The Different Networks of Firms Implied by the News

V. Hilt & G. Schwenkler*

September 3, 2024†

Abstract

The interconnectedness of firms through various networks – such as production, credit, and competition – plays a critical role in determining firm-level and aggregate outcomes. However, data on these connections are often limited. This paper introduces a novel artificial intelligence methodology that extracts explicit firm relationship networks from financial news articles, providing comprehensive and interpretable data across multiple dimensions. Applying this methodology to New York Times articles since 1981, we generate extensive networks that predict key macroeconomic indicators. Our publicly accessible dataset offers valuable insights for future research on firm networks and aggregate fluctuations.

Keywords: firm networks, predictability, aggregate fluctuations, natural language processing, large language models. JEL codes: C45, C82, D20, E32, G00.

*Hilt & Schwenkler are at the Department of Finance, Santa Clara University Leavey School of Business. Schwenkler is corresponding author. Email: gschwenkler@scu.edu, website.

†We are grateful to Sanjiv Das; and seminar participants at Stanford University, and Santa Clara University for useful comments and suggestions.

1 Introduction

How firms are connected to each other is a critical determinant of firm-level and aggregate outcomes in an economy. The way firms are connected in the production network that captures the input and output of goods has been shown to be a driver of aggregate risks by Acemoglu et al. (2012), Acemoglu et al. (2017), Carvalho (2010), Gabaix (2011), Herskovic (2018), and Herskovic et al. (2020), among others. Production network linkages also drive stock returns and contagion across assets, as shown by Cohen and Frazzini (2008), Hertzel et al. (2008), Jorion and Zhang (2009), and others. The network of credit links is an important determinant of firm-level and aggregate default risk, as demonstrated by Azizpour et al. (2018), Eisenberg and Noe (2001), Elliott et al. (2014), and others. Non-credit financing relations in the form of parent-subsidiary connections, common ownership, or strategic partnerships also have an impact on asset prices, as highlighted by Antón and Polk (2014), Aragon and Strahan (2012), Boone and Ivanov (2012), and others.

Most of the existing work either focuses on theoretical results or only indirectly estimates the effects of network linkages. This is because data on different types of links across firms are unusually sparse. While there are some commercial data vendors, like Compustat, that offer data on supply chain relationships, the existing datasets are updated infrequently and only cover a narrow cross section of firms. Some researchers have developed indirect approaches to obtain more extensive data on peer linkages (Hoberg and Phillips (2016), Lee et al. (2015)), common use of technology (Lee et al. (2019)), and correlated shocks (Diebold and Yılmaz (2014), Duan et al. (2012), Hautsch et al. (2014)). However, indirect approaches often result in networks that are hard to interpret because the links they contain are not explicit (see Abbassi et al. (2017) and Craig et al. (2023) for recent analyses). Furthermore, there are simply no accessible datasets for some type of firm relationships, like credit.

Our paper closes this gap. We exploit the rich information contained in financial news reporting to identify different types of relationships that firms may share with each other.

We develop a novel artificial intelligence methodology that reads financial news articles and outputs five different networks of firm relationships: a credit network, a non-credit financing network, a peer network, a supply chain network, and a network that captures all other possible links. We apply our methodology to New York Times articles to obtain monthly networks going back to 1981. Our networks are exhaustive. They contain nearly 40,000 unique links among more than 8,300 distinct firms. Our networks are also explicit because each link can be backed by reporting in a financial news article. We show that summary statistics of the different types of networks predict aggregate outcomes, like declines in industrial production, stock market contractions, and increases in equity volatilities and credit spreads. We make all of our network data freely accessible in an online repository that can be accessed at http://www.news-networks.net. This enables the open usage of our data. All in all, the results of our paper facilitate the development of new research and applications that rely on firm network data that had previously been intractable.

The central contribution of our paper is an artificial intelligence technology that can read news articles and output labelled networks of firms. Our methodology extends the recently developed natural language processing approach of Schwenkler and Zheng (2024). That methodology carries out the first step: It reads news articles and outputs unlabelled links between two firms when they are co-mentioned in the same sentence of an article. We take all the sentences in which the methodology of Schwenkler and Zheng (2024) establishes two firm mentions, transform them into a numerical representation using a modern embedding algorithm, and run the numerical representations through a deep learning neural network model that classifies each sentence as describing one of the following link types: *(i)* credit, *(ii)* financing which is not of the credit type, *(iii)* peer or competitive, and *(iv)* supply chain. We assign the label of the link type that obtains the highest likelihood of being the true representation of the sentence, as long as the likelihood is above 50%. If no relationship achieves a likelihood of more than 50%, then we assign a label of "other."

The deep learning model that underpins our methodology is high dimensional, with

more than 460,000 free parameters. In order to ensure that our model is properly trained, we generate synthetic training data using the GPT 4o large language model of OpenAI. For each of the five firm relationships we consider, we ask ChatGPT to generate 500 sentences that clearly describe such a relationship between two artificial firms. We intentionally ask for artificial firms to not confuse the learning with actual firm names. We also ask to roughly mirror the writing style that would be used in a financial news outlet, like the Wall Street Journal, Financial Times, or the New York Times. We make all of the prompts that were used to generate our synthetic training data, as well as the synthetic data itself, openly available. Because our synthetic dataset contains clear sentences that unambiguously describe different types of firm links, our model is able to learn to accurately distinguish the link types. In out-of-sample tests with synthetic data that were never used for training or cross-validation purposes, our methodology achieve a precision score of 90%, a recall score of 88%, and an F1 score of 89%. Even though the methodology we introduce in this paper only identifies five different types of firms links, it is easy to extend our methodology to identify additional types of linkages. We show how to extend our approach in Section 2.6. A related extension of our methodology is used by Schwenkler and Zheng (2023).

We apply our methodology to 455,583 business news articles we scrape from the New York Times online archive. We combine the news sample with firm data from CRSP, Compustat, and I/B/E/S. Our data cover the time from January 1, 1981, through December 31, 2023. Over this period, we extract a credit network that contains 14,419 links among 2,823 unique firms, a financing network that contains 34,527 links among 4,898 unique firms, a peer network that has 30,958 links among 2,942 firms, and a supply chain network that has 19,930 links among 3,422 firms. The network of "other" links is the densest, containing 76,359 links among 5,789. This shows that our methodology is selective when assigning a link label, which only shrinks our link samples but does not confuse them. We run a keyword analysis to understand the informational content of the sentences in which we identify the different types of firm links. We find that credit relationships are often identified in sentences

that mention words like *bond*, *borrow*, *credit*, *debt*, and *yield*, validating the nature of this type of firm link. Non-credit financing relationship are identified in sentences that contain words like *acquire*, *buy*, *own*, *share*, and *stock*, suggesting that this type of linkages captures common equity across firms. Because of this, we interpret the non-credit financing network as containing M&A, parent-subsidiary, and common ownership links. Sentences that contain peer links often mention words like *competitor*, *dominate*, *industry*, and *rival*. Supply chain links are often extracted from sentences that contain the words *client*, *customer*, *develop*, *make*, *partnership*, and *work*. Note that these keywords are identified in actual sentences published in the New York Times in which we identify firm links. This shows that our methodology is able to properly identify firm links in real news data outside of the artificially clean synthetic dataset we used to train our models.

We apply our newly obtain networks to test the theories of Acemoglu et al. (2012), Acemoglu et al. (2017), Carvalho (2010), Gabaix (2011), Herskovic (2018), and Herskovic et al. (2020), who show that characteristics of the production network of an economy drive aggregate fluctuations. For each link type, we implement a vector autoregressive model for several macroeconomic variables, including consumption growth, industrial production, the level and slope of the yield curve, corporate credit spreads, as well as the S&P 500 return and the VIX, together with three key connectivity metrics that summarize the architecture of the network. One is the density of the network, another is the degree of centralization of the network, and a third one is an interconnectivity measure that captures whether firms are strongly or weakly interconnected through intermediary firms. Validating the existing theoretical work, we find that it is primarily shocks to the supply chain network that predict macroeconomic outcomes. We find that periods of time when the news-implied supply chain network becomes more interconnected are followed by persistent declines in industrial production and transient declines in consumption growth. Extending the existing literature, we also find that shocks to the architecture of the peer and financing networks predict deteriorations in asset prices. Shocks to the interconnectivity measure of the peer network

(as well as the supply chain network) are followed by increases in credit spreads that persist for 12 months. Shocks to the density of the peer and financing networks predict persistent increases in the VIX, while shocks to the density of the financing network also predict short-term declines in the S&P 500. Furthermore, shocks to the interconnectivity of the peer and financing networks forecast declines in the level and increases in the slope of the US yield curve. Our results show that non-production networks, such as peer and financing networks, also play a role in driving aggregate fluctuations. They complement existing research that highlights the influence of peer and non-credit financial linkages for firm-level outcomes, such as Antón et al. (2024), Antón and Polk (2014), Azar et al. (2022), Hou (2007), Hou and Moskowitz (2005), Jorion and Zhang (2007), Lee et al. (2015), and many others. Our findings show that the different types of networks we extract from the news facilitate the accurate prediction of macroeconomic risks.

This paper is organized as follows. Section 2 introduces our methodology and data sources. Section 3 presents the different types of firm networks that are extracted from the data by our methodology. Section 4 contains our predictive analysis of macroeconomic indicators. We conclude in Section 5. Our data repository is available at `http://www.news-networks.net`. There is an Online Appendix, which is also available in our data repository.

2 Data & methodology

We develop a methodology to identify firm connections from news data and label each connection as one of the following five types: 1) credit, 2) financing (not credit), 3) supply chain, 4) peer (competitive), and 5) other. Our approach extends the methodology of Schwenkler and Zheng (2024) that builds networks of firm interconnections without labelling the types of relationships firms may share in the network. We apply our methodology to composite data from The New York Times, CRSP, Compustat, and I/B/E/S. All network data we generate with our methodology is openly accessible in our data repository.

We provide a summary of our methodology in this section and offer details in Online Appendix B. While we only apply our methodology to identify five different types of firm links in this paper, the methodology can be easily extended to identify additional types. We outline how this can be done in Section 2.6. An extension of our methodology is applied by Schwenkler and Zheng (2023).

2.1 Firm population

We construct our population of firms from CRSP, Compustat, and I/B/E/S. Online Appendix A provides details of our data collection approach. We only consider US corporations with equity securities and otherwise apply no filters. We obtain a sample consisting of 26,374 unique firms. Table 1 reports summary statistics and Figure 1 display the industry distribution in our firm population.

2.2 News data

We scrape news articles published in the New York Times after January 1, 1981.[1] We only keep articles for which the names of the published section or the contributing news desk include the words "business" or "financial." We filter out any company announcements, which we identify as those articles that either are tagged with the keyword "Company Reports" by The New York Times or include the string "* COMPANY REPORTS *" in their bodies. For the purpose of this paper, we end our news sample on December 31, 2023, because our CRSP data is only updated annually. However, we will update our sample every year in our online data repository.

We collect 455,583 full-text articles over the sample time period. Figure 2 shows the number of articles per month, together with the average length of an article and the number of contributing news desks in a month. We see that The New York Times published less

[1] We developed our own scraper to collect these articles from the New York Times API.

business-relevant articles in the later stages of our sample than it did at the beginning of our sample. However, the articles have become longer as measured by the number of words they include. There are also more news desks contributing business articles now.

2.3 Unlabelled news-implied network

We apply the methodology of Schwenkler and Zheng (2024) to identify firm mentions in our news data and construct news-implied networks. The methodology proceeds in four steps. It first uses a Named Entity Recognition algorithm (NER) to identify and classify any entities named in the articles. For those entities labelled as organizations, it follows heuristics to separate firms from non-firms. In the third step, it matches the sample of possible firms with the CRSP database. In the fourth step, the methodology establishes a link between two firms when they are co-mentioned in a sentence and the sentence contains only two firm mentions.[2] The methodology of Schwenkler and Zheng (2024) is highly accurate, achieving a precision score of 96.5% and an F1 score of 65%.

