Paper*.
- Zou, H. and Zhang, H. H. (2009). On the Adaptive Elastic-Net with a Diverging Number of Parameters. *The Annals of Statistics*, 37(4):1733–1751.

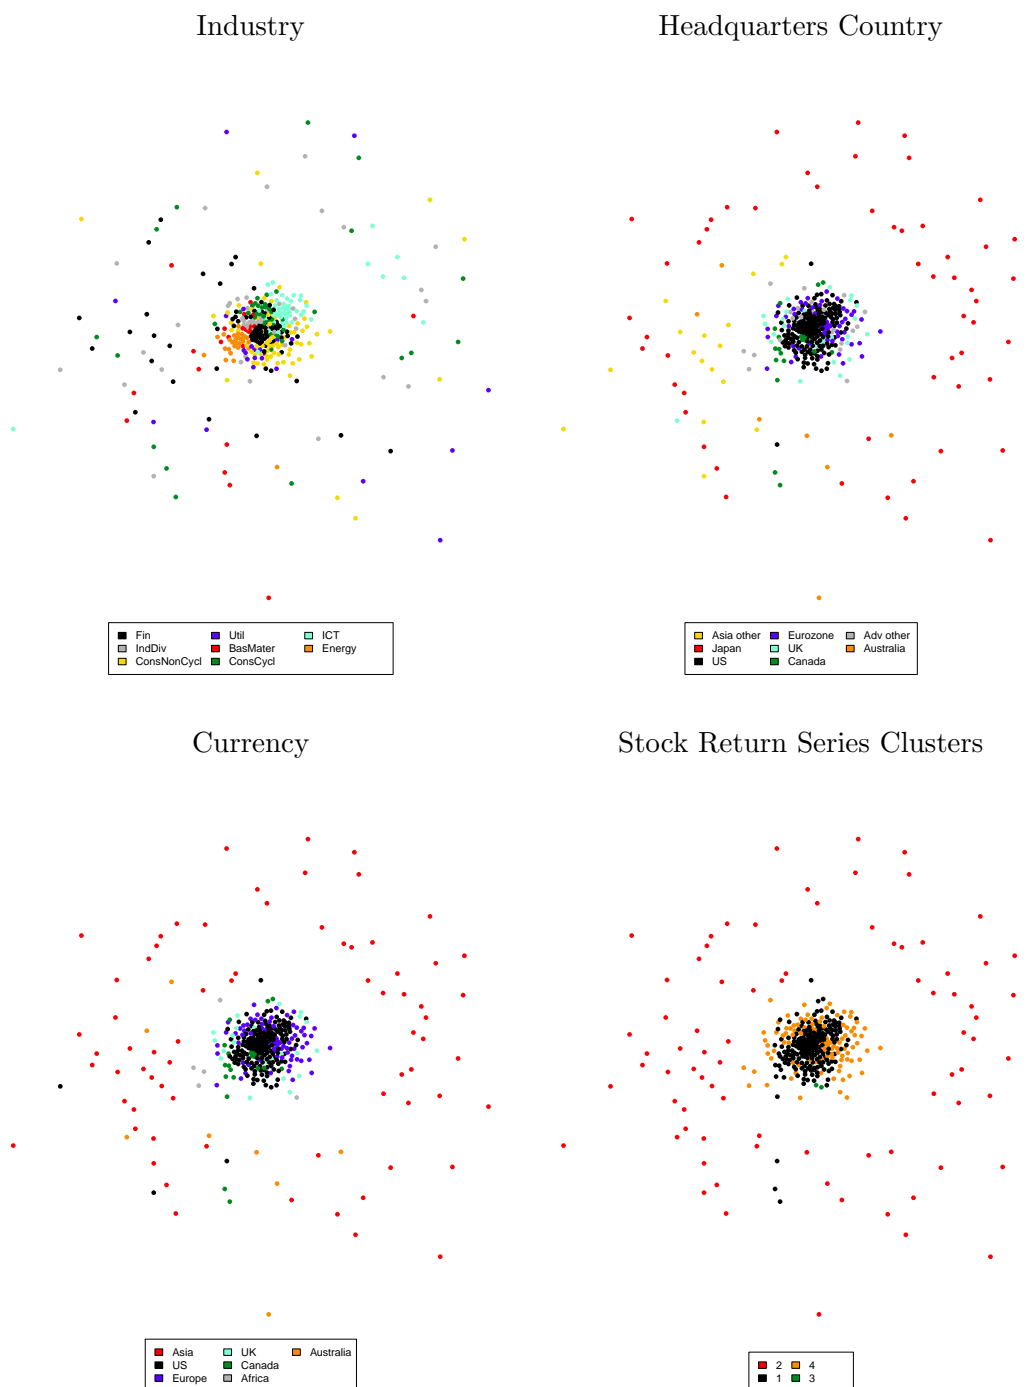
Figure 1: Robust-yet-Fragile Contagion in Firm Networks



Note: Shaded areas indicate the contagion window. Results of model simulations assuming a network of 1,000 firms, with 1,000 repetitions for each level of connectedness. Plots are the average results when one firm is randomly selected to receive a shock that forces it into bankruptcy. A contagious default outbreak is defined as at least 5% of firms becoming insolvent.

