The Different Networks Of Firms Implied By The News

V. Hilt & G. Schwenkler*

October 15, 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, risk management, natural language processing, large language models. JEL codes: C80, D20, E32, G12, G32.

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

†We are grateful to Will Cong, Sanjiv Das, Daniel Giamouridis, Gerard Hoberg, Christos Makridis, and Alireza Tahbaz-Salehi; as well as 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 Factset and Compustat, that offer data on supply chain relationships, the existing datasets are updated infrequently and only cover a subset of firms. Some researchers have developed indirect approaches to obtain 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`, enabling 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 achieves a precision score of 90%, a recall score of 88%, and an F1 score of 89%. Our methodology does not suffer under model collapse, which is common among generative AI models as recently highlighted by Shumailov et al. (2024), because it is not generative in nature: we only classify the firms links that are published in financial news. Even though the methodology we introduce in this paper only identifies five different types of firms links, we present two extensions of our methodology: one that identifies additional types of linkages and another one that provides confidence-weighted links. A related extension 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*. Non-credit financing links are identified in sentences that contain words like *acquire*, *buy*, *own*, *share*, and *stock*, suggesting that this type of link captures direct and indirect equity relationships across firms. 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*. These keywords are identified in actual sentences published in the New York Times in which we identify firm links, highlighting the interpretable nature of our firm links.[1] We further validate our networks by comparing them to existing network databases. We find that our peer network closely correlates with the TNIC peer network of Hoberg and Phillips (2016) while our supply chain network positively correlates with the customer segments network from Compustat. Compared to these existing networks, however, our networks are available at arbitrary frequencies. We also find that a variance decomposition network due to Demirer et al. (2018), which primarily captures stock return correlations, is positively associated with our peer and other networks. Our results show that our methodology is able to identify actual firm links in real news data, going beyond the artificially clean and synthetic dataset we used to train our models.

We apply our newly obtained 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 and financial 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 shocks to the supply chain network primarily predict macroeconomic fluctuations. We find that shocks to the degree of interconnectivity of our supply chain network predict transient declines in consumption growth, persistent declines in industrial production, and persistent increases in corporate credit spreads. Extending the existing literature, we find that shocks to the architecture of the peer and financing networks forecast deteriorations in aggregate asset prices. Shocks to the degree of interconnectivity of our peer network predict persistent increases in credit spreads as well as a steepening of the US yield curve. Shocks to the degree of interconnectivity of the financing network

[1] Interpretability is a desirable feature of textual analyses in economic applications, as highlighted by Bybee et al. (2024), Cong et al. (2020), Cong et al. (2024), Gentzkow et al. (2019), and others.

predict declines in the level of the US yield curve while shocks to its density predict short-term declines in the S&P 500 and persistent increases in the VIX. These 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 links 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 aggregate risks. They complement new research that highlights the power of news to forecast aggregate outcomes; see Bybee et al. (2024), Chahrour et al. (2021), Engle et al. (2020), Larsen et al. (2021), and others.

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 and compares them to existing network databases. 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`. The data repository also contains our Online Appendix.

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. 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.

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. 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 collect news articles published in the New York Times after January 1, 1981, by scraping the New York Times API. We do this because the New York Times archive is freely accessible while commercial news data bases, like Factiva, are only accessible in exchange for large fees.[2] However, our methodology can be applied to any text data set. Schwenkler and Zheng (2023) apply an extension of our methodology to crypto news articles from CryptoCompare.

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 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

[2]When we first started this research, we did not have access to a commercial news data vendor.

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.[3]

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 than the average firm in the population, 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 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 network over our sample period is showcased in Figure 3. Similar to Schwenkler and Zheng (2024) albeit using a different news data source, we find that the news-implied network has a star architecture as in Acemoglu et al. (2012). We observe that big banks, like Goldman Sachs (ticker: GS), Morgan Stanley (MS), and UBS, are 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

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

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 links. 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." Our model can be viewed as one big, segmented neural network. Figure 4 showcases the architecture of our model. Online Appendix B provides details.

Table 2 provides two examples of how our methodology works, where we mark in parentheses the tickers of the two firms that were identified in the sentence. Our methodology uniquely labels the first sentence as describing a supply chain relationship, which is reasonable given the information contained in the sentence. On the other hand, our approach does not assign more than 50% likelihood to any of the four firm links for the second sentence. As a result, we label it "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 may result in unbalanced panels that may confound the learning of our model. For example, peer linkages are more frequently mentioned in news articles (when, e.g., the market performance of two competitors is compared) than credit relationships, for which information might be hard to collect by news publishers. Second, manual labelling is costly as it requires human supervision and expertise. It is also subjective because it is based on how individual 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 the ChatGPT 4o model. For each firm link type (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. 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 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. Note that our methodology does not suffer under model collapse as pinned by Shumailov et al. (2024) because we do not generate language. Instead, our model classifies language published in actual news articles. All in all, the results this section show that our methodology is highly accurate.

[4]We choose to stop the training after 10 epochs with batch sizes of 32 because we observe that the cross-validation error declines with more than 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 Extensions

We present two extensions of our methodology. The first one relaxes the fact that we only assign binary links in our networks. Our methodology is based on probabilities for each link type (see Table 2). If a user wanted to construct confidence-based links that are stronger the more confident the model is in the assigned link type, then the user could weight the links with the link type probability determined by our neural network model. Our data repository contains all link type probabilities assigned by our methodology to the sentences in our news data.

Another extension allows for the identification of additional link types. We choose to restrict our model to five different link types to limit the scope of the paper. But, in principle, our methodology can identify as many firm relationships as it is fed with training data. Suppose that a user 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, a user can follow these steps:

(1) 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, ask the LLM to include exactly two fictional firms and a fictional location in the sentences. Also ask to write the sentences in a way that would appear in common financial news outlets. Generate the same number of sentences as had previously been generated for all other link types in order to keep the panel balanced.

(2) Run sentences through the same embedding used by our model.

(3) Construct an additional neural network 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.

(4) Re-train all other classifiers with the new, extended training set replicas.

(5) Run unlabelled link sentences through the newly trained classifiers and assign the link label that achieves the highest probability, as long as the probability is above 50%. Otherwise,

assign the label "other."

As this example shows, 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. But there are some caveats. 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. Figure 7 shows the four named firm networks, aggregated over the whole sample period. Figure A.1 in the Online Appendix shows the network of other links.

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*, *credit*, *yield*, *borrow*, and *subordinate* are among the most frequently mentioned keywords in sentences in which we identify credit relationships. In contrast, we find that *rival*, *industry*, *competitor*, *compete*, and *dominate* are among the most frequent keywords in sentences in which we identify peer links.

Going into more detail, we observe that the credit network has the typical star architecture posited by Acemoglu et al. (2012). Big banks like Citibank (ticker: C), Goldman Sachs (GS), and JP Morgan (JPM) are highly central and interconnected, surrounded by smaller banks and industry-specific clusters. A star architecture for the credit network is consistent with Bernanke et al. (1999) and Carvalho and Gabaix (2013), who highlight the central role of the banking industry. The non-credit financing network is more dispersed. Table 4 shows that many of the links in the financing network are equity links, which can be of the form of parent-subsidiary, direct equity investments, or indirect common ownership links. This is also validated by Figures 5 and 6, which show that keywords like *acquire*, *buy*, *own*, *share*, and *stock* are frequently mentioned in sentences in which we identify financing links. Because of this, we interpret the non-credit financing network as an equity network.

The peer network showcases strong bilateral links, e.g., between AT&T (ticker: T.1) and Verizon (VZ), Merck (MRK) and Johnson & Johnson (JNJ), Bank of America (BAC) and Wells Fargo (WFC), or Mastercard (MA) and Visa (V). The sentences of Table 5 indicate that the supply chain network includes direct sale linkages as well as indirect linkages through common suppliers. This network showcases some clusters. For example, in the center of the network we observe a cluster between Coca-Cola (ticker: KO), Pepsi (PEP), Procter & Gamble (PG), Johnson & Johnson (JNJ), Walmart (WMT), and CVS. Finally, Figure A.1 in the Online Appendix shows that the network of other links showcases more of a mesh architecture. Table 5 shows that the this network captures alternative links, like common geographic locations or common executives.

3.1 Comparison to alternative networks

We compare our networks to the networks extracted from three alternative databases to which we have access: *(i)* a peer network based on the Text-based Network Industry Classifications (TNIC) measure of Hoberg and Phillips (2016), *(ii)* a supply chain network based on the Compustat Segments data, and *(iii)* an equity correlation network based on the variance decomposition methodology of Demirer et al. (2018). We aggregate links over sample periods that cover the range of each of the alternative networks as well as our networks, and we restrict ourselves to the 500 largest firms over these samples. Online Appendix E summarizes how we construct these networks.

We run two regression analyses for each of the alternative networks. First, we consider the extensive margin and run logit regressions for the likelihood of observing a link in the alternative networks based on the strength of the link in our networks. Second, we consider the intensive margin and restrict ourselves to the subset of links that are observed in the alternative network. Then, we regress the strength of the link in the alternative network on the strength of the link in our networks. We measure the strength of a link in our networks through the logarithm of one plus the number of the times the links is observed over the sample. We measure the strength of a link in one of the alternative networks as follows:

- For the TNIC network, we take the average baseline TNIC score over the sample from the Hoberg-Phillips Data Library.
- For the Segments network, we take the logarithm of one plus the sum of all sales among a firm pair, neglecting the direction of the sales.
- For the equity correlation network, we take the average directional connectedness between the two firms, with computations as proposed by Demirer et al. (2018). We censor values at zero whenever an absolute value is below 0.5%.

In all regressions, we include fixed effects for each firm in a link. We also cluster standard errors by firm. Table 6 summarizes our regression estimates.

Starting with the TNIC network, we find that the link likelihood and link strength correlate with the link strength in our peer network. These findings indicate that our peer network captures similar links as those contained in the TNIC dataset of Hoberg and Phillips (2016), validating this link type. However, our peer network is sparser than the TNIC network as measured by the number of links among the 500 largest firms in the sample. This occurs even though we cover a longer time period with higher sampling frequencies. Our estimates also show that the strength of a link in the TNIC network is negatively correlated with the strength of the same link in the credit and supply chain networks. These findings validate that our credit and supply chain networks are unlikely to contain peer linkages. Finally, we find that the other network also correlates with the TNIC network. This finding suggests that the other network may contain information about peer linkages that we neglect.

Turning to the Segments network, we find that both the link indicator and link strength positively correlate with the strength of a link in our supply chain network, though this is only statistically significant for the link strength regression. Our supply chain network is more extensive than the Segments network according to the number of links among the 500 largest firms recorded in both datasets. These findings provide validation for our supply chain network. The strength of a link in the other network positively correlates with the link indicator and negatively correlates with the link strength in the Segments network, which suggests that the other network only contains ambiguous information about supply chain links. We find that the link likelihood in the Segments network negatively correlates with the link strength in the peer network. This suggests that our peer network is unlikely to contain information about supply chain relationships, further validating the peer network.

The estimates of the variance decomposition network suggest that it correlates most closely with our network of other links. They also show that the variance decomposition network weakly correlates with our peer and supply chain networks. These findings indicate that equity correlations may contain information that goes beyond credit, equity, peer, and supply chain relationships, complementing recent findings by Abbassi et al. (2017) and Craig et al. (2023). All in all, the results of this section provide external validation for our approach.

4 Predictability

Schwenkler and Zheng (2024) show that the architecture of the network of firm links extracted from financial news reporting predicts periods of aggregate distress. In this section, we extend these results by dissecting the type of linkages that drive this predictability.

We follow Acemoglu et al. (2012) and Herskovic (2018) 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 8 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.[6] We estimate separate VARs for each of the networks. We follow 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 G 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. But we obtain similar results if we restrict ourselves to the original, unadjusted VIX series. Putting everything together, 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.2–A.16 in the Online Appendix show all impulse response functions. Here, we summarize our findings in Figures 9–11. We neglect the impulse

[6] A similar approach is followed by Baker et al. (2016) and Bybee et al. (2024).

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.

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 network 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 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 and increases in corporate credit spreads. Finally, shocks to the density of the financing and peer networks predict persistent increases in the VIX, while shocks to the density of the financing network predict short-term declines in stock returns.

4.1 Discussion

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 forecast Treasury and corporate bond yields.

Several papers have demonstrated that competitive links enable 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 show 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 our credit network 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 8 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 papers that study 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. We validate our methodology by measuring its out-of-sample accuracy and by comparing the resulting networks to existing databases. 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.

Baker, Scott R., Nicholas Bloom and Steven J. Davis (2016), 'Measuring Economic Policy Uncertainty*', *The Quarterly Journal of Economics* **131**(4), 1593–1636.

Bernanke, Ben S., Mark Gertler and Simon Gilchrist (1999), The financial accelerator in a quantitative business cycle framework, *in* '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.

Bybee, Leland, Bryan T. Kelly, Asaf Manela and Dacheng Xiu (2024), 'Business news and business cycles', *The Journal of Finance* . Forthcoming.

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.

Chahrour, Ryan, Kristoffer Nimark and Stefan Pitschner (2021), 'Sectoral media focus and aggregate fluctuations', *American Economic Review* **111**(12).

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

Cong, Lin William, Tengyuan Liang, Baozhong Yang and Xiao Zhang (2020), *Analyzing Textual Information at Scale*, World Scientific, chapter 10, pp. 239–271.

Cong, Lin William, Tengyuan Liang, Xiao Zhang and Wu Zhu (2024), Textual factors: A scalable, interpretable, and data-driven approach to analyzing unstructured information. Working Paper.

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

Demirer, Mert, Francis X. Diebold, Laura Liu and Kamil Yılmaz (2018), 'Estimating global bank network connectedness', *Journal of Applied Econometrics* **33**(1), 1–15.

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.

Engle, Robert F, Stefano Giglio, Bryan Kelly, Heebum Lee and Johannes Stroebel (2020), 'Hedging Climate Change News', *The Review of Financial Studies* **33**(3), 1184–1216.

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.

Gentzkow, Matthew, Bryan T. Kelly and Matt Taddy (2019), 'Text as data', *Journal of Economic Literature* **57**(3), 535–74.