Table 1 provides summary statistics and Figure 1 shows the industry distribution for the firms that are identified in our news sample using the methodology of Schwenkler and Zheng (2024). Over the course of our 43-year-long sample, we identify a total of 2,112,140 mentions of 13,526 distinct firms. We observe that the firms that are covered in the New York Times tend to be larger, with higher net incomes, sales, and costs of goods sold. They also have higher monthly returns and are covered by more earnings analysts. We observe that the manufacturing, retail, and transportation industries are overrepresented in the New York Times sample. Notably, the finance industry is underrepresented in the New York Times. This stands in contrast to Schwenkler and Zheng (2024), who find that the finance industry is overrepresented in a Reuters news sample.

Moving to firm links, the methodology of Schwenkler and Zheng (2024) finds 176,193

[2]Analyses in Schwenkler and Zheng (2024) suggest that it suffices to identify firm connections by co-mentions in sentences versus whole body.

sentences with two firm mentions in our sample. It establishes a total of 39,629 unique links among 8,322 distinct firms. The implied firm networks are showcased in Figure 3, where we break the time series into decades. Similarly to Schwenkler and Zheng (2024) albeit using a different news data source, we find that the news-implied networks have a star architecture as in Acemoglu et al. (2012). We observe that some of the biggest banks tend to be in the center of the network, surrounded by clusters of non-financial firms.

2.4 Network labelling

We extend the methodology of Schwenkler and Zheng (2024) to label the types of relationships firms may share in the news-implied network. For each link identified by the approach of Schwenkler and Zheng (2024), our methodology assigns one of five possible labels: 1) credit, 2) financing (ex credit), 3) supply chain, 4) peer, and 5) other.

Our methodology proceeds as follows. We take all sentences that the methodology of Schwenkler and Zheng (2024) establishes as having two firm mentions and run them through an open-source sentence transformer algorithm called "all-mpnet-base-v2" provided by Hugging Face in Python. The transformer converts the text of the sentence into a numerical representation in the form of a 768-dimensional vector embedding that captures its most relevant semantic features. The transformer model was trained for the purpose of information retrieval and sentence similarity tasks using over 1 billion paired sentences from Reddit, S2ORC, WikiAnswers, Yahoo Answers, and others. Because our goal is to identify sentences that describe similar types of business relationships, we believe that this transformer model works well for our purposes.

Once we obtain the embeddings, we run them through independent deep learning algorithms for each firm relationship type. We construct four separate neural-network-based binary classifiers, one for each named firm relationship: credit, financing, supply chain, and peer. The classifiers compute probabilities that a sentence, as represented by its embedding,

characterizes the four different types of firm relationships. We run each sentence through the four classification models to obtain link type probabilities. Then, we assign the relationship label with the highest probability, as long as the probability is larger than 50%. If no link type obtains a probability above 50%, then we assign the label of "other." All in all, our model can be viewed as one big, segmented neural network. Figure 4 showcases the architecture of our neural network model. Online Appendix B provides details.

Table 2 provides two examples of how our methodology works. The first sentence states: *"United (UAL) is the only American airline that operates Boeing (BA) 777s equipped with the PW4000 engine series, and the airline said on Sunday that it was grounding those 24 planes in its active fleet while it awaited F.A.A. guidance."* Here, we marked in parentheses the tickers of the two firms that were identified in the sentence. Our methodology uniquely labels this sentence as describing a supply chain relationship, which is reasonable given the information contained in the sentence. The second sentence reads: *"The many local Xinjiang police bureaus on the list buy commercial American technology like Intel (INTC) microchips and Microsoft (MSFT) Windows software, according to procurement documents."* Our approach does not assign more than 50% likelihood to any of the four firm relationships, so we label it as "other." This is again reasonable as the sentence does not clearly outline the relationship between Intel and Microsoft, other than these two being American firms.

2.5 Training & testing

Our methodology is high dimensional. In total, there are 460,808 parameters to fit. Fitting such a complex model requires good training data. One approach we could take to construct a training dataset is to hand-label several sentences that contain two firms as describing one of the five different types of firm relationships we consider. However, such an approach faces several challenges. First, it is costly as it requires costly human supervision and expertise. Second, it is subjective and therefore also possibly ad-hoc because it is based on how individ-

ual people interpret different sentences. Third, it is prone to errors because many sentences do not clearly describe unique firm relationships. For example, the sentence *"In April, Kodak demanded that Cetus goes to arbitration to determine Kodak's rights to certain production methods and technology the companies jointly developed"* could reasonably be classified as *peer*, *supply chain*, or *other*.

We follow an alternative approach that levers the power of modern large language models (LLMs). We create synthetic sets of sentences that clearly describe the different types of firm relationships using ChatGPT. For each one of the five firm links (credit, financing, supply chain, peer, and other), we ask ChatGPT to generate 500 sentences that describe such a relationship between two fictional firms. We intentionally ask ChatGPT to construct sentences for fictional firms to not contaminate the training of our models with the names of firms that may normally be associated with a certain type of relationship. For example, we do not want "Citi" to always be associated with a credit relationship just because it is a bank or "Morningstar" to always be associated with the "other" label just because it provides market research. We ask ChatGPT to construct diverse sentence structures that resemble the way The Wall Street Journal, Financial Times, or The New York Times would describe relationships between firms. We also ask ChatGPT to consider different types of relationships for each label. For example, for a credit relationship, ChatGPT could describe one firm getting loan from a bank or it could describe a customer getting trade credit from a supplying firm. We provide the prompts we used to construct the synthetic datasets in Online Appendix C.[3] We also uploaded the synthetic dataset to our data repository for full transparency. In total, our synthetic training dataset contains 2,500 sentences that are equally distributed across the five firm relationships.

We proceed as follows to train our model. For each relationship other than "other," we create a replica of our synthetic training set and assign a numerical category value of "1" if a synthetic sentence describes the given relationship and "0" otherwise. Each synthetic

[3]We used the GPT 4o model to generate our synthetic datasets.

training data replica has 2,500 rows (one row for each sentence) and 769 columns (768 columns for the embedding and one column for the link category indicator). We split each replica the same way into a training set (60% of the sample or 1,500 sentences), a cross-validation set (20% of the sample or 500 sentences), and a test set (20% of the sample or 500 sentences). We then fit the four neural network models for each named firm relationship to the corresponding synthetic data replica and measure the cross-validation errors in the process.[4] We measure the out-of-sample accuracy of our approach using the test data.

Table 3 displays the accuracy metrics as well as the out-of-sample confusion matrix of our methodology for each firm relationship and in the aggregate. We observe that our methodology is highly accurate. On average across all link types, we achieve a precision score of 90%, a recall score of 88%, and an F1 score of 89%.[5] The in-sample, cross-validation, and out-of-sample metrics for the individual link models are similar, suggesting that our methodology does not overfit the data. Our methodology is most precise for identifying credit relationships, for which all accuracy metrics are above 95%. Our methodology is least accurate for identifying "other" links, which is acceptable because the sentences that have this type of label do not describe a clear relationship. The high recall scores for the four named firm relationships (credit, financing, supply chain, and peer) indicate that our methodology rarely assigns a label erroneously. The high precision scores indicate that our methodology is generally correct when it assigns a label. The confusion matrix shows that, when errors occur, they generally occur because our methodology assigns the label "other" to a named relationship. This mainly reduces the category sample size but does not confound it. The results of Table 3 show that our methodology to label firm relationships extracted from news data is highly accurate.

[4]We choose to stop the training after 10 epochs with batch sizes of 32 as we observe that the cross-validation error peaks at around 10 epochs.

[5]Here, precision (recall) measures the ratio of true positive over all predicted (actual) positives, while the F1 score gives the harmonic mean of precision and recall. These metrics are commonly used in the machine learning literature to measure the performance of labeling algorithms; see El Mrabet et al. (2021), Shobha and Rangaswamy (2018), and others.

2.6 Extension

One downside of our methodology is that it only identifies five different types of firm links. However, this is by choice to limit the scope of the paper rather than by an actual limitation of the methodology. In principle, our methodology is able to identify as many firm relationships as it is fed with training data. We provide a heuristic to extend our methodology to identify additional types of firm linkages.

Suppose that we wanted to also build a network of firms with common geographic locations (the second example in Table 2 suggests that the network of "other" links may include such links). To extend our methodology to also identify common location links, we would do the following. First, we would use a large language model (LLM) like ChatGPT to construct a synthetic set of sentences that describe two firms that share a common geographic location. To not confuse the learning, we would ask the LLM to include exactly two fictional firms and a fictional location in the sentences. We would also ask to write the sentences in a way that would appear in common financial news outlets. And we would generate the same number of sentences as we had previously generated for all other link types in order to keep the panel balanced. Then, we would run these sentences through the same embedding we run all other synthetic training sentences. In the next step, we would construct an additional deep learning classifier for common location links and train it with the synthetic dataset replica that includes the new common location links as well as all other firm links. We would also re-train all other classifiers with the new, extended training set replicas. Once all models have been trained, we run the unlabelled links extracted by the methodology of Schwenkler and Zheng (2024) from actual news articles through our classifiers and assign the link label that achieves the highest probability, as long as the probability is above 50%. Otherwise, we assign the label "other." In short, extending our methodology only requires generating more synthetic training data using an LLM and expanding the model of Figure 4 by adding additional neural networks.

Even though extending our methodology is straightforward, including additional types of firm links may impact links that had been previously labeled. For example, once we include common location links, some of the "other" links may be re-classified as common location. There are trade-offs between including a large number of link types and identifying expansive networks of firms. With a large number of link types, there will be more synthetic training data, which should translate into higher accuracy. However, with a limited set of actually published news articles, each link category would cover a smaller fraction of the news, resulting in sparser networks.

3 Output of the methodology

We analyze the output of our methodology. Tables 4 and 5 shows the five most frequent links for each of the different relationships. Figures 5 and 6 displays the most frequently mentioned verbs and nouns in the sentences that contain the different types of firm links. In the word clouds, we exclude the words *"say"* and *"company"* because they are the most frequently mentioned words but are irrelevant for our purposes. Figures 7–11 shows the time series of the five different networks broken down by decades.

Compared to the network of Schwenkler and Zheng (2024) in Figure 3, our networks are interpretable and economically meaningful. We observe that the credit network is the sparsest while the "other" network is the densest as measured by the total number of identified links and unique firms. The keywords of Figures 5 and 6 show that each network captures information that is related to the link type they represent. For example, words like *debt*, *bond*, and *credit* are among the most frequently mentioned nouns in sentences in which we identify credit relationships, while *yield*, *borrow*, and *subordinate* are among the most frequently used verbs in those sentences. *Acquire*, *buy*, and *own* are among the most frequent verbs for non-credit financing relationships, while *share* and *stock* are among the most frequent nouns. For peer relationships, we find that *rival*, *competitor*, and *industry* are some

of the most frequent nouns and *compete* and *dominate* are among the most frequent verbs. *Work*, *make*, and *develop* are among the most frequent verbs for supply chain relationships, and *client*, *customer*, and *partnership* are among the most frequent nouns.

Going into more detail, we observe that the credit network has the typical star architecture posited by Acemoglu et al. (2012), whereby the big banks are highly central and interconnected. They are surrounded by smaller banks and industry-specific clusters. A star architecture for the credit network is consistent with the theoretical works of Bernanke et al. (1999) and Carvalho and Gabaix (2013), who highlight the central role of the banking industry. Table 4 suggests that our credit network captures both direct credit relationships (like between Citigroup and Enron) as well as indirect exposures. For example, some firms have common holdings of credit assets (like Chrysler, Ford, and General Motors) while other firms have common liabilities (like Chrysler & General Motors or Bank of America & Citigroup). These results complement recent findings by Abbassi et al. (2017) and Craig et al. (2023), who study the degree to which networks extract from asset pricing data contain information about direct and indirect credit linkages. We show that the news reports about such linkages.