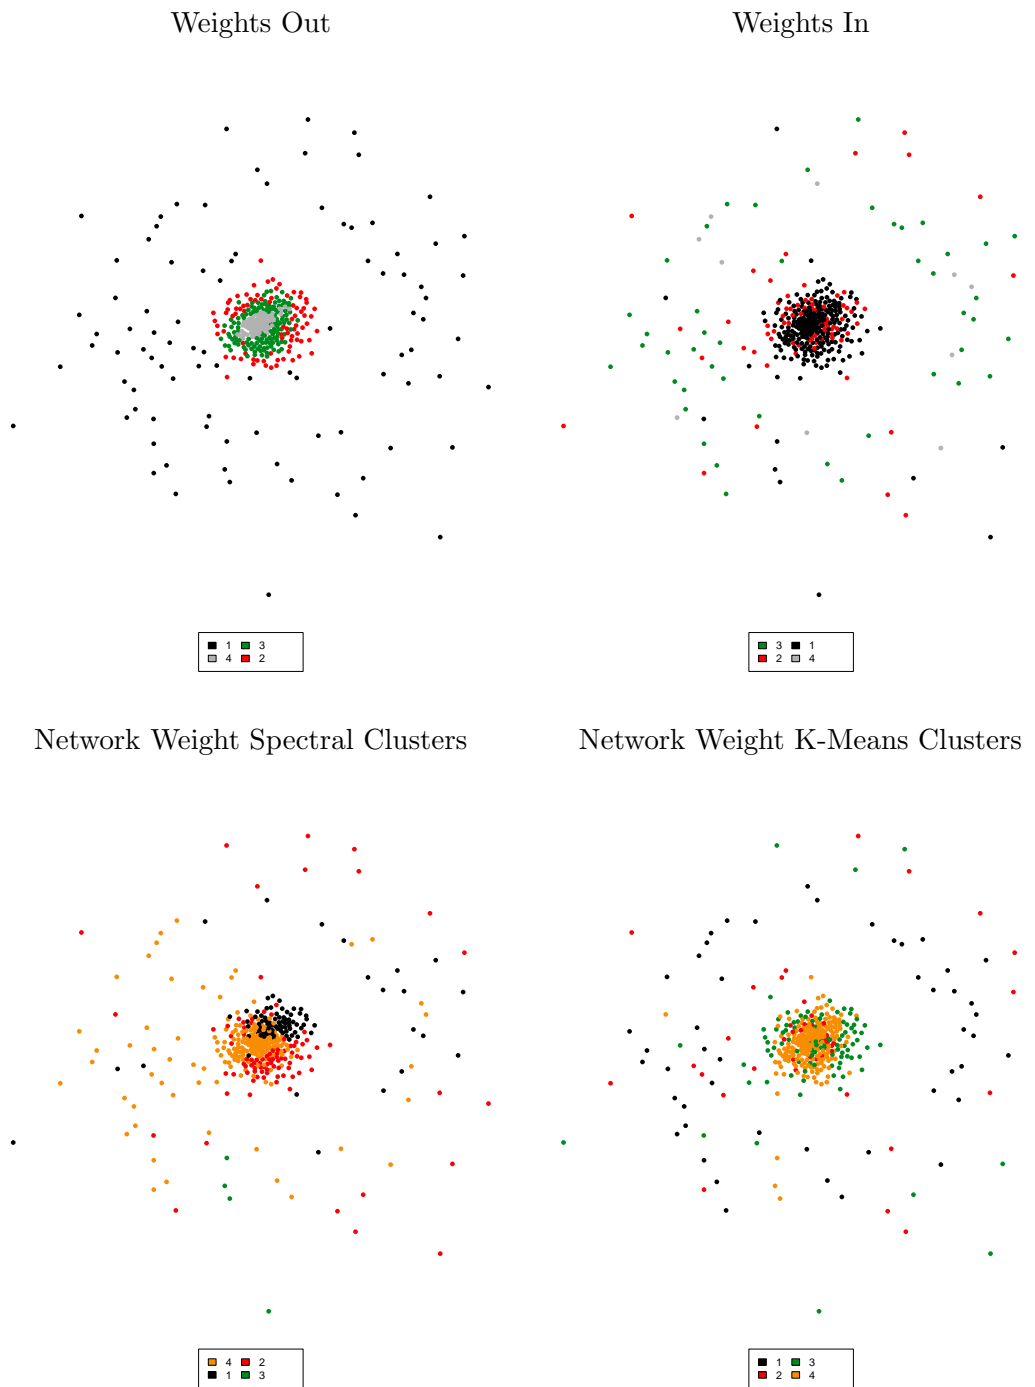
Source: Gai and Kapadia (2010).

Figure 2: Global Network, 1991-2016, by Firm Categories



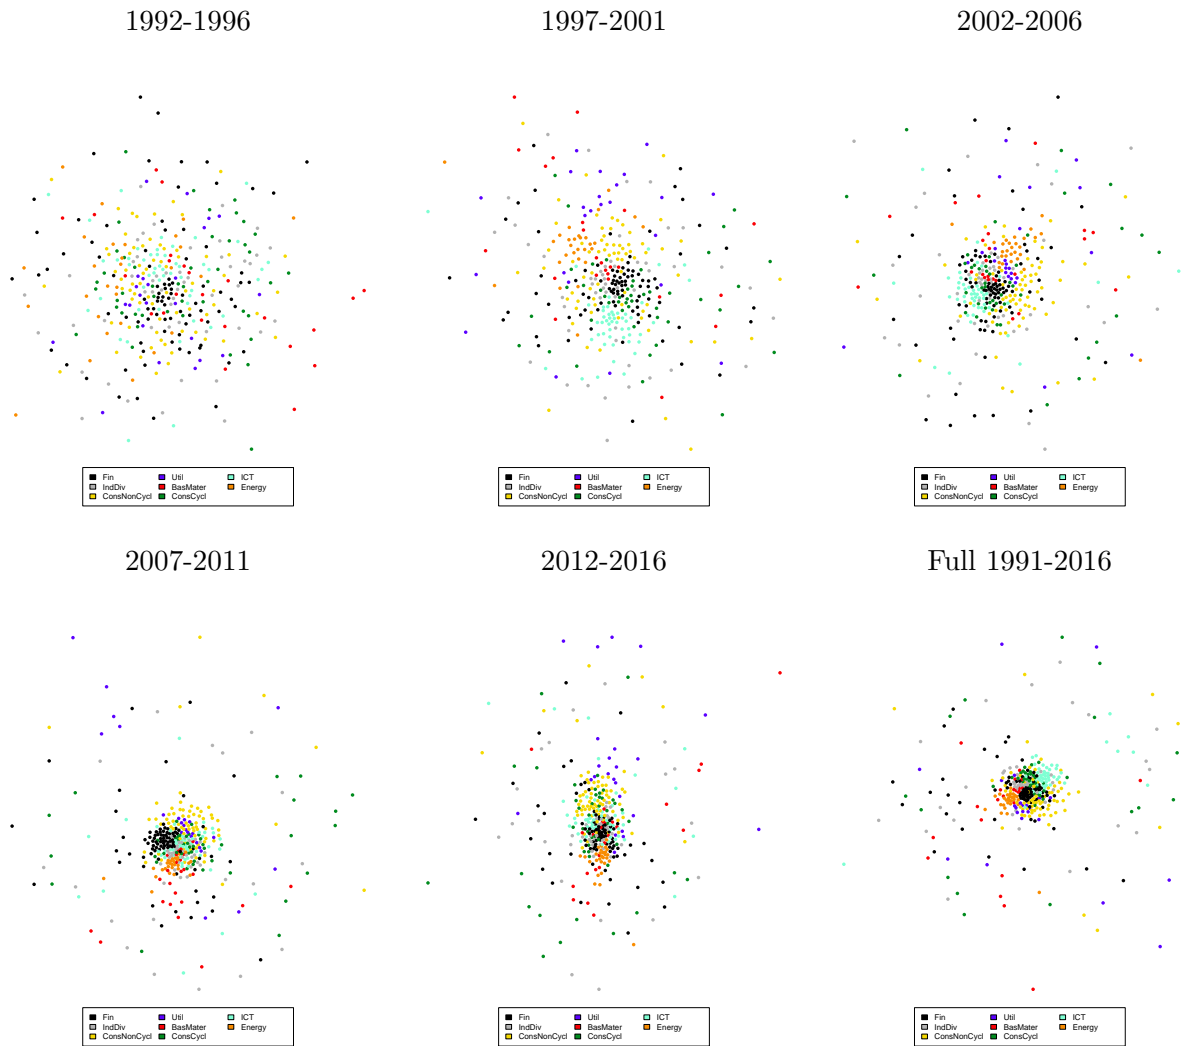
Note: Daily equity return based network for all firms available continuously from January 1991 through September 2016. Estimated using 1 lag in the VAR and GIRF horizon=1. This sample includes 382 firms across 18 countries, with equities issued in 13 currencies.