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.

Larsen, Vegard H., Leif Anders Thorsrud and Julia Zhulanova (2021), 'News-driven inflation expectations and information rigidities', *Journal of Monetary Economics* **117**, 507–520.

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.

Shumailov, Ilia, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson and Yarin Gal (2024), 'Ai models collapse when trained on recursively generated data', *Nature* **631**, 755–759.

	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 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	Out-of-sample Recall	Out-of-sample 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:	Assigned label: 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 the GPT 4o model. 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 total 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 total 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 total 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 total 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 total 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.

	TNIC network		Segments network		Correlation network	
	Likelihood	Strength	Likelihood	Strength	Likelihood	Strength
Credit link strength	−0.058	⋄ −0.082	−1.025	0.692	0.104	0.030
	(−0.341)	(−1.726)	(−1.595)	(1.353)	(0.745)	(0.038)
Financing link strength	** 0.495	0.071	0.196	0.203	0.067	0.016
	(3.092)	(1.453)	(0.527)	(0.645)	(0.707)	(0.026)
Peer link strength	*** 0.969	⋄ 0.088	*** −1.304	−0.105	* 0.285	⋄ 0.040
	(6.103)	(1.823)	(−3.464)	(−0.421)	(2.524)	(0.024)
Supply chain link strength	−0.354	*** −0.262	0.669	* 0.822	* −0.422	⋄ −0.073
	(−1.482)	(−4.692)	(1.488)	(2.019)	(−2.498)	(0.038)
Other link strength	*** 0.960	*** 0.331	*** 1.199	* −0.535	*** 0.641	*** 0.097
	(7.233)	(6.012)	(5.334)	(−1.998)	(5.791)	(0.019)
Total firm pairs	124,750		124,750		124,750	
Active links	5,938		201		8,984	
Start date	1988-01-01		1994-12-01		1981-01-01	
End date	2021-12-31		2023-12-31		2023-12-31	

Table 6: *Regression analysis of alternative networks.* We regress the link indicator and strength of an alternative network on the link strength of the different news-implied networks we extract. We fit a logit model for the link likelihood regressions and a Gamma model with a logarithmic link function for the link strength regressions. We include firm fixed effects and also cluster standard errors at the firm-level. We aggregate the data over sample periods that cover the largest range contained in both an alternative network and our networks. We restrict ourselves to the largest 500 firms over the corresponding sample periods. Online Appendix E summarizes how we construct the alternative networks. The values in parentheses give the z-statistic for a Wald test for the significance of the estimates. ***, **, *, and ⋄ denote significance on the 99.9%, 99%, 95%, and 90% confidence levels, respectively.

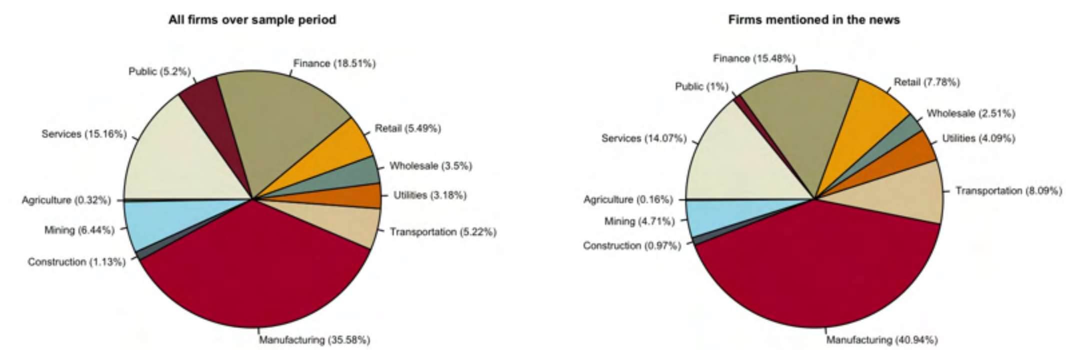

Figure 1: *Industry distribution in population and news-mentioned sample.* The population includes 26,374 distinct firms and the sample of firms mentioned in the news covers 13,526 unique firms. The data cover the time period between January 1, 1981, and December 31, 2023.

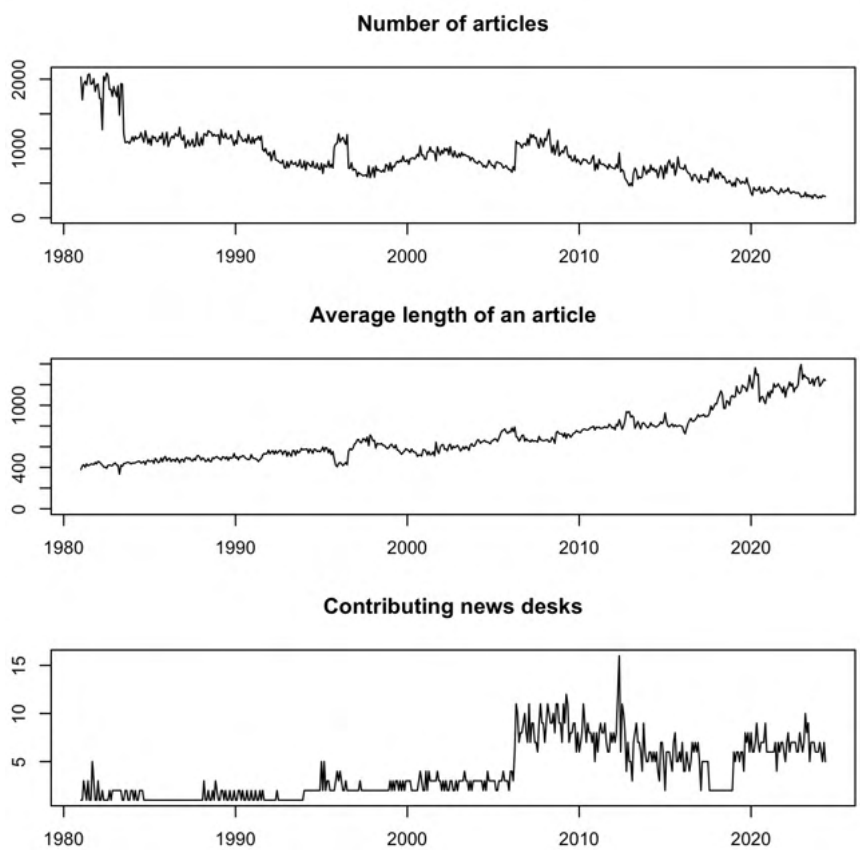

Figure 2: *New York Times sample.* The plot shows the monthly number of articles contained in our news sample, together with the average length of an article and the number of contributing news desks. We measure the length of an article through the number of words it contains.

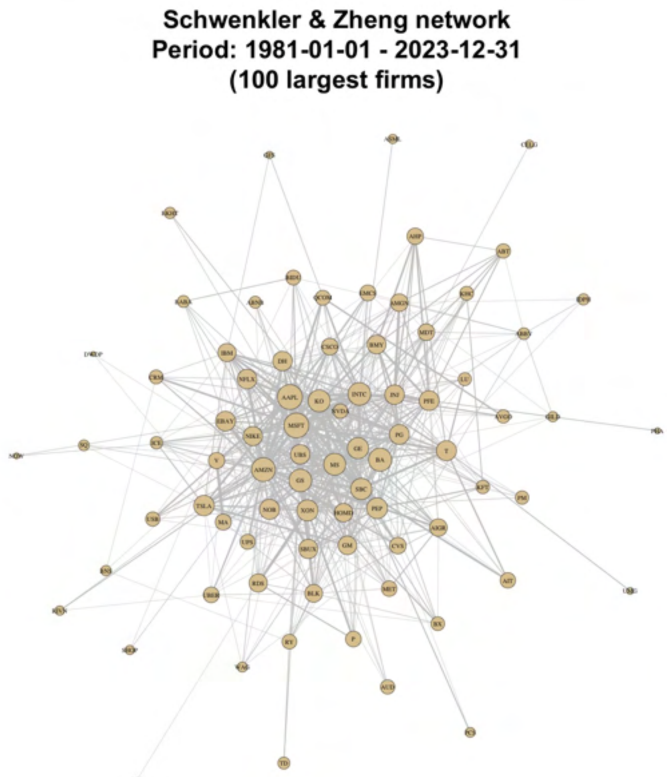

Figure 3: *News-implied network in our sample.* This graph shows the network extracted from the news using the methodology of Schwenkler and Zheng (2024). The size of a node (link) is proportional to the logarithm of one plus the number of firm (link) mentions in the news.

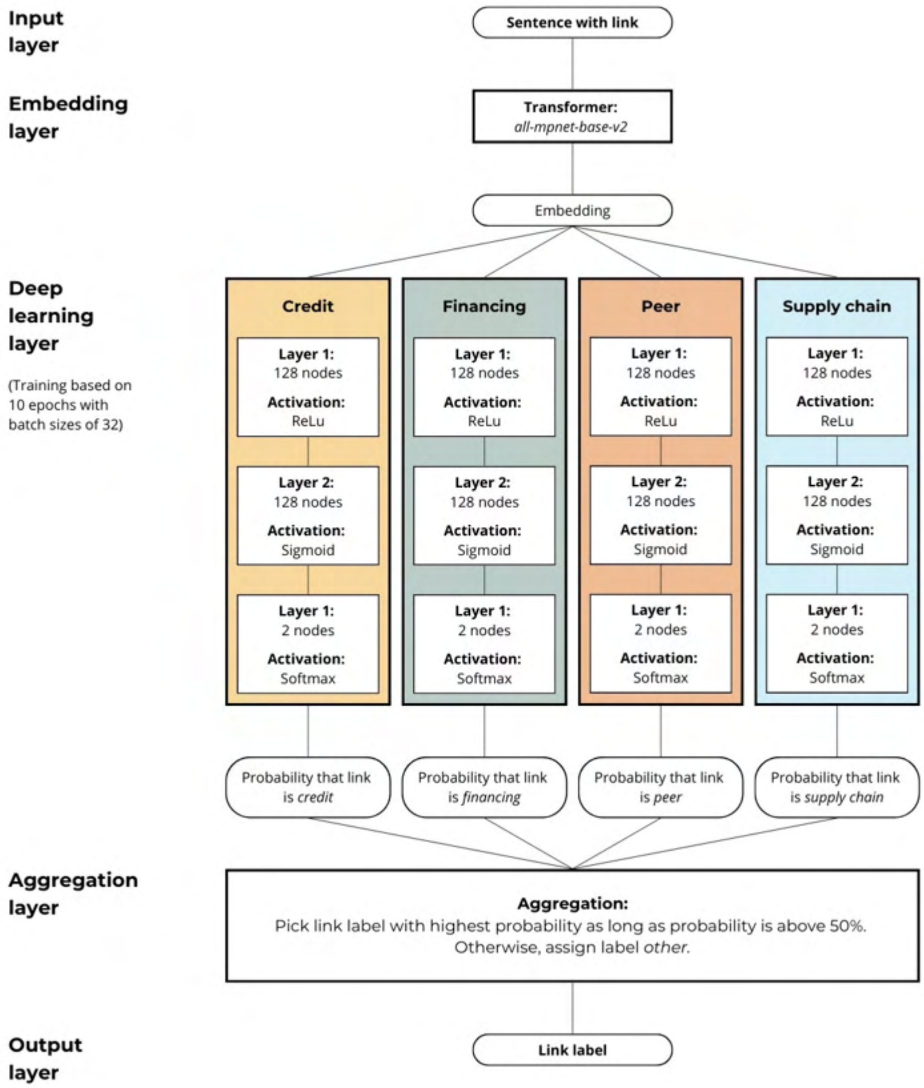

Figure 4: *Neural network architecture.* This figure shows the architectures of the model that underpins our link labelling methodology.

(a) Credit.

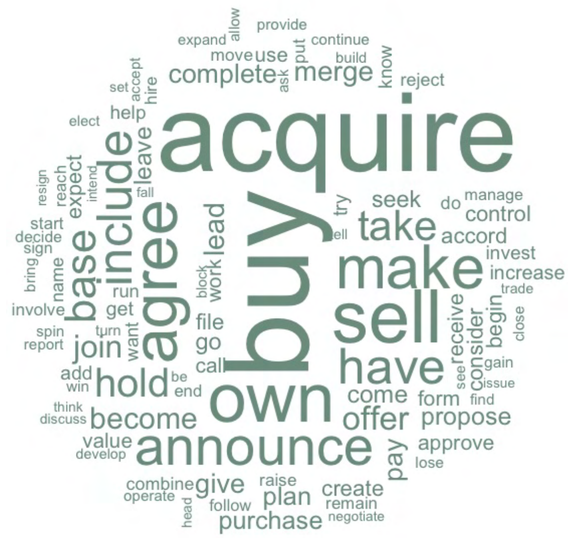

(b) Financing (ex credit).

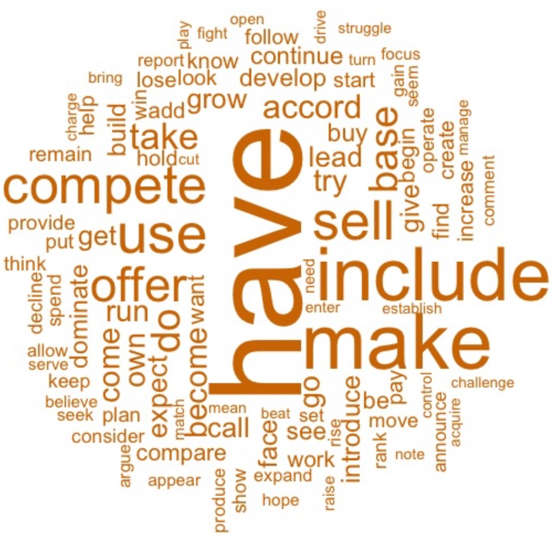

(c) Peer.

(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) Credit.

