News and Networks: Using Financial News Coverage to Measure Bank Interconnectedness*

Sophia Kazinnik Cooper Killen Daniela Scidá John Wu †

This version: July 6, 2023

Abstract

This paper examines the interconnectedness of stress-tested banks using financial news coverage. We use the COVID-19 pandemic as an exogenous shock to investigate the behavior of bank networks during periods of stress. We then propose a new measure of systemic risk using text-based eigenvector centrality, a relative metric of influence within a network. We show that this measure provides a valuable complement to existing systemic risk measures. Our findings highlight the importance of soft information in the context of financial stability. Our approach offers a novel tool to study the financial system's architecture and complements insights from traditional structured data.

Keywords: Banks, networks, financial news, systemic risk, text analysis, COVID-19
JEL Classification: G1, D85, L14, G21, G28, G32

*The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Richmond or the Federal Reserve System. We thank David Aldama-Navarrete, Celso Brunetti, Erik Heitfield, Xin Huang, Filippo Curti, Jung-Eun Kim, Fabrizio Marodin, Ping McLemore, Seung J. Lee, Filip Zikes, and Matthew Plosser for their valuable comments and suggestions. We thank Federal Reserve Bank of Richmond QSR Brown Bag, RiskLab/BoF/ESRB Conference on Systemic Risk Analytics, AEA/ASSA Annual Meeting, FRS Interconnectedness Brown Bag, WEAI Annual Conference, FRS Econometrics Conference, FRS Conference on Financial Institutions, Regulations, and Markets, and WEA Annual Conference participants for valuable feedback. We thank Enya Truong, Caroline Norris, Bryson Alexander, and Ethan Butler for excellent research assistance.

†Sophia Kazinnik, Daniela Scidá, and John Wu are affiliated with the Quantitative Supervision and Research (QSR) division of the Federal Reserve Bank of Richmond, 530 E Trade St, Charlotte, NC 28202. Cooper Killen is affiliated with the Wholesale Credit Risk Center (WCRC) division of the Federal Reserve Bank of Chicago, 230 S La Salle St, Chicago, IL 60604. Emails: Sophia.Kazinnik@rich.frb.org, Cooper.Killen@chi.frb.org, Daniela.Scida@rich.frb.org, John.Wu@rich.frb.org. Corresponding author: Daniela Scidá.

1 Introduction

In 2008, the Global Financial Crisis (GFC) showed that financial interconnectedness is fundamentally tied to systemic risk. In 2023, the failure of Silicon Valley Bank and Signature Bank demonstrated that interconnectedness between financial institutions is not only strong but also multifaceted. Moreover, these recent events showcase the role that digital media plays in today's technology-driven world. In this paper we introduce a novel measure of financial interconnectedness rooted in financial news narrative, and apply it to examine the behavior of banks under extreme stress in real time.

Denser financial networks can lead to shared vulnerabilities, exposing banks to simultaneous spread of illiquidity, insolvency, and losses during periods of financial distress. Popular methods of capturing financial interconnectedness rely on either data that explicitly documents bank connectivity, like inter-bank lending data (e.g., Gofman (2011), Schoar, Afonso, and Kovner (2014)), or on econometric models that infer connectivity from co-movements in market data (e.g., Billio, Getmansky, Lo, and Pelizzon (2012), Diebold and Yilmaz (2014), Härdle, Wang, and Yu (2016)). In the former case, such detailed data is not always available; in the latter case, the data is available in real-time, but it is often difficult to interpret each observed connection.

Depending on the data and methodology, the existing interconnectedness metrics are bound to reflect different economic phenomena (e.g., Brunetti, Harris, Mankad, and Michailids (2019)). Coupled with data availability and interpretability trade-offs, this creates a methodological challenge for constructing a real-time interconnectedness measure. To overcome these issues, our paper exploits an alternative way of measuring financial interconnections in real time by studying banks' relationships in the context of financial news. We leverage financial news from major news outlets and construct weekly network matrices based on co-mentioning of banks in the news (i.e., two banks being mentioned in the same article). This approach allows us to capture the *soft* aspects of financial interconnectedness (e.g., market sentiment, trust, reputation, competitive dynamics, etc.) that are not captured by traditional structured data sources. It also provides an underlying narrative to each observed connection derived from the associated news articles.

Generally, a text-to-network methodology involves transforming textual content into a network of interconnected entities. This approach has become increasingly popular in recent years as advances in natural language processing and network analysis techniques have made it possible to extract valuable information from unstructured text data. For example, Rönnqvist and Sarlin (2015) study the interconnectedness among European large and complex banking groups surrounding the GFC.[1] We focus instead on large U.S. financial institutions that fall under the Dodd-Frank Act Stress Test (DFAST) umbrella and events surrounding the stress period related to the COVID-19 pandemic. We center our attention on DFAST banks precisely because these institutions are the largest in the U.S. and, thus, are arguably the most systemically important.[2]

We contribute to the literature both empirically and methodologically. On the financial networks front, we are the first to study the network among U.S.-based stress tested banks, and, in particular, the impact of a stress event such as the COVID-19 crisis on their network topology. Furthermore, we deliver an interconnectedness approach in real-time with a clear associated narrative.

Our results provide an intuitive characterization of the connection patterns of banks that are subject to DFAST stress testing, as observed through media narratives. We find evidence of clustering among similar types of banks and a core-periphery topology where the largest banks are clustered together at the center with smaller banks at the periphery. We also find a rise in the number of connections across regional banks during the pandemic, which

[1]In general, there have been multiple studies in finance involving networks and unstructured data. For example, Hoberg and Phillips (2016) use manager discussions of competing products to determine interconnectedness of firms.

[2]DFAST banks are institutions with at least $100 billion in total consolidated assets. DFAST banks include LISCC firms, large and complex firms, and large and non-complex firms. Background and further details on the DFAST program can be found at www.federalreserve.gov.

reflects the important role these banks played in facilitating PPP loans.

On the systemic risk front, we introduce a novel (text-based) systemic risk measure relying upon negative sentiment eigenvector centrality. We rank systemic importance of financial institutions according to our metric, and compare it with structured-data-based systemic risk measures (i.e., DIPS, SRISK, and CoVaR). We also conduct principal component analysis to further assess how these measures compare. Our results show that our measure provides more robust and intuitive rankings across the time series while capturing soft information that is embedded in news articles, thereby adding to the available information set of systemic risk measures. In addition, we demonstrate the usefulness of eigenvector centrality in explaining movements in financial variables, such as banks' abnormal returns.

What does it mean when two banks are mentioned in the same article? Financial news reflects market dynamics and information flows. The fact that two banks are mentioned together in financial news suggests some level of interdependence or correlation between their activities, even if the correlation is inverse. The co-occurrence of banks in news articles is therefore indicative of shared economic factors or events that influence their operations. In the paper, we show that the analysis of bank co-occurrences can provide important macro-level insights.

Overall, our contributions offer several key insights into the measurement and analysis of systemic risk, with important implications for policymakers, regulators, and market participants. By utilizing an innovative approach that leverages the informational content of financial news articles, we provide a public discourse picture of the interbank network dynamics during both normal and stress times. This can ultimately help inform regulators on effective risk management and decision-making. Furthermore, text-based networks could be a useful tool for monitoring interconnectedness in the financial system in real time. Our network visualizations can also serve as a powerful tool for regulators, enabling the portrayal of complex sets of connections in a more comprehensible manner.

In summary, our proposed approach to study financial interconnectedness has several advantages. First, it allows for the analysis of both cross-sectional and temporal elements of systemic risk. Second, our method allows for frequent and granular updating of both the network topology and the systemic risk measure. Third, text data has the advantage of providing a richer narrative that can be used to understand the observed connection patterns. Our framework enables real-time analysis of a financial system's architecture based on a clear underlying narrative.

The rest of the paper is organized as follows. Section 2 covers relevant literature, Section 3 discusses data, and Section 4 discusses methodology. Section 5 presents general results on the network topology, and Section 6 discusses the use of eigenvector centrality to pin down systemic importance. Section 7 discusses how eigenvector centrality could be useful to help explain banks' abnormal returns. Section 8 presents robustness checks, and Section 9 concludes.[3]

2 Relevant Literature

Our paper draws from the literature on text analysis, financial networks, and systemic risk. In the next two subsections, we discuss relevant work related to these fields.

2.1 Financial Interconnectedness

Previous works have used interbank lending data (e.g., Gofman (2011), Schoar et al. (2014)) or balance sheet measures (e.g., Greenwood, Augustin, and Thesmar (2015)) to capture bank interconnectedness. Other papers, including Elliott, Golub, and Jackson (2014), examined interconnectedness through shared holdings of equities, debt, and liabilities. There are also papers that measure systemic risk and financial interconnectedness according to the contagion channel and shared default risk through shared liabilities and assets (e.g., Jackson and

[3]Main tables and figures are gathered at the end of the paper, and supplementary materials are provided in the online appendix.

Pernoud (2021)). These data sources provide valuable information that allows for the identification of links based on fundamentals which can be linked back to business characteristics. However, this type of data, and lending data in particular, is not always available.

In the financial econometrics literature, a large set of papers have instead used market co-movement data to capture financial interconnectedness (e.g., Billio, Getmansky, Lo, and Pelizzon (2012), Diebold and Yilmaz (2014), related papers by these authors, and Härdle, Wang, and Yu (2016)); while other papers have used VaR and other market measures to determine financial connectivity (e.g., Hautsch, Schaumburg, and Schienle (2015) and Barigozzi and Brownlees (2019)). The use of market data (e.g., stock return volatility) is appealing since it is available at a high frequency; yet, it is often challenging to relate the observed connections to fundamentals.

In this paper, we are capturing bank interconnectedness by utilizing co-occurrence matrices based on news data. By adopting this approach, we aim to overcome common data limitations that arise with other data sources, such as availability and interpretability. Furthermore, news data provides several advantages over market data, such as its real-time availability and the ability to gain a deeper understanding of the nature of observed connections through the context of the news. We next discuss how text data is currently used in the financial literature.

2.2 News Data in Finance

There is an established body of literature in finance that uses text analysis to convert qualitative information contained in news stories and corporate announcements into quantifiable financial metrics.[4] A subset of this literature relies on financial news as a text source.[5]

[4] For example, Loughran and McDonald (2016) review textual analysis literature in accounting and finance. Gentzkow, Kelly, and Taddy (2019) provide a comprehensive survey of historical advances and recent innovations in text analysis in the social sciences with emphasis on finance

[5] Earlier work by Tetlock (2007), for example, quantitatively measures the interactions between the media and the stock market using daily content from a popular Wall Street Journal column. The paper finds that higher media pessimism predicts downward pressure on market prices and reversion to fundamentals.