The non-credit financing network is more dispersed. Often, it is not banks which are at the center of the networks but industrial firms, like General Motors and General Electric. We attribute this to the fact that non-credit financing relationships may involve equity in M&A transactions in which the counterparties are not necessarily banks. In line with this interpretation, Table 4 shows that many of the links in the financing network are due to mergers or acquisitions (like between Microsoft & Yahoo or Time Warner & AT&T) or direct equity holdings (like between Yahoo & Google). There are also indirect financing links if, for example, two acquirers pursue the same target (like Microsoft & Google) or two firms have common investors (like Chrysler & Ford). We conclude that the non-credit financing network contains equity links, which can be of the form of parent-subsidiary, direct equity investments, or indirect common ownership links.

The peer network showcases strong bilateral relationships like, for example, between Ebay and Walmart, Coca Cola and Pepsi, Apple and Microsoft, and Ford and Chrysler. Strong bilateral links would naturally arise if, say, the news compared one peer to another. The sample sentences of Table 4 show that this often occurs in our data. Our methodology identifies peer linkages when the news explicitly reports about competitors (like between Coca Cola & Pepsi, Microsoft & Google, or Yahoo & Google). It also identifies peer linkages when the news reports about peers indirectly through comparisons (like between Apple & Google or Apple & Microsoft).

The sentences of Table 5 indicate that the supply chain network includes direct sale linkages (like between Microsoft, Apple, Amazon, and Intel) as well as indirect linkages through common suppliers (like between Chrysler and Ford). The supply chain network showcases pronounced clusters. For example, throughout the decades it is visible that Boeing (ticker: BA) is in a customer-supplier relationship with American Airlines (AMR), Delta (DAL), and United (UAL). The airlines are also linked in the supply chain network because they often share strategic partnerships or cooperations. For example, The New York Times reported on September 14, 1982: *"United Airlines and Delta Air Lines said they had agreed that each would market the other's computerized passenger reservations and accounting systems to travel agencies, commercial accounts and governmental agencies."*

The network of other links showcases more of a mesh architecture. Table 5 shows that the "other" network captures alternative links, like common geographic locations (as between Apple & Google) or common executives (Chrysler & Ford). Some "other" links are uninformative and arise when the news randomly co-mentions two firms in passing, say, if it is providing a market update (like between Ford & General Motors or Intel & Microsoft).

All in all, the results of this section show that the output of our methodology is vast. It includes five different types of firm links for which there are currently only sparse datasets. Our methodology identifies links among a large number of firms. It can construct networks over arbitrary time frequencies. While we only display networks over decades here, our data

repository contains all networks at the monthly frequency. These data are freely available for download. Our methodology is flexible and can be easily adjusted and extended to recognize more types of firms links. This only requires generating additional synthetic training datasets using an LLM and training additional deep learning classifiers, as highlighted in Section 2.6.

4 Predictability

The literature shows that the architecture of the network of firm interconnections in the economy is a key driver of aggregate fluctuations. Most of the evidence is theoretical in nature. Empirical analyses are challenging because data on firm networks are notoriously sparse. In new work, Schwenkler and Zheng (2024) show that the architecture of the network of firm links extracted from financial news reporting indeed predicts periods of aggregate distress, validating the existing theoretical results. However, the work of Schwenkler and Zheng (2024) is silent about what kinds of firm links drive their results. In this section, we extend the results of Schwenkler and Zheng (2024) by dissecting the type of linkages that drive this predictability.

We follow Schwenkler and Zheng (2024) and measure three connectivity metrics that summarize the architecture of a network:

- The average degree, which counts the average number of links of a node and is inversely related to the network sparsity measure of Herskovic (2018). It indicates how dense a network is.

- The first-order interconnectivity measure of Acemoglu et al. (2012) given by the coefficient of variation of the degree distribution in the network. This metric indicates how centralized a network is.

- The second-order interconnectivity measure of Acemoglu et al. (2012), which is the weighted covariance of the degree of two nodes that are indirectly connected through

a third node. It indicates whether cluster of firms are strongly interconnected through intermediary firms.

Figure 12 shows the time series of the connectivity metrics for the four named firm networks. In the caption, we report details of how we compute them. We observe that the peer network has become denser over time. We also observe that both the peer and the supply chain network have become more interconnected as measured by the second-order interconnectivity metric. The credit network is the sparsest network. The financing network has the most volatile connectivity metrics.

Next, we estimate a one-lag vector autoregressive (VAR) model with an intercept for the joint dynamics of consumption growth, industrial production growth, the level and slope of the yield curve, the AAA and BAA corporate credit spreads, S&P 500 returns, the VIX, and the news-implied connectivity measures. We estimate separate VARs for each of the networks. We follow Schwenkler and Zheng (2024) and base the identification on a Cholesky decomposition of the residual variance-covariance matrix in which consumption and industrial production growth are the most exogenous variables and the connectivity measures are the most endogenous variables; Online Appendix E provides details. We differentiate all non-stationary variables first and then standardize all variables prior to running the VAR. Note that the time series for the VIX only begins in January 1990 but the time series for the VXO (the option-implied volatility index for the S&P 100) starts in January 1986. To obtain a longer time series for the VIX, we complement early observations with those of the VXO. Because of this and because we take first differences of some variables, the sample for the VAR analyses is limited to the time after February 1986.

We evaluate cumulative impulse response functions to orthogonal shocks for the connectivity measures of the different networks. Figures A.1–A.15 in the Online Appendix show all the impulse response functions. Here, we summarize our findings in Figures 13 and 14. We neglect the impulse response functions for the first-order interconnectivity metric as they

are statistically insignificant. We also neglect the impulse response functions for shocks to the connectivity metrics of the network of other links as they are mostly insignificant, other than having a minor predictive role for the VIX. These observations validate the secondary nature of the links in the "other" network.

Consistent with Schwenkler and Zheng (2024) albeit using a different and more extensive news dataset, we find that shocks to the connectivity measures of the news-implied networks predict periods of heightened macroeconomic risks. Also consistent with Schwenkler and Zheng (2024), we find that the average degree (i.e., the *density* of the network) generally predicts deterioration in asset prices while second-order interconnectivity (i.e., the *interconnectivity* of the network) predicts aggregate distress. Extending Schwenkler and Zheng (2024), we show which networks drive the predictability. We find that shocks to the interconnectivity of the supply chain networks predict persistent declines in industrial production and temporary declines in consumption. They also predict drops in the level of the yield curve. Shocks the interconnectivity of the non-credit financing network predict declines in the level of the yield curve, while shocks to the interconnectivity of the peer network predict a steepening of the yield curve. In addition, we find that shocks to the density of the financing and peer networks predict persistent increases in the VIX. Shock to the density of the financing network also predict short-term declines in stock returns, while shocks to the density of the peer network predict increases in credit spreads.

The results of our VAR analyses expand our understanding of the drivers of aggregate risks. The models of Acemoglu et al. (2012), Acemoglu et al. (2017), Carvalho (2010), and Gabaix (2011) show that the architecture of the production network of an economy drives aggregate output fluctuations. Herskovic (2018) and Herskovic et al. (2020) show that the architecture of the production network also drives equity prices and volatilities. The network of ours that most closely reflects production links is the supply chain network. The impulse response functions for the supply chain network provide support for the theories of Acemoglu et al. (2012), Acemoglu et al. (2017), Carvalho (2010), and Gabaix (2011) in that

they show that shocks to the degree of interconnectivity of the supply chain network forecast declines in industrial production and consumption growth. We do not find evidence that the architecture of the supply chain network impacts aggregate stock returns and volatilities, even though it does influence Government and corporate bond yields.

Several papers have demonstrated that peer linkages enable competition and contagion effects that yield cross-momentum and cross-reversal among peer asset prices; see Hou (2007), Hou and Moskowitz (2005), Jorion and Zhang (2007), Lang and Stulz (1992), and Lee et al. (2015), among others. However, few papers have considered the aggregate implications of the network of peer linkages. We close this gap by showing that shocks to the density and interconnectivity of the peer network predict increases aggregate stock market volatilities as well as corporate credit spreads. A more interconnected peer network also forecasts a steeping of the US yield curve.

Somewhat surprising, we find that the credit network we extract from our New York Times dataset has no statistically significant predictive power for macroeconomic outcomes. This stands in contrast to established results that show that the network of credit connections across firms is a key driver of aggregate risks; see Azizpour et al. (2018), Eisenberg and Noe (2001), and Farboodi (2023), among others. We attribute the lack of predictive power to the sparseness of our credit network. Figure 12 shows that the credit network has the fewest amount of links on average.

However, we find that the non-credit financing network has predictive power for aggregate outcomes. Shocks to the density of the financing network predict declines in aggregate stock returns and increases in aggregate stock volatilities. Shocks to the degree of interconnectivity of the financing network also predict declines in the level of the yield curve. These results complement several works that have studied how non-credit financing relationship impact asset prices (see Aragon and Strahan (2012), Boone and Ivanov (2012), Fernando et al. (2012), and others) by highlighting the aggregate implications of such linkages. To the extent that our financing network captures equity holdings that are shared across firms, our

findings suggest that common ownership can exacerbate aggregate measures of risks. Our findings align with Ederer and Pellegrino (2024), who show that common ownership can impact aggregate welfare. They complement papers that study the implications of common ownership on stock returns (Antón and Polk (2014)), product market competition (Azar et al. (2018), Azar et al. (2022), Koch et al. (2021)), corporate governance (Edmans et al. (2018)), and innovation (Antón et al. (2024), Li et al. (2023)), among other aspects.

5 Conclusion

We introduce an artificial intelligence methodology to extract labelled networks of firm interconnections from news data. Our methodology is based on modern natural language processing technologies and large language models. We use our methodology to construct extensive credit, financing, peer, and supply chain networks that span the time from 1981 to 2023. The networks we extract contain information that predicts declines in industrial production, consumption, and aggregate stock returns, as well as increases in aggregate equity volatility and credit spreads. We make all of our network data freely accessible, opening up avenues for new research that was previously intractable.

References

Abbassi, Puriya, Christian Brownlees, Christina Hans and Natalia Podlich (2017), 'Credit risk interconnectedness: What does the market really know?', *Journal of Financial Stability* **29**, 1–12.

Acemoglu, Daron, Asuman Ozdaglar and Alireza Tahbaz-Salehi (2017), 'Microeconomic origins of macroeconomic tail risks', *American Economic Review* **107**(1), 54–108.

Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar and Alireza Tahbaz-Salehi (2012), 'The network origins of aggregate fluctuations', *Econometrica* **80**(5), 1977–2016.

Antón, Miguel and Christopher Polk (2014), 'Connected stocks', *The Journal of Finance* **69**(3), 1099–1127.

Antón, Miguel, Florian Ederer, Mireia Giné and Martin Schmalz (2024), 'Innovation: The bright side of common ownership?', *Management Science* . Forthcoming.

Aragon, George O. and Philip E. Strahan (2012), 'Hedge funds as liquidity providers: evidence from the Lehman bankruptcy', *Journal of Financial Economics* **103**(3), 570–587.

Azar, José, Martin C. Schmalz and Isabel Tecu (2018), 'Anticompetitive effects of common ownership', *The Journal of Finance* **73**(4), 1513–1565.

Azar, José, Sahil Raina and Martin Schmalz (2022), 'Ultimate ownership and bank competition', *Financial Management* **51**(1), 227–269.

Azizpour, S., K. Giesecke and G. Schwenkler (2018), 'Exploring the sources of default clustering', *Journal of Financial Economics* **129**(1), 154–183.

Bernanke, Ben S., Mark Gertler and Simon Gilchrist (1999), The financial accelerator in a quantitative business cycle framework, Handbook of Macroeconomics, Elsevier, Amsterdam, pp. 1341–1393.