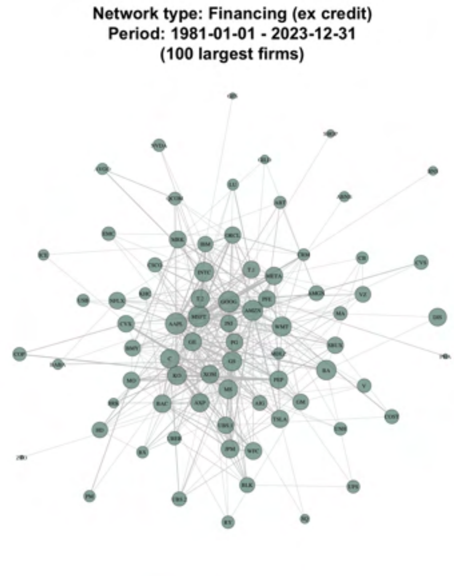

(b) Financing (ex credit).

(c) Peer.

(d) Supply chain.

Figure 7: *News-implied networks.* These graphs show the different networks extracted from the New York Times sample using our methodology, aggregated over the sample period. The size of a node (link) is proportional to the logarithm of one plus the number of firm (link) mentions in the news. We restrict ourselves to the 100 largest firms by average market cap over lifetimes that overlapped with our sample period.

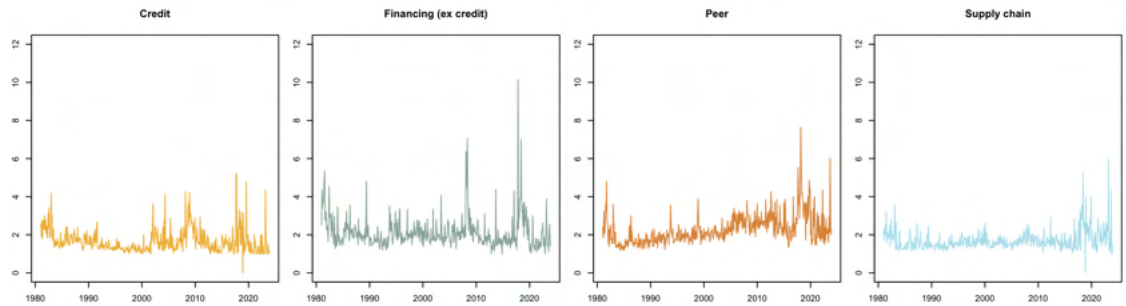

(a) Average degree.

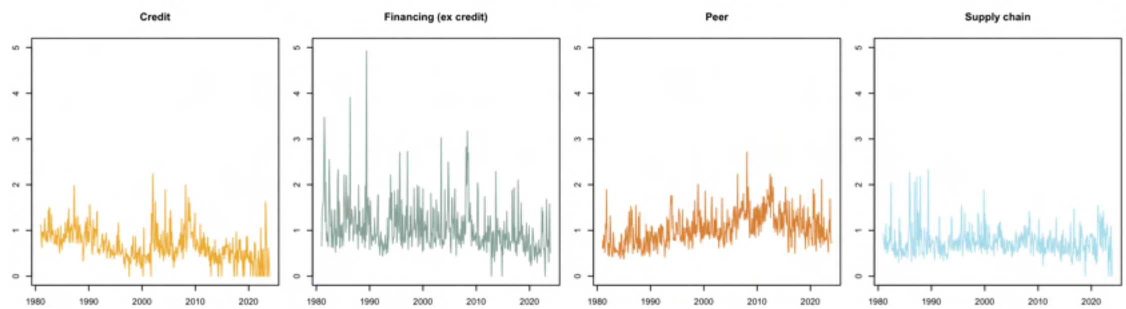

(b) First-order interconnectivity.

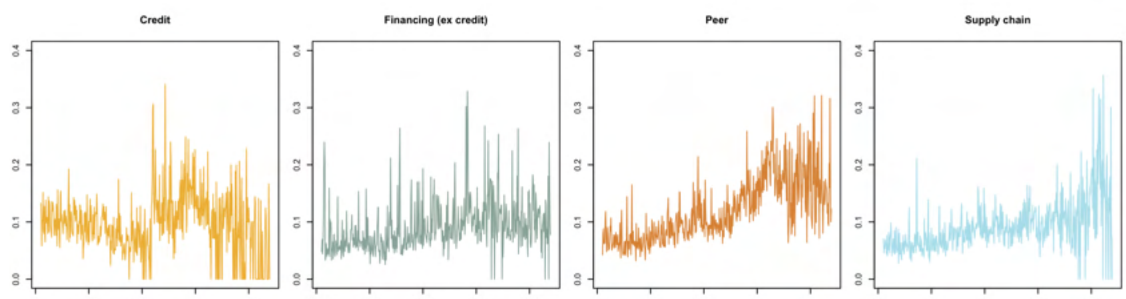

(c) Second-order interconnectivity.

Figure 8: *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 $\left(\sum_{n=1}^{N}\sum_{j\neq n}\sum_{k\neq j,n}\frac{w_t^{j,n}}{N_t}\frac{w_t^{k,n}}{N_t}\right)^{1/2}$.

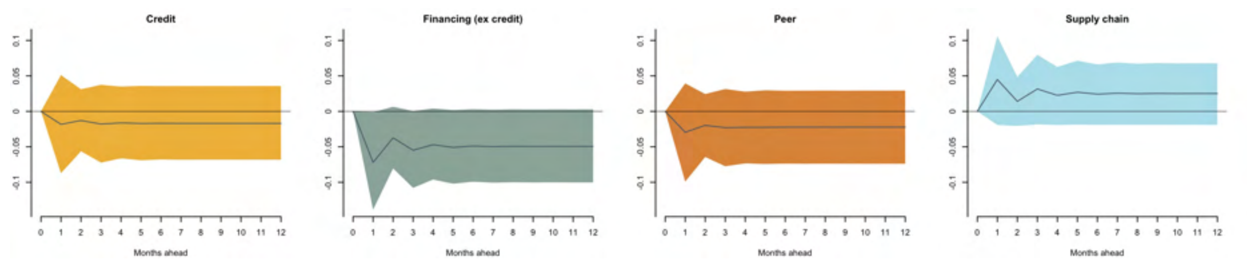

(a) S&P 500.

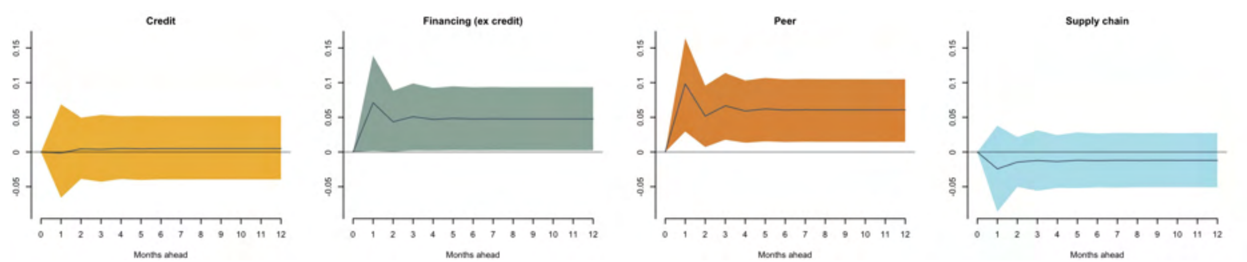

(b) VIX*.

Figure 9: *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.

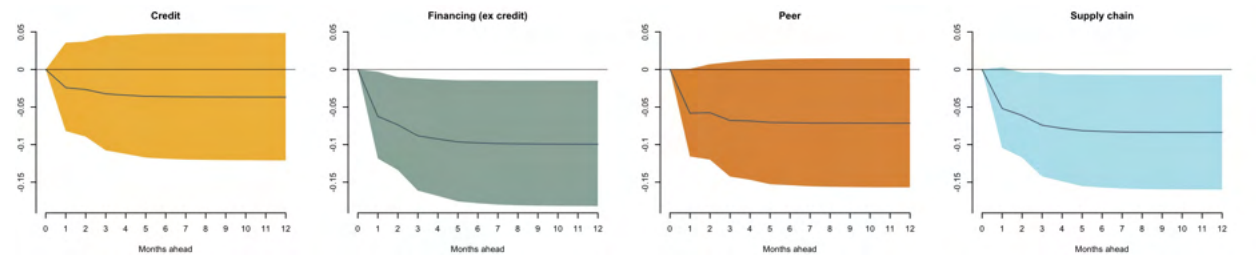

(a) Yield curve level.

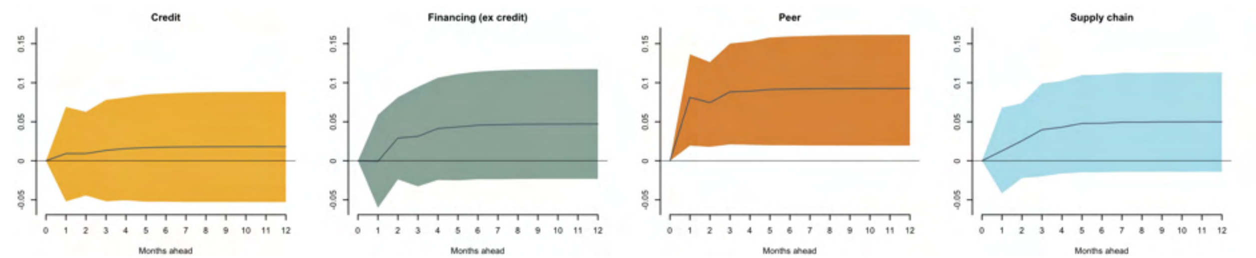

(b) Yield curve slope.

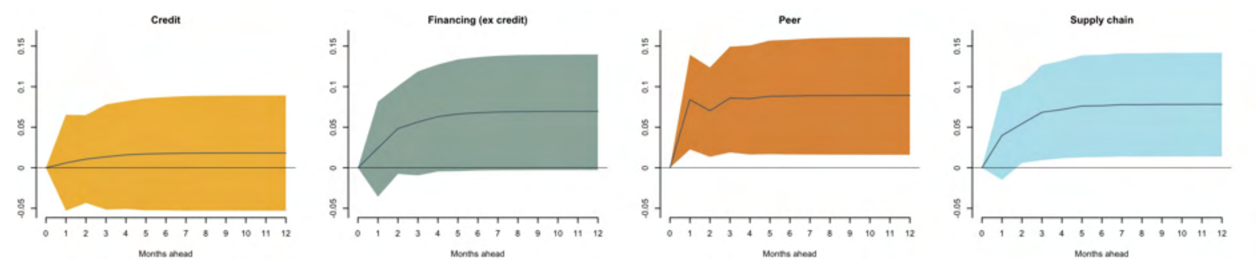

(c) AAA credit spread.

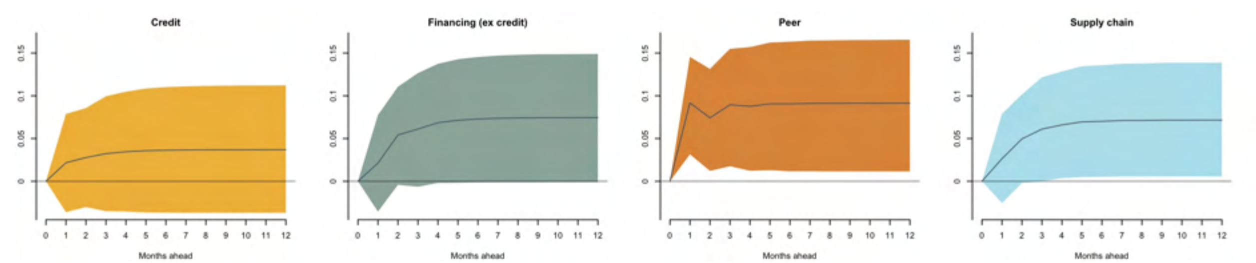

(d) BAA credit spread.

Figure 10: *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.

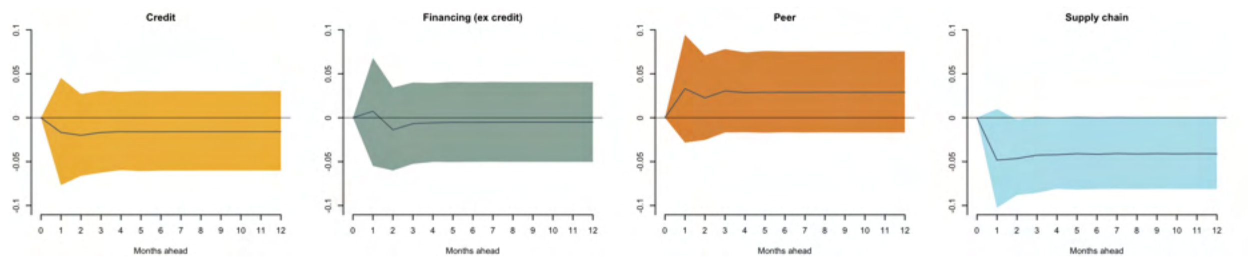

(a) Consumption growth.

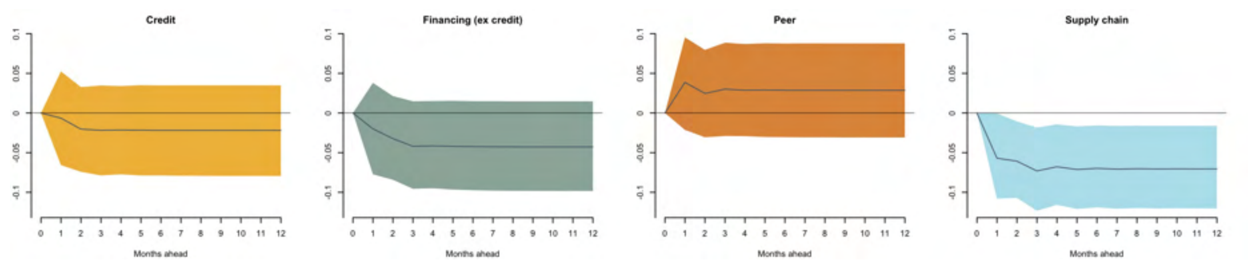

(b) Industrial production growth.

Figure 11: *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.