For example, Boudoukh, Feldman, Kogan, and Richardson (2013) use text analysis to identify fundamental information in news and show that firm-level public news is a meaningful component of stock return variance. Baker, Bloom, and Davis (2016) develop an index of economic policy uncertainty (EPU) based on newspaper coverage frequency. The authors look at news articles in 10 leading U.S. newspapers and show how specific keywords within newspaper text can yield useful proxies for economic conditions. Shapiro, Sudhof, and Wilson (2020) develop a new time-series measure of economic sentiment derived from economic and financial newspaper articles from January 1980 to April 2015 and aggregate these scores into a daily time-series index. Calomiris and Mamaysky (2019) classify news articles and use that classification to predict risk and return in stock markets in developed and emerging economies.

Some recent literature looks at the relationship between news article sentiment and market reaction during the COVID-19 period. Costola, Hinz, Nofer, and Pelizzon (2023) investigate COVID-19 related news and show that there is a significant and positive relationship between sentiment score and market returns. Mamaysky (2020) studies how financial markets interact with news related to the COVID-19 pandemic and shows that markets were hypersensitive to news during its early onset. Baker, Bloom, Davis, Kost, Sammon, and Viratyosin (2020) examine news articles going back to 1900 and find that no previous infectious disease episode led to stock market volatility even remotely resembling the response to COVID-19. Overall, there is a considerable interest and ever-increasing reliance on news articles as an alternative source of quantitative information.

3 News Data

In this paper, we study financial networks by comparing inter-bank network topology snapshots among large U.S. based financial institutions. This topology is based on the sample of financial news articles described in what follows. We retrieve news articles from the Dow

Jones Factiva Analytics product (Factiva hereafter), and we focus on the top financial news sources. We gather all articles from these sources starting from July 1st, 2019 to September 30, 2020. Overall, this gives us about 49,000 news articles with bank mentions and about 11,400 news articles with co-mentions after cleaning the data.

Our sample includes news articles starting in July 2019 to make sure we capture some of the pre-pandemic news articles. We consider the pandemic period starts in March 2020.[6] The months that capture the peak of pandemic related stress (March and April, 2020) are identified based on the following facts. First, during these months uncertainty and economic stress were the highest as the pandemic was declared by the World Health Organization (WHO) in mid-March, quickly leading to the first lock-down of the U.S. economy. Moreover, the unemployment rate escalated to unprecedented levels during these two months, achieving a peak of 14.8% in April. Furthermore, based on the Chicago Fed National Financial Conditions Index (NFCI), these two months exhibited tighter than average financial conditions.[7]

We limit our sample to only those articles that contain mentions of our entities of interest: DFAST banks. We classify DFAST banks into five types:

- **Big 6** - the largest four U.S. banks in terms of asset size (BofA, WFC, Citi, JPMC) plus the two largest trading firms (GS and MS);
- **Trusts** - custodian banks which are principally involved in the trust business (BNY, NTRS, and STT);
- **Credit Card** - banks with credit cards as their primary line of business (Amex, COF, and DFS);

[6] Although the NBER identifies the beginning of the recession in the U.S. as February 2020, the pandemic was declared by the World Health Organization (WHO) in March 2020. Moreover, in February, travel, restaurants, and shops in the U.S. were functioning without restrictions as the first lock-down took place in mid-March.

[7] Positive values of the NFCI indicate financial conditions that are tighter than average, while negative values indicate financial conditions that are looser than average.

- **Regionals** - depository institutions larger than community banks but which generally operate below the state level;
- **IHC** - U.S. intermediate holding companies for foreign banks with over $50B in U.S. non-branch/agency assets.

We rely on this classification to better understand the observed network patterns, and account for heterogeneity in terms asset size and business profile. A list of these entities, along with their ticker symbols and group classifications, is provided in Table 1.

[Insert Table 1 about here]

Factiva has the capability to assign entity codes to each news article.[8] Specifically, if a news article discusses a bank in sufficient detail, Factiva matches this article to the bank's entity code. This enables us to identify relevant articles based on a bank's identity rather than a less precise text match. Factiva's text processing addresses the complication that DFAST banks are often mentioned in the news in several nominal variations, including known short names and acronyms.[9] Capital and lower case variations as well as M&A activity lead to additional name variations accounted for by Factiva's entity tagging.[10] We present the source groups and specific news outlets in the online appendix, but our range of sources is diverse and includes Dow Jones, Reuters, SNL Financial, Business Insider, WSJ, and others.

[8]Factiva leverages its own proprietary machine learning algorithm to assign entity codes to each news article. This is done in combination with manual quality control procedures.

[9]For example, American Express Company can be called American Express, Amex, AMEX, or by its ticker symbol, AXP.

[10]Since we are dealing with bank holding companies (BHCs), we have to consider their subsidiaries. Dow Jones integrates the Dun & Bradstreet corporate family tree such that news about BHCs' subsidiaries would also be captured as news about the parent company. In the case of mergers, the news articles are backdated to be captured at the current resulting merged entity (e.g., Truist would capture historical BB&T and SunTrust news).

4 Network Methodology

4.1 Text-to-Network

This section outlines the text-to-network approach we employ to analyze financial news data. Our approach is based on the concept of co-occurrence which involves identifying pairs of entities that appear together in the same text. In our setup, two banks are considered to co-occur if their names appear in the same news article (i.e., if they are co-mentioned). We treat each news article as a separate observation. Mathematically, the co-occurrence matrix is defined by:

$$C = M^T \times M \tag{1}$$

where M is the term-document matrix in which the rows correspond to documents and the columns correspond to list of terms. The elements of a term-document matrix corresponds to the number of times a term appears within an article.

We cap the elements of the term-document matrix at 1 to avoid over-weighting due to writer bias. For example, if "Bank of America" is mentioned in an article once while "Wells Fargo" is mentioned 5 times, this would be counted as one co-occurrence instance. Below, we provide an illustration of the process using three quotes from our data:

- Article 1: JP Morgan has hired a financial advisor from UBS for its private banking business in Miami.
- Article 2: Bank of America, JPMorgan Chase, Citigroup, Wells Fargo and US Bancorp have set aside an additional 35$ billion during the first quarter to cushion against loans that go bust, according to a tally by Edward Jones.
- Article 3: In Q3, WFC was fined \$1.5B for its dealings with JPMC. WFC plans to appeal.

Then, the term-document matrix can be represented as:

	JPM	UBS	BofA	C	WFC	USB
Article 1	1	1	0	0	0	0
Article 2	1	0	1	1	1	1
Article 3	1	0	0	0	1	0

For Wells Fargo, the two occurrences in Article 3 are capped at 1 to avoid over-weighting due to writer bias. Using Equation 1, the co-occurrence matrix is then:

	JPM	UBS	BofA	C	WFC	USB
JPM	3	1	1	1	2	1
UBS	1	1	0	0	0	0
BofA	1	0	1	1	1	1
C	1	0	1	1	1	1
WFC	2	0	1	1	2	1
USB	1	0	1	1	1	1

The interpretation of this matrix from a network perspective is the following. Each DFAST bank listed in the matrix corresponds to a node, and each non-zero cell $\{i, j\}$ entry indicates the number of articles in which banks i, j. The value in cell $\{i, j\}$ represents the weight of the link between i and j. A zero entry means that the two banks are not connected (i.e, do not co-occur). For instance, cell $\{5, 1\}$ in the co-occurrence matrix indicates that WFC and JPM co-occur 2 times (i.e., are mentioned together in 2 articles). Links are weighted using the co-occurrence count to account for variation across binary relations and provide a sense of the relative importance of the observed relationships. We aggregate co-occurrences over time to compute links for a period of interest (e.g., a week or month), so a value of 5 in a cell would represent five co-occurrences between those banks during the time period. It is important to note that there is no direction in the co-occurrence relationship, which means that the network matrix is symmetric (i.e, undirected).[11]

In the next subsection, we discuss how we analyze the co-occurrence matrices to study the network topology.

4.2 Network Analysis

Using the text-to-network methodology, we build co-occurrence network matrices at a weekly frequency. We use a weekly frequency since news tend to preserve some relationship at that frequency.[12] As explained in the previous Section, we use non-zero co-occurrences to represent the links between every bank pair and the corresponding co-occurrence value to measure the importance of the connection (i.e., network weights). We conduct our analysis in several steps.

We start by looking at our co-occurrence measure across time by bank type to broadly understand the behavior of co-occurrence through time. Next, since DFAST networks have not been previously studied in the literature, we study the network topology during normal times. Focusing first on normal times allows us to better understand our baseline. Finally, we investigate the changes to the network topology that the pandemic stress may have introduced by comparing "normal times" and "peak of stress" networks.

We extract further information contained in the co-occurrence matrices by calculating the nodes' eigenvector centrality. Eigenvector centrality generalizes degree centrality (i.e., number of connections) by also taking into account the quality of connections (i.e., a node is more central if it is connected to well-connected nodes). We argue that by focusing on direct and indirect connections, eigenvector centrality is a good centrality measure to proxy for the systemic importance of a node (i.e., a bank) in a given network.

Formally, the eigenvector centrality of a node i is proportional to the sum of the centrality

[11] There is no obvious or natural way to add direction to the links in these networks. Working with directed co-occurrence is left for future research.

[12] We also investigate using monthly co-occurrence in the Robustness section.

of all the nodes that i is connected to:

$$C_i = \frac{1}{\lambda} \sum_{j:j \neq i} a_{i,j} C_j \tag{2}$$

where C_i (C_j) is the eigenvector centrality of node i (j); $j \in \mathcal{M}(i)$, $\mathcal{M}(i)$ is the set of neighbors of i, $a_{i,j}$ is the weight (i.e., co-occurrence) of node i and j with $a_{i,j} = 0$ if i and j are not connected; and λ is a constant.[13] We use eigenvector centrality primarily to study the systemic ranking among DFAST banks and to compare it to other known measures of systemic risk importance (see Section 6). Additionally, we incorporate centrality into our visualizations (i.e., node size) to allow for further analyses of centrality and connectivity, both on system and individual bank levels.

5 Network Topology Results

In this section, we first discuss the economic meaning of connections and then present the results on our network topology analysis. Subsection 5.1 explores the question of what does it mean that two banks are mentioned in an article from an economic point of view. Subsection 5.2 describes the temporal progression of bank co-occurrences and emphasizes the importance of visualizing such data. Subsection 5.3 discusses the narrative associated with these networks.

5.1 The Economic Meaning of Co-mentions

What does it mean that two banks are mentioned in an article? Financial news reflects market dynamics and information flows. The fact that two banks are mentioned together in financial news suggests some level of interdependence or correlation between their activities, even if the correlation is inverse. The co-occurrence of banks in news articles is therefore

[13] Note that λ is the greatest eigenvalue of the associated eigenvector equation.

indicative of shared economic factors or events that influence their operations.

Our notion of interconnectedness is similar in spirit to those based on market data correlations but with the advantage of providing narratives to alleviate the black-box critique. Indeed, interconnectedness measures based on market data correlations deliver links among two banks based on the existence of some association between them (positive or negative). Similarly, two banks are mentioned in the same article when there is some interdependence (i.e., positive or negative association) between them that engenders a comparison in the financial news. While there is no direct explanation from market data, besides correlations, on why two banks are connected, there is a rich narrative that can be extracted from text data to understand the links derived in text-based networks. This is an important distinction because the additional dimension provided by text-data can fill knowledge gaps and help improve the economic intuition behind bank networks.

