Tell me a story: Quantifying economic narratives and their role during COVID-19*

Daniel Borup[†] Jorge W. Hansen[††]

Benjamin Liengaard[‡] Erik C. M. Schütte[**]

Abstract

We retrieve the most salient COVID-19 narratives from daily and real-time open-ended questionnaires presented to a large number of US stockholders prior to, during, and after the first wave of the pandemic (February through June, 2020). We elicit thirteen narratives (e.g. *supply disruption*, *investor fear*, *stay at home*, *infection worry*, and *fiscal and monetary policy intervention*) using textual analysis from the survey responses and quantify their propagation over time with individual time series. We document that narratives have strong behavioral implications for human mobility during the first wave of the pandemic. We assess the economic implications via network analysis, which allows for bi-directional effects, showing that (i) narratives significantly drive unexpected fluctuations in the real economy and financial markets and that (ii) narratives are themselves shaped by the real economy and financial markets. The two directions are generally equally strong. These effects exist at a daily horizon, and cumulate over a weekly and monthly horizon. Finally, we show that narratives carry significant risk premia in the financial markets.

Keywords: Narrative economics, COVID-19, textual analysis, network analysis, Google Mobility, latent Dirichlet allocation, risk premia.

JEL Classification: C55, D91, E44, E71.

This version: February 1, 2021.

---

*We are grateful to Robert Shiller, Tony Cookson, Leif Anders Thorsrud (discussant), Christoph Merkle, Olga Balakina, Jonas Nygaard Eriksen, Umut Akovalı, and participants at the Nordic Finance Network Finance Webinar Series hosted by the Norwegian School of Economics, and the joint Econometrics-Finance seminar series at Aarhus University. Support from CREATES and Danish Finance Institute (DFI) is gratefully acknowledged.

[†]Corresponding author. CREATES, Department of Economics and Business Economics, Aarhus University, Fuglesangs Allé 4, DK-8210 Aarhus V, Denmark, and the Danish Finance Institute (DFI). Email: dborup@econ.au.dk.

[††]CREATES, Department of Economics and Business Economics, Aarhus University, Fuglesangs Allé 4, DK-8210 Aarhus V, Denmark, and the Danish Finance Institute (DFI). Email: jh@econ.au.dk.

[‡]CREATES, Department of Economics and Business Economics, Aarhus University, Fuglesangs Allé 4, DK-8210 Aarhus V, Denmark. Email: benlien@econ.au.dk.

[**]CREATES, Department of Economics and Business Economics, Aarhus University, Fuglesangs Allé 4, DK-8210 Aarhus V, Denmark, and the Danish Finance Institute (DFI). Email: christianms@econ.au.dk.

I. Introduction

Storytelling is a human universal. Evidence for this dates back to 1800 BC when the oldest known story in the world "The Epic of Gilgamesh" was written on clay tablets in ancient Mesopotamia. Recent evidence from evolutionary biology and neuroscience links our cooperative problem-solving ability to deep structures in our brains that hardwire us to be naturally predisposed to think and communicate in narratives (Zak, 2015; Smith et al., 2017; Bietti et al., 2019). Despite the ubiquity of narratives and stories in human life, economics has given little attention to their impact on economic phenomena. This has recently changed as a result of the novel work by Shiller (2017, 2019, 2020) who coined the term "narrative economics" to describe the study of how the stories, explanations, and justifications of events that we tell ourselves and others shape individual's decision-making and drive economic fluctuations. In other words, the prevailing narratives, which are those that go viral and spread throughout society, have the ability to determine our economic actions since they function as scripts to guide behavior in times of uncertainty (Schank and Abelson, 1977). This includes the propensity to spend or invest, the likelihood of starting new, possibly risky business, or hiring new employees, for instance. It is thus important to understand and characterize narratives since their propagation may constitute important means to forming policy decisions or anticipating economic activity.