Figure 3: Global Network, 1991-2016, by Network Weight Categories



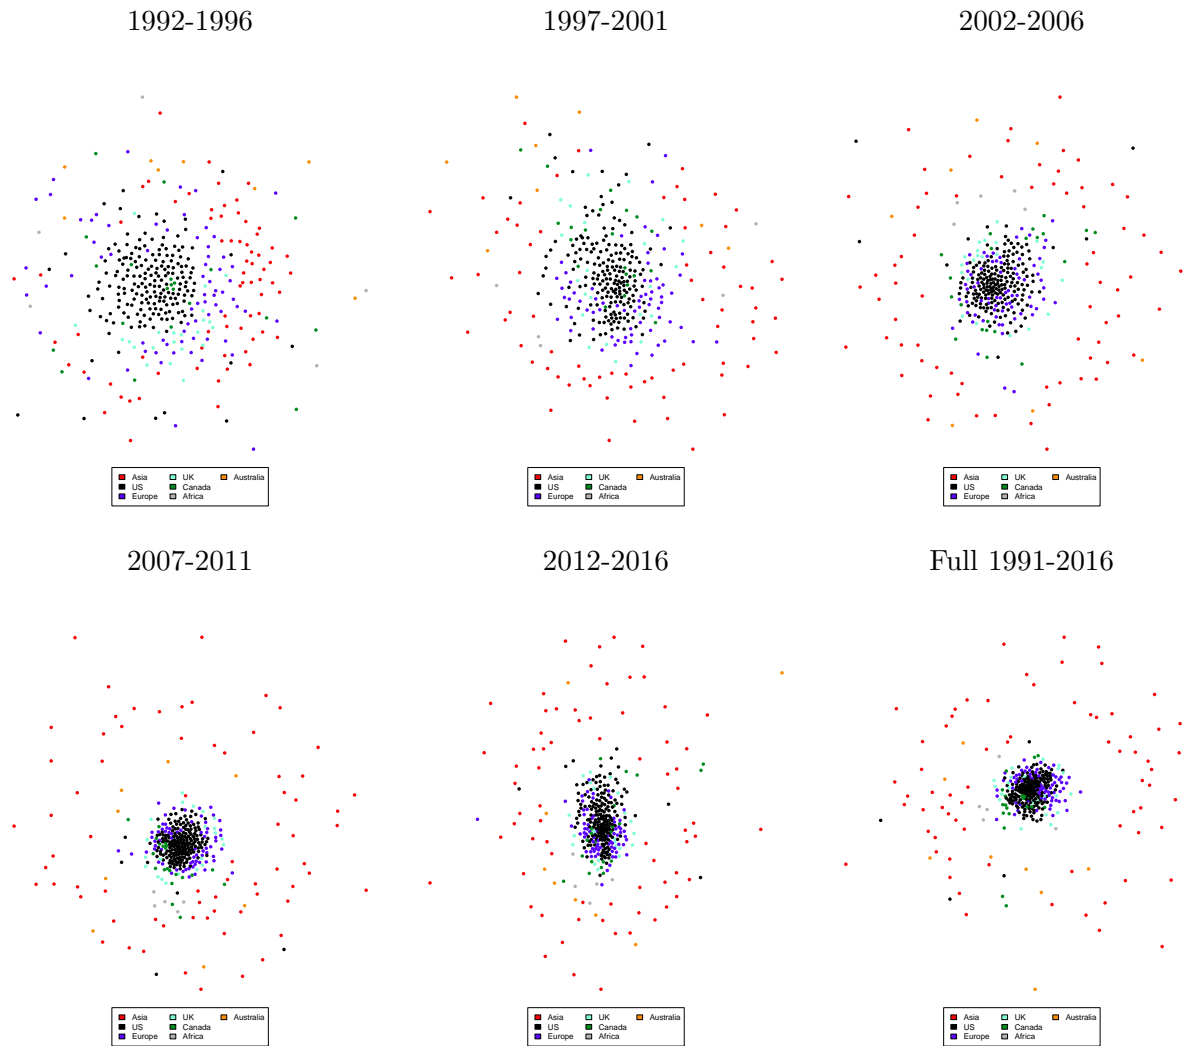
Note: Daily equity return based network for all firms available continuously from January 1991 through September 2016. Estimated using 1 lag in the VAR and GIRF horizon=1. This sample includes 382 firms across 18 countries, with equities issued in 13 currencies. For the top two plots, higher quartile numbers indicate greater weights.