The recent liquidity crisis in March 2023 provides clear examples to illustrate the economic meaning of co-mentions. For instance, during the Silicon Valley Bank (SVB) collapse, Signature Bank was prominently mentioned in relation to SVB across most news and social media outlets. This was in part attributable to their common exposure to the crypto sector as well as significant levels of uninsured deposits in their balance sheets. The example of SVB and Signature Bank illustrates that, under stress conditions, a high co-mention value between two institutions can be linked to real economic consequences.

A different example involves an article mentioning SVB and some of the largest banks (e.g., JP Morgan, Bank of America, etc.). A priori, we may not think this link has much economic meaning, or perhaps less than the link between SVB and Signature Bank. To that point, the relative size (i.e., frequency) of co-occurrence between institutions provides a metric for weighing the relative importance of news-derived connections, and, intuitively, higher co-occurrence values indicate connections with more economic significance. In other words, we expect that more meaningful relationships between institutions will be more prevalent in the news. Therefore, links with low co-mention values will lose importance in the aggregate, analogous to a one-off stock market movement. In summary, while similar in spirit to market correlation-based interconnectedness measures, we believe our measure has the advantage of carrying a more direct economic interpretation.

Our measure also has advantages over those that directly report links in that we pull from a more diverse dataset and can therefore draw conclusions based on a wider array of connections. While interconnectedness measures based on specific data segments, such as loan data, are easy to interpret because, by construction, they capture a known and concrete economic relationship (e.g., a link represents a lending relationship), they only provide insight into a narrow slice of the complex relationships among financial entities. Our news coverage-based measure is not subject to this data limitation precisely because it does not limit ex-ante the nature of the links to a specific transaction, relationship, and/or aspect of interconnectedness. This is an important distinction as co-mention networks could also help highlight the relationships which are most publicly visible and important at any given time.

Furthermore, a researcher analyzing co-mention networks at a given point in time is equipped with a corpus of text (i.e., the news article), which allows her to build a clear narrative surrounding each connection. Based on our manual reading exercise, for high co-mention links, one could easily grasp the larger narrative driving the connection by reading a small sample of articles. In addition, a researcher could potentially automate this "manual reading" analysis by applying additional NLP techniques, such as topic analysis, on the text of interest. While we do not explore topic analysis in conjunction with building text-based networks, this is a potential expansion of the methodology presented in this paper.

5.2 Visualizing Network Topology

Monitoring and analyzing bank interconnectedness data is a challenging task given its multi-dimensional nature. One approach to overcoming this challenge is to use data visualizations that aid in understanding complex data and identifying patterns and trends. For example, a study by Sarlin (2016) highlights the value of effective visualizations as a macroprudential tool for detecting potential financial instability and identifying systemic risks. Flood, Lemieux, Varga, and Wong (2016) argue that leveraging visual analytics for dynamic and heterogeneous data helps increase supervisors' comprehension of such data, transforming it into actionable knowledge to support informed decision-making.

In the spirit of this research, we provide network topology plots to represent complex interconnections among banks in a clear and intuitive way. Network topology visualizations allow us to identify patterns of interconnectivity and dependencies among financial entities more easily. This idea is largely based on "cognitive load" theory (Sweller (2011)), which suggests that our working memory has a limited capacity and that visualizations can help reduce cognitive load by presenting information in a more structured and organized way. Visualizations make it easier for us to process and analyze the information, helping us identify patterns and relationships that might be harder to discern in a purely textual or numerical representation.

We first examine the average number of bank co-occurences over time for the entire sample. Figure 1 shows the average number of bank co-occurrences over time by bank type. The large peaks correspond to quarterly earnings releases, which suggests that banks are discussed more heavily together during earnings releases. Earnings release months follow a clear pattern of focus on different types of banks in different weeks, which explains why the observed peaks follow a clear cyclical pattern each quarter.

We next move to visualizing the network topology. The cyclical earnings releases are very apparent in the visualizations of the network across each quarter in our sample.[14] We use this feature to compare the pre-crisis and crisis periods, attributing differences to COVID-19 stress based on the stable pattern of connectivity in earning release periods.

We compare the pre-crisis and crisis periods using the January 2020 and April 2020 earning release weeks, respectively. We choose April 2020 as it corresponds to the peak of

[14]Weekly network graphs for the full sample are available upon request.

the pandemic.[15] We focus on the second week of the month during earnings releases, which is the most important week for earnings and has the highest co-occurrences and discussions.[16] We use this week to improve the comparison between pre-crisis and crisis periods and to capture the peak of stress in April.

We examine the network topology during normal times (January 2020) in Figure 2, Panel A. The visualization shows clusters of similar types of banks (e.g. Big 6, Trusts, Credit Cards) and a core-periphery topology where larger banks are more central.[17] There is also significant connectivity across bank types (gray links):

[Insert Figure 2 about here]

Figure 2, Panel B adds the co-occurrence measure to connections (shown with color and thickness) to demonstrate their importance. Banks are sorted by bank-type and asset size for readability. Overall, we observe that a large number of weak connections (i.e., 1-2 co-mentions) interconnect the whole network. Connections with medium weight levels either involve one of the Big 6 or take place within a given bank-type. The strongest links are observed amongst the Big 6 with co-mentions above 41.[18]

In Figure 3, shows the network graphs for the peak of stress in April 2020. We observe stronger ties across DFAST bank networks during periods of stress, which is consistent with previous literature. However, the broader topology remains similar to normal times, with the largest banks still at the center of the network. IHCs appear less interconnected in April compared to January, likely due to their lesser role in supporting the US economy during the crisis.[19]

[15]Of the three possible earning release months (July 2019, October 2019, and January 2020) for the pre-pandemic period, we discard July 2019 due to potential confounding from the release of DFAST Stress Testing results around that time.

[16]The vast majority of U.S.-based banks (including all *Big 6* banks) release their earnings in the second week of the month (around the 15^{th}).

[17]It is worth noticing that NTRS releases its earnings a week later. Therefore, some additional connections might be revealed for the Trusts if one were to include the third week of earnings releases.

[18]Note that the co-mention buckets in the figure are based on the entire sample quantiles. As a result, the maximum number of co-mentions in a given week might be below the upper bound of 415.

[19]This can also be seen in the heatmap comparison in Figure 4, which is particularly useful to quickly identify clusters of banks.

[Insert Figures 3 & 4 about here]

In Table 2, we show the difference in the number of connections and co-occurrences by bank-type during the COVID-19 crisis. While the total number of connections decreased (-18.1%), the number of co-occurrences increased (1.7%). This was mainly driven by decreased connectivity within IHCs and between Trusts and non-Trusts, compensated by an increase in connectivity from Regionals, which played a key role in facilitating PPP loans. The Big 6 continued to be fully interconnected, and the strength of connections increased (16.7%), while connectivity from Big 6 to the rest of the network remained strong (only a slight decline of 6 connections).

[Insert Table 2 about here]

During stress, the "clustering coefficient" increases by 10%, and the average path length is shorter (1.41) than pre-crisis (1.51) week, indicating tighter connections.[20,21] These results are consistent with the literature.

Overall, the stressed network exhibits a more pronounced core-periphery topology than in January, with the Big 6 banks at the core. Though the number of connections is slightly lower, the crisis period emphasizes the vital role played by Regional banks in supporting the US economy and underscores the importance of the Big 6 banks.

5.3 Network Topology Narrative

One advantage of deriving network connections from text is that we can leverage the text narrative to better understand the underlying connections being captured. To that end, we

[20]The "clustering coefficient" is also known as transitivity or the probability that adjacent vertices of a vertex (i.e., triangles) are connected. It is calculated as ${c^w}_i = \frac{1}{s_i(k_i-1)} \sum_{j,h} \frac{w_{ij}+w_{ih}}{2} a_{ij}a_{ih}a_{jh}$ where s_i is the strength of vertex i (i.e., sum of edge weights of adjacent edges to vertex i), a_{ij} are elements of the adjacency matrix, k_i is the vertex degree, w_{ij} are the weights (adapted from Barrat, Barthelemy, Pastor-Satorras, and Vespignani (2004))

[21]The average path length is defined as the mean number of shortest paths (or going from one node to another) between all nodes in the network.

manually review a 100 articles from January and April (17.4% and 15.3% of total articles in each week, respectively) to understand the underlying connections in the network. The observed links can be classified into two categories: Earnings news and pandemic news, which is in line with our prior as both weeks are earnings release weeks with the addition of the pandemic environment in April.

Earnings news dominated the headlines in both weeks, but there were distinct differences between the pre- and during-pandemic seasons. In January, articles discussed the Fed's rate hikes, US/China trade tensions, and the Big 6's re-orientation toward wealth management, as well as conventional metrics for measuring bank performance. In April, nearly 85% of the sample mentioned COVID-19, with focus on banks' responses, the Federal Reserve's steps to bolster the economy, and the Paycheck Protection Program (PPP).

In April, there is significantly more news regarding loan issuance, loan relief, and fundraising in the form of drawing corporate and consumer credit lines and new rounds of public stock offerings. The major airlines, Airbnb, Marriott, and even the state of Rhode Island turned to banks to either provide or help them raise capital at the height of the pandemic. This activity indicates a deepening of activity between banks and with their partners and clients. Also, consistent with findings in earlier sections, we note that Regionals are discussed more frequently and in a more substantive manner in the pandemic period, with attention paid to their relative importance and success in administering PPP loans.[22]

Our reading exercise suggests that news coverage based networks contain valuable information that deepens our understanding of bank connectivity. It not only portrays the public view on such connectivity, but also, by aggregating information from a rich set of news sources, brings unknown interconnectedness information to surface.

[22]Performing topic analysis could be a natural way to automate some of the information found during manual reading. We leave this potential expansion of the methodology for future research.

6 Systemic Risk and Network Centrality

This section establishes eigenvector centrality as a useful measure for ranking the systemic importance of banks. By incorporating indirect connections in addition to direct ones, eigenvector centrality offers a more nuanced understanding of systemic risk. To explore its value, we compare our text-based eigenvector centrality measure to three traditional systemic risk measures (i.e., DIP, SRISK, and CoVaR). In the following subsections, we aim to answer two main questions: (1) Do these measures rank systemic importance of financial institutions similarly? and (2) Do these measures capture similar sources of information?

In order to construct our text-based eigenvector centrality measure, we first focus on news articles that convey negative sentiment. We postulate that by subsetting our sample of articles in this way, we are able to derive the elements of risk embedded in those articles in a more robust fashion.[23]

To isolate news articles with negative sentiment, we conduct sentiment analysis on our body of texts. Sentiment analysis, in short, is a computational treatment of opinions expressed in written texts. We rely on dictionary-based text analysis in combination with polarity scoring using Loughran and McDonald (2011)'s financial sentiment dictionary. We also take into account valence shifters.[24] We measure overall sentiment on the sentence level, as opposed to word level, as this allows us a higher degree of precision. This approach includes incorporating features such as negation (e.g., "good" to "not good"), amplification/de-amplification (e.g., "very" to "barely"), and adversative conjunctions (e.g., "not too smart").[25] It's been shown that extending dictionary methods with contextual va-

[23]Box-plots of negative sentiment eigenvector centrality by bank, for both pre-pandemic and pandemic periods, are provided in the online appendix.