Online Appendix for:

The Different Networks of Firms Implied by the News

V. Hilt & G. Schwenkler*

October 15, 2024

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

A Firm data

We collect data from CRSP, Compustat, and I/B/E/S through the Wharton Research Data Service (WRDS), which we access through the R package "tidyfinance" (see Scheuch et al. (2023)). We obtain monthly pricing data from CRSP. We only download data for US corporations with equity securities. We take the total returns of these stocks (CRSP item "MTHRET"). We compute monthly market capitalizations as the product of the closing price (CRSP item "MTHPRC") and the number of shares outstanding (CRSP item "SHROUT"). We follow Bali et al. (2017) and cluster industries based on their SIC codes.

- SIC ≥ 1 & SIC ≤ 999: Agriculture
- SIC ≥ 1000 & SIC ≤ 1499: Mining
- SIC ≥ 1500 & SIC ≤ 1799: Construction
- SIC ≥ 2000 & SIC ≤ 3999: Manufacturing
- SIC ≥ 4000 & SIC ≤ 4899: Transportation
- SIC ≥ 4900 & SIC ≤ 4999: Utilities
- SIC ≥ 5000 & SIC ≤ 5199: Wholesale
- SIC ≥ 5200 & SIC ≤ 5999: Retail
- SIC ≥ 6000 & SIC ≤ 6799: Finance
- SIC ≥ 7000 & SIC ≤ 8999: Services
- SIC ≥ 9000 & SIC ≤ 9999: Public
- All other SIC: Other

We compute monthly volatilities using the daily stock file from CRSP. We collect daily shares outstanding (CRSP item "SHROUT"), prices (CRSP item "PRC"), and returns (CRSP item "RET") for all firms. For each stock series (as indicated by the PERMNO id), we compute market caps as the product of price and shares outstanding. We aggregate daily returns at the firm level (using the PERMCO id) in a value-weighted way. We then compute monthly volatilities as the standard deviation of daily returns and merge with the original monthly data from CRSP.

We obtain quarterly fundamentals data from Compustat. We measure debt as quarterly total liabilities (Compustat item "LTQ"). We define leverage as the ratio of debt over total assets. Whenever data are missing, we interpolate with the most recent annual observation. If data are still missing, we interpolate with the last available observation. We merge CRSP and Compustat using the merge file on WRDS.

We collect earnings analyst and surprise data from I/B/E/S and match firms using the I/B/E/S and CRSP merge file from WRDS. For each firm, we collect all EPS analyst forecasts for the current and next three fiscal quarters. We compute the number of analysts that track a firm in a month as the average number of earnings analysts that reported forecasts for each of the considered fiscal quarters in one month. We compute the number of earnings analysts upgrades (downgrades) as the average number of forecast upgrades (downgrades) across any of the considered quarters. We assign a value of zero when unavailable. Earnings surprises are defined as instances in which the actual quarterly EPS disclosed in the month is higher or lower than the consensus analyst estimate.

Once all data has been collected, we aggregate at the firm level as suggested by the CRSP "PERMCO" identifier. When more than one stock series is available for a given PERMCO, we aggregate market caps across series by summing them up. We aggregate returns and earning surprises data by taking value-weighted averages across stock series. We aggregate book and analyst data by summing up across stock series. We keep the firm name, ticker, exchange code, and industry of the stock series with the largest market cap.

B Firm link labelling method

We collect all sentences that contain exactly two firm mentions. We construct embeddings for these sentences using the "all-mpnet-base-v2" transformer model of Hugging Face. This transformer model is an extension of the well-known MPNet model of Song et al. (2020) that is fine-tuned to extract information and detect similarities at the sentence level. It is trained on over 1 billion paired sentences in order to detect which one out of a randomly selected subset of sentences was paired with a given sentence based on the informational content of the sentences. The transformer maps sentences of at most 384 words into a 768-dimensional vector that captures the most relevant semantic features of the sentence. We have implemented the transformer model in Python using the "sentence_transformers" package. The script is called "HS_news_sentence_embedding.ipynb" and is available in our data repository.

We take all the sentence embeddings and run them through independent deep learning classification models. We build four different models for each one of the following firm relationships: credit, financing (ex credit), supply chain, and peer. Each model consists of a three-layer neural network. The first and second layers contain 128 nodes. The third layer only has 2 nodes. The first layer uses a ReLu activation function, the second layer uses sigmoid activation function, and the third layer uses a softmax activation function. We minimize the cross entropy of each model using the adaptive learning rate algorithm "rmsprop" in Python. To prevent overiftting, we train each model using 10 epochs with a batch size of 32. We determine the maximum number of epochs by maximizing the cross-validation error of each model. To ensure that our algorithm are near-replicable, we fix all seeds and we prevent the training algorithm from shuffling the data. The models were implemented in Python. The script is called "HS_news_sentence_labelling.ipynb" and is available in our data repository.

C Synthetic link data

We use ChatGPT 4o to generate sample sentences that describe one of the five firm relationships (credit, financing, supply chain, peer, and other) between two artificial firms. We use the following prompts to generate these sentences.

Credit:

- "Can you give me some of examples of how financial news would report about a credit relationship between two firms? I want the sentences to be reflective of how the news reports about corporations and I only want to refer to credit relationships, that could be loans, convertible debt, trade credit, subscription lines, credit lines, or similar instruments. I want the firms to be fictional and not real."

- "Can you also include any other type of credit relationship you could think of that may occur between two firms?"
- "I would like the sentence structure to be a bit more complex and more realistic of, say, how the Wall Street Journal would report about credit relationships between two firms."
- "Can you create different sentences now in the tone of the New York Times?"

Financing (ex credit):

- "I would like to create a synthetic training sample of sentences that would be used in financial news articles to report about financing relationships between two firms. I would like these sentences to be similar to what would appear in the Wall Street Journal, the New York Times, or Financial Times. I would like there to be only two firms mentioned in the sentence. And I would like the financing relationship to be anything that is not debt or credit. Would you be able to generate a few examples?"
- "I would also like to consider possible reports of M&A transactions and any other type of financing activity that could occur between two firms that is not based on credit. Can you come up with some examples?"
- "Are there any other non-credit financing relationships between two firms that you can think of?"
- "These are good, but the sentence structure appears to be simple. Can you create a new sample of sentences in a tone as would appear in the Wall Street Journal?"
- "OK, now create sentences describing any possible non-credit financing relationships between two firms, in the tone as it would appear in the New York Times."
- "Can you now create new sentences describing any possible non-credit financing relationships between two firms, including M&A, in the tone as it would appear in Financial Times?"

Supply chain:

- "Create a list of different news sentences that vary in structure and describe two businesses, each sentence with different business names and industries in a supply chain relationship."
- "I would like to create a synthetic training sample of sentences that would be used in financial news articles to report about supply chain relationships between two firms. I would like these sentences to be similar to what would appear in the Wall Street Journal, the New York Times, or Financial Times. I would like there to be exactly two firms mentioned in the sentence. The firms should be fictional and not real. Would you be able to generate a few examples?"
- "Can you generate additional examples with more complex sentence structures as they would be used in the Wall Street Journal?"
- "Can you change up the industries and firms and generate examples with complex sentence structures as they would be used in the New York Times?"

Peer:

- "I would like to create a synthetic training sample of sentences that would be used in financial news articles to report about peer or competitive relationships between two firms. I would like these sentences to be similar to what would appear in the Wall Street Journal, the New York Times, or Financial Times. I would like there to be exactly two firms mentioned in the sentence. Would you be able to generate a few examples? The firms should be fictional and not real."
- "Can you add more variety for the sentences, including positive and negative relationships between peer firms?"
- "Can you add sentences that compare the stock market performance of two industry rivals as it would be described in the Wall Street Journal?"
- "Can you generate additional examples of sentences with complex sentence structures as they would be mentioned in Financial Times?"

Other:

- "In order to train an LLM, I want to create a synthetic data sample of sentences that mention two artificial firms in passing. The two mentioned firms should have no credit, financing, peer, or customer-supplier relationship – they should just be mentioned randomly. But I do want the sentences to mirror how professional news outlets like the New York Times or the Wall Street Journal would write about two firms. Can you create a sample of sentences, formatted in a table without any quotations marks, that match my requirements?"

- "I would also like more complex sentence structures and different types of firms mentioned. But only two firms at a time in a sentence."

- "Can you mix up the informational content of the sentences. I don't just want sentences that say that the firms made announcements."

- "Can you change up the sentence structure so that it's not just comparing the two companies?"

- "Can you also come up with sentences that describe how two firms performed in the stock market with the firms being unrelated to each other?"

- "Can you create additional sentences that mention two firms in one sentence for which the two firms have no relationship whatsoever with each other? The information described in the sentences can be anything you think would appear in a news article. But it is important the the firms do not share a credit, financing, supply chain, or competitive relationship."

- "Can you create sentences now that list the stock market performance of two stocks without the stocks having any relationship to each other as it would be reported in the Wall Street Journal?"

D Output of the methodology

Figure A.1 showcases the network of other links, aggregated over the whole sample period. We restrict ourselves to the 75 largest firms rather than the 100 largest firms as the graph otherwise becomes illegible.

E Alternative networks

We collect Text-based Network Industry Classifications (TNIC) data from the Hoberg-Phillips Data Library. We use the baseline database that is calibrated to match three-digit SIC codes. These data cover the years 1988 through 2021. We assume that coverage begins in January 1988 and ends in December 2021. We match the GVKEY identifiers in these data to CRSP PERMCO identifiers using our own matching methodology. We then aggregate links across years and construct a network in which the strength of a link is given by the average TNIC score over the sample. We consider the 500 largest firms by average market cap over the sample period.

We follow an analogous approach for the Compustat Segments network, in which we measure the strength of a link as the logarithm of one plus the sum of all sales among a firm pair. We restrict ourselves to business and operating segments only. We obtain these data through WRDS. The sample period covers December 1994 through December 2023.

For the variance decomposition network, which we label equity correlations, we follow the approach proposed by Demirer et al. (2018). This approach requires that we compute daily volatility measures using intraday data, which we collect from CRSP through WRDS for the 500 largest firms by average market cap throughout our sample period (January 1981 through December 2023). We measure the strength of a link in this network as the average of the two directional connectedness measures of the firms in a link. This network is usually very dense. So we censor links for which the absolute value of the strength is below 0.5%.

When comparing one of the alternative networks to our networks, we restrict ourselves to the same sample period and subset of firms to ensure consistency.

F Aggregate data

This section provides a summary of our aggregate economic variables. Table A.1 provides summary statistics as well as coverage dates for all series. Below, we provide details on how we collect these data.

The BAA and AAA credit spreads correspond to the Moody's Seasoned corporate bond yields minus the Federal Funds Rate. The level of the yield curve is measured via the 3-month Treasury Bill constant-maturity yield, while the slope is constructed as the spread between the constant-maturity yields of the 10-year and the 1-year Treasury Bills. All rates are evaluated as averages within month. Industrial production growth is given by the month-to-month growth rate in the industrial production index of the Board of Governors of the Federal Reserve System. We measure consumption growth as the percent monthly change in the Real Personal Consumption Expenditures Index of the U.S. Bureau of Economic Analysis. We obtain all of the above data from the St. Louis FRED system.

We collect monthly S&P 500 returns from CRSP. We collect daily VIX and VXO data from CRSP and aggregate these to monthly time series by taking averages in a month.

G VAR analyses

We estimate one-lag vector autoregressive (VAR) models with intercepts 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 except the "other" network. We base our identification on a Cholesky decomposition of the residual variance-covariance matrix, where the variables are ordered as follows: (1) consumption growth, (2) industrial production growth, (3) yield curve level, (4) yield curve slope, (5) AAA credit spread, (6) BAA credit spread, (7) S&P 500 returns, (8) VIX, (9) average degree, (10) first-order interconnectivity, and (11) second-order interconnectivity. Because the VIX level only became available in January 1990, we complement the VIX series with VXO data, and use this VIX-like series in our analyses. However, in unreported experiments, we obtain similar results if we restrict ourselves to the original, unadjusted VIX series. To ensure that all variables are stationary and comparable, we take first-differences of all time series except consumption growth, industrial production growth, and the S&P 500 return, and then standardize all variables prior to running the VAR. The sample period over which we fit the VARs is February 1986 through December 2023.

Figure 8 of the main body of the paper shows the time series of the connectivity metrics for each link type. The captions report details of how we compute them. Figures A.2–A.16 show the cumulative impulse response functions for one-standard-deviation orthogonal shocks to the average degree, the first-order interconnectivity, and the second-order interconnectivity of each of the networks.

References

Bali, Turan, Robert Engle and Scott Murray (2017), *Empirical Asset Pricing: The Cross Section of Stock Returns: An Overview*, John Wiley & Sons, Ltd, pp. 1–8.

Demirer, Mert, Francis X. Diebold, Laura Liu and Kamil Yılmaz (2018), 'Estimating global bank network connectedness', *Journal of Applied Econometrics* **33**(1), 1–15.

Scheuch, Christoph, Stefan Voigt and Patrick Weiss (2023), *Tidy Finance with R*, 1st edn, Chapman and Hall/CRC.
URL: *https://tidy-finance.org*

Song, Kaitao, Xu Tan, Tao Qin, Jianfeng Lu and Tie-Yan Liu (2020), Mpnet: masked and permuted pre-training for language understanding, NIPS '20.

Variable	Start	End	Mean	Median	Std dev.	Min.	Max.
S&P 500 return	1981-01	2023-12	0.44%	1.26%	4.86%	-16.94%	10.77%
VIX	1990-01	2023-12	22.74	20.37	10.32	10.83	62.67
VXO	1986-01	2021-01	22.55	19.53	11.00	10.49	65.45
AAA credit spread	1983-01	2023-12	3.67%	3.91%	1.62%	0.04%	5.73%
BAA credit spread	1986-01	2023-12	4.95%	5.19%	1.91%	0.95%	8.82%
Level of yield curve	1981-09	2023-12	1.01%	0.13%	1.69%	0.01%	5.16%
Slope of yield curve	1981-01	2023-12	1.95%	1.97%	1.04%	-0.42%	3.40%
Industrial production	1981-01	2023-12	0.01%	0.20%	0.86%	-4.30%	1.40%
Consumption growth	1981-01	2023-12	0.08%	0.10%	0.28%	-0.90%	0.90%

Table A.1: *Summary statistics of our aggregate variables.* All factors are sampled at the monthly frequency.