In this paper, we present a new framework that allows us to quantify the interplay between prevailing economic narratives and the real economy and financial markets. The advent of sophisticated algorithms for textual and semantic analyses enables credible quantitative analysis. We proceed in three steps. First, we obtain survey responses from US stockholders on their stories about the impact of COVID-19 on the economy and financial markets. Second, we use textual analysis to retrieve narratives from the survey responses and quantify their propagation over time. Third, we assess the narrative basis for economic fluctuations by means of three analyses. The first analysis begins by establishing the validity of the extracted narratives by testing their natural implications for human mobility using data from Google Maps. We then use network analysis on the extracted time series of narrative prevalence in conjunction with a large set of daily macro-financial time series to examine the role narratives have in driving the fluctuations within this system over time. Our approach is agnostic, allows for bi-directional effects, and does not impose any a-priori structure on the relationship between narratives and the macro-financial system. We find strong evidence for the narrative basis for economic fluctuations, yet also find that narratives are themselves driven by those. The third analysis examines asset pricing implications by estimating the risk premia for each of the narratives (and the macro-finance variables). This is motivated by the Intertemporal Capital Asset Pricing Model (ICAPM) of Merton (1973) which predicts that variables affecting investors' marginal utility of consumption constitute systematic risk factors. Our analysis reveals significant risk premia for narratives, supporting their role in financial markets.

The COVID-19 pandemic and the global economic recession that resulted provide an ideal testing ground to examine the narrative basis for economic fluctuations. It is the largest economic shock since the Great Depression of the 1930s and has resulted in huge volatility swings in financial markets.[1] Large shocks make it easier to measure and isolate any links between the macro-financial system and popular narratives. The rarity of global pandemics means that there is a large degree of uncertainty associated with COVID-19, particularly during its early stages. This has led to a plethora of heterogeneous and distinct narratives as people tried to comprehend the severity and impact of something they have never experienced in their lifetimes. Finally, in contrast to other historical economic shocks, this time is indeed different since social media and the internet allow for a much more rapid spread of narratives. For example, the number of adult people in the US that used social media during the advent of the 2008-2009 recession was only 21%; as of 2020, this number has increased to approximately 80%.[2] This represents a shift from traditional news and word of mouth as primary mediums of transmission.[3]

To obtain data on narratives, we rely on daily direct surveys sent to US stockholders in the period from February 29, 2020, to June 26, 2020. This sample period allows us to capture narratives that developed in the early stages of the pandemic timely and in real time. This information is likely to be difficult to recover using ex-post surveys due to contamination by a cognitive recency bias. The survey asks the following open-ended question:

> *"Please describe what, in your opinion, are the main reasons that the spread of the coronavirus has a negative (or positive) effect on the financial markets."*

[1]The S&P500 Index lost 30% of its value in only 22 trading days starting on February 19, 2020. This was followed by a rally that started on March 23 that brought the index back to its record high in five months.

[2]For a direct link to the source see `https://ourworldindata.org/` under technology adoption in US households.

[3]Cinelli et al. (2020) analyze the impact of social media on the transmission of narratives and discourse about the COVID-19 pandemic. Taking an epidemiological approach, they calculate a basic reproduction number R_0 for different social media platforms and find that narratives spread in a manner similar to viral pandemics.

We used Amazon's Mechanical Turk (MTurk) to distribute our survey. MTurk is a large crowd-sourcing platform that provides researchers with the opportunity to obtain survey data by paying workers to complete discrete on-demand tasks (e.g., questionnaires). This platform is previously used by Goetzmann et al. (2017), but is otherwise novel to the economics literature. The advantage of this approach is that we can directly elicit investor narratives and beliefs. This direct approach has several advantages over proxy methods such as using newspaper data to extract narratives. The first is that news aim to report facts, often in a distilled and objective way, whereas popular narratives usually contain an emotional and subjective component. The second is that, in contrast to news, narratives do not necessarily need to be true to have a large effect on the economy or financial markets, they only need to have certain degree of prevalence and spread such that enough people believe in them.[4] This distinction is crucial in an age of misinformation, conspiracies, and "fake news".[5] Additionally, our medium and the format of the question allows investors to express expansively about suggestive causes and effects as strongly advocated by Shiller (2019).

Our final and cleaned data set contains about 80,000 words, approximately equivalent to two fiction novels, across 1,812 survey responses. To retrieve the economic narratives from the survey responses, we rely on the latent Dirichlet allocation (LDA) topic modeling approach of Blei et al. (2003). LDA is an unsupervised learning algorithm that extracts a fixed number of topics from a high-dimensional corpus of text. LDA achieves this by imposing a factor structure on the text corpus, representing each survey response as a mixture of topics, where each topic is itself composed of words or terms with a probabilistic weight to each.[6] We represent a narrative by each of these topics. The advantage of this method is that it allows us to compress the high-dimensional text corpus of all survey responses into a manageable number of comprehensible narratives. Since the method is unsupervised, there is no subjective researcher bias involved in the process of extracting the narratives. This contrasts a common approach to textual analysis in economic research which relies on dictionary methods in which the researcher pre-defines a set of terms of interest and then computes their counts across documents. This appealing feature has motivated some

[4]A recent example of narratives influencing economic behavior are the extreme price movements of the GameStop Corp. stock during the first weeks of January, 2021 as a result of narratives going viral on Reddit about the potential upside of the stock.