[24]Valence shifters are supplemental words that can enhance, subdue, or flip the degree of positivity or negativity within the given unit of text.

[25]We consider each article as a unit of observation and decompose each article into a set of sentences. We then run our algorithm to derive a sentiment score for each sentence in that article by considering both the presence of positive and negative terms as well as amplifiers and negations in the neighborhood of each relevant term. Once we assign a sentiment score to each sentence, we aggregate the sentiment score to the article level.

lence shifters improves the accuracy of the classification (e.g., Kennedy and Inkpen (2006)).

6.1 Traditional Systemic Risk Measures

To facilitate the analysis that follows, we briefly discuss each of the four systemic risk measures involved.[26] Eigenvector centrality (Eigen) measures a firm's network importance using financial news text and captures both traditional and soft financial data. Distress Insurance Premium (DIP) measures the hypothetical insurance premium against systemic financial distress based on size, default probability, and asset return correlations. Both Eigen and DIP are relative to a given portfolio of firms, while SRISK and CoVaR focus on broader market factors. SRISK measures a bank's systemic vulnerability as expected capital shortfall conditional on a large market downturn based on size, leverage, and risk; and CoVaR measures spillovers to the whole financial network from a distressed bank using stock returns.

We use readily available data for DIP, SRISK, and CoVaR from the Research and Statistics Department at the Board of Governors for the 12 Large Institution Supervision Coordinating Committee (LISCC) firms as of 2020.[27]. These firms are a subset of the DFAST banks and correspond to large financial institutions that pose the greatest risk to U.S. financial stability.[28] The data is available on daily frequency, and we compute weekly averages to match the frequency of eigenvector centrality. The time period expands 1-year from October 2019 to September 2020.

These measures capture different aspects of the relationship between institution-level and systemic risk. SRISK and CoVaR focus on broad market-wide risk factors, while eigenvector

[26]See Huang, Zhou, and Zhu (2009, 2012), Brownlees and Engle (2017), and Adrian and Brunnermeier (2016) for further details on DIP, SRISK, and CoVaR, respectively.

[27]The LISCC portfolio includes 8 U.S. banks (BofA, Citi, JPMC, WFC, GS, MS, BNY, and STT) and 4 IHCs (BCS, CS, DB, and UBS). As of 2021, the IHCs are no longer part of the LISCC portfolio.

[28]The Federal Reserve created the Large Institution Supervision Coordinating Committee (LISCC) supervisory program in 2010 to coordinate its supervisory oversight of these systemically important firms. LISCC firms include (i) any firm subject to Category I standards under the regulatory tailoring framework, (ii) any non-commercial, non-insurance savings and loan holding company that would be identified for Category I standards if it were a bank holding company, and (iii) any nonbank financial institution designated as systemically important by the FSOC.

centrality focuses on relationships between a small number of banks. We expect DIP to be closer to eigenvector centrality since it tries to measure the cost of insuring tail risk for a portfolio of a small number of institutions.

6.2 Systemic Risk Rankings

We first look at how these four measures rank the twelve LISCC firms from 1 to 12 using monthly averages of each risk measure (with 1 being "most systemically important" and 12 being "least systemically important"). We use monthly averages instead of weekly to acknowledge that weekly measures are often volatile and can lead to noisy rankings.

Figure 5 showcases how the Big 6 firms rank among the 12 LISCC firms by measure and across time.[29] Since the Big 6 are arguably the most systemically important firms, focusing on them has the advantage of helping us understand if the rankings appear intuitive. In particular, we expect the Big 6 to occupy the rankings between 1-6 among the 12 LISCC firms if the measures are working as intended.

[Insert Figure 5 about here]

Based on these ranking we see that SRISK and CoVaR, both of which rely on market data, are much more volatile relative to the other two measures, even with monthly smoothing. In fact, SRISK has received criticism as a policy-making tool for its dependence on market measures which makes it reactive to current market analysis rather than underlying systemic risk (see Danielsson et al. (2016b); Danielsson et al. (2016a)).[30] A clear example of this shortcoming in SRISK is the ranking of JPMC in Figure 5. JPMC, being arguably one of the most systemically important firms, is ranked last (out of the 12 LISCC firms) according to SRISK prior to the pandemic lock-down. Then, as the pandemic unfolds, JPMC is ranked

[29]We omit the results for the remaining LISCC firms for graph clarity given the graph gets quickly crowded when adding more firms.

[30]At the same time, Nucera et al. (2016) notes that SRISK is actually one of the more stable risk measures compared to VaR, ΔCoVaR, and MES because it is based on book value rather than price.

first when the peak of stress hits in April 2020. We clearly see that as the stress gets some relief, JPMC ranking goes up again. A similar shortcoming is observed in CoVaR (e.g., WFC), which is also market based.

6.3 Principal Component Analysis of Systemic Risk Measures

To understand how these four measures relate to each other, we first look at raw correlations across these four measures in Table 3. Eigen and DIP are the most correlated (0.61), while CoVaR is correlated near-zero or negatively with all three other measures and is especially negatively associated with SRISK (-0.35, -0.06, and 0.01 with SRISK, Eigen, and DIP respectively). Also, Eigen's correlation with SRISK is small but non-negligible (0.16).[31]

[Insert Table 3 about here]

Next, we perform a Principal Component Analysis (PCA) in order to investigate whether the four systemic risk measures capture similar sources of information or variation. The results are shown in Table 4. The three first principal components are able to explain 92% of the variation, with PC1, PC2, and PC3 explaining 46%, 30%, and 16% respectively. Notice that PC2 and PC3 help explain almost half of the variation, suggesting that all three PCs should be taken into consideration. Eigen and DIP load more heavily in the first principal component (-0.57 and -0.62), while SRISK and CoVaR load relatively more heavily in the second (-0.47 and +0.74) and third component (+0.65 and +0.58).

This is consistent with the observed patterns in the raw correlations. Note also that Eigen's contribution to the second and third principal component is sizeable (+0.37 and -0.47) and that Eigen and DIP diverge in the third principal component. Furthermore, none of the measures clearly dominate a given principal component suggesting that each of these measures carry important pieces of information that complement each other.

[Insert Table 4 about here]

[31] Results are qualitatively similar using a rank correlation.

Figure 6, a PCA bi-plot, provides a clear visualization of these patterns. Consistent with our prior, we see that Eigen and DIP tend to be the closest (except for PC3, where they diverge), yet distinct enough. Overall, these figures suggest more clearly that the four measures are complement of each other rather than substitutes.

[Insert Figure 6 about here]

6.4 Eigenvector Centrality and Firm's Fundamentals

Finally, to further assess Eigen as a potential measure of systemic risk, we perform a regression analysis that relates Eigen to firms' fundamentals and performance metrics. This helps us understand what sources of variation eigenvector centrality is primarily correlated with. We include variables that are known drivers of traditional systemic risk measures like downside risk and size, but also variables that relate to the business mix of a firm. In particular, we include variables related to a significant presence in trading activities (e.g., large broker-dealer arm) captured through the size of a firm's trading assets as well as their trading income.[32]

This analysis is based on all DFAST banks except BNP Paribas due to data limitations. As explanatory variables in our baseline we include: Log of Total Assets (log of TA), log of Trading Assets, Return on Assets (ROA), firm level Value at Risk (VaR). In addition, we include several sources of income (i.e., interest income (II), non-interest income (NII) excluding trading income, and trading income) as a ratio to total assets and a dummy equal 1 for the top 5 largest trading firms (i.e., BofA, Citi, JPMC, GS, and MS).[33]

All variables (except VaR) are sourced from the FR Y-9C and are quarterly. VaR is also quarterly and represents the 95% Value at Risk derived from the distribution of firms' daily

[32]Arguably, some lines of business, like trading, are more important from a systemic point of view and, thus, it is useful to account for the business mix when studying a firm's systemic importance.We focus on the business mix as opposed to the capital structure since a firm can make strategic decisions over the former while the latter is largely driven by regulations.

[33]Note that WFC is not included given its trading activities are not at par with the top 5 firms.

stock returns from CRSP. Since all firm fundamentals and performance metrics are available at a quarterly frequency, we average Eigen at the quarterly level for consistency.[34] Our main regression specification is:

$$Eigen_{i,t} = log(TotalAssets)_{i,t} + log(TradingAssets)_{i,t} + ROA_{i,t} + VaR_{i,t} + TradingIncome/TA_{i,t} + NII/TA_{i,t} + II/TA_{i,t} + \epsilon_{i,t} \tag{3}$$

where i identifies each DFAST bank in our sample and t captures our time component.

Table 5 presents the results of our regression models. Unsurprisingly, firms' asset and trading asset size are the most correlated with Eigen. Both are positive and significant across all specifications. This suggests that larger firms are associated with larger values of eigenvector centrality and that firms with higher trading assets are also more central. In specification (2), we add several sources of income as a ratio to total assets to further understand how the business mix relates to eigenvector centrality.

[Insert Table 5 about here]

We observe a large and highly significant (1% level) positive coefficient for the trading income ratio, while the remaining non-interest income (i.e., excluding trading income) is negative and of smaller magnitude (5% level significant). The interest income ratio coefficient is also small in magnitude yet positive (5% level significant). This suggests that firms with larger trading and interest income ratios tend to have higher centrality values, while non-trading, non-interest income firms have lower centrality values, indicating that strong trading or certain interest income activities make firms more interconnected and central, while non-trading non-interest income is less likely to have systemic implications.

In specification (3), we add our top 5 largest trading firms dummy to further explore the observation on trading. When adding this dummy, the trading income ratio is no longer

[34] We also tried including the income variables as a percentage of total income as opposed to total assets. Results were qualitatively similar.

significant, suggesting that the size of this activity is of particular relevance and that the large positive effect was driven by the largest trading firms.

In summary, these fundamentals and performance metrics jointly account for about 72% to 79% of the variation in Eigen. The key takeaway from Table 5 is that Eigen is significantly correlated with firms' size and their business mix, with trading activities being particularly relevant when sizeable. This further suggests that Eigen could be a useful complement to existing traditional systemic risk measures as it could helps capture systemic risk implications from the business mix.

7 Eigenvector Centrality and Returns

To further demonstrate the potential usefulness of the measure, we conduct a regression analysis exercise that models weekly Cumulative Abnormal Returns using eigenvector centrality and other firm fundamentals. While we do not claim causation, we are not the first to link measures of systemic risk to cumulative abnormal returns. Sabri, Gilder, and Onali (2019) study market reactions to FOMC decisions about the Fed funds target rate and conclude that SRISK tends to increase abnormal returns as rates fall and vice versa. Cotter et al. (2018) study abnormal returns in the wake of firm decisions to initiate or omit dividends and include change in systemic risk (ΔSRISK) in their regression specifications. We also leverage work done by Hoffman (2010) and Cooper, Gulen, and Schill (2008) when developing our regression specifications.