Boone, Audra L. and Vladimir I. Ivanov (2012), 'Bankruptcy spillover effects on strategic alliance partners', *Journal of Financial Economics* **103**(3), 551–569.

Carvalho, Vasco (2010), Aggregate fluctuations and the network structure of intersectoral trade, Technical Report 1206, Department of Economics and Business, Universitat Pompeu Fabra.

Carvalho, Vasco and Xavier Gabaix (2013), 'The great diversification and its undoing', *American Economic Review* **103**(5), 1697–1727.

Cohen, Lauren and Andrea Frazzini (2008), 'Economic links and predictable returns', *The Journal of Finance* **63**(4), 1977–2011.

Craig, Ben R., Madina Karamysheva and Dilyara Salakhova (2023), Do market-based networks reflect true exposures between banks? ECB Working Paper No. 2023/2867.

Diebold, Francis X. and Kamil Yılmaz (2014), 'On the network topology of variance decompositions: Measuring the connectedness of financial firms', *Journal of Econometrics* **182**(1), 119 – 134.

Duan, Jin-Chuan, Jie Sun and Tao Wang (2012), 'Multiperiod corporate default prediction—a forward intensity approach', *Journal of Econometrics* **170**(1), 191–209.

Ederer, Florian and Bruno Pellegrino (2024), A tale of two networks: Common ownership and product market rivalry. Working Paper.

Edmans, Alex, Doron Levit and Devin Reilly (2018), 'Governance Under Common Ownership', *The Review of Financial Studies* **32**(7), 2673–2719.

Eisenberg, Larry and Thomas H. Noe (2001), 'Systemic risk in financial systems', *Management Science* **47**(2), 236–249.

El Mrabet, Mohammed Amine, Khalid El Makkaoui and Ahmed Faize (2021), Supervised machine learning: A survey, *in* '2021 4th International Conference on Advanced Communication Technologies and Networking (CommNet)', pp. 1–10.

Elliott, Matthew, Benjamin Golub and Matthew O. Jackson (2014), 'Financial networks and contagion', *American Economic Review* **104**(10), 3115–3153.

Farboodi, Maryam (2023), 'Intermediation and voluntary exposure to counterparty risk', *Journal of Political Economy* **131**(12), 3267–3309.

Fernando, Chitru S., Anthony D. May and William L. Megginson (2012), 'The value of investment banking relationships: evidence from the collapse of Lehman Brothers', *The Journal of Finance* **67**(1), 235–270.

Gabaix, Xavier (2011), 'The granular origins of aggregate fluctuations', *Econometrica* **79**(3), 733–772.

Hautsch, Nikolaus, Julia Schaumburg and Melanie Schienle (2014), 'Financial Network Systemic Risk Contributions*', *Review of Finance* **19**(2), 685–738.

Herskovic, Bernard (2018), 'Networks in production: Asset pricing implications', *The Journal of Finance* **73**(4), 1785–1818.

Herskovic, Bernard, Bryan Kelly, Hanno Lustig and Stijn Van Nieuwerburgh (2020), 'Firm volatility in granular networks', *Journal of Political Economy* **128**(11), 4097–4162.

Hertzel, Michael G., Zhi Li, Micah S. Officer and Kimberly J. Rodgers (2008), 'Inter-firm linkages and the wealth effects of financial distress along the supply chain', *Journal of Financial Economics* **87**(2), 374–387.

Hoberg, Gerard and Gordon Phillips (2016), 'Text-based network industries and endogenous product differentiation', *Journal of Political Economy* **124**(5), 1423–1465.

Hou, Kewei (2007), 'Industry Information Diffusion and the Lead-lag Effect in Stock Returns', *The Review of Financial Studies* **20**(4), 1113–1138.

Hou, Kewei and Tobias J. Moskowitz (2005), 'Market Frictions, Price Delay, and the Cross-Section of Expected Returns', *The Review of Financial Studies* **18**(3), 981–1020.

Jorion, Philippe and Gaiyan Zhang (2007), 'Good and bad credit contagion: evidence from credit default swaps', *Journal of Financial Economics* **84**(3), 860–883.

Jorion, Philippe and Gaiyan Zhang (2009), 'Credit contagion from counterparty risk', *The Journal of Finance* **64**(5), 2053–2087.

Koch, Andrew, Marios Panayides and Shawn Thomas (2021), 'Common ownership and competition in product markets', *Journal of Financial Economics* **139**(1), 109–137.

Lang, Larry and Rene Stulz (1992), 'Contagion and competitive intra-industry effects of bankruptcy announcements', *Journal of Financial Economics* **32**, 45–60.

Lee, Charles M.C., Paul Ma and Charles C.Y. Wang (2015), 'Search-based peer firms: Aggregating investor perceptions through internet co-searches', *Journal of Financial Economics* **116**(2), 410 – 431.

Lee, Charles M.C., Stephen Teng Sun, Rongfei Wang and Ran Zhang (2019), 'Technological links and predictable returns', *Journal of Financial Economics* **132**(3), 76–96.

Li, Xuelin, Tong Liu and Lucian A. Taylor (2023), 'Common ownership and innovation efficiency', *Journal of Financial Economics* **147**(3), 475–497.

Schwenkler, Gustavo and Hannan Zheng (2023), News-driven peer co-movement in crypto markets. Working Paper.

Schwenkler, Gustavo and Hannan Zheng (2024), The network of firms implies by the news. Working Paper.

Shobha, Gangadhar and Shanta Rangaswamy (2018), Chapter 8 - machine learning, *in* V. N.Gudivada and C.Rao, eds, 'Computational Analysis and Understanding of Natural Languages: Principles, Methods and Applications', Vol. 38 of *Handbook of Statistics*, Elsevier, pp. 197–228.

	Mean	Median	Std dev.	Min.	Max.
Whole population					
Market capitalization (million USD):	1,229.03	102.98	8,327.52	0.04	504,183.54
Total assets (million USD):	3,864.60	144.52	43,794.00	0.03	2,168,173.11
Total debt (million USD):	3,120.66	61.67	40,932.84	0.00	2,021,128.24
Book leverage:	56.97%	52.81%	89.27%	0.00%	7,418.05%
Cash holdings (million USD):	224.55	11.16	3,561.09	-37.10	361,034.90
Quarterly net income (million USD):	16.47	-0.03	177.03	-5,829.59	5,980.48
Quarterly sales (million USD):	367.03	20.60	2,340.85	-191.45	99,938.93
Quarterly cost of goods sold (million USD):	262.62	12.62	1,910.59	-10.64	89,752.48
Average monthly return:	0.27%	1.02%	4.73%	-89.62%	224.48%
Average monthly volatility:	4.13%	3.60%	4.23%	0.00%	346.81%
Quarterly dividends per share (USD):	0.06	0.00	0.23	0.00	24.56
Monthly analyst coverage:	3.14	2.06	2.93	1.00	32.24
Sample mentioned in news data					
Market capitalization (million USD):	1,924.11	139.15	11,990.35	0.09	504,183.54
Total assets (million USD):	5,901.75	188.15	60,704.43	0.03	2,646,347.38
Total debt (million USD):	4,818.93	83.73	56,664.74	0.00	2,454,577.08
Book leverage:	54.22%	52.78%	43.51%	0.00%	2,916.68%
Cash holdings (million USD):	333.39	11.00	4,233.33	-36.90	287,893.14
Quarterly net income (million USD):	25.79	0.58	230.84	-2,691.67	6,534.50
Quarterly sales (million USD):	548.66	36.93	2,977.04	-516.75	86,445.53
Quarterly cost of goods sold (million USD):	394.46	22.70	2,437.82	-44.96	82,866.47
Average monthly return:	1.98%	1.11%	20.70%	-87.37%	988.97%
Average monthly volatility:	4.15%	3.31%	4.55%	0.00%	223.39%
Quarterly dividends per share (USD):	0.07	0.00	0.20	0.00	6.00
Monthly analyst coverage:	3.76	2.50	3.60	1.00	30.68
Monthly news mentions:	3.05	1.50	8.21	1.00	543.00
Monthly news links:	4.97	3.54	4.45	1.00	35.64

Table 1: *Summary statistics of firm population & sample covered in news data.* The population includes 26,374 distinct firms while the news-covered sample covers 13,526 distinct firms. The data span the time period between January 1, 1981, and December 31, 2023. The above statistics are time series moments over months in which a firm was alive during our sample. Market data come from CRSP, fundamentals from Compustat, and analyst following from I/B/E/S, all of which we access through the WRDS API. Total debt is total liabilities. Book leverage is the ratio of total debt over total assets. Analyst coverage counts the number of analysts following a firm.

"United (UAL) is the only American airline that operates Boeing (BA) 777s equipped with the PW4000 engine series, and the airline said on Sunday that it was grounding those 24 planes in its active fleet while it awaited F.A.A. guidance."			
Credit:	0.2%	Financing:	0.0%
Supply chain:	50.2%	Peer:	0.0%

"The many local Xinjiang police bureaus on the list buy commercial American technology like Intel (INTC) microchips and Microsoft (MSFT) Windows software, according to procurement documents."			
Credit:	0.0%	Financing:	1.2%
Supply chain:	13.3%	Peer:	0.1%

Table 2: *Examples of firm link labelling.* We provide two example sentences that are labelled by our methodology, the top one as "supply chain" and the bottom one as "other." We mark in yellow a relationship that is discarded by our methodology because it does not achieve more than 50% link likelihood based on the corresponding neural network. We mark in green a relationship that passes this test.

Accuracy rates	In-sample Accuracy	Cross-validation Accuracy	Out-of-sample Precision	Recall	F1 Score
Credit	99.53%	98.40%	100%	95%	97%
Financing (ex credit)	97.33%	93.60%	94%	92%	93%
Supply chain	97.40%	93.20%	90%	80%	85%
Peer	98.27%	96.80%	93%	80%	86%
Other			72%	94%	81%
Overall			90%	88%	89%
Data points	1,500	500	500		

Confusion matrix	Credit	Fin. (ex credit)	Supply chain	Peer	Other
Credit	95	0	1	0	4
Financing (ex credit)	0	92	0	1	7
Supply chain	0	5	81	1	13
Peer	0	0	7	80	13
Other	0	1	1	4	94

Table 3: *Accuracy rates & confusion matrix of our methodology.* We construct a synthetic dataset of sentences that describe one of the five different relationships between two fictional firms using ChatGPT. We train our model using 60% of this sample as training data, which we also use to measure the in-sample error. We determine the number of epochs to train our models based on the minimum cross-validation error, which we evaluate on 20% of the sample. Finally, we measure the out-of-sample error of our methodology on 20% of our synthetic data that was not used for training or cross validation. In this table, "Precision" measures the ratio of true positives over all predicted positives, "Recall" measures the ratio of true positives out of all actual positives, and "F1 Score" gives the harmonic mean between precision and recall. The confusion matrix counts the number of instances in the test data in which our category assigned a label given the true label.