Figure 4: Global Dynamic Networks by Industry



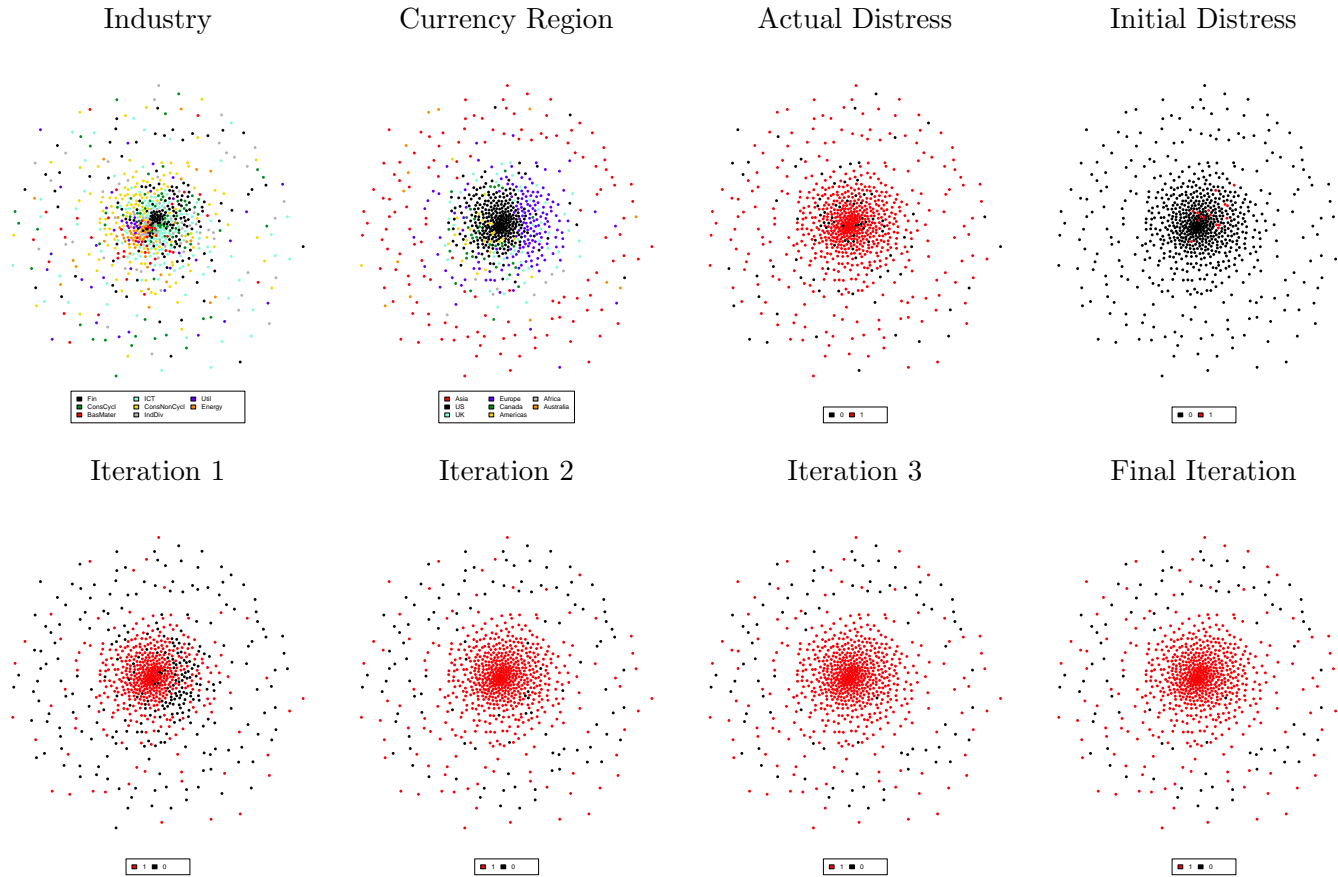
Note: Daily equity return based network for all firms available continuously during each period. Estimated using 1 lag in the VAR and GIRF horizon=1. This sample includes 382 firms across 18 countries, with equities issued in 13 currencies.

Figure 5: Global Dynamic Networks by Currency




Note: Daily equity return based network for all firms available continuously during each period. Estimated using 1 lag in the VAR and GIRF horizon=1. This sample includes 382 firms across 18 countries, with equities issued in 13 currencies.

Figure 6: Fall 2008 Contagion — Actual Data vs Model Simulations



Note: Spring plots are based on the rolling global firm sample network from 2003–2007. Simulations are based off of quarterly predictions from the main sample using the latent linear model estimates from a probit model estimated from 1996–2016, with firm distress as the dependent variable and the five robust and fragile terms from our main regressions with the level of the VIX as the network crisis variable, in addition to the TED spread as its own term. VIX and TED spread levels from the end of 2008 and the above initially distressed firms are entered into the model, and it is simulated until reaching a steady state. See Online Appendix Table A24 for details on the initially distressed firms, and Table A25 for the estimated latent probit model used. This sample includes 756 firms across 40 countries, with equities issued in 25 currencies.

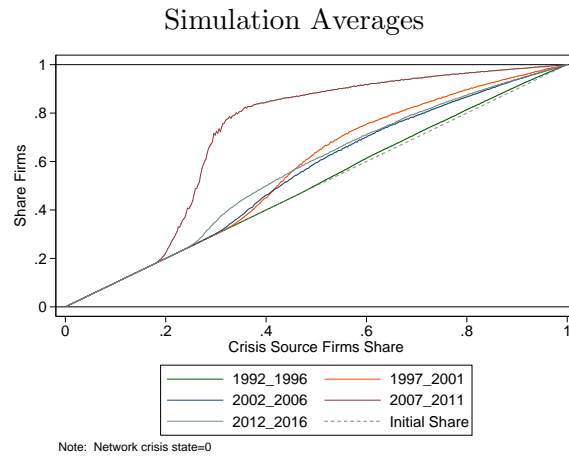
Figure 7: Spread of Local Sector Shocks vs Sector Firm Share



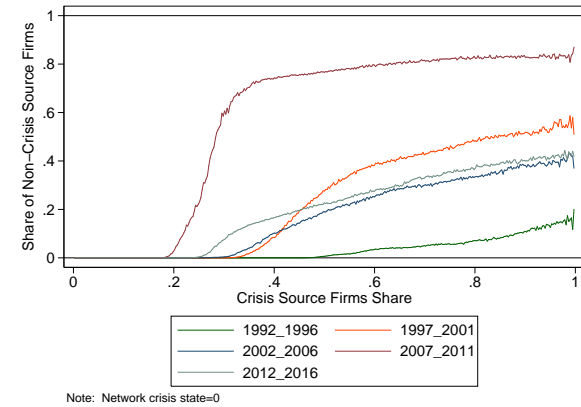
Note: This figure plots the final extent of contagion from a shock to each currency-industry locality pair. The simulations use a firm network of our long, continuous global firm sample from 1991-2016, using lag and GIRF orders of one. Simulations are based off of the latent linear model estimates of the probit model from column (5) of Table A5. The TED spread is assumed to be 2.9, the value of the top percentile of its sample distribution, in order to consider the worst case extent of contagion from a shock to each local sector. This sample includes 382 firms across 18 countries, with equities issued in 13 currencies.