Specifically, we model CAR (aggregated weekly by BHC) against BHC quarterly asset growth, weekly market capitalization scaled by assets, weekly retained earnings (RE) scaled by assets, quarterly Return on Assets (ROA), weekly stock price, weekly share of trading volume, and the Chicago Board of Options Exchange's measure of market volatility (VIX, weekly). Our data are sourced from CRSP, the FR Y-9C, and the Chicago Board of Options Exchange. The time period covers our full sample, and our main regression specification is give by:

$$CAR_{i,t} = \alpha_{i,t} + Eigen_{i,t} + AssetGrowth_{i,t} + MarketCap_{i,t} + RE_{i,t} \tag{4}$$
$$+ ROA_{i,t} + Price_{i,t} + Vol_{i,t} + VIX\_h_{i,t}$$

where i identifies each DFAST bank in our sample, and t captures our time component. Eigen captures our weekly measure of bank eigenvector centrality, and the other variables are as defined in the preceding paragraph. We believe a weekly CAR is an appropriate level of temporal aggregation for our analysis given that our primary explanatory variable, Eigen, while a measure of systemic risk, is derived from news, and therefore its effects are contemporaneously reflected in asset prices pursuant to the efficient market hypothesis.

We expect Eigen to have the largest impact on CAR at the tails of the distribution, and as such, we run quantile regressions at the 10th and 90th percentiles. Our results are presented in Tables 6 and 7.

[Insert Tables 6 & 7 about here]

We find that Eigen is correlated with CAR at the 10th percentile (negative abnormal returns) but not at the 90th percentile (positive abnormal returns). More specifically, Eigen is positively correlated with CAR and significant at the 5% level in the lower tail of the distribution, after controlling for firm fundamentals. In contrast, Eigen is not significant in any specification for the upper tail of the CAR distribution.[35] This result makes sense if we think of Eigen as a measure of risk and that, as such, it is more likely to help explain negative abnormal returns than positive ones.

Taking into account that at the lower tail CAR is negative, an increase in Eigen reduces CAR towards 0. This is in line with a priori expectations. Given that our data sample includes only the largest financial institutions, the argument can be made that, even with

[35]This results also holds for Eigen based on all financial news as opposed to only negative sentiment news. Results are available in the online appendix.

investors' expectations that firms are underwritten by the U.S. government, the most highly connected and largest of the financial institutions in our sample cannot escape from prevailing market trends nor the expectation that larger and more capitalized firms experience lower cumulative abnormal returns. More broadly, our results suggest that Eigen could be a useful explanatory variable in studies focused on popular financial variables such as CAR.

8 Robustness Checks

This section presents robustness checks to two methodology choices that impact the network topology. In subsection 8.1, we use monthly instead of weekly frequency for our co-occurrence measure to check whether the key topology features hold. In subsection 8.2, we limit our news sources to only Reuters, a news source that is typically used in the finance literature, and again check whether the observed patterns in our network topology are preserved.

8.1 Monthly vs Weekly Co-occurrence

In our comparison of pre-crisis (January 2020) vs crisis (April 2020) network matrices, we use the second week of the earnings release month to perform the analysis since in that week most firms release their earnings. Also, the second week of April coincides with the peak of COVID-19 induced financial stress, making it an ideal candidate for our exercise. However, given that some of the Trusts release their earnings in the third week of the month and the IHCs in the fourth week, as a robustness check, we repeat the analysis constructing co-occurrence matrices for the full months.

The monthly co-occurrence network graphs for the comparison of January vs April 2020 earnings releases are available in the online appendix. While by construction the figures based upon monthly co-occurrence matrices have a higher number of co-occurrences in comparison to the weekly ones, the general pattern is preserved. Overall, we find our main conclusions to be robust to the frequency change.

8.2 Co-occurrence Using Select Publications: Reuters

In this subsection, we rerun our analysis using one of the most popular sources of news employed in the finance literature: Reuters. Using Reuters instead of Factiva Analytics' top 10 sources for financial news produces a less dense network.[36] This is expected since the amount of news articles and, thus, the potential number of co-occurrences are much smaller when using only one source. Key clusters of connections are captured, but cross-cluster connectivity is limited. There is also less variation in co-occurrences, potentially indicating a single source bias.

Overall, this robustness exercise suggests that using Factiva Analytics has the advantage of providing a richer mix of sources, thus allowing us to capture a wider range of connections and delivering more stable eigenvector centrality based systemic risk rankings.[37]

9 Conclusion

In this paper, we investigate the interconnectedness of U.S. based bank holding companies by analyzing their co-mentions in financial news articles. News data provide a timely and relevant source of information on financial risks and stress factors that is not represented in structured numerical data, making news a valuable source of soft information for analysis.

We show that these text-based bank networks exhibit a core-periphery topology and network ties becomes stronger during times of financial stress. We then propose an alternative systemic risk measure called eigenvector centrality. We find it to be more robust in ranking banks according to their systemic importance during normal and stress periods, as demonstrated by a comparison with popular systemic risk measures, using the COVID-19 pandemic shock as an exogenous period of stress.

Importantly, we show that our proposed measure captures soft information that is not

[36]The heatmap and circle graphs based on this analysis are avalilable in the online appendix.
[37]Additional graphs using only Reuters news are available upon request.

contained in numerical data and is correlated with a firm's size and business mix, with trading activities being particularly relevant when sizable. Additionally, our approach allows for frequent and granular updating of both the network topology and the systemic risk measure, and provides a narrative that can be used to understand observed movements in systemic risk.

Overall, our contributions offer a rich real-time analysis of the financial system's architecture which has important implications for policymakers and regulators. By introducing an innovative methodology and demonstrating its usefulness in capturing systemic risk, our study provides valuable insights into financial risk management and decision-making.

References

Adrian, T. and M. K. Brunnermeier (2016): "CoVaR," *The American Economic Review*, 106, 1705.

Baker, S. R., N. Bloom, and S. J. Davis (2016): "Measuring economic policy uncertainty," *The Quarterly Journal of Economics*, 131, 1593–1636.

Baker, S. R., N. Bloom, S. J. Davis, K. Kost, M. Sammon, and T. Viratyosin (2020): "The Unprecedented Stock Market Reaction to COVID-19," *The Review of Asset Pricing Studies*, 10, 742–758.

Barigozzi, M. and C. Brownlees (2019): "Nets: Network estimation for time series," *Journal of Applied Econometrics*, 34, 347–364.

Barrat, A., M. Barthelemy, R. Pastor-Satorras, and A. Vespignani (2004): "The architecture of complex weighted networks," *Proceedings of the national academy of sciences*, 3747–3752.

Billio, M., M. Getmansky, A. W. Lo, and L. Pelizzon (2012): "Econometric measures of connectedness and systemic risk in the finance and insurance sectors," *Journal of Financial Economics*, 104, 535–559.

Boudoukh, J., R. Feldman, S. Kogan, and M. Richardson (2013): "Which news moves stock prices?" *National Bureau of Economic Research*.

Brownlees, C. and R. F. Engle (2017): "SRISK: A Conditional Capital Shortfall Measure of Systemic Risk," *The Review of Financial Studies*, 30, 48–79.

Brunetti, C., J. Harris, S. Mankad, and G. Michailids (2019): "Interconnectedness in the interbank market," *Journal of Financial Economics*, 133, 520–538.

Calomiris, C. W. and H. Mamaysky (2019): "How news and its context drive risk and returns around the world," *Journal of Financial Economics*, 133, 299–336.

Cooper, M. J., H. Gulen, and M. J. Schill (2008): "Asset Growth and the Cross-Section of Stock Returns," *The Journal of Finance*, 63, 1609–1651.

Costola, M., O. Hinz, M. Nofer, and L. Pelizzon (2023): "Machine learning sentiment analysis, COVID-19 news and stock market reactions," *Research in International Business and Finance*, 64, 101881.

Cotter, J., H. von Eije, M. Faour, and C. Muckley (2018): "Investor anticipation and the stock price reaction to dividend initiations and omissions," *Economics*.

Danielsson, J., K. James, M. Valezuela, and I. Zer (2016a): "Can we prove a bank guilty of creating systemic risk? A minority report," *Journal of Money, Credit, and Banking*, 48, 795–812.

——— (2016b): "Model Risk of Risk Models," *Journal of Financial Stability*, 23, 79–91.

Diebold, F. X. and K. Yilmaz (2014): "On the network topology of variance decompositions: Measuring the connectedness of financial firms," *Journal of Econometrics*, 182, 119–134.

Elliott, M., B. Golub, and M. O. Jackson (2014): "Financial Networks and Contagion," *American Economic Review*, 104, 3115–3153.

Flood, M. D., V. L. Lemieux, M. Varga, and B. W. Wong (2016): "The application of Visual Analytics to Financial Stability Monitoring," *Journal of Financial Stability*, 27, 180–197.

Gentzkow, M., B. Kelly, and M. Taddy (2019): "Text as Data," *Journal of Economic Literature*, 57, 535–574.

Gofman, M. (2011): "A Network-based Analysis of Over-the-counter Markets," *AFA 2012 Chicago Meetings Paper*.

Greenwood, R., L. Augustin, and D. Thesmar (2015): "Vulnerable banks," *Journal of Financial Economics*, 115, 471–485.

Hautsch, N., J. Schaumburg, and M. Schienle (2015): "Financial network systemic risk contributions," *Review of Finance*, 685–738.

Hoberg, G. and G. Phillips (2016): "Text-based network industries and endogenous product differentiation," *Journal of Political Economy*, 124, 1423–1465.

Hoffman, A. J. (2010): "Developing multivariate models to predict abnormal stock returns - Using cross-sectional differences to identify stocks with above average return expectations," in *Proceedings of the International Conference on Fuzzy Computation and 2nd International Conference on Neural Computation*, 411–419.

Huang, X., H. Zhou, and H. Zhu (2009): "A framework for assessing the systemic risk of major financial institutions," *Journal of Banking Finance*, 33, 2036–2049.

——— (2012): "Systemic risk contributions," *Journal of Financial Services Research*, 42, 55–83.

Härdle, W. K., W. Wang, and L. Yu (2016): "TENET: Tail-Event driven NETwork risk," *Journal of Econometrics*, 192, 499–513.

Jackson, M. O. and A. Pernoud (2021): "Systemic Risk in Financial Networks: A Survey," *Annual Review of Economics*, 13, 171–202.

Kennedy, A. and D. Inkpen (2006): "Sentiment Classification of Movie Reviews Using Contextual Valence Shifters," *Computational Intelligence*, 22, 110–125.

Loughran, T. and B. McDonald (2011): "When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks," *The Journal of Finance*, 66, 35–65.

——— (2016): "Textual Analysis in Accounting and Finance: A Survey," *Journal of Accounting Research*, 54, 1187–1230.

Mamaysky, H. (2020): "Financial markets and news about the coronavirus," *COVID Economics*, 38, 68–128.