Firm 1	Firm 2	Count	Sample sentence
Credit links: 14,419 links among 2,823 unique firms.			
Chrysler	Ford	295	"*Ford and Chrysler will offer five-year zero percent deals on many of their vehicles as well.*"
Chrysler	General Motors	104	"*The bleakest numbers will most likely come from General Motors and Chrysler, which both received billions of dollars in loans from the federal government at the end of December to help them remain solvent.*"
Bank of America	Citigroup	100	"*In a case where a lead underwriter, Citigroup, settled for $2.65 billion, perhaps Bank of America, a participating underwriter, felt the wise course was to settle and cut off any future exposure to liability.*"
Citigroup	Enron	98	"*Citigroup has put its Enron loan exposure at $1.2 billion, though it also has some insurance-related obligations.*"
Ford	General Motors	91	"*For both Ford and General Motors, lending money to car buyers has been more lucrative than selling the cars.*"
Financing (ex credit) network: 34,527 links among 4,898 unique firms.			
Microsoft	Yahoo	466	"*After a months-long standoff, Microsoft and Yahoo are now engaged in active merger talks, people involved in the discussions said Friday.*"
Time Warner	AT&T	218	"*Until last week , AT&T 's pending $85.4 billion acquisition of Time Warner seemed destined to close by the end of the year.*"
Yahoo	Google	173	"*Before the settlement, Yahoo owned 5.5 million shares of Google and had planned to sell 550,000 in the auction.*"
Microsoft	Google	157	"*According to one executive, Mr. Parsons called Steven A. Ballmer, Microsoft 's chief executive, at 10:30 a.m. Friday to tell him that the deal that Microsoft had so eagerly sought – and had thought it had won – was going to Google.*"
Chrysler	Ford	155	"*It bought Ford in 1981 and Chrysler in 1982 and held them for twelvefold gains.*"
Peer network: 30,958 links among 2,942 unique firms.			
Coca Cola	Pepsi	972	"*Coke, after all, started as a fountain business, whereas Pepsi began selling cola in bottles, avoiding a fight with its rival at the drugstore lunch counter.*"
Microsoft	Google	793	"*Competition is heating up between Google, the world's dominant search engine, and Microsoft, which recently entered the Web search market.*"
Apple	Google	767	"*Apple's strategy carries risks, however, especially in developing countries where smartphone sales are growing briskly but its market share is a blip compared with devices running Google's Android software.*"
Apple	Microsoft	475	"*Apple partisans have long asserted that Microsoft stole the basic look and feel of Windows from Apple's operating system, Mac OS.*"
Yahoo	Google	452	"*Google and Yahoo each dominate one segment of the online advertising market.*"

Table 4: *Most frequently mentioned links (Part I).* This table lists the five most frequently identified links for each link category. It also provides a sample sentence for each link.

Firm 1	Firm 2	Count	Sample sentence
Supply chain network: 19,930 links among 3,422 unique firms.			
Chrysler	Ford	192	"*He manages the factory, which forges tie rods and other steering links for trucks, including pickup trucks assembled by Ford and Chrysler.*"
Microsoft	Amazon	152	"*Microsoft announced a plan for Amazon.com to open an electronic book store on its Web site that would distribute free copies of Microsoft's Reader software.*"
Apple	Microsoft	149	"*During the trial in 1998, Apple's lead software designer, Avi Tevanian, described his company's efforts to persuade PC makers to bundle Apple's QuickTime media software with their machines and how Microsoft demanded that Apple 'knife the baby' – in other words , drop the QuickTime bundling effort.*"
Apple	Amazon	136	"*Millions of people unknowingly interact with A.W.S. every day when they stream movies on Netflix or store photos on Apple's iCloud, services that run off Amazon's machines.*
Intel	Microsoft	101	"*So inextricably linked are the dual monopolies of the Intel Corporation's microprocessors and the Microsoft Corporation's Windows operating systems in desktop computing that industry watchers commonly describe them in a single word : Wintel.*"
Other network: 76,359 links among 5,789 unique firms.			
Chrysler	Ford	1713	"*Mr. Lutz, who was born in Switzerland, was also a top executive at Ford Motor and Chrysler during his 47-year automotive career.*"
Apple	Microsoft	949	"*In the 1990s, Apple was teetering close to bankruptcy and Microsoft dominated the personal computer industry.*"
Ford	General Motors	772	"*The pricing action follows a price increase by General Motors on April 3 by an average of 3.5 percent and by Ford yesterday by an average of 2.1 percent.*"
Apple	Google	630	"*Mr. Obama arrived in the heart of Silicon Valley at a time of great tension with companies here, including Apple and Google, both represented at the event.*"
Intel	Microsoft	572	"*Even Intel and Microsoft were down for much of the day, but they rallied late, to recover.*"

Table 5: *Most frequently mentioned links (Part II)*. This table lists the five most frequently identified links for each link category. It also provides a sample sentence for each link.

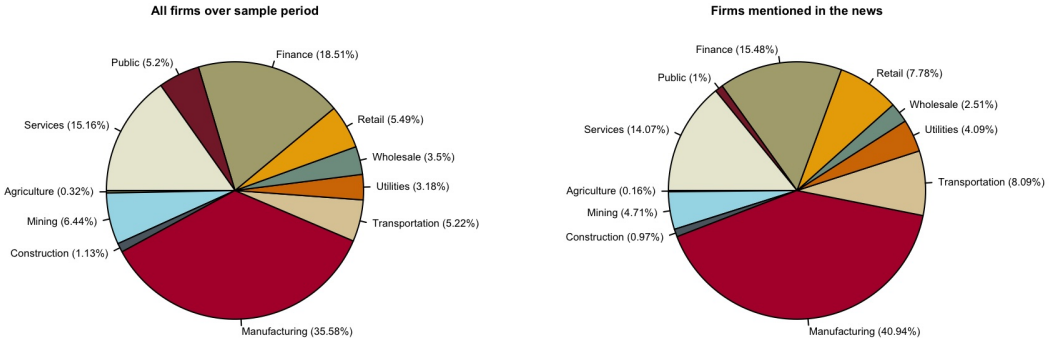


Figure 1: *Industry distribution in population and news-mentioned sample.* The population includes 26,374 distinct firms and the sample of firms mentioned in the news covers 13,526 unique firms. The data cover the time period between January 1, 1981, and December 31, 2023.

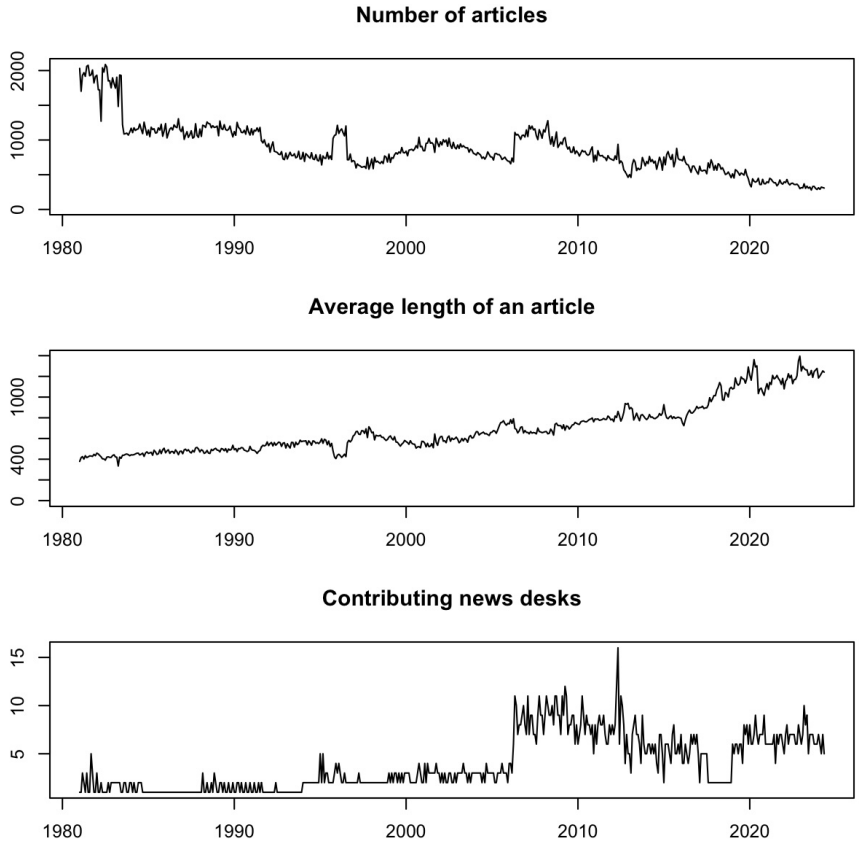
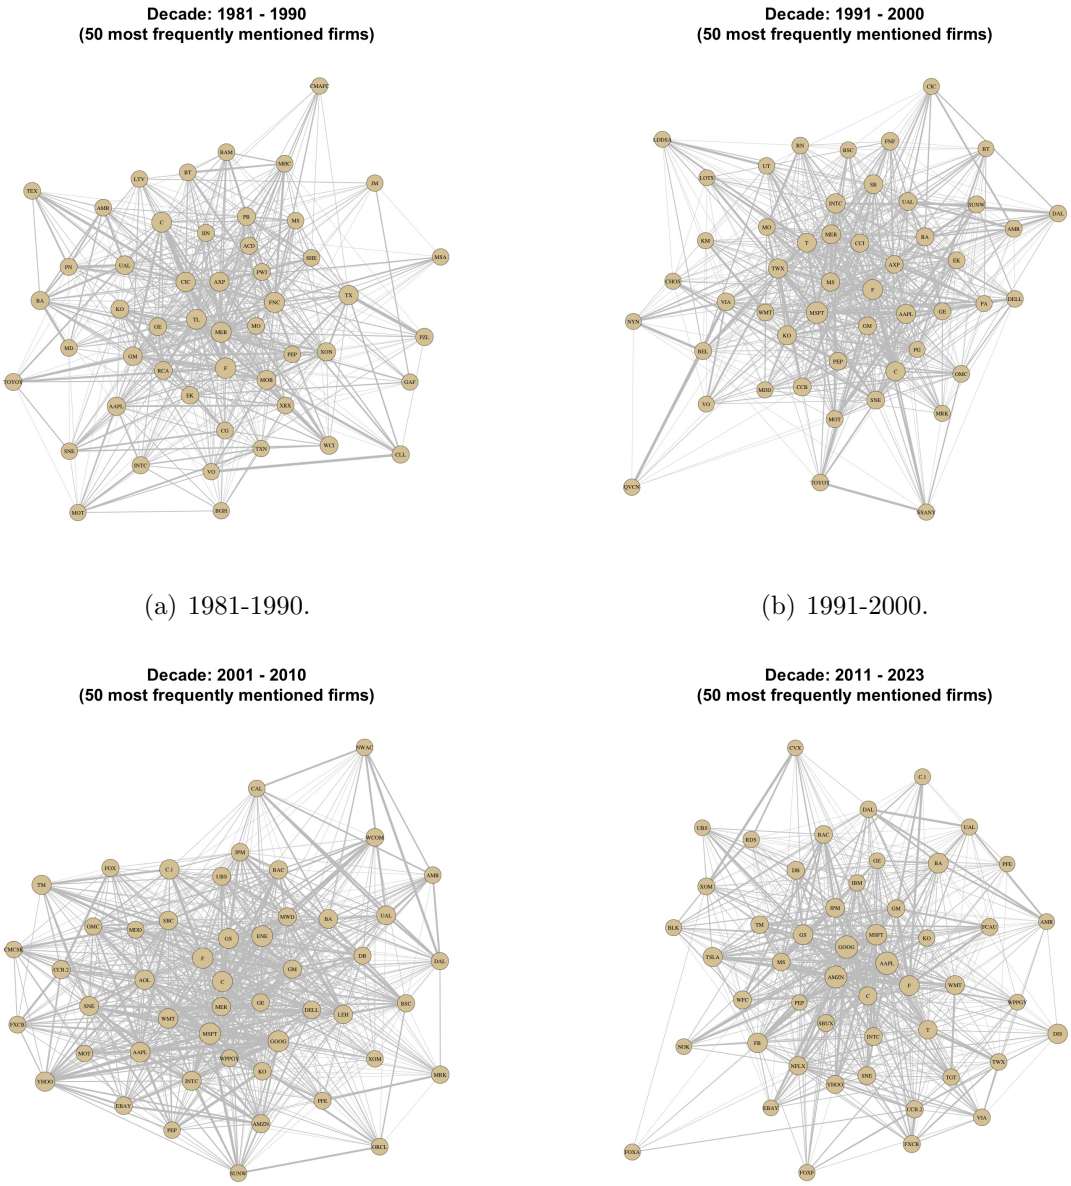


Figure 2: *New York Times sample.* The plot shows the monthly number of articles contained in our news sample, together with the average length of an article and the number of contributing news desks. We measure the length of an article through the number of words it contains.