Figure 8: Network Contagion Simulations Using Various Time Period Networks

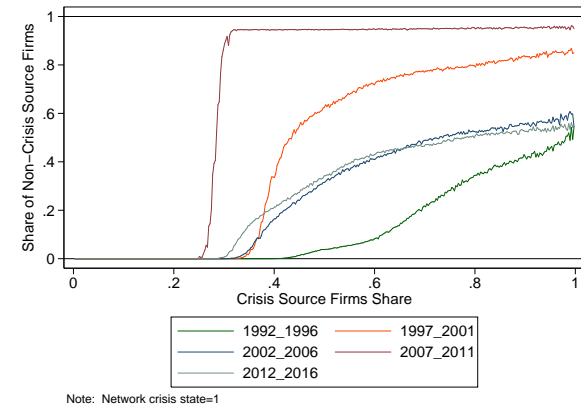
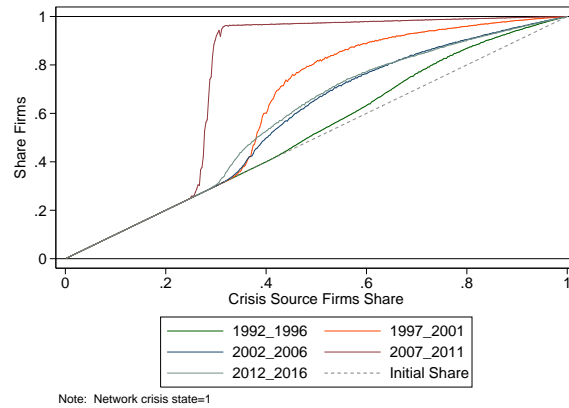
Without Network Level Crisis State