Nucera, F., B. Schwaab, S. J. Koopman, and A. Lucas (2016): "The information in systemic risk rankings," *Journal of Empirical Finance*, 38, 461–475.

Rönnqvist, S. and P. Sarlin (2015): "Bank networks from text: interrelations, centrality and determinants," *Quantitative Finance*, 15, 1619–1635.

Sabri, A., D. Gilder, and E. Onali (2019): "Monetary Policy and Systemic Risk," *SSRN Working Paper*.

Sarlin, P. (2016): "Macroprudential oversight, risk communication and visualization," *Journal of Financial Stability*, 27, 160–179.

Schoar, A., G. M. Afonso, and A. Kovner (2014): "Trading Partners in the Interbank Lending Market," *Federal Reserve Bank of New York Staff Report*.

Shapiro, A. H., M. Sudhof, and D. Wilson (2020): "Measuring news sentiment," *Federal Reserve Bank of San Francisco*.

Sweller, J. (2011): "Cognitive Load Theory," 55, 37–76.

Tetlock, P. C. (2007): "Giving content to investor sentiment: The role of media in the stock market," *The Journal of Finance*, 62.

Tables

Table 1. List of DFAST Bank Holding Companies (BHC)
We classify DFAST banks into five types: *Big 6*, which corresponds to the largest four U.S. banks in terms of asset size (BofA, WFC, Citi, JPMC) plus the two largest trading firms (GS and MS); *Trusts*, which are custodian banks with main activity in the trust business; *Credit Card*, which correspond to banks with credit cards as their primary line of business; *Regionals*, which are depository institutions larger than community banks but which generally operate below the state level, and *IHC*, which are U.S. intermediate holding companies for foreign banks with over $50B in U.S. non-branch/agency assets.

Bank Type	N	Bank Name	Symbol
Big 6	1	Bank of America	BofA
	2	Citigroup	Citi
	3	Goldman Sachs	GS
	4	JPMorgan Chase	JPMC
	5	Morgan Stanley	MS
	6	Wells Fargo	WFC
Trusts	7	BNY Mellon	BNY
	8	Northern Trust	NTRS
	9	State Street Corp	STT
Credit Card	10	American Express	Amex
	11	Capital One	COF
	12	Discover Financial	DFS

Bank Type	N	Bank Name	Symbol
Regionals	13	Ally Financial	Ally
	14	Citizens Financial Group	CFG
	15	Fifth Third Bank	FITB
	16	Huntington Bank	HBAN
	17	KeyCorp	KEY
	18	M&T Bank	MTB
	19	PNC Group	PNC
	20	Regions Financial	RF
	21	Truist	TFC
	22	US Bancorp	USBC
IHC	23	BBVA Compass	BBVA
	24	Bank of Montreal	BMO
	25	BNP Paribas	BNP
	26	Barclays Bank	BCS
	27	Credit Suisse	CS
	28	Deutsche Bank	DB
	29	HSBC Bank	HSBC
	30	MUFG Union	MUFG
	31	Royal Bank of Canada	RBC
	32	Santander Bank	SAN
	33	TD Bank	TD
	34	UBS Group	UBS

Table 2. Summary statistics: January vs April network matrices

January 2020 earnings releases (pre-crisis) extend from January 13^{th}-19^{th}, 2020; and April 2020 earnings releases (crisis) extend from April 13^{th}-19^{th}, 2020. Connections is the number of links and co-occurrences is the number of co-mentions in articles (weight of connections). Clustering coefficient is calculated as the transitivity or connectivity of a network, and average path length is the mean shortest path between two nodes.

	Connections			**Co-occurrences**		
Type	**Jan**	**Apr**	**%Δ**	**Jan**	**Apr**	**%Δ**
Within *Big 6*	15	15	0%	504	588	16.7%
Between *Big 6* and Non-*Big 6*	103	97	-5.8%	357	327	-8.4%
Within *Regionals*	9	12	33.3%	11	49	345%
Between *Regionals* and Non-*Reg*	58	69	19.0%	103	226	119%
Within *Trusts*	1	0	-100%	3	0	-100%
Between *Trusts* and Non-*Trusts*	34	6	-82.4%	61	11	-82.0%
Within *IHC*	30	12	-60.0%	103	26	-74.8%
Between *IHC* and Non-*IHC*	77	63	-18.2%	266	186	-30.1%
Within *CC*	3	1	-66.7%	6	3	-50%
Between *CC* and Non-*CC*	32	29	9.4%	73	68	-6.8%
Within All Non-*Big 6*	184	120	-34.8%	392	320	-18.4%
Total	210	172	-18.1%	1057	1075	1.7%
Average Path Length	1.51	1.41	-6.6%			
Clustering Coefficient	0.70	0.77	10.0%			

Table 3. Systemic risk measures: Correlation matrix

This correlation is based on the 12 LISCC banks as of 2020. The analysis is performed on three traditional systemic risk measures (DIP, SRISK, and CoVaR) and our text-based eigenvector centrality from negative news. All measures are at a weekly frequency (weekly average for DIP, SRISK, and CoVaR) over a 1-year period from October 2019 to September 2020.

	Eigen	DIP	SRISK	CoVaR
Eigen	1.00	0.61	0.16	-0.06
DIP	0.61	1.00	0.34	0.01
SRISK	0.16	0.34	1.00	-0.35
CoVaR	-0.06	0.01	-0.35	1.00

Table 4. Systemic risk measures: Principal Component Analysis (PCA)

This Principal Component Analysis is based on the 12 LISCC banks as of 2020. The analysis is performed on three traditional systemic risk measures (DIP, SRISK, and CoVaR) and our text-based eigenvector centrality from negative news. All measures are at a weekly frequency (weekly average for DIP, SRISK, and CoVaR) over a 1-year period from October 2019 to September 2020.

Factor Loadings	PC1	PC2	PC3	PC4
Eigen	-0.57	0.37	-0.47	0.57
DIP	-0.62	0.31	0.15	-0.70
SRISK	-0.48	-0.47	0.65	0.36
CoVaR	0.24	0.74	0.58	0.24

Variance Explained	PC1	PC2	PC3	PC4
Standard Deviation	1.35	1.10	0.80	0.57
Proportion of Variance	0.46	0.30	0.16	0.08
Cumulative Proportion	0.46	0.76	0.92	1.00

Table 5. Eigenvector centrality and firm characteristics
This analysis is based on all DFAST banks (except BNP Paribas due to data limitations). We relate eigenvector centrality (Eigen) generated from "negative sentiment" news to salient firm fundamentals and performance metrics. This helps understand what firm characteristics could explain movements in Eigen. We include: log of Total Assets (log of TA), log of trading assets, Return on Assets (ROA), firm level 95% Value at Risk (VaR), and several sources of income (i.e., interest income, non-interest income excluding trading income, and trading income) as a ratio to total assets. The dummy "Top 5 trading firms" equals 1 for the top 5 largest trading firms (i.e., BofA, Citi, JPMC, GS, and MS). All variables are quarterly and Eigen is averaged at that frequency level for consistency. Our data are sourced from CRSP and the FR Y-9C. The time period covers our full sample.

	Eigen (quarterly)		
	(1)	(2)	(3)
Log(Total Assets)	0.1040***	0.0921***	0.0698***
	(0.0070)	(0.0084)	(0.0141)
Log(Trading Assets)	0.0030**	0.0041***	0.0044***
	(0.0012)	(0.0012)	(0.0012)
ROA	-0.0021*	-0.0022*	-0.0020
	(0.0012)	(0.0012)	(0.0013)
Value at Risk	0.1856	-0.3264	-0.2034
	(0.2001)	(0.2686)	(0.2401)
Trading income/TA		11.58***	3.241
		(2.826)	(3.767)
Non-interest income/TA		-1.626**	-1.073
		(0.7160)	(0.6692)
Interest income/TA		1.053**	1.002**
		(0.4576)	(0.4261)
Top 5 trading firms			0.1086**
			(0.0452)
Constant	-1.961***	-1.803***	-1.374***
	(0.1292)	(0.1581)	(0.2691)
Adj. R^2	0.72	0.77	0.79
Num. obs.	164	164	164

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$.
Heteroskedasticity-robust standard-errors in parentheses.

Table 6. Weekly CAR quantile regression (10^{th} percentile)
This analysis is based on all DFAST banks. We model weekly Cumulative Abnormal Returns (CAR) using eigenvector centrality (Eigen) from "negative sentiment" news and salient firm fundamentals to help explain movements in CAR. CAR and Eigen are weekly; asset growth and Return on Assets (ROA) are quarterly; market capitalization and retained earnings are weekly and scaled by asset size; stock price, trading volume (number of shares traded in billions), and the CBOE's volatility index (VIX) are weekly. Our data are sourced from CRSP, the FR Y-9C, and the Chicago Board of Options Exchange. The time period covers our full sample.

	Cumulative Abnormal Return (CAR)		
	(1)	(2)	(3)
Eigen	0.029**	0.014**	0.015**
	(0.012)	(0.006)	(0.008)
Asset Growth	−0.001**	−0.0004***	−0.0004**
	(0.0003)	(0.0001)	(0.0002)
Market Cap	0.052***	0.031***	0.026***
	(0.005)	(0.004)	(0.008)
Retained Earnings	0.022	0.005	0.008
	(0.028)	(0.020)	(0.022)
ROA	0.002***	0.0003	0.0002
	(0.0005)	(0.0004)	(0.0004)
Stock Price		0.00005**	0.00004**
		(0.00002)	(0.00002)
Trading Volume ($B)		−0.041***	−0.014***
		(0.003)	(0.005)
VIX			−0.001***
			(0.0002)
Constant	−0.065***	−0.005	−0.002
	(0.003)	(0.004)	(0.004)
Num. obs.	2,122	2,122	2,122

Note: *p<0.1; **p<0.05; ***p<0.01

Table 7. Weekly CAR quantile regression (90^{th} percentile)
This analysis is based on all DFAST banks. We model weekly Cumulative Abnormal Returns (CAR) using eigenvector centrality (Eigen) from "negative sentiment" news and salient firm fundamentals to help explain movements in CAR. CAR and Eigen are weekly; asset growth and Return on Assets (ROA) are quarterly; market capitalization and retained earnings are weekly and scaled by asset size; stock price, trading volume (number of shares traded in billions), and the CBOE's volatility index (VIX) are weekly. Our data are sourced from CRSP, the FR Y-9C, and the Chicago Board of Options Exchange. The time period covers our full sample.