(a) 1981-1990.

(b) 1991-2000.

(c) 2001-2010.

(d) 2011-2023.

Figure 3: *Time series of news-implied networks in our news data sample.* These graphs show the networks extracted from the news using the methodology of Schwenkler and Zheng (2024), broken down by decades. The size of a node (link) is proportional to the logarithm of one plus the number of firm (link) mentions in the news.

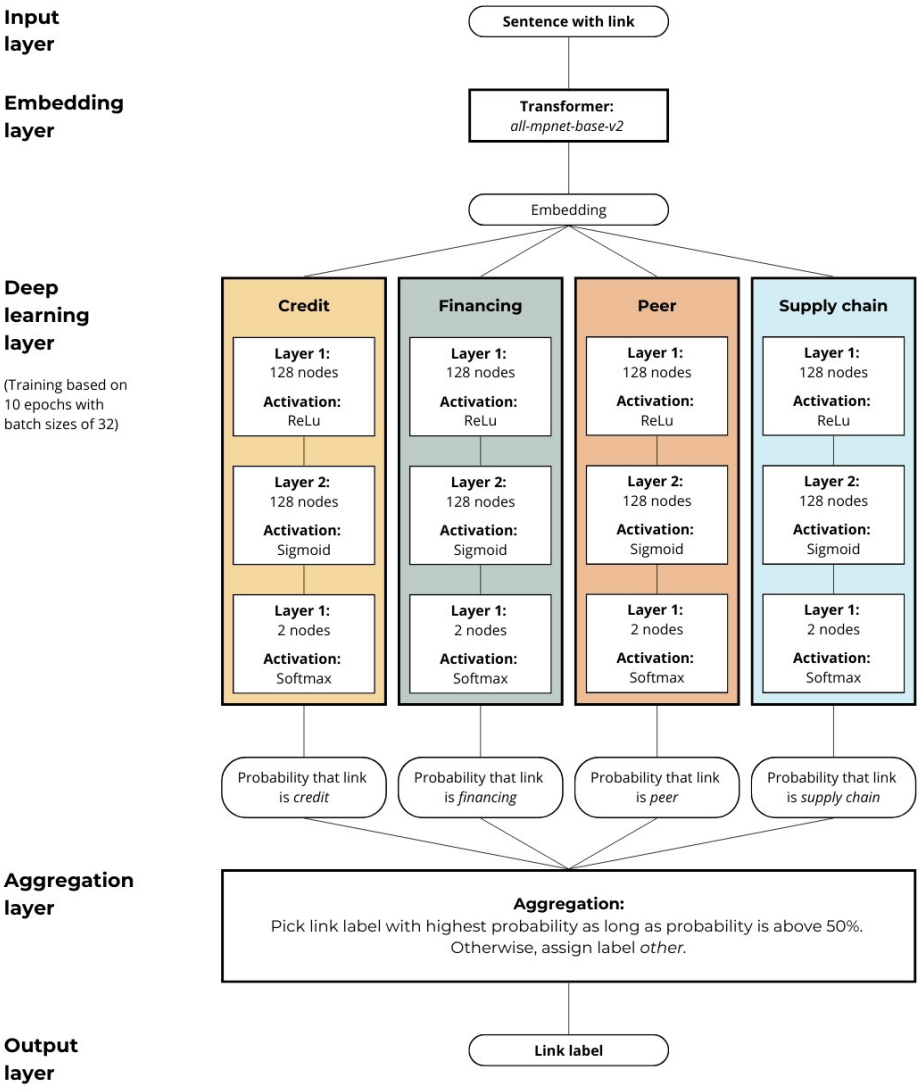


Figure 4: *Neural network architecture.* This figure shows the architectures of the model that underpins our link labelling methodology.



(a) Credit.

(b) Financing (ex credit).

(c) Peer.

(d) Supply chain.

Figure 5: *Most frequent verbs.* These word clouds display the 100 most frequent verbs in sentences that our methodology tags as describing one of the four named firm relationships. We use the universal point-of-speech tagger from the cleanNLP package available in R. We restrict ourselves to the English language model from spaCy (*"en_core_web_sm"*). We lemmatize all verbs prior to displaying them in a word cloud. We exclude the verb "say," which is the most frequently mentioned verb for all link types.



(a) Credit.



(b) Financing (ex credit).

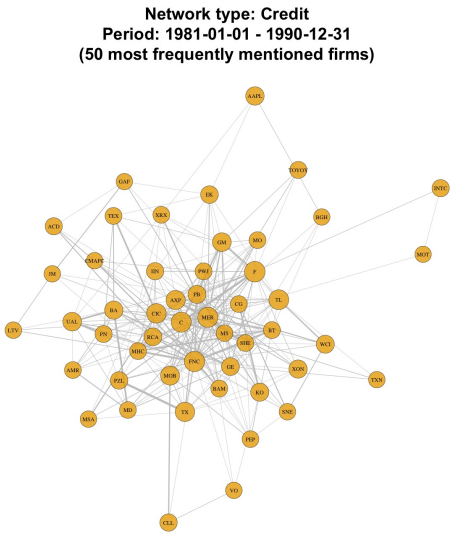


(c) Peer.

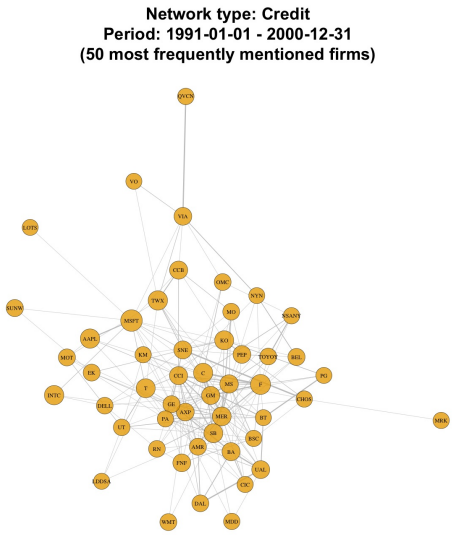


(d) Supply chain.

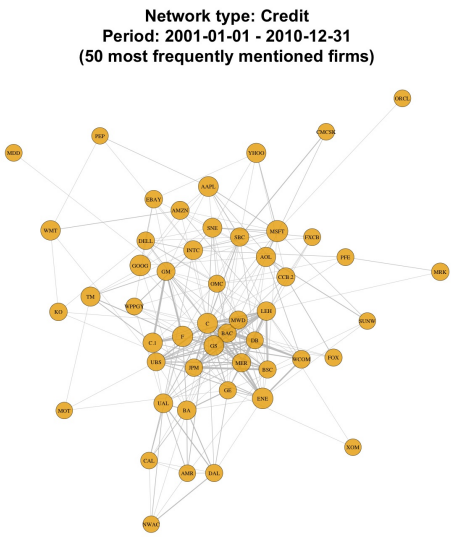
Figure 6: *Most frequent nouns.* These word clouds display the 100 most frequent nouns in sentences that our methodology tags as describing one of the four named firm relationships. We use the universal point-of-speech tagger from the cleanNLP package available in R. We restrict ourselves to the English language model from spaCy (*"en_core_web_sm"*). We lemmatize all nouns prior to displaying them in a word cloud. We exclude proper nouns as well as the noun "company," which is the most frequently mentioned noun for all link types.



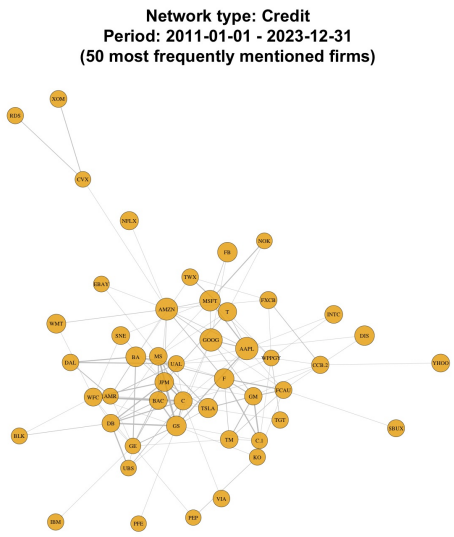
(a) 1981-1990.



(b) 1991-2000.

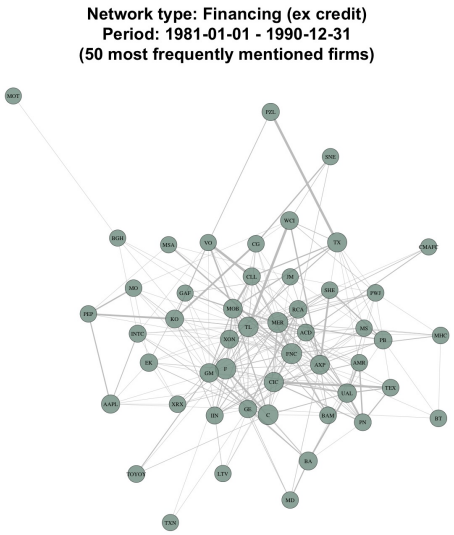


(c) 2001-2010.



(d) 2011-2023.

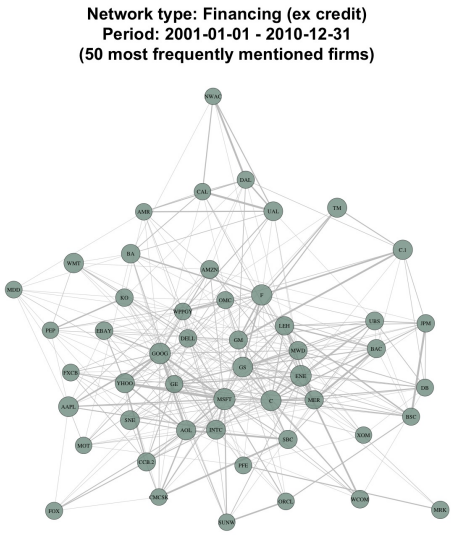
Figure 7: *Time series of news-implied credit network.* These graphs show the credit networks extracted from the New York Times sample using our methodology, broken down by decades. The size of a node (link) is proportional to the logarithm of one plus the number of firm (link) mentions in the news.



(a) 1981-1990.



(b) 1991-2000.