Contagion Spread Fraction Averages

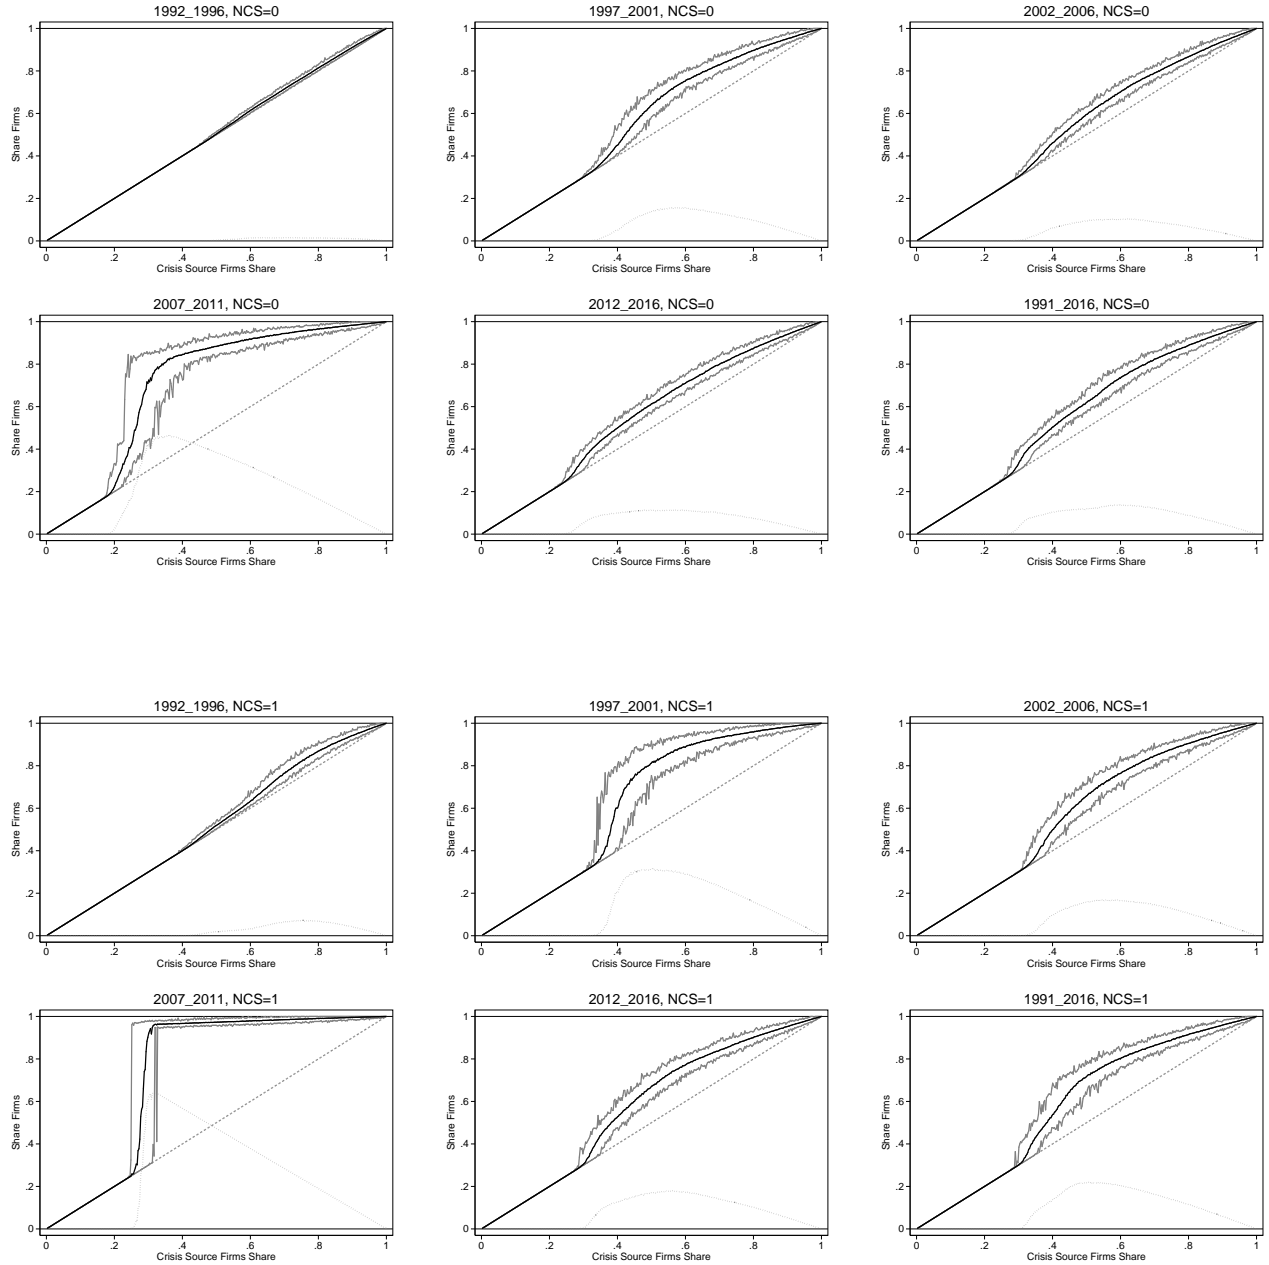


With Network Level Crisis State



Note: Network contagion simulations using daily equity return based networks for all firms available continuously during each period. Estimated using 1 lag in the VAR and GIRF horizon=1. Contagion determined using simulations based upon estimates from the latent equation of Column (5) in Table 3. Crisis state is equivalent to a TED spread crisis state. The contagion fractions are the average contagion minus the initial crisis share divided by one minus the initial crisis share.

Figure 9: Details of Network Contagion Simulations Using Various Time Period Networks



Note: Network contagion simulations using various global continuous sample estimated monthly networks based on regression estimates from Column (5) in Table 3. Crisis state is equivalent to a TED spread crisis state. Black lines are simulation averages, gray lines are minimum and maximum levels, and the dotted lines are the spread of contagion (the distance between the averages and the 45-degree line that marks the initial share of firms in crisis). This sample includes 382 firms across 18 countries, with equities issued in 13 currencies.

Table 1: Global Top Firms by Network Out Weights

Rank	Ticker	Name	Industry	Currency	Measure
1	jpm-us	JPMORGAN CHASE	Fin	USD	29.42
2	ge-us	GENERAL ELECTRIC	IndDiv	USD	29.4
3	ben-us	FRANKLIN RES INC	Fin	USD	29.36
4	c-us	CITIGROUP INC	Fin	USD	29.04
5	axp-us	AMERICAN EXPRESS	Fin	USD	28.75
6	bac-us	BANK OF AMERICA	Fin	USD	28.27
7	bk-us	BANK NY MELLON	Fin	USD	28.22
8	ntrs-us	NORTHERN TRUST	Fin	USD	28.2
9	ppg-us	PPG INDS INC	BasMater	USD	27.93
10	emr-us	EMERSON ELEC CO	IndDiv	USD	27.82
11	pcar-us	PACCAR INC	ConsCycl	USD	27.16
12	wfc-us	WELLS FARGO AND CO	Fin	USD	27.11
13	sti-us	SUNTRUST BANKS	Fin	USD	27.06
14	dd-us	DU PONT (EI)	BasMater	USD	26.98
15	itw-us	ILLINOIS TOOL WO	IndDiv	USD	26.96
16	key-us	KEYCORP	Fin	USD	26.93
17	l-us	LOEWS CORP	Fin	USD	26.8
18	pnc-us	PNC FINANCIAL SE	Fin	USD	26.78
19	stt-us	STATE ST CORP	Fin	USD	26.71
20	etn-us	EATON CORP PLC	IndDiv	USD	26.62
21	cat-us	CATERPILLAR INC	IndDiv	USD	26.46
22	aa-us	ALCOA CORP	BasMater	USD	26.46
23	utx-us	UNITED TECH CORP	IndDiv	USD	26.28
24	hd-us	HOME DEPOT INC	ConsCycl	USD	26.23
25	schw-us	SCHWAB (CHARLES)	Fin	USD	26.07

Note: Daily equity return based network for all firms available continuously from January 1991 through September 2016. Estimated using 1 lag in the VAR and GIRF horizon=1. Self-loops not included. This sample includes 382 firms across 18 countries, with equities issued in 13 currencies.

Table 2: Global Top Currency & Industry Aggregates by Network Out Weights

Rank	Name	Measure
1	USD-Fin	760.26
2	USD-ConsNonCycl	728.06
3	USD-ICT	622.75
4	USD-IndDiv	515
5	USD-ConsCycl	445.44
6	USD-Energy	366.96
7	EUR-Fin	221.46
8	USD-BasMater	212.46
9	USD-Util	190.03
10	CAD-Fin	147.84
11	EUR-ConsNonCycl	143.07
12	GBP-Fin	118.74
13	EUR-ConsCycl	91.96
14	EUR-IndDiv	82.55
15	EUR-ICT	81.63
16	CAD-Energy	80.71
17	CHF-Fin	78.79
18	GBP-ConsNonCycl	59.89
19	EUR-Util	57.42
20	JPY-ConsCycl	54.57
21	JPY-IndDiv	51.15
22	CAD-BasMater	48.13
23	EUR-Energy	43.53
24	EUR-BasMater	43.15
25	CAD-ICT	42.58

Note: Daily equity return based network for all firms available continuously from January 1991 through September 2016. Estimated using 1 lag in the VAR and GIRF horizon=1. Self-loops not included. This sample includes 382 firms across 18 countries, with equities issued in 13 currencies.

Table 3: Global Network Monthly Equity Return Crises, 1996-2016

Panel A: Regression Estimates						
	(1)	(2)	(3)	(4)	(5)	
Robustness: Diversification	0.000208*		-0.00248***	-0.00232***	-0.00226***	-
$\phi \sum_{j \neq i} w_{jit}$	(0.000108)		(0.000205)	(0.000200)	(0.000197)	
Fragility: Direct Contagion		0.0170***	0.0217***	0.0210***	0.0211***	+
$\gamma \sum_{j \neq i} w_{jit} D_{jt}$		(0.000904)	(0.000936)	(0.000927)	(0.000952)	
Fragility: Network Vulnerability				0.0811***	0.137***	+
λN_t				(0.00684)	(0.0117)	
Robustness: Network Resistance					-0.00269***	-
$\omega N_t \sum_{j \neq i} w_{jit}$					(0.000602)	
Fragility: Network Crisis Reinforced Contagion					0.000388	+
$\theta N_t \sum_{j \neq i} w_{jit} D_{jt}$					(0.00139)	
Observations	95,118	95,118	95,118	95,118	95,118	

Panel B: Standardized Coefficients			
	$\Delta \text{Pr}(\text{Crisis})$ for one StdDev increase in explanatory variable for regression:		
	(4)	(5)	
Robustness: Diversification	-3.7%	-3.6%	-
Fragility: Direct Contagion	5.9%	5.9%	+
Fragility: Network Vulnerability	8.1%	13.7%	+
Robustness: Network Resistance		-4.3%	-
Fragility: Network Crisis Reinforced Contagion		0.1%	+