[5]For example, Kogan et al. (2019) find that fake articles induce abnormal trading activity and increase price volatility in financial markets.

[6]The procedure resembles Principal Component Analysis (PCA), but for text data, since each survey response is composed of a mixture of topics (similar to the components of PCA), each of which is comprised by a set of terms with assigned weights (similar to the weights on variables in the components).

recent uses in other areas of economics and finance, e.g., Hoberg and Lewis (2017), Hansen et al. (2018), Huang et al. (2018), Larsen and Thorsrud (2019), Adämmer and Schüssler (2020), Lowry et al. (2020), Bellstam et al. (2020), and Bybee et al. (2020). Once the narrative weights have been extracted, it is possible to aggregate across survey responses on a daily basis to obtain a time series that quantifies the narrative prevalence of each narrative over time.

The narratives obtained from the LDA estimation are highly interpretable and can easily be linked to plausible popular narratives during the pandemic's early and later stages. To facilitate interpretability, we label these narratives using the words with the highest probabilistic loading and our reading of the survey responses with the purest topic loading. We obtain the following 13 narratives in no particular order: *consumer confidence*, *personal spending*, *stock market crash*, *monetary policy intervention*, *supply disruption*, *business closure*, *job loss*, *infection worry*, *financial market impact*, *fiscal policy intervention*, *investor fear*, *stay at home*, and *COVID-19 status*.

The time series dynamics of these narratives tell an interesting story about their prevalence and importance during the first wave of the pandemic. For example, the *investor fear* and *COVID-19 status* narratives have a notable increase in mid-March 2020 when the national state of emergency was announced and a wave of lockdown and stay-at-home orders were imposed across the US territory. The *supply disruption* narrative has a much higher prevalence in the beginning of our sample when the concern for the majority of the population centered around potential shortage of goods and supply chain disruptions in China as well as on the effects of travel restrictions. These worries faded away as infections started increasing in the US and restrictions were imposed. This latter development is captured by the increased prevalence of the *business closure* narrative which increases sharply after compulsory closure of non-essential retail establishments were introduced across the US territory in the second half of March 2020 and beginning of April 2020. The *job loss* narrative has an increasing prevalence throughout the sample period, reflecting the increasing worry about the labor market and an historically unprecedented and elevated level of initial unemployment claims throughout the sample period.

Some of the narratives have clear implications for human mobility, which we utilize to formally verify that they capture relevant dynamics for individuals' behavior. Using data from Google Maps on people's mobility (daily time series of visits and length of stay at different places), we find that the *stay at home* narrative relates positively to mobility at people's own residence and negatively to mobility at their workplace, the *business closure* narrative relates positively to mobility at people's own residence and negatively to mobility at their workplace, and that the *supply disruption narrative* relates positively to mobility at retail and groceries consistent with runs to supermarkets in fear of shortage of goods. These relations are statistically significant and often very strong; for example, more than half of humans' mobility at their homes and workplaces is related to changes in the prevalence of the *business closure* narrative.

We then shift our focus to quantifying the interplay between narratives and the macro-financial systems. Since the relationship is hypothesized to be bi-directional (Shiller, 2019, 2020) we cannot rely on simple regression analysis since that is likely to face endogeneity issues. Further, since we lack a formal economic theory to guide us on the structural form of the model and on which variables to include, we are at risk of an omitted variable bias. The latter can result in the incorrect identification of links between narratives and the macro-financial systems. To account for the lack of economic theory to guide us, endogeneity, and omitted variable bias issues, we rely on an agnostic data-rich approach. Specifically, we consider a high-dimensional Vector AutoRegressive (VAR) system that contains numerous macro-finance variables and allows each variable to have a bi-directional role. We include 17 economic indicators that are available at a daily frequency, including some related to the real economy (e.g., the ADS index (Aruoba et al., 2009) and inflation expectations) and financial variables such as equity, bonds, and commodity market returns. We also include two news media variables; the economic policy uncertainty index (Baker et al., 2016) and a recently developed infectious disease tracker of the equity market volatility (Baker et al., 2020). The inclusion of these indices allows us to measure the effect of the news on narratives, the macro-financial indicators, and vice versa.