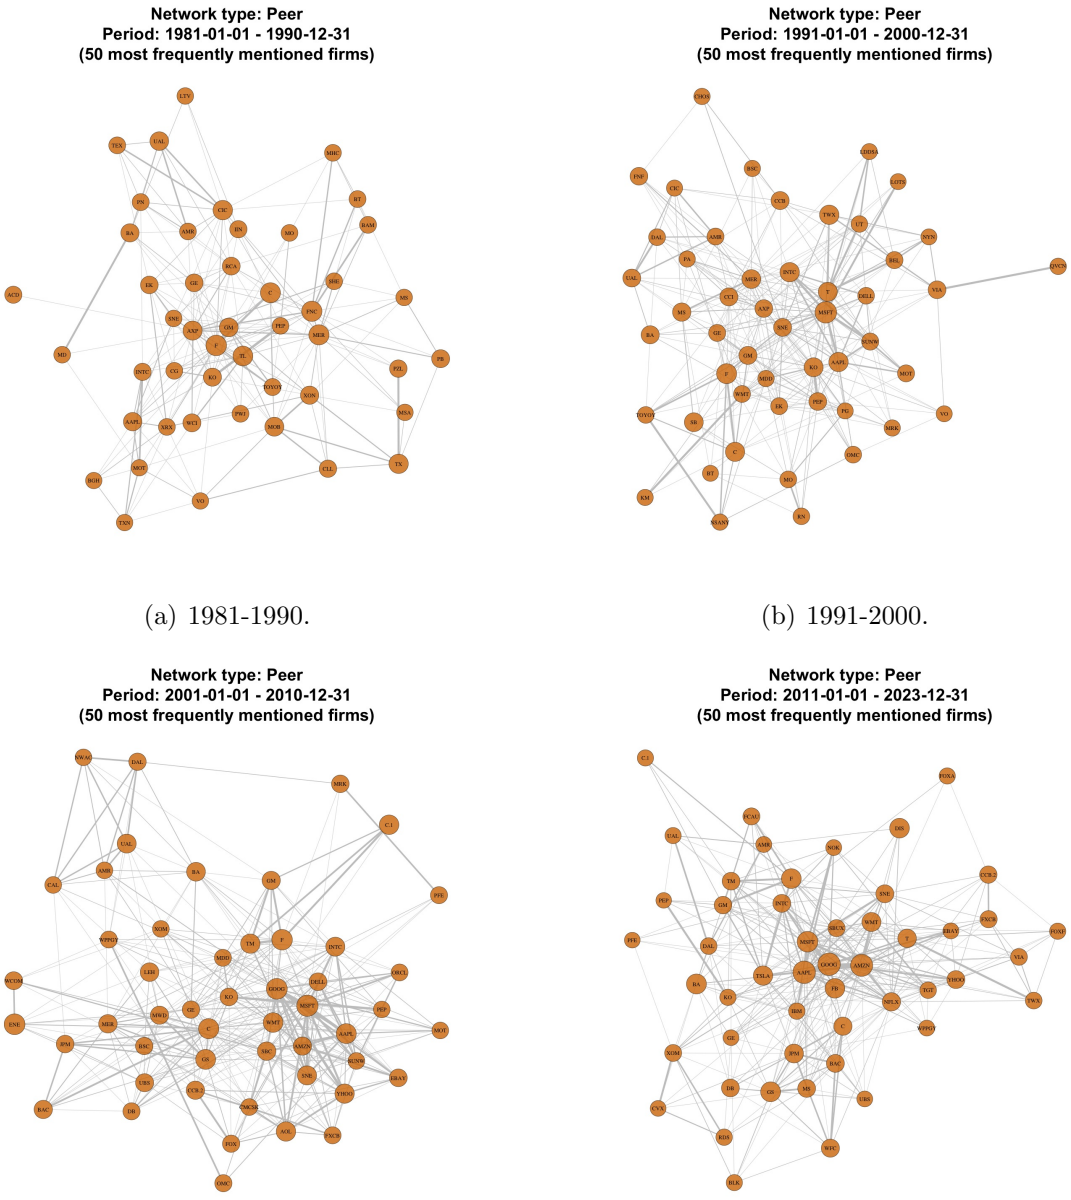


(c) 2001-2010.



(d) 2011-2023.

Figure 8: *Time series of news-implied financing (non credit) network.* These graphs show the networks of financing relationships outside of credit extracted from the New York Times sample using our methodology, broken down by decades. The size of a node (link) is proportional to the logarithm of one plus the number of firm (link) mentions in the news.



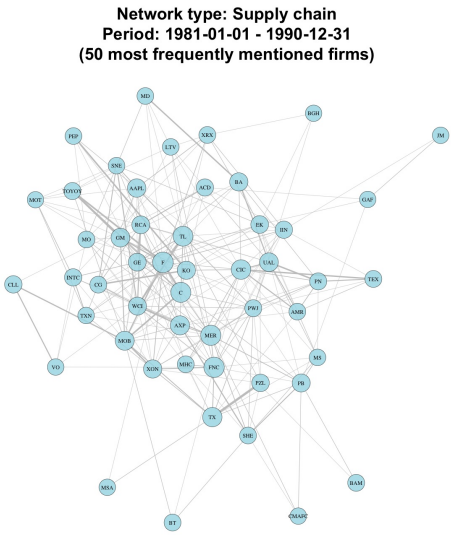
(a) 1981-1990.

(b) 1991-2000.

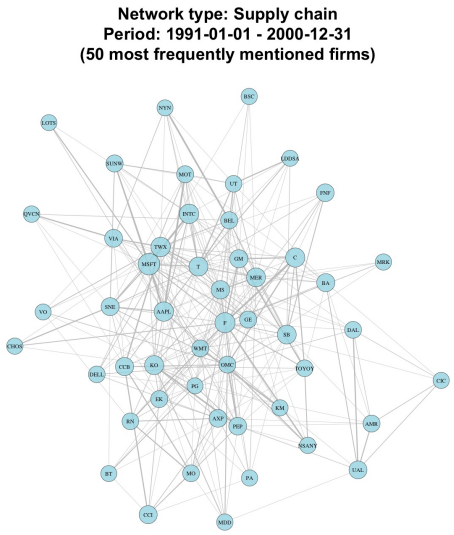
(c) 2001-2010.

(d) 2011-2023.

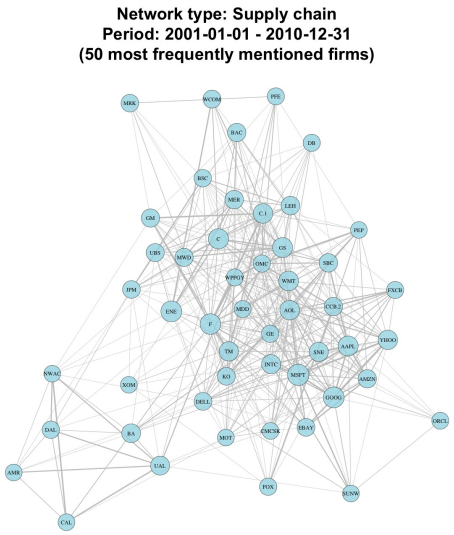
Figure 9: *Time series of news-implied peer network.* These graphs show the peer networks extracted from the New York Times sample using our methodology, broken down by decades. The size of a node (link) is proportional to the logarithm of one plus the number of firm (link) mentions in the news.



(a) 1981-1990.



(b) 1991-2000.

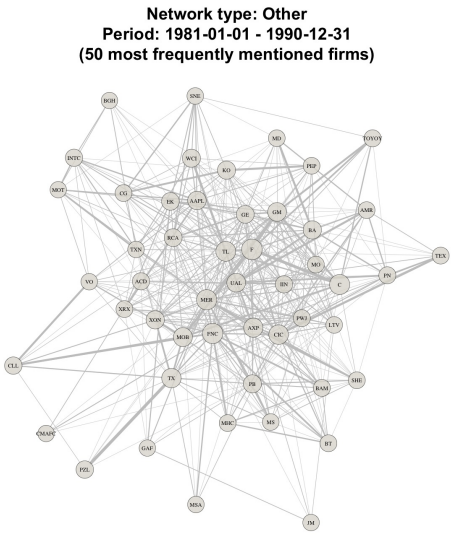


(c) 2001-2010.

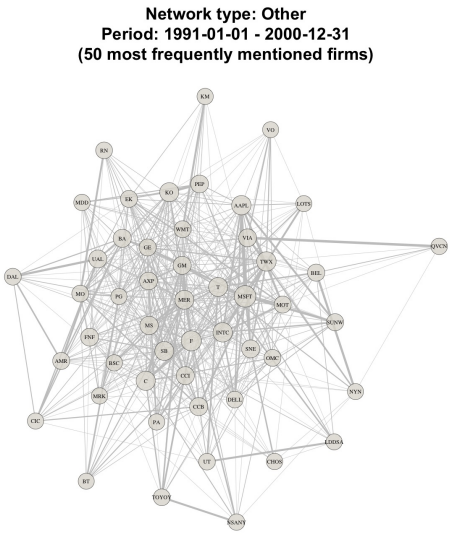


(d) 2011-2023.

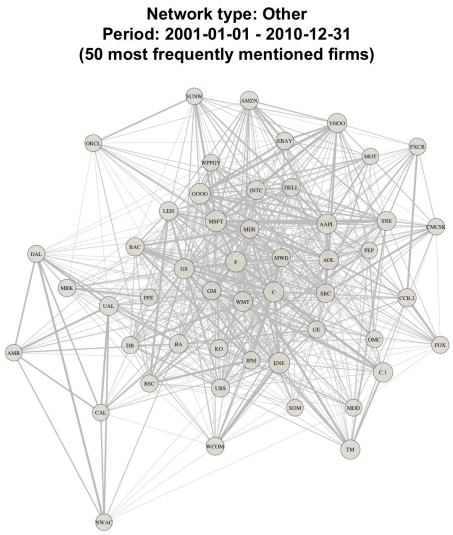
Figure 10: *Time series of news-implied supply chain network.* These graphs show the supply chain networks extracted from the New York Times sample using our methodology, broken down by decades. The size of a node (link) is proportional to the logarithm of one plus the number of firm (link) mentions in the news.



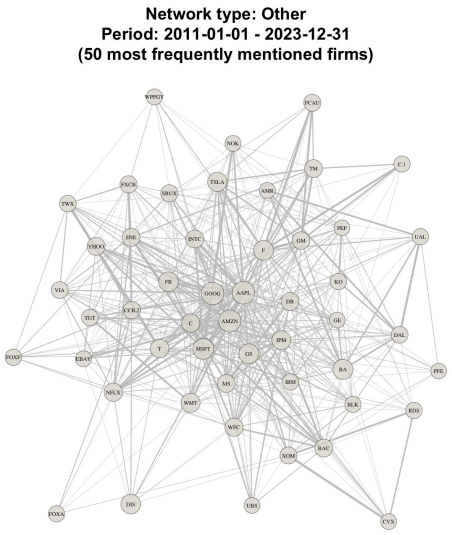
(a) 1981-1990.



(b) 1991-2000.

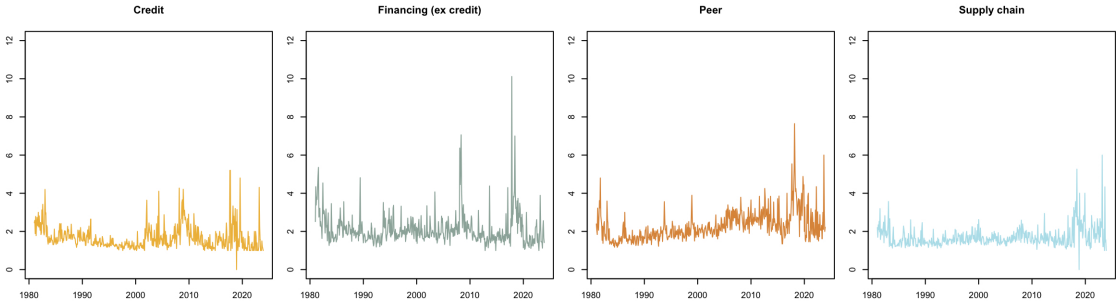


(c) 2001-2010.

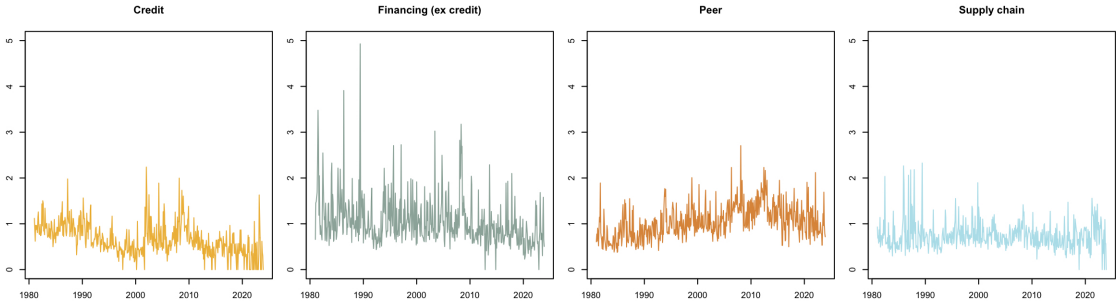


(d) 2011-2023.