Note: The dependent variable is the monthly equity return distress indicator for firm i , the neighboring firm health variable (D_{jt}) is the equity return distress indicator for firm j , and the network state variable (N_t) is an indicator for the average monthly TED spread being in the top 10%. Effects of a network crisis state indicate the effect of the indicator variable going from zero to one, and the interactions with the weight (w_{jit}) and crisis sums are this times the standard deviation of the interacted sum term. This sample includes 382 firms across 18 countries, with equities issued in 13 currencies. Robust standard errors clustered at the firm level in parentheses. The marginal effects for the probit regressions are provided in place of the latent regression coefficient estimates. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 4: Global Network Continuous Monthly Equity Returns, 1996-2016

Panel A: Regression Estimates						
	(1)	(2)	(3)	(4)	(5)	
Robustness: Diversification	-1.48e-05		0.000821***	0.000771***	0.000749***	+
$\phi \sum_{j \neq i} w_{jit}$	(1.49e-05)		(4.12e-05)	(3.89e-05)	(3.81e-05)	
Fragility: Direct Contagion		-0.00962***	-0.0116***	-0.0111***	-0.0112***	-
$\gamma \sum_{j \neq i} w_{jit} D_{jt}$		(0.000432)	(0.000503)	(0.000492)	(0.000516)	
Fragility: Network Vulnerability				-0.0388***	-0.0602***	-
λN_t				(0.00266)	(0.00466)	
Robustness: Network Resistance					0.00218***	+
$\omega N_t \sum_{j \neq i} w_{jit}$					(0.000268)	
Fragility: Network Crisis Reinforced Contagion					-0.00196***	-
$\theta N_t \sum_{j \neq i} w_{jit} D_{jt}$					(0.000610)	
Observations	95,118	95,118	95,118	95,118	95,118	
R-squared	0.000	0.079	0.094	0.099	0.101	

Panel B: Standardized Coefficients			
	Standardized coefficients for regression:		
	(4)	(5)	
Robustness: Diversification	12.3%	12.0%	+
Fragility: Direct Contagion	-31.0%	-31.2%	-
Fragility: Network Vulnerability	-38.8%	-60.2%	-
Robustness: Network Resistance		34.9%	+
Fragility: Network Crisis Reinforced Contagion		-5.5%	-

Note: The dependent variable is the monthly equity return for firm i , the neighboring firm health variable (D_{jt}) is the equity return distress indicator for firm j , and the network state variable (N_t) is an indicator for the average monthly TED spread being in the top 10%. Effects of a network crisis state indicate the effect of the indicator variable going from zero to one, and the interactions with the weight (w_{jit}) and crisis sums are this times the standard deviation of the interacted sum term. This sample includes 382 firms across 18 countries, with equities issued in 13 currencies. Robust standard errors clustered at the firm level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 5: Global Network RyF Regressions with Various Continuous Dependent Variables, 1996-2016

Panel A: Regression Estimates								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	
	EQ Return	Change 5yr CDS Spread†	5yr CDS Spread†	Change EBITDA	Change Revenue	Change RoE	RoE	
Robustness: Diversification	0.000771***	-0.479***	-1.114***	0.00229***	0.000990***	0.0984***	-0.0859*	+
$\phi \sum_{j \neq i} w_{jit}$	(3.89e-05)	(0.0872)	(0.218)	(0.000453)	(0.000183)	(0.0350)	(0.0496)	
Fragility: Direct Contagion	-0.0111***	7.823***	14.38***	-0.0310***	-0.0120***	-1.120**	-1.572***	-
$\gamma \sum_{j \neq i} w_{jit} D_{jt}$	(0.000492)	(0.739)	(2.078)	(0.00556)	(0.00198)	(0.462)	(0.464)	
Fragility: Network Vulnerability	-0.0388***			-0.0989***	-0.0151	-5.289***	-2.910**	-
λN_t	(0.00266)			(0.0201)	(0.0133)	(1.553)	(1.447)	
Observations	95,118	7,620	7,976	21,479	27,913	6,375	6,571	
R-squared	0.099	0.028	0.010	0.007	0.004	0.005	0.009	
Period	M	M	M	Q	Q	Y	Y	
Panel B: Standardized Coefficients								
	Standardized coefficients for regression:							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	
Robustness: Diversification	12.3%	-10.4%	-31.1%	9.2%	7.2%	5.6%	-3.6%	+
Fragility: Direct Contagion	-31.0%	46.0%	30.2%	-11.9%	-10.4%	-6.8%	-6.3%	-
Fragility: Network Vulnerability	-38.8%	-	-	-23.0%	-6.6%	-17.5%	-7.3%	-

Note: The dependent variables are the monthly averages of each variable for firm i , the neighboring firm health variable (D_{jt}) indicates the most extreme 10% of distress as measured by the corresponding dependent variable for firm j , and the network state variable (N_t) is an indicator for the average monthly TED spread being in the top 10%. CDS data covers October, 2013 - November, 2015 during which there were no network distress events so that term is excluded from those regression results. Effects of a network crisis state indicate the effect of the indicator variable going from zero to one, and the interactions with the weight (w_{jit}) and crisis sums are this times the standard deviation of the interacted sum term. †Expected to have coefficients with signs opposite the other regressions. This sample includes 382 firms across 18 countries, with equities issued in 13 currencies. Robust standard errors clustered at the firm level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 6: Global Network Monthly Equity Return Crises, 1996-2016 with Various Network Crisis Measures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	
	Crisis TedSprdAvg	TedSprdAvg	log TedSprdAvg	Crisis VIXAvg	VIXAvg	Crisis RetSP500	RetSP500†	Fraction of Firms in Distress	Crisis Fraction of Firms in Distress	
Robustness: Diversification	-0.00232***	-0.00204***	-0.00210***	-0.00211***	-0.00165***	-0.00129***	-0.000883***	4.29e-05	-0.00127***	-
$\phi \sum_{j \neq i} w_{jit}$	(0.000200)	(0.000197)	(0.000203)	(0.000193)	(0.000181)	(0.000159)	(0.000138)	(0.000115)	(0.000163)	
Fragility: Direct Contagion	0.0210***	0.0199***	0.0204***	0.0187***	0.0156***	0.0120***	0.00912***	0.00245***	0.0112***	+
$\gamma \sum_{j \neq i} w_{jit} D_{jt}$	(0.000927)	(0.000933)	(0.000951)	(0.000874)	(0.000796)	(0.000684)	(0.000568)	(0.000419)	(0.000605)	
Fragility: Network Vulnerability	0.0811***	0.0406***	0.0225***	0.0607***	0.00350***	0.148***	-1.079***	0.518***	0.120***	+
λN_t	(0.00684)	(0.00239)	(0.00172)	(0.00369)	(0.000133)	(0.00758)	(0.0414)	(0.0163)	(0.00395)	
Observations	95,118	95,118	95,118	95,118	95,118	95,118	95,118	95,118	95,118	