We follow Diebold and Yilmaz (2009, 2012, 2014) and use a variance decomposition of unexpected fluctuations for each variable in the high-dimensional VAR framework to obtain a network representation of our system. Under this framework, we obtain a measure of connectedness by analyzing the magnitude of unexpected fluctuations over a horizon from time t to $t+h$, $h \geq 1$, of a given variable due to shocks arising elsewhere in the VAR or network at time t. This measure has a direct relation to impulse response functions. These effects can be obtained at different horizons (h), which permits an analysis on the speed of transmission in shocks across series. We consider effects spanning one day up to one month. To address the high-dimensionality of the underlying VAR model in an agnostic manner, we use regularization via Elastic Net estimation (Zou and Hastie, 2005). This removes superfluous VAR relations, and only leaves in those relations that are potentially Granger causal, retaining economic interpretation. We tailor a residuals-based parametric bootstrap for statistical inference.

We find strong evidence that narratives are connected with macro-financial variables. Starting with the daily horizon, 12% of the unexpected fluctuations in the macro-financial set of variables are directly attributable to fluctuations in narratives. These numbers are statistically (at the 1% level) and economically significant. The effects accumulate over time, increasing to 15% at a weekly horizon and topping at over 20% at a monthly horizon. The share of unexpected fluctuations in narratives attributable to macro-financial variables is moderately larger at approximately 17% at the daily horizon and increasing to 32% over a monthly horizon. We find that the net effect, that is, the difference between the fluctuations of macro-finance driven by narratives and the fluctuations of narratives driven by macro-finance is generally statistically insignificant, enforcing the bi-directional nature of narratives' role for economic fluctuations.

Visualizing the connectedness measures in a network representation, we discern several important patterns. First, networks become denser as the horizon increases. This indicates that some transmission between narratives and the macro-financial system is not immediate but might span at least one month. The network as a whole is, nonetheless, significantly connected at the daily horizon. Second, the network clusters variables around three tightly integrated groups that remain relatively stable over the time horizons we analyze. Two of these groups contain a mix of narratives and macro-financial variables, while the last group is mainly composed of the latter, though the groups strongly interact with one another.

We then look at individual links in the network and identify multiple interesting relations. The *supply disruption* narrative plays a significant role in the network at the monthly horizon, driving a sizeable 15.0% of the unexpected fluctuations of VIX and 14.2% of the TED spread. The reverse direction is negligible. We interpret this as evidence that supply chain shortage was an important driver of investor fear in the early stages of the pandemic.[7] There is also a strong one-sided connection between the ADS index and the *fiscal policy intervention* narrative at the weekly and monthly horizon. A notable 20.9% of the unexpected variation in narratives on fiscal policy and stimulus packages are driven solely by unexpected changes to the ADS index at the monthly horizon. On the other hand, a mere 0.7% of unexpected fluctuations in the ADS is attributable to *fiscal policy intervention*. This indicates that individuals' stories on government actions to mitigate the economic impact of COVID-19 are driven by their observations of the economic environment. The Fed funds rate is always at the center of the network. It is strongly shaping the

[7]The VIX and TED spread are commonly used as proxies for investor fear.

COVID-19 status narrative. This means that unexpected changes in the Fed funds rate shape how people talk about the societal impact of COVID-19.[8]

In our final analysis, we estimate risk premia for each of the narratives as well as the macro-finance variables. The aim is to quantify how narratives manifest in asset prices and expected returns. Guided by the ICAPM of Merton (1973), we use innovations from the VAR used in the network analysis as state variables that constitute candidate risk factors. To account for the possibility of omitted risk factors (a relevant concern given the plethora of factors available (Harvey et al., 2016)) we use the novel methodology of Giglio and Xiu (2021). Our tests assets span a large number of equity portfolios formed on the basis of a variety of firm characteristics. We identify multiple significant risk premia, two of which are associated with narratives. The *supply discruption* narratives, which was also dominant in the network analysis, carries a significant negative risk premium. This suggests that stories spreading on, e.g., goods shortage are associated by investors with a bad state and high marginal utility. The *monetary policy intervention* narrative also carries a negative and significant risk premium, suggesting that stories spreading on the interest rate cuts by the Federal Reserve are related to periods of high marginal utility states. This is intuitive since investors are more likely to talk about rate cuts when the economic situation deteriorates. These findings support the role of narratives in the financial markets and strengthens the conclusions from the network analysis. Altogether, we document quantitative evidence in support of narrative economics posed by Shiller (2017, 2019, 2020).