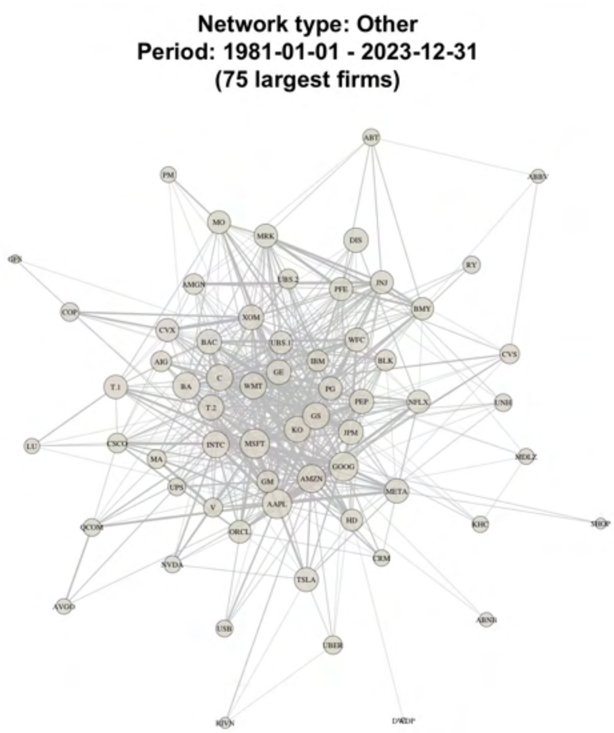

Figure A.1: *News-implied other network.* These graph shows the network of other links extracted from the New York Times sample using our methodology, aggregated over the sample period. The size of a node (link) is proportional to the logarithm of one plus the number of firm (link) mentions in the news. We restrict ourselves to the 75 largest firms by average market cap over lifetimes that overlapped with our sample period.

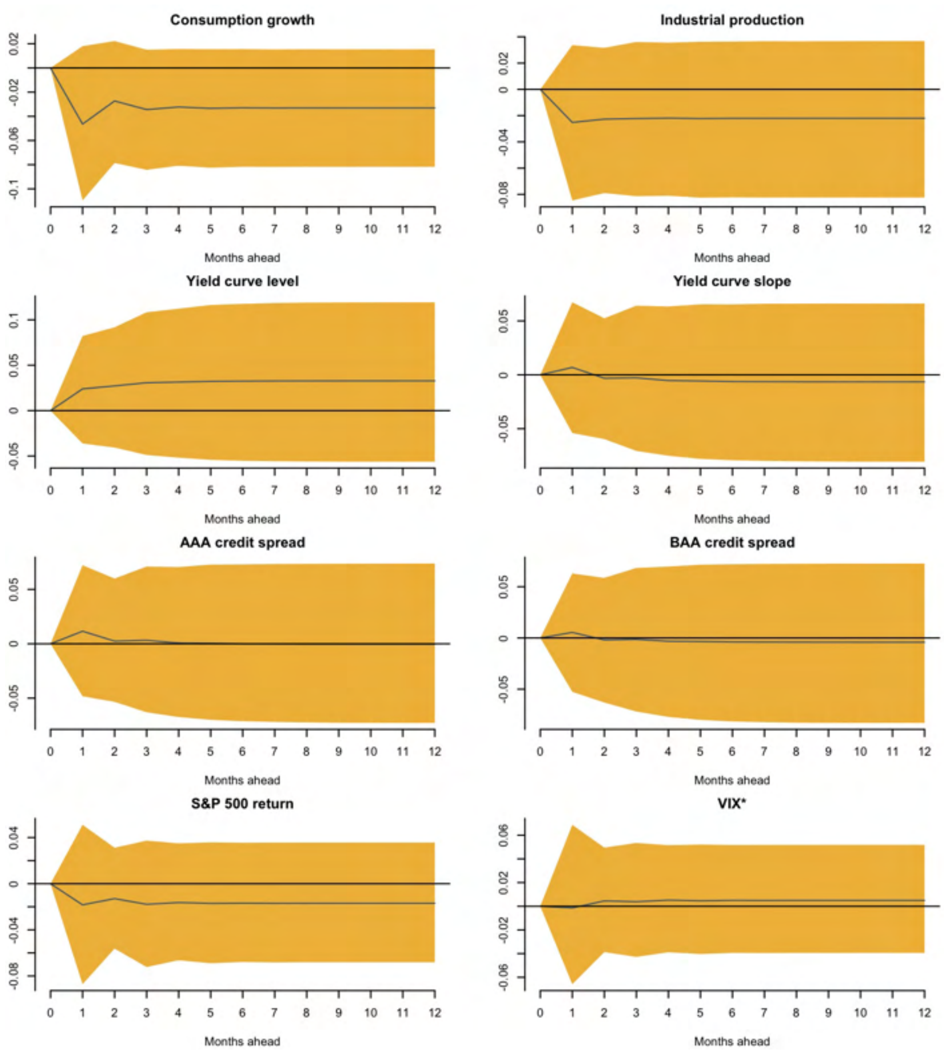

Figure A.2: *Cumulative impulse response functions for orthogonal shocks to the average degree of the credit network.* We consider a one-lag vector autoregressive model that includes a constant intercept. The grey lines give the impulse responses and the colored areas give 90% confidence bands computed via bootstrap with 10,000 bootstrap samples. We consider one-standard-deviation shocks to the orthogonal component of the impulse variable. Online Appendix F provides summary statistics of these macroeconomic variables. The * indicates that this applies for an adjusted VIX time series that includes VXO observations between 1986 and 1990.

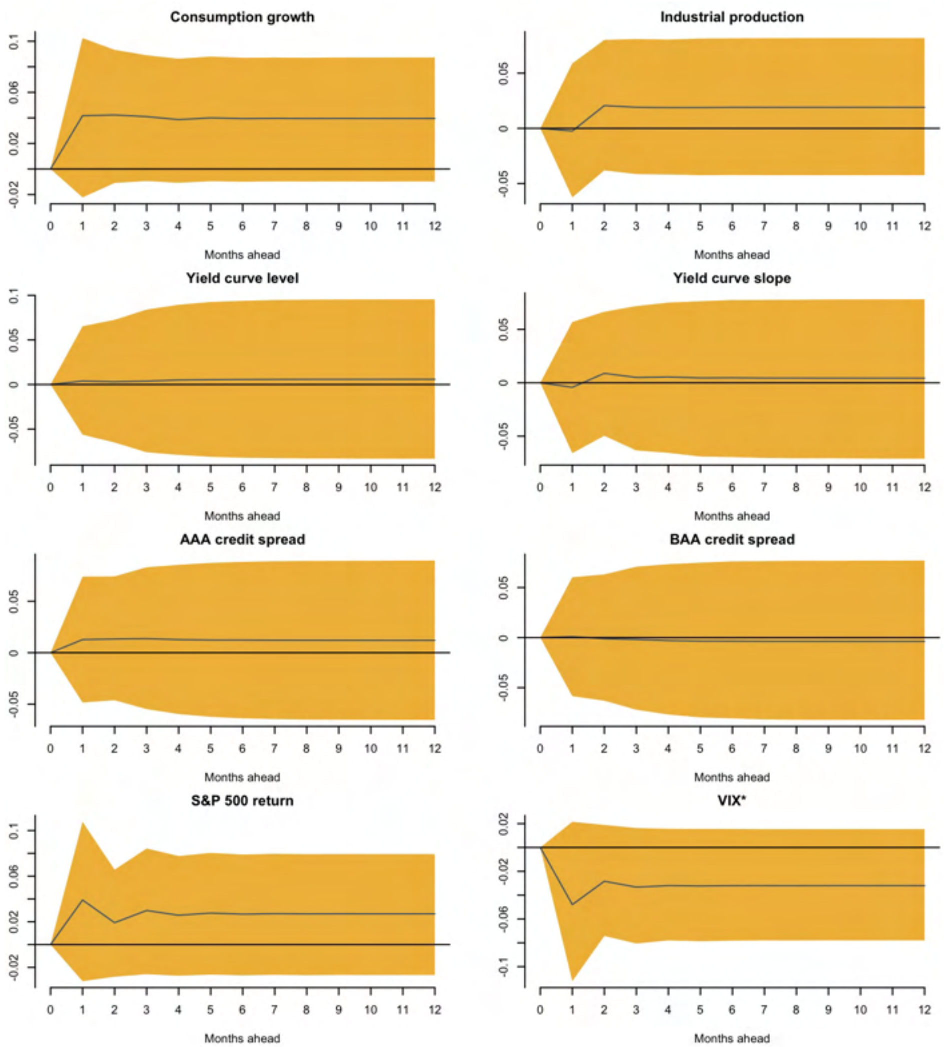

Figure A.3: *Cumulative impulse response functions for orthogonal shocks to the first-order interconnectivity measure of the credit network*. We consider a one-lag vector autoregressive model that includes a constant intercept. The grey lines give the impulse responses and the colored areas give 90% 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.

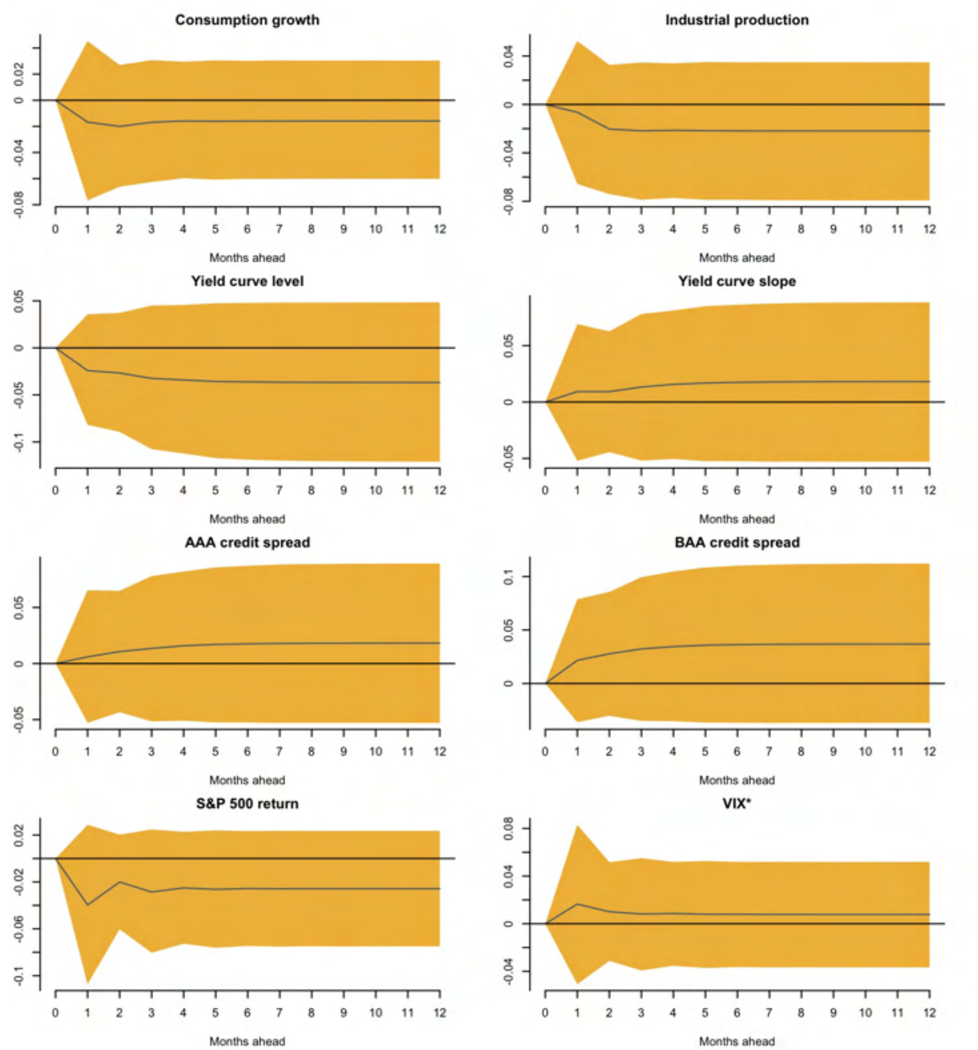

Figure A.4: *Cumulative impulse response functions for orthogonal shocks to the second-order interconnectivity measure of the credit network.* We consider a one-lag vector autoregressive model that includes a constant intercept. The grey lines give the impulse responses and the colored areas give 90% confidence bands computed via bootstrap with 10,000 bootstrap samples. We consider one-standard-deviation shocks to the orthogonal component of the impulse variable. Online Appendix F provides summary statistics of these macroeconomic variables. The * indicates that this applies for an adjusted VIX time series that includes VXO observations between 1986 and 1990.