Note: The dependent variable is the monthly equity return distress indicator for firm i , the neighboring firm health variable (D_{jt}) is the equity return distress indicator for firm j , and the network state variable (N_t) varies across the columns. †Expected to have a negative network vulnerability coefficient. This sample includes 382 firms across 18 countries, with equities issued in 13 currencies. Robust standard errors clustered at the firm level in parentheses. The marginal effects for the probit regressions are provided in place of the latent regression coefficient estimates. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 7: Global Network Continuous Monthly Equity Returns, 1996-2016 with Various Network Crisis Measures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	
	Crisis TedSprdAvg	TedSprdAvg	log TedSprdAvg	Crisis VIXAvg	VIXAvg	Crisis RetSP500	RetSP500†	Fraction of Firms in Distress	Crisis Fraction of Firms in Distress	
Robustness: Diversification	0.000771***	0.000771***	0.000834***	0.000802***	0.000794***	0.000484***	0.000167***	3.35e-05	0.000483***	+
$\phi \sum_{j \neq i} w_{jit}$	(3.89e-05)	(4.00e-05)	(4.28e-05)	(4.15e-05)	(4.30e-05)	(3.09e-05)	(2.15e-05)	(2.34e-05)	(3.16e-05)	
Fragility: Direct Contagion	-0.0111***	-0.0113***	-0.0116***	-0.0113***	-0.0113***	-0.00691***	-0.00312***	-0.00174***	-0.00729***	-
$\gamma \sum_{j \neq i} w_{jit} D_{jt}$	(0.000492)	(0.000501)	(0.000512)	(0.000513)	(0.000533)	(0.000383)	(0.000272)	(0.000320)	(0.000386)	
Fragility: Network Vulnerability	-0.0388***	-0.00752***	0.00114*	-0.00547***	-0.000209***	-0.0592***	0.823***	-0.303***	-0.0432***	-
λN_t	(0.00266)	(0.00123)	(0.000616)	(0.00138)	(6.47e-05)	(0.00274)	(0.0256)	(0.00978)	(0.00156)	
Observations	95,118	95,118	95,118	95,118	95,118	95,118	95,118	95,118	95,118	
R-squared	0.099	0.095	0.094	0.094	0.094	0.120	0.188	0.158	0.125	

Note: The dependent variable is the monthly equity return distress indicator for firm i , the neighboring firm health variable (D_{jt}) is the equity return distress indicator for firm j , and the network state variable (N_t) varies across the columns. †Expected to have a positive network vulnerability coefficient. This sample includes 382 firms across 18 countries, with equities issued in 13 currencies. Robust standard errors clustered at the firm level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 8: Global Network Monthly Equity Return Crises, 1996-2016 with Various Networks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	
	Bilateral Return Correl	H=0	H=1	H=2	H=3	H=4	H=5	
Robustness:								
Diversification	-0.00117***	-0.00126***	-0.00232***	-0.0122***	-0.0307***	-0.0801***	-0.145***	-
$\phi \sum_{j \neq i} w_{jit}$	(5.17e-05)	(6.12e-05)	(0.000200)	(0.000659)	(0.00247)	(0.00596)	(0.0104)	
Fragility: Direct Contagion	0.00736***	0.00899***	0.0210***	0.0842***	0.222***	0.563***	0.919***	+
$\gamma \sum_{j \neq i} w_{jit} D_{jt}$	(0.000204)	(0.000268)	(0.000927)	(0.00321)	(0.0103)	(0.0251)	(0.0399)	
Fragility: Network Vulnerability	0.000109	0.0214***	0.0811***	0.0844***	0.104***	0.121***	0.139***	+
λN_t	(0.00415)	(0.00509)	(0.00684)	(0.00654)	(0.00807)	(0.00797)	(0.00771)	
Observations	95,118	95,118	95,118	95,118	95,118	95,118	95,118	

Note: The dependent variable is the monthly equity return distress indicator for firm i , the neighboring firm health variable (D_{jt}) is the equity return distress indicator for firm j , and the network state variable (N_t) is an indicator for the average monthly TED spread being in the top 10%. This sample includes 382 firms across 18 countries, with equities issued in 13 currencies. Robust standard errors clustered at the firm level in parentheses. The marginal effects for the probit regressions are provided in place of the latent regression coefficient estimates. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 9: Global Network Continuous Monthly Equity Returns, 1996-2016 with Various Networks

	(1) Bilateral Return Correl	(2) H=0	(3) H=1	(4) H=2	(5) H=3	(6) H=4	(7) H=5	
Robustness: Diversification	0.000256*** (8.87e-06)	0.000298*** (9.98e-06)	0.000771*** (3.89e-05)	0.00340*** (0.000166)	0.00895*** (0.000502)	0.0224*** (0.00112)	0.0483*** (0.00239)	+
Fragility: Direct Contagion	-0.00446*** (9.21e-05)	-0.00538*** (0.000120)	-0.0111*** (0.000492)	-0.0456*** (0.00184)	-0.120*** (0.00609)	-0.299*** (0.0124)	-0.503*** (0.0224)	-
Fragility: Network Vulnerability	-0.00950*** (0.00216)	-0.0219*** (0.00243)	-0.0388*** (0.00266)	-0.0436*** (0.00254)	-0.0485*** (0.00269)	-0.0538*** (0.00262)	-0.0567*** (0.00259)	-
Observations	95,118	95,118	95,118	95,118	95,118	95,118	95,118	
R-squared	0.166	0.157	0.099	0.100	0.093	0.084	0.071	

Note: The dependent variable is the monthly equity return distress indicator for firm i , the neighboring firm health variable (D_{jt}) is the equity return distress indicator for firm j , and the network state variable (N_t) is an indicator for the average monthly TED spread being in the top 10%. This sample includes 382 firms across 18 countries, with equities issued in 13 currencies. Robust standard errors clustered at the firm level in parentheses. *** p<0.01, ** p<0.05, * p<0.1