The rest of the paper is laid out as follows. In Section II, we describe the open-ended questionnaires aimed at eliciting the popular narratives and provide relevant summary statistics and examples. Section III covers the LDA methodology and retrieves the survey responses' narratives. It also examines their implications for people's mobility. Section IV describes the empirical results from the network analysis and our graphical interpretation as well as the asset pricing implications. Section V concludes.

II. Open-ended questionnaire

This section describes our data set of open-ended questionnaires aimed at eliciting investor beliefs central to the financial market's behaviour during the COVID-19 pandemic. It also provides some summary statistics and examples.

We design a survey that asks the following question:

[8]The Federal Reserve lowered the rate twice during our sample period.

> *"Please describe what, in your opinion, are the main reasons that the spread of the coronavirus has a negative (or positive) effect on the financial markets."*

On a daily basis, the survey was randomly presented to respondents using Amazon's Mechanical Turk (MTurk). MTurk is a large crowd-sourcing platform that provides researchers with the opportunity to obtain survey data by paying workers to complete human intelligence tasks (e.g., questionnaires). This platform is also used in Goetzmann et al. (2017). It allows us to obtain real-time and high-frequency survey responses from a broad sample of US investors. Our findings thus inductively infer prevailing narratives from this representative sample of the population. The survey respondents at MTurk have been found to be representative of the US population as a whole based on gender balance, racial composition, and income. Yet, the average person on MTurk is somewhat younger than the US as a whole. This justification for MTurk for data collection has recently been summarized in Lowry et al. (2016). If the investor beliefs elicited in the surveys are skewed in some direction, e.g., due to so-called echo chambers (Brock and Balloun, 1967; Gentzkow and Shapiro, 2011; Cookson et al., 2020), our retrieved narratives, cf. Section III, are not necessarily representative for the entire population. In that case, any identified relationship among macro-financial variables and narratives, cf. Section C, may be viewed as conservative since we are only capturing a subset of popular narratives. That is, our findings of the narrative basis for economic fluctuations are likely stronger than what is presented.

The survey was presented to US stockholders through daily waves with the first wave sent on February 29.[9] The outbreak in China had already occurred at this point, yet its global spread had not been widely reported nor commonly recognized. As such, our first wave was sent before the outbreak of the COVID-19 pandemic in the US, well ahead of the first steep increase in reported cases and deaths in the US, cf. Figure 1.[10] We continued presenting the survey to respondents on a daily basis, ending the survey on June 26, 2020, covering a total of 119 days. This elicits changes in respondents' beliefs about the stock market's behaviour in a very frequent manner, capturing the beginning, peak, and dampening of the first wave of the COVID-19 pandemic in the US. The survey was initiated shortly after the US stock market realized a sudden crash starting on February 24, 2020, amounting

[9]The MTurk service allows you to restrict survey respondents by *premium qualifications*. In our case, we restricted our sample to be stockholders currently living in the US. The MTurk service validates these restrictions by requiring respondents to provide documentation that shows that they own financial assets.

[10]Further, the first state-level declaration of emergency in the US happened 29 February in the state of Washington, matching the start date of our survey.

Figure 1: Survey start and end dates versus COVID-19 spreading

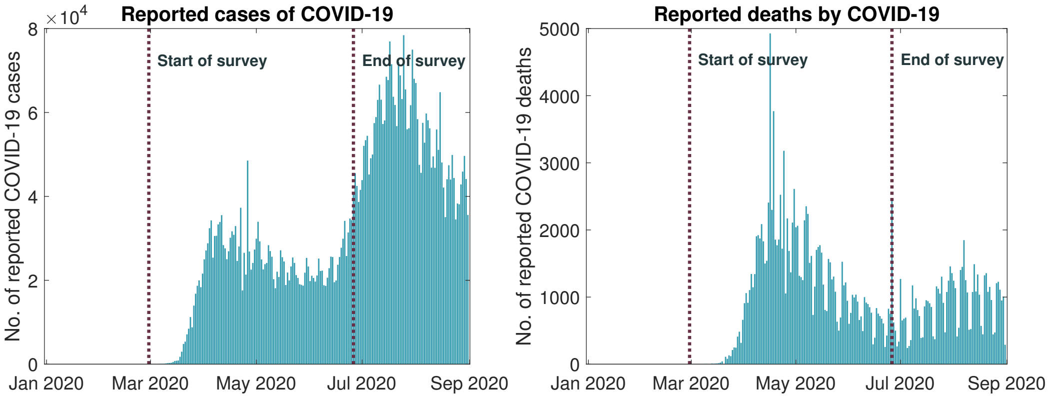

This figure depicts reported cases (left panel) and deaths (right panel) of COVID-19 in the US. The horizontal dashed lines in each figure indicate the start and end of our daily survey send to US stockholders.

to the largest one-week decline since the 2008 Great Recession. In this way, our collection of daily survey responses offers an extraordinary and frequent insight into the changing beliefs throughout the outbreak of the pandemic in the US, and one of the most sudden stock market crashes in the last century.