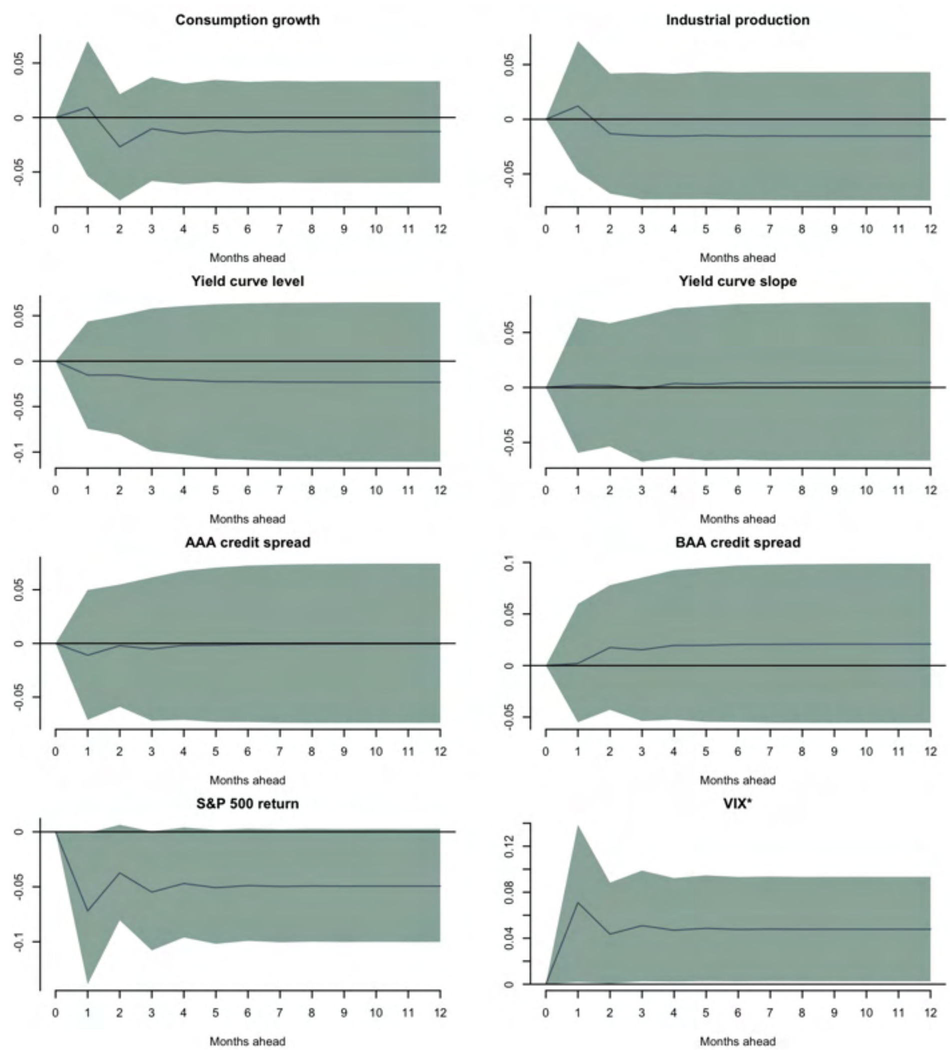

Figure A.5: *Cumulative impulse response functions for orthogonal shocks to the average degree of the financing (ex credit) network*. We consider a one-lag vector autoregressive model that includes a constant intercept. The grey lines give the impulse responses and the colored areas give 90% confidence bands computed via bootstrap with 10,000 bootstrap samples. We consider one-standard-deviation shocks to the orthogonal component of the impulse variable. Online Appendix F provides summary statistics of these macroeconomic variables. The * indicates that this applies for an adjusted VIX time series that includes VXO observations between 1986 and 1990.

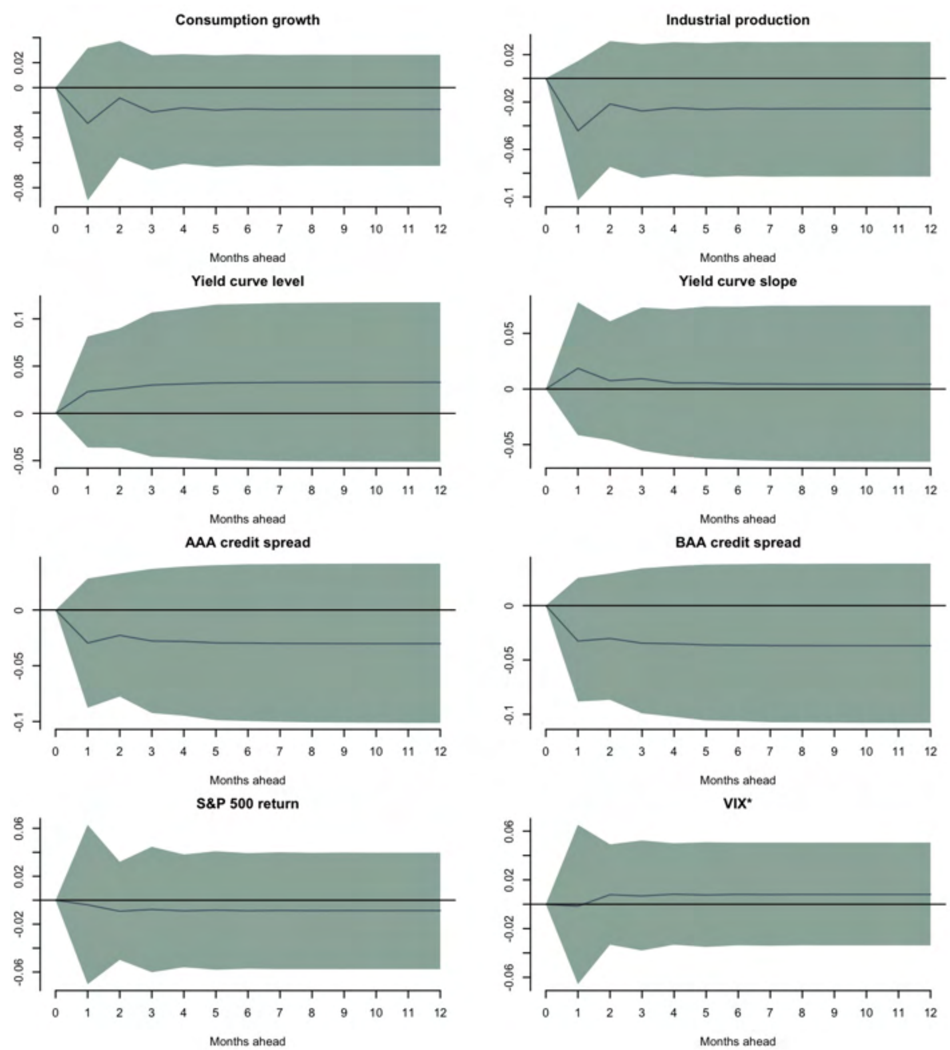

Figure A.6: *Cumulative impulse response functions for orthogonal shocks to the first-order interconnectivity measure of the financing (ex credit) network.* We consider a one-lag vector autoregressive model that includes a constant intercept. The grey lines give the impulse responses and the colored 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.

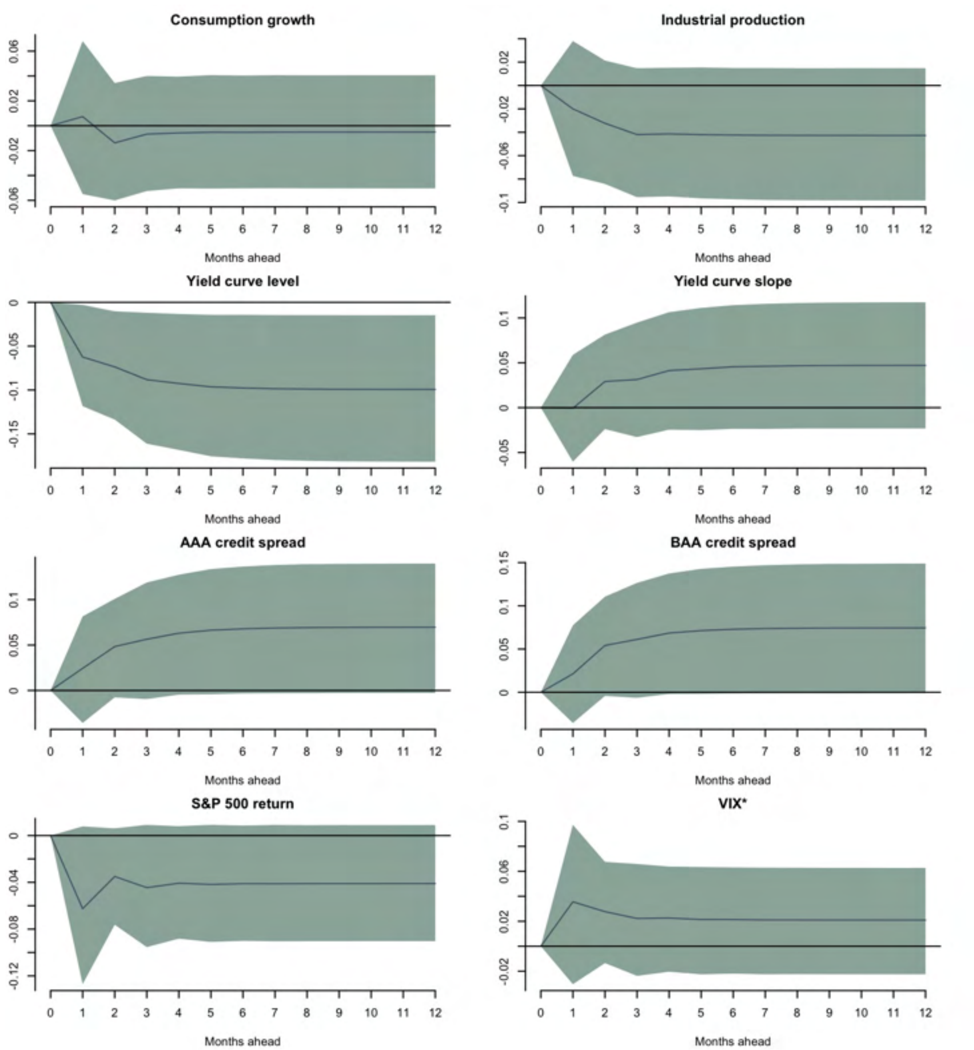

Figure A.7: *Cumulative impulse response functions for orthogonal shocks to the second-order interconnectivity measure of the financing (ex credit) network.* We consider a one-lag vector autoregressive model that includes a constant intercept. The grey lines give the impulse responses and the colored 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. Online Appendix F provides summary statistics of these macroeconomic variables. The * indicates that this applies for an adjusted VIX time series that includes VXO observations between 1986 and 1990.

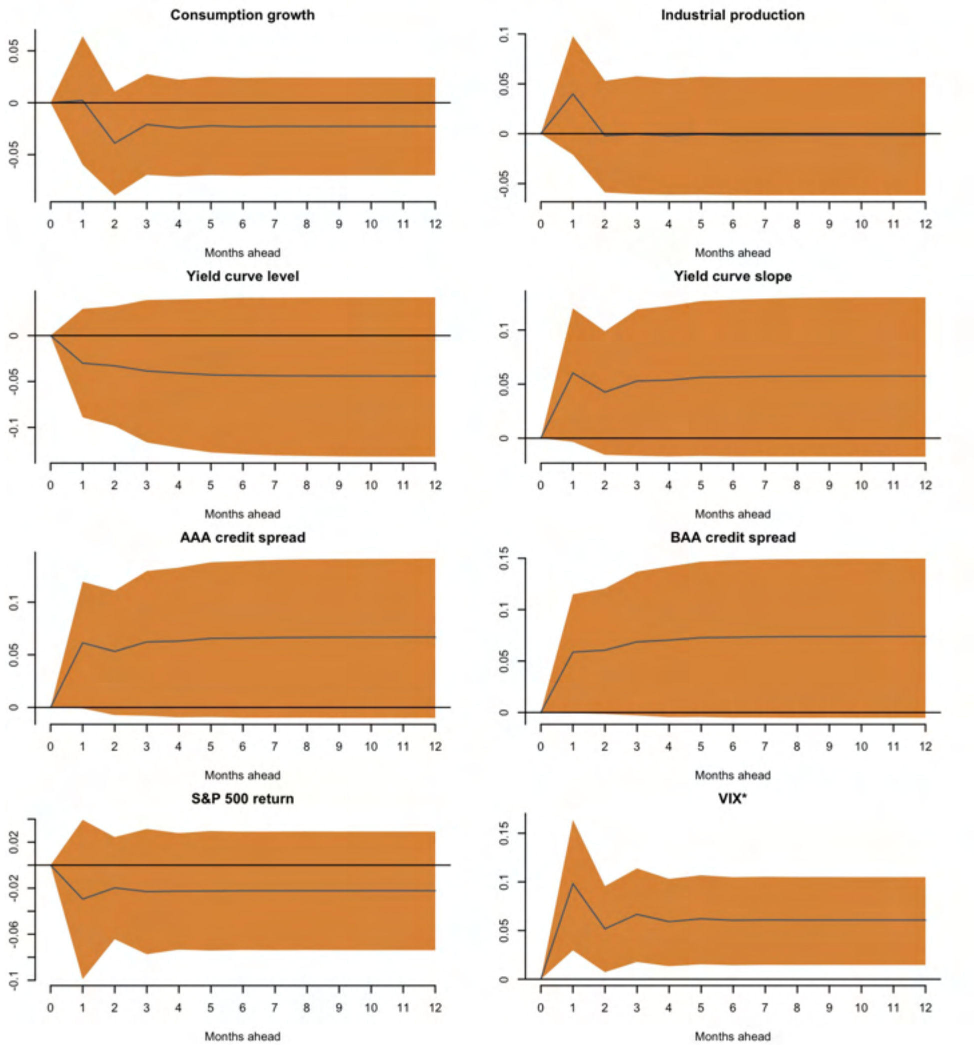

Figure A.8: *Cumulative impulse response functions for orthogonal shocks to the average degree of the peer network.* We consider a one-lag vector autoregressive model that includes a constant intercept. The grey lines give the impulse responses and the colored areas give 90% confidence bands computed via bootstrap with 10,000 bootstrap samples. We consider one-standard-deviation shocks to the orthogonal component of the impulse variable. Online Appendix F provides summary statistics of these macroeconomic variables. The * indicates that this applies for an adjusted VIX time series that includes VXO observations between 1986 and 1990.