	Cumulative Abnormal Return (CAR)		
	(1)	(2)	(3)
Eigen	−0.011	−0.007	−0.006
	(0.016)	(0.012)	(0.012)
Asset Growth	−0.0004	−0.0004*	−0.0004
	(0.0003)	(0.0002)	(0.0002)
Market Cap	−0.044***	−0.023**	−0.023**
	(0.009)	(0.011)	(0.011)
Retained Earnings	−0.018	−0.001	−0.0003
	(0.029)	(0.031)	(0.031)
ROA	−0.001	−0.001	−0.001
	(0.001)	(0.001)	(0.001)
Stock Price		−0.00002	−0.00002
		(0.00003)	(0.00004)
Trading Volume ($B)		0.028***	0.028***
		(0.004)	(0.007)
VIX			0.00003
			(0.0003)
Constant	0.054***	0.013**	0.013*
	(0.004)	(0.006)	(0.006)
Num. obs.	2,122	2,122	2,122

Note: *p<0.1; **p<0.05; ***p<0.01

Figures

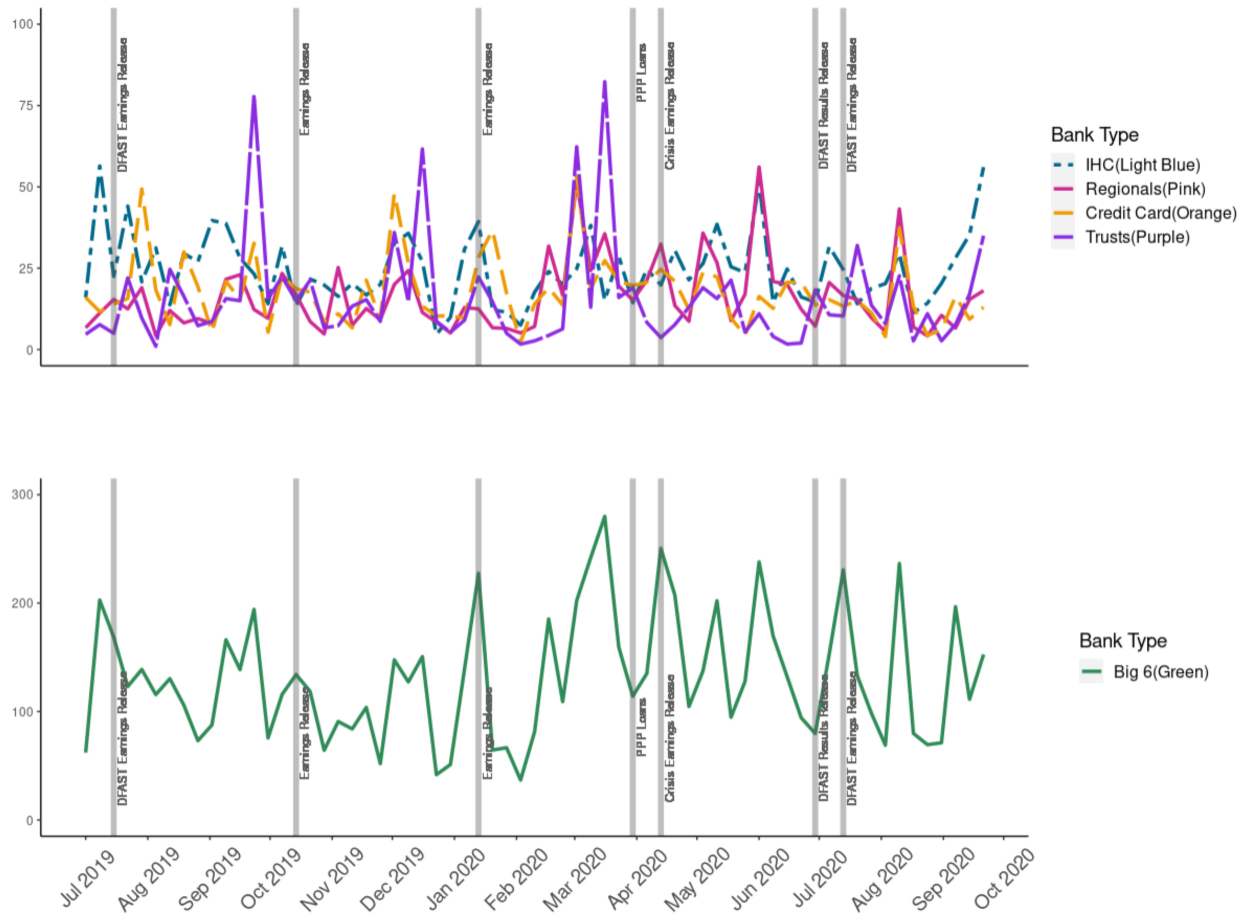

Figure 1. Time series of bank co-occurrences, by bank type

Each axis represents the weekly number of co-occurrences by bank type, from July 2019 - September 2020. *Big 6* banks are plotted in the bottom panel due to scale difference. *Big 6* corresponds to the largest four U.S. banks in terms of asset size (BofA, WFC, Citi, JPMC) plus the two largest trading firms (GS and MS); *Trusts* are custodian banks principally involved in the trust business (BNY, NTRS, and STT); *Credit Card* corresponds to banks with credit cards as their primary line of business (Amex, COF, and DFS); *Regionals* are depository institutions larger than community banks but which generally operate below the state level; and *IHC* are U.S. intermediate holding companies for foreign banks with over \$50B in U.S. non-branch/agency assets. Earnings release events are marked as follows: "DFAST earnings release" correspond to "earnings release" following the publications of DFAST results, and "Crisis earnings release" corresponds to "earnings release" that coincided with the pandemic peak of stress. "PPP Loans" mark the beginning of the program.

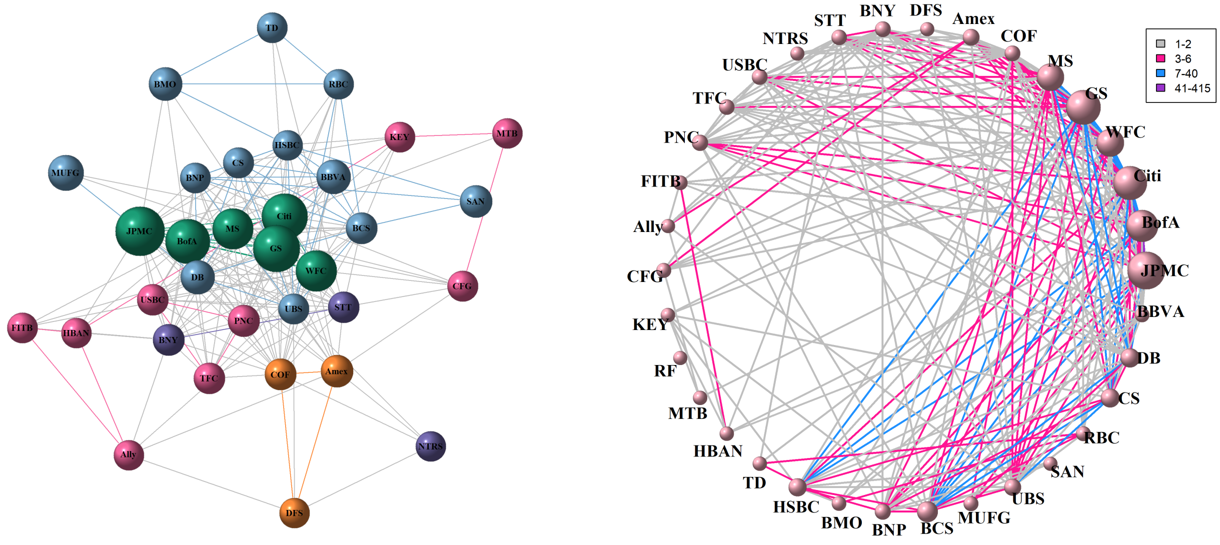

Panel A. Connections & clusters

Panel B. Co-occurrences

Figure 2. Network graphs: January 2020
Panel A: Nodes are colored by bank-type (Big 6 (green), CC (orange), Trusts (purple), Regionals (Pink), and IHCs (light blue)), and link colors correspond to within bank-type connections. Panel B: Nodes are sorted by bank-type and asset size, and link colors correspond to co-occurrence counts (i.e., connections' weights). In both panels, node size represents eigenvector centrality (i.e., larger nodes are more central). January 2020 earnings releases (pre-crisis) extends from January 13^{th}-19^{th}, 2020.

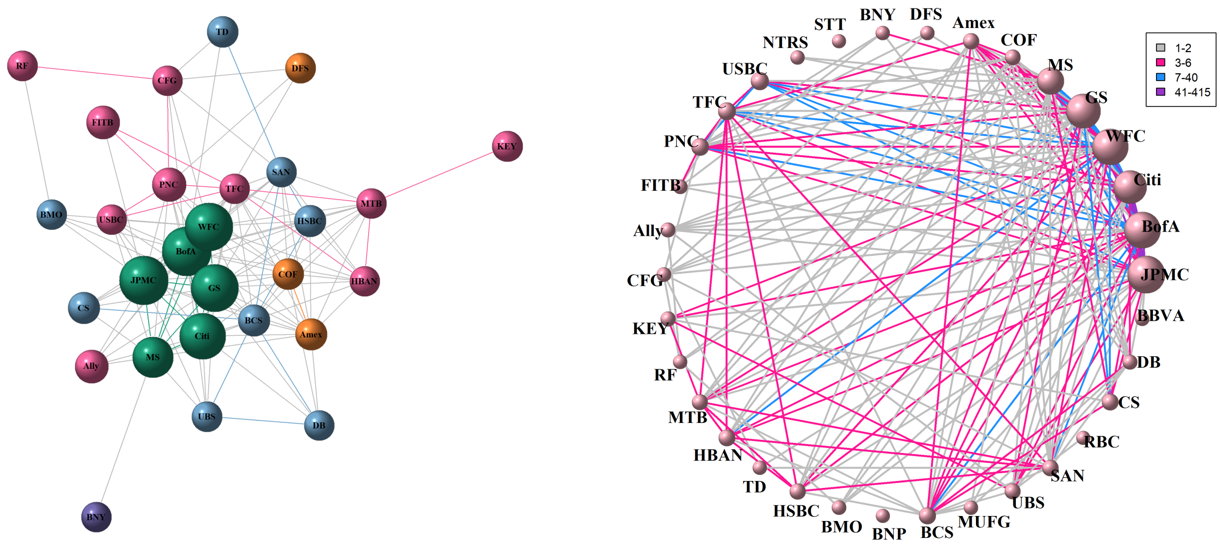

Panel A. Connections & clusters

Panel B. Co-occurrences

Figure 3. Network graphs: April 2020
Panel A: Nodes are colored by bank-type (Big 6 (green), CC (orange), Trusts (purple), Regionals (Pink), and IHCs (light blue)), and link colors correspond to within bank-type connections. Panel B: Nodes are sorted by bank-type and asset size, and link colors correspond to co-occurrence counts (i.e., connections' weights). In both panels, node size represents eigenvector centrality (i.e., larger nodes are more central). April 2020 earnings releases (crisis) extend from April 13^{th}-19^{th}, 2020.