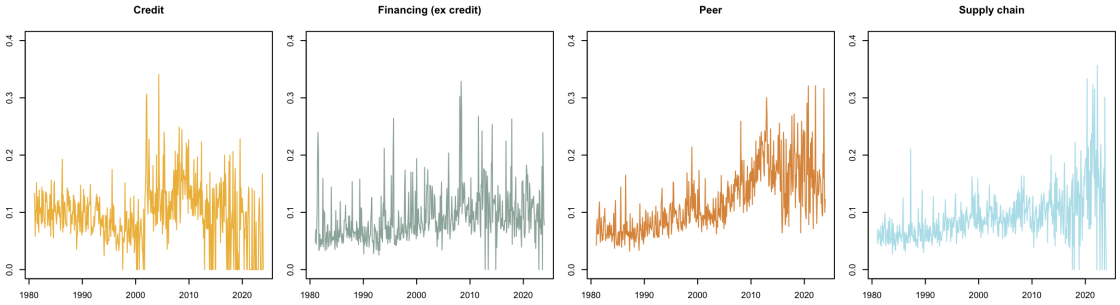
Figure 11: *Time series of news-implied other network.* These graphs show the networks of other links extracted from the New York Times sample using our methodology, broken down by decades. The size of a node (link) is proportional to the logarithm of one plus the number of firm (link) mentions in the news.



(a) Average degree.

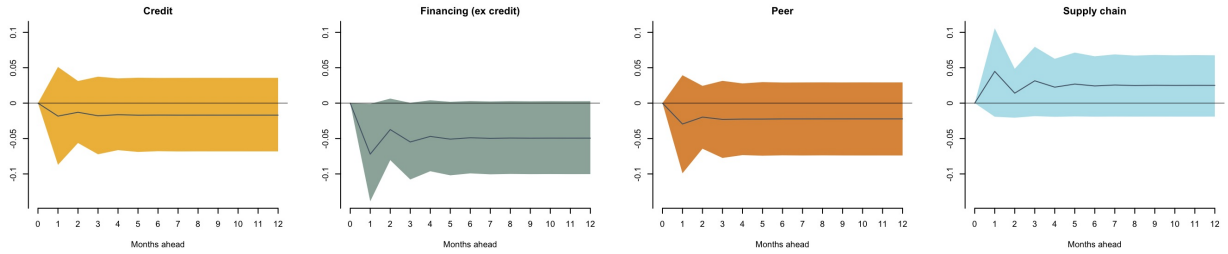


(b) First-order interconnectivity.

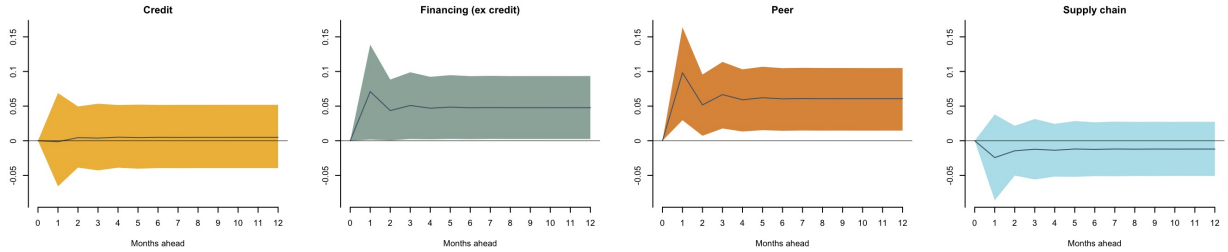


(c) Second-order interconnectivity.

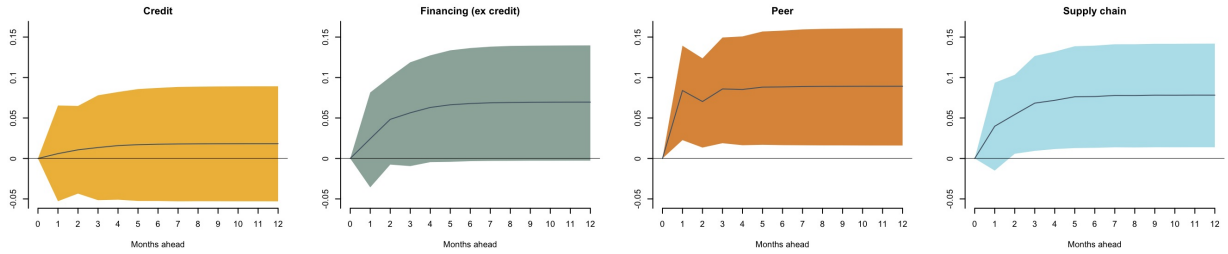
Figure 12: *Time series of connectivity metrics.* Given a monthly network, we follow Schwenkler and Zheng (2024) and evaluate the average degree as $\bar{d}_t = \frac{1}{N_t}\sum_{n=1}^{N_t} d_t^n$, where N_t is the number of nodes in the network of Month t and $d_t^n = \sum_{j=1}^{N} w_t^{j,n}$ is the number of links of Node n in Month t. Here, $w_t^{j,n}$ is the number of identified links that connect Nodes j and n in Month t. We evaluate first-order interconnectivity measure as $\frac{1}{\bar{d}_t}\left(\frac{1}{N-1}\sum_{n=1}^{N_t}(d_t^n - \bar{d}_t)^2\right)^{1/2}$ and second-order interconnectivity as $\left(\sum_{n=1}^{N}\sum_{j\neq n}\sum_{k\neq j,n}\frac{w_t^{j,n}}{N_t}\frac{w_t^{k,n}}{N_t}\right)^{1/2}$.



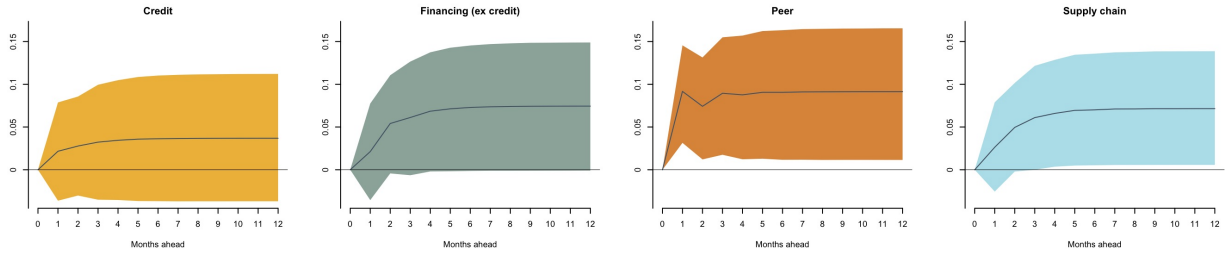
(a) S&P 500.



(b) VIX*.

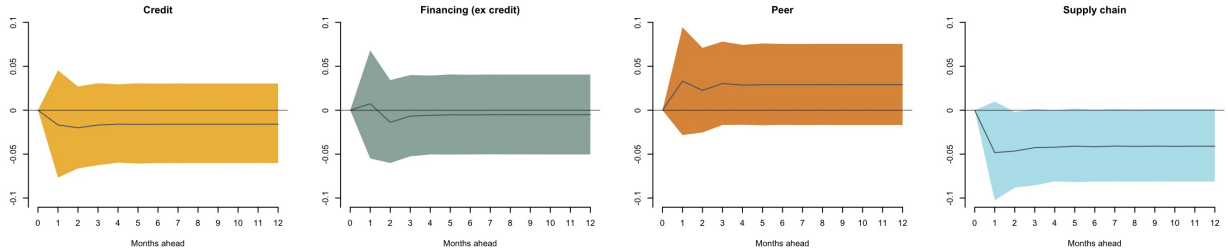


(c) AAA credit spread.

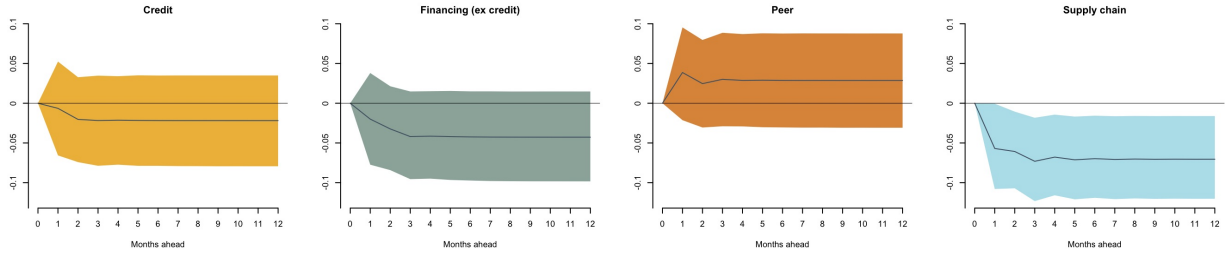


(d) BAA credit spread.

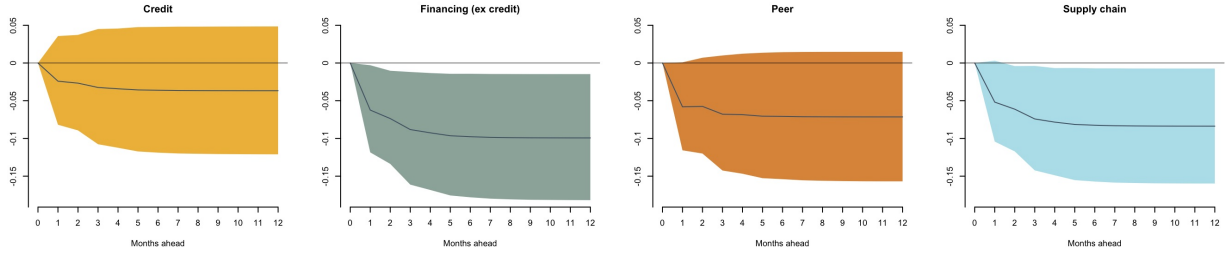
Figure 13: *Cumulative impulse response functions for orthogonal shocks to the average degree of the different networks.* The grey lines give the impulse responses and the colored areas give 90% confidence bands computed via bootstrap with 10,000 bootstrap samples. We consider one-standard-deviation shocks to the orthogonal component of the impulse variable. The * indicates that this applies for an adjusted VIX time series that includes VXO observations between 1986 and 1990.



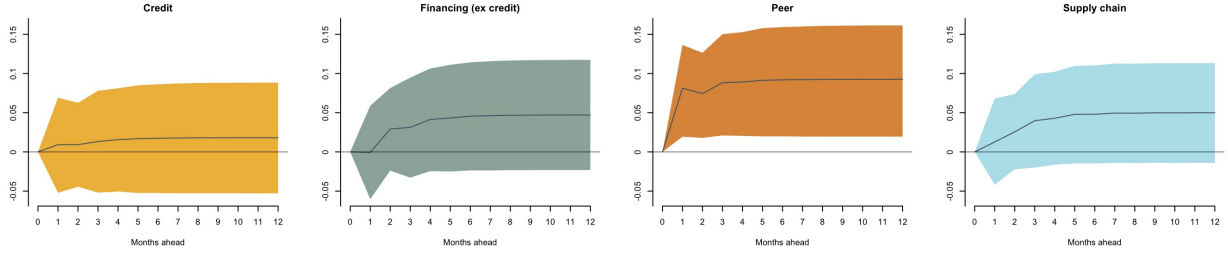
(a) Consumption growth.



(b) Industrial production growth.



(c) Yield curve level.



(d) Yield curve slope.

Figure 14: *Cumulative impulse response functions for orthogonal shocks to the second-order interconnectivity measure of the different networks.* The grey lines give the impulse responses and the colored areas give 90% confidence bands computed via bootstrap with 10,000 bootstrap samples. We consider one-standard-deviation shocks to the orthogonal component of the impulse variable.