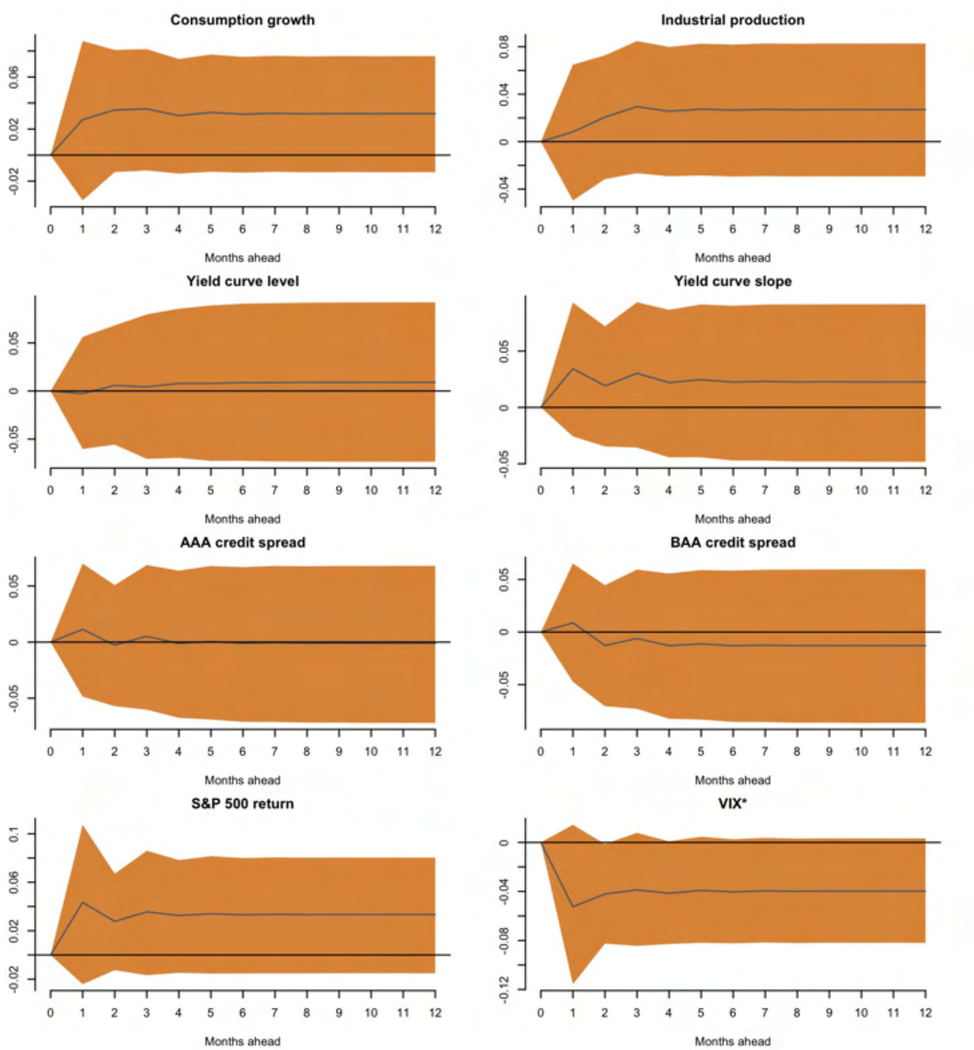

Figure A.9: *Cumulative impulse response functions for orthogonal shocks to the first-order interconnectivity measure of the peer network*. We consider a one-lag vector autoregressive model that includes a constant intercept. The grey lines give the impulse responses and the colored areas give 90% confidence 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.

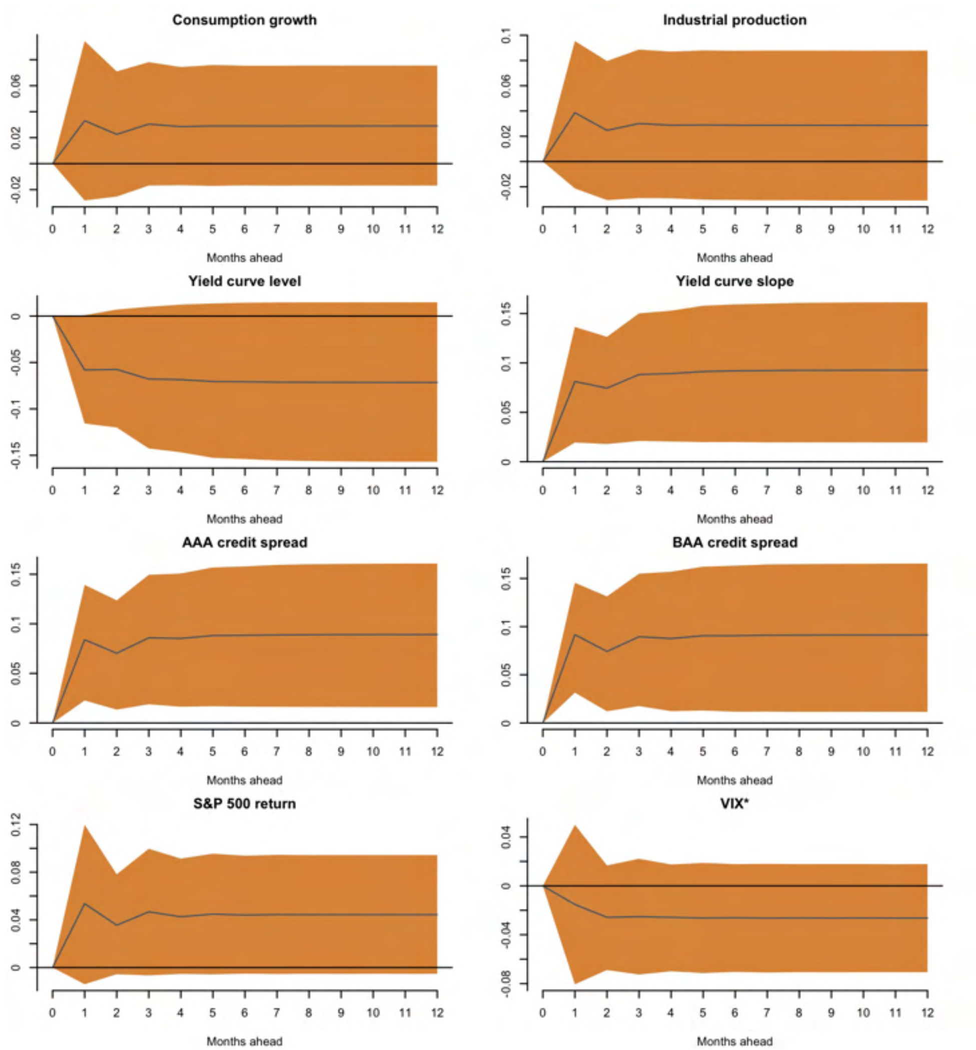

Figure A.10: *Cumulative impulse response functions for orthogonal shocks to the second-order interconnectivity measure of the peer network.* We consider a one-lag vector autoregressive model that includes a constant intercept. The grey lines give the impulse responses and the colored areas give 90% confidence bands computed via bootstrap with 10,000 bootstrap samples. We consider one-standard-deviation shocks to the orthogonal component of the impulse variable. Online Appendix F provides summary statistics of these macroeconomic variables. The * indicates that this applies for an adjusted VIX time series that includes VXO observations between 1986 and 1990.

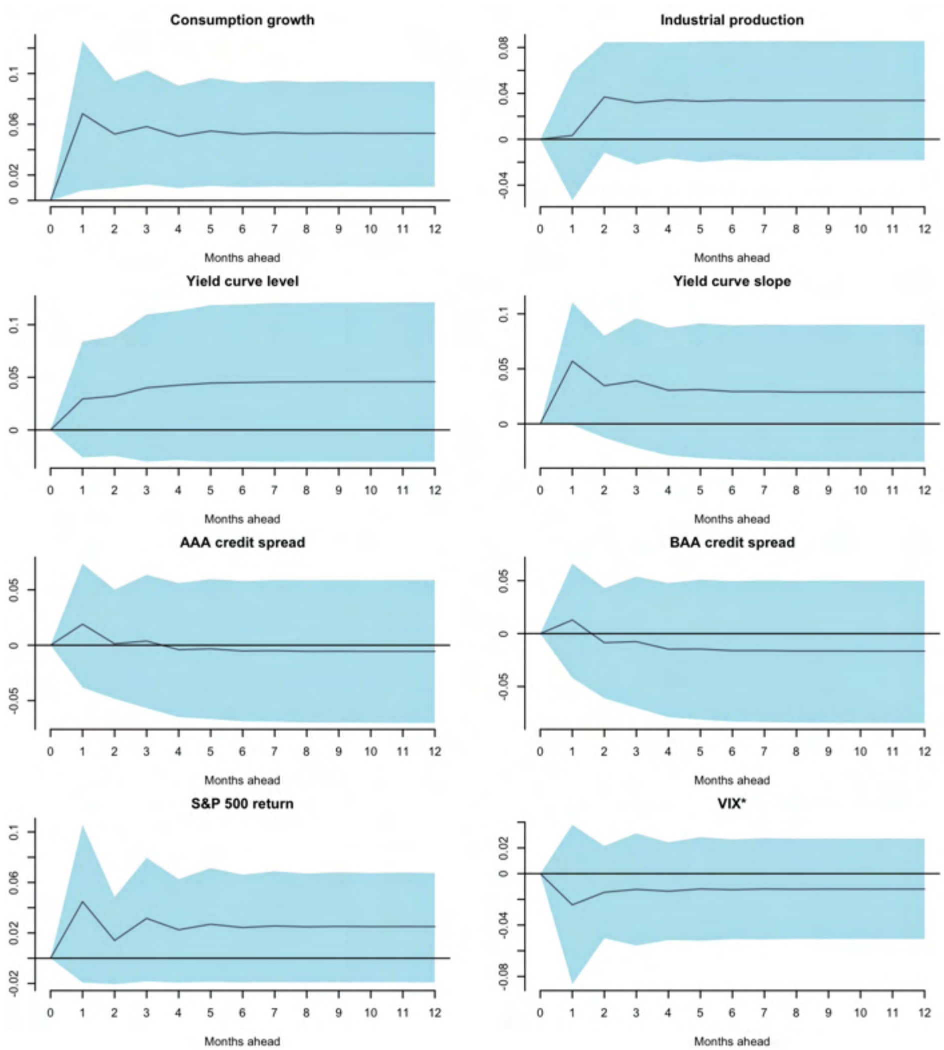

Figure A.11: *Cumulative impulse response functions for orthogonal shocks to the average degree of the supply chain network.* We consider a one-lag vector autoregressive model that includes a constant intercept. The grey lines give the impulse responses and the colored areas give 90% confidence bands computed via bootstrap with 10,000 bootstrap samples. We consider one-standard-deviation shocks to the orthogonal component of the impulse variable. Online Appendix F provides summary statistics of these macroeconomic variables. The * indicates that this applies for an adjusted VIX time series that includes VXO observations between 1986 and 1990.

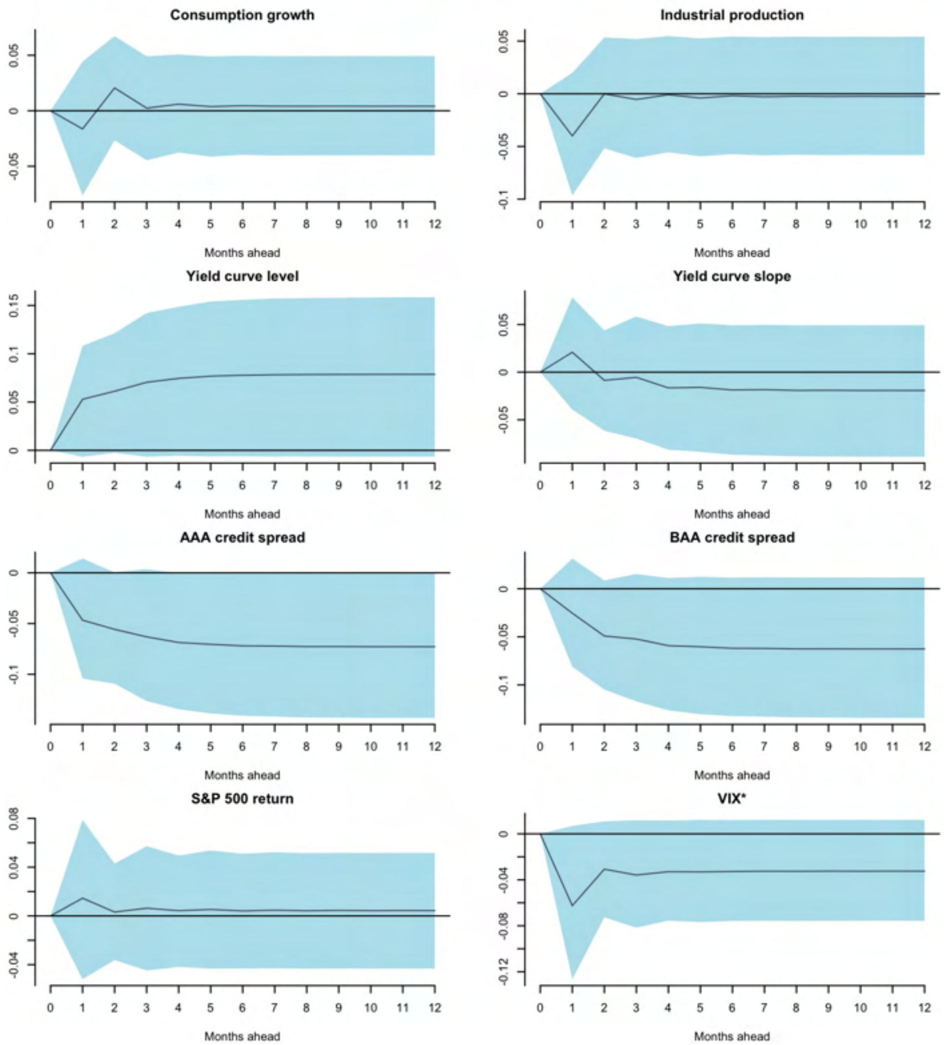

Figure A.12: *Cumulative impulse response functions for orthogonal shocks to the first-order interconnectivity measure of the supply chain network.* We consider a one-lag vector autoregressive model that includes a constant intercept. The grey lines give the impulse responses and the colored areas give 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.