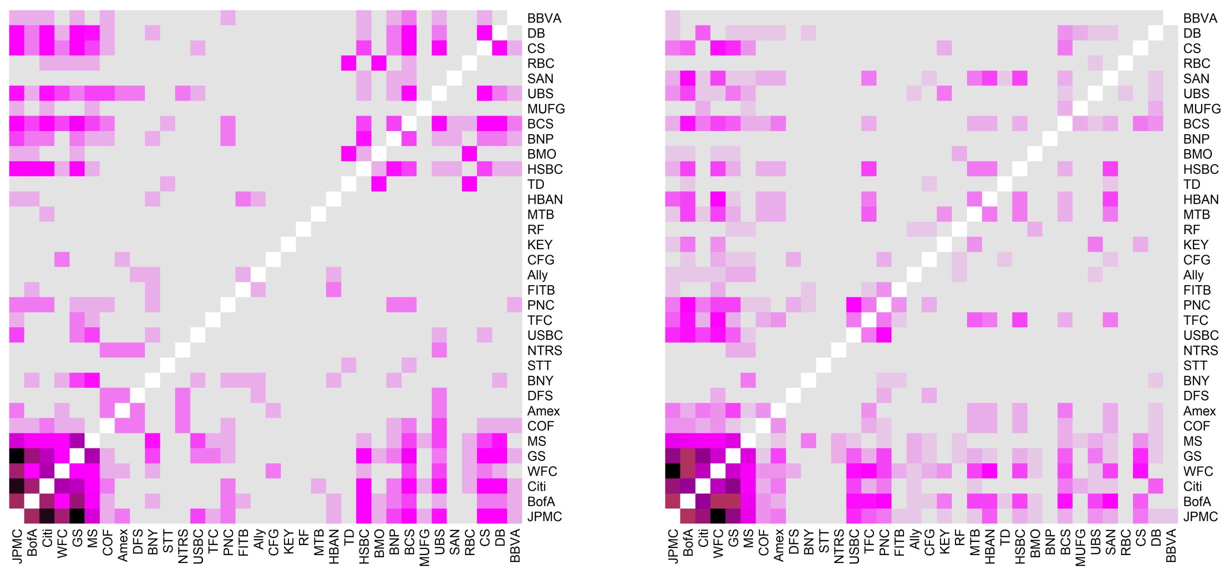

Panel A. January 2020

Panel B. April 2020

Figure 4. Heatmaps: Pre-crisis vs crisis periods
Number of co-occurrences between two banks is represented by the darkness of the corresponding square. More (less) co-occurrences corresponds to darker (lighter) squares. Banks are sorted by bank-type and asset size. January 2020 earnings releases (pre-crisis) extend from January 13^{th}-19^{th}, 2020; and April 2020 earnings releases (crisis) extend from April 13^{th}-19^{th}, 2020.

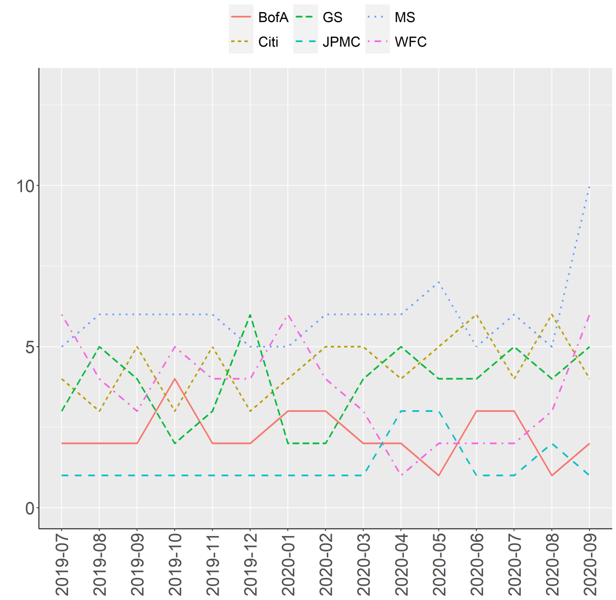

a. Eigenvector Centrality

b. DIP

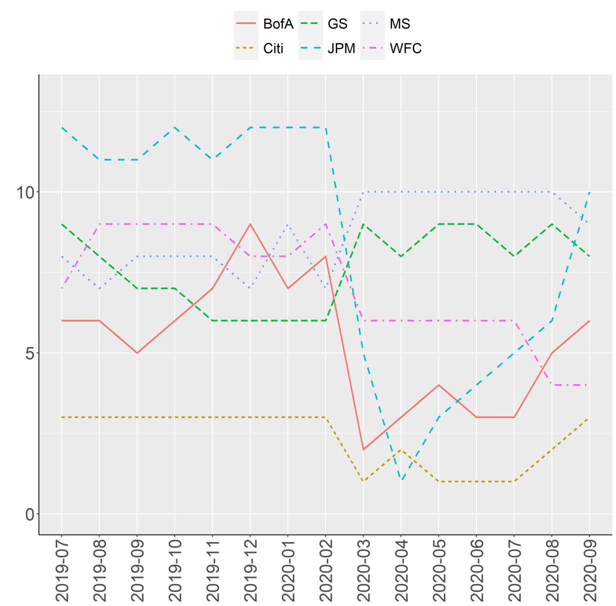

c. SRISK

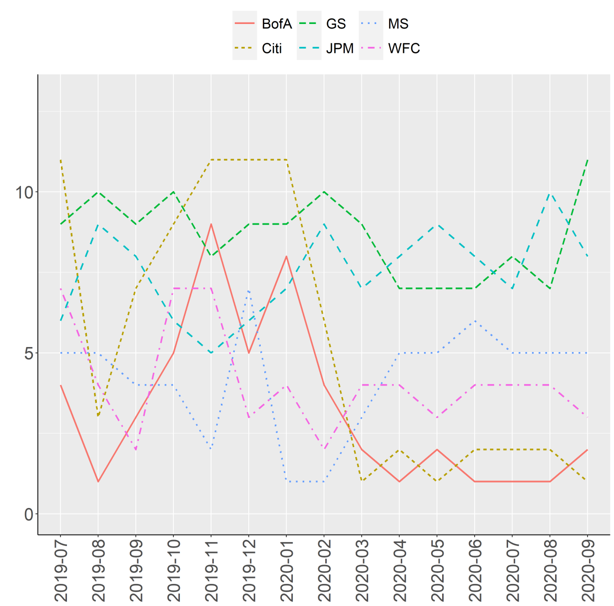

d. CoVaR

Figure 5. Ranking comparison: Traditional systemic risk measures vs Eigenvector centrality
This rank comparison is based on the 12 LISCC banks as of 2020. Only the Big 6 banks are plotted for graph clarity. All banks are ranked 1-12 in terms of monthly average risk measure (with 1 being "most systemically important" and 12 being "least systemically important"), where the y-axis represents the bank ranking. A tie in value between banks in a given measure is given the lowest value/highest rank. Rankings are plotted by month over a 1-year period from October 2019 to September 2020. Negative sentiment news articles are used for text-based Eigenvector centrality.

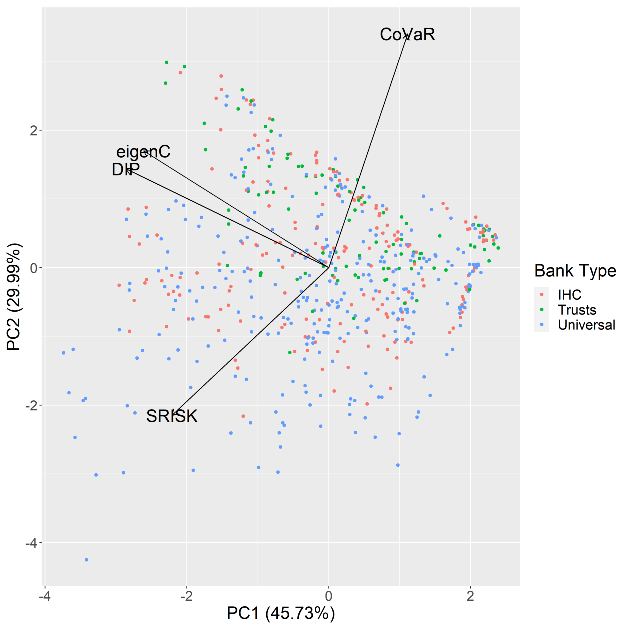

a. PC1 vs PC2

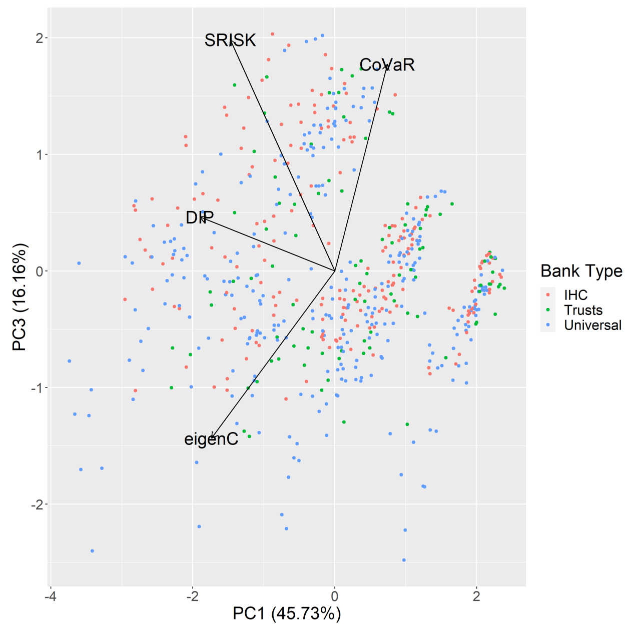

b. PC1 vs PC3

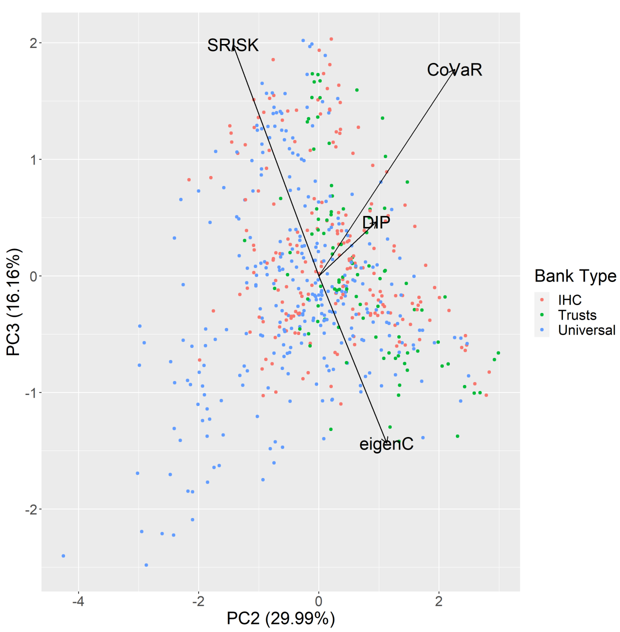

c. PC2 vs PC3

Figure 6. Systemic risk measures: Principal Component Analysis (PCA)
This Principal Component Analysis (PCA) is based on the 12 LISCC banks as of 2020. The analysis is performed on three traditional systemic risk measures (DIP, SRISK, and CoVaR) and our text-based eigenvector centrality. All measures are at a weekly frequency (weekly average for DIP, SRISK, and CoVaR) over a 1-year period from October 2019 to September 2020. Negative sentiment news articles are used for text-based Eigenvector centrality.