It is important to emphasize that our survey is *real-time*, allowing us to obtain the narratives that prevailed when each daily survey wave was sent to respondents. This is crucial, since any ex-post reconstruction of history by asking people, say half a year later, will be severely impacted by a recency bias. That is, individuals' recollection of a particular event will be influenced to a large extent by their current, most recent knowledge and memories.[11] To get an accurate representation of the extant narratives at each point in time, frequent and real-time collection of responses is thus essential. Moreover, our survey question is formed such that it instructs respondents to tell a story that is "*suggestive of causes in the current environment*" (Shiller, 2019) and it is "*inviting them to talk expansively*" (Shiller, 2019) by leaving the question open-ended. This renders the responses elaborate and focuses on the individual's story, facilitating an understanding of the prevailing narratives that may influence their behavior. The approach is motivated by and, thus, consistent with the recommendations in Shiller (2019).

[11]The empirical evidence of recency bias roots in the psychology literature, e.g., the seminal paper of Murdock (1962) in which it is shown that individuals presented to a list of words of irrelevant orders tend to recall the words presented the latest. Its relevance in economics and finance has since been documented in, for instance, the context of investor trading behavior (Nofsinger and Varma, 2013), financial statements auditing (Tubbs et al., 1990), mutual fund selection (Gruber, 1996), and stock price momentum (Bhootra and Hur, 2013).

Figure 2: Survey characteristics over time

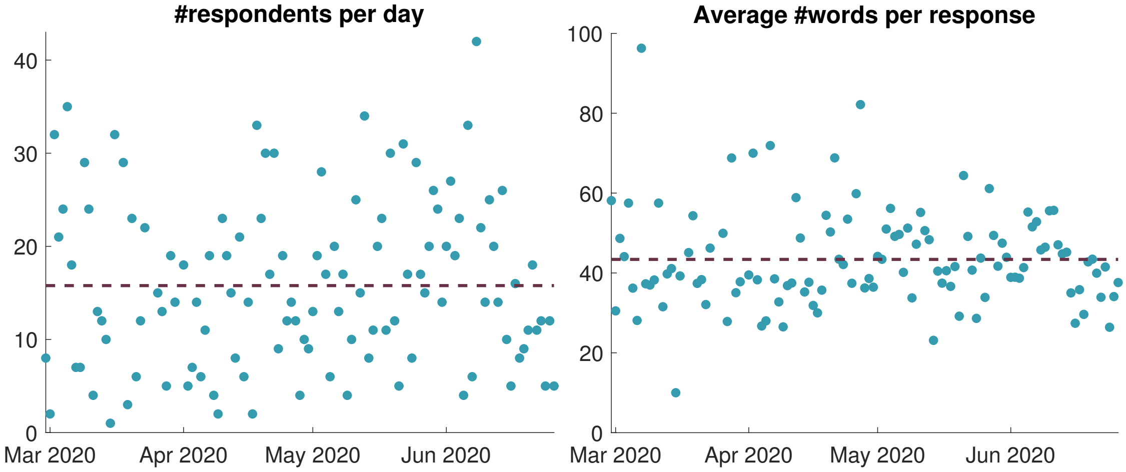

This figure depicts the number of respondents per day (left panel) from the beginning of the first survey date of February 29, 2020, through the end date of June 26, 2020. The right panel depicts the average number of words per response per day seen over the same time period. The dashed purple line measures the full sample average.

A. *Survey characteristics and examples*

The survey contains a total of 2,076 survey respondents. We remove duplicate responses (128 in total), days with no responses (four in total), responses with fewer than five words (93 in total), and responses that are deemed invalid by manual reading (40 in total) due to respondents misunderstanding of the question. This leaves 1,815 replies, corresponding to an average of 16 replies each day. Figure 5 shows that the number of respondents fluctuates over time, with some days having a total of 42 subjects and other days about a handful, yet there is no systematic pattern over time. It also shows that the average number of words per response fluctuates over the sample period. Yet, the average length per response of 44 words (equivalent to about 700 words per day) appears representative for the full sample period and, again, is without a systematic pattern. The