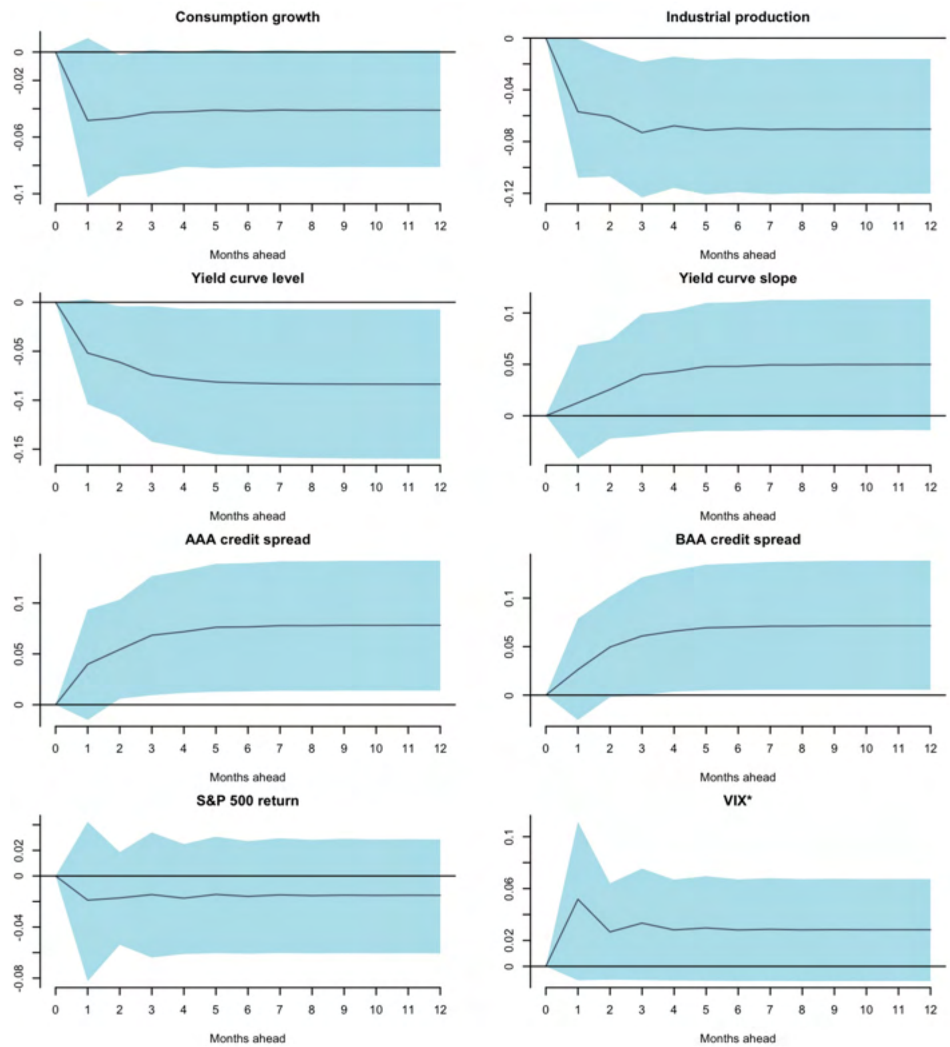

Figure A.13: *Cumulative impulse response functions for orthogonal shocks to the second-order interconnectivity measure of the supply chain network.* We consider a one-lag vector autoregressive model that includes a constant intercept. The grey lines give the impulse responses and the colored areas give 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. Online Appendix F provides summary statistics of these macroeconomic variables. All variables except consumption growth, industrial production growth, and the S&P 500 return are differentiated before running the VAR to ensure that all variables are stationary. All variables are standardized with their time-series mean and standard deviation prior to running the VAR.

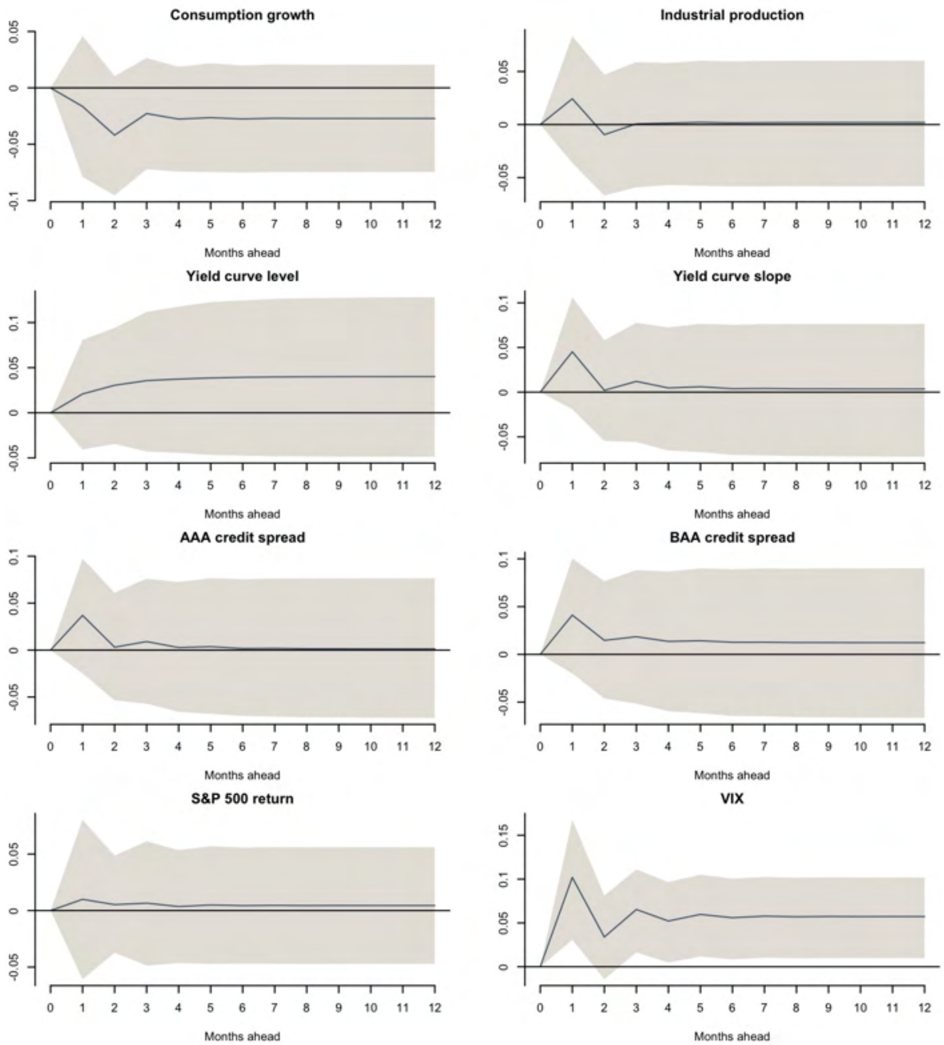

Figure A.14: *Cumulative impulse response functions for orthogonal shocks to the average degree of the other network*. We consider a one-lag vector autoregressive model that includes a constant intercept. The grey lines give the impulse responses and the colored areas give 90% confidence bands computed via bootstrap with 10,000 bootstrap samples. We consider one-standard-deviation shocks to the orthogonal component of the impulse variable. Online Appendix F provides summary statistics of these macroeconomic variables. The * indicates that this applies for an adjusted VIX time series that includes VXO observations between 1986 and 1990.

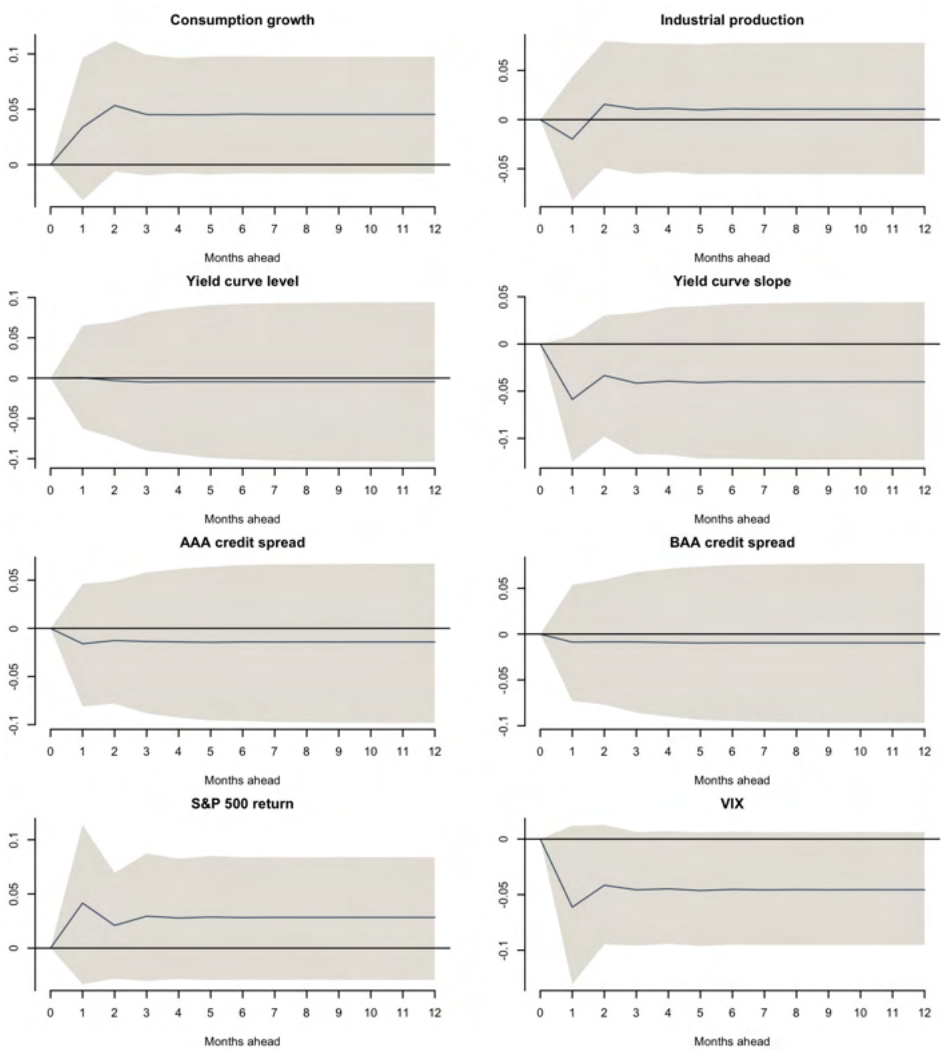

Figure A.15: *Cumulative impulse response functions for orthogonal shocks to the first-order interconnectivity measure of the other network.* We consider a one-lag vector autoregressive model that includes a constant intercept. The grey lines give the impulse responses and the colored areas give 90% 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.

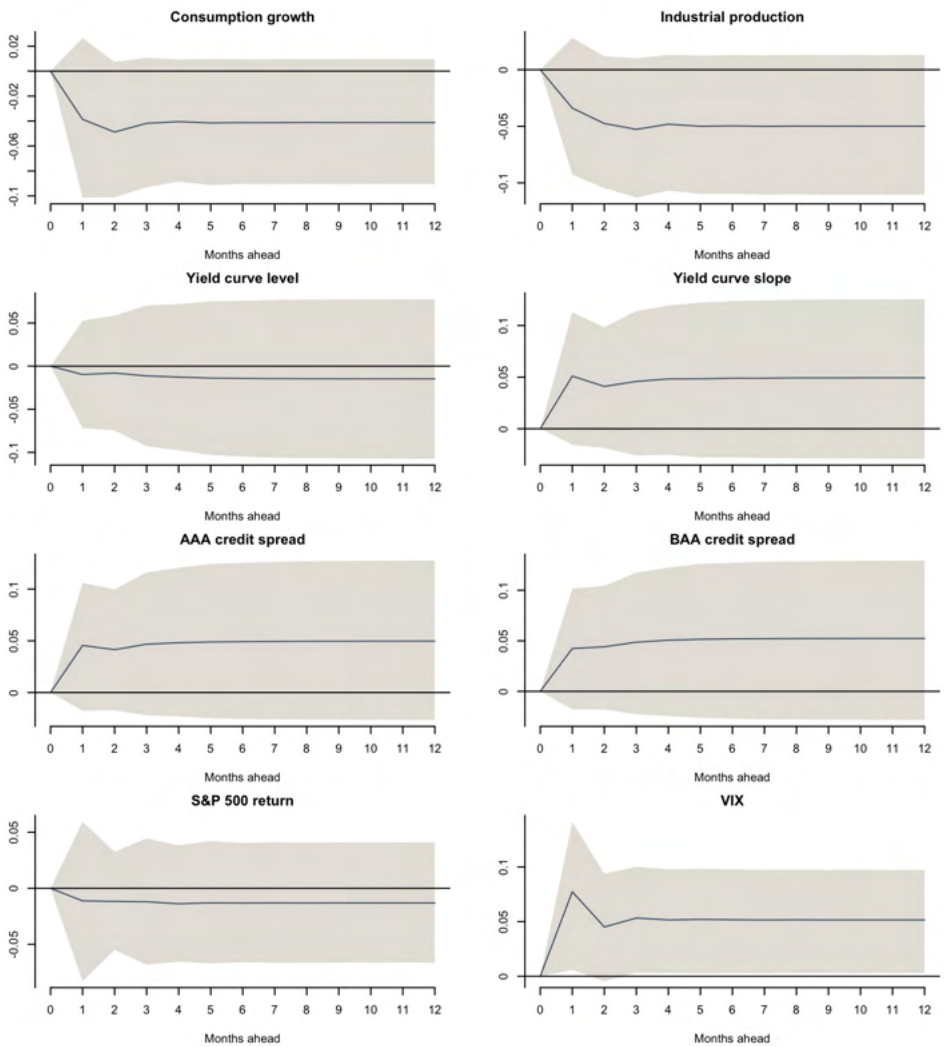

Figure A.16: *Cumulative impulse response functions for orthogonal shocks to the second-order interconnectivity measure of the other network.* We consider a one-lag vector autoregressive model that includes a constant intercept. The grey lines give the impulse responses and the colored areas give 90% confidence bands computed via bootstrap with 10,000 bootstrap samples. We consider one-standard-deviation shocks to the orthogonal component of the impulse variable. Online Appendix F provides summary statistics of these macroeconomic variables. The * indicates that this applies for an adjusted VIX time series that includes VXO observations between 1986 and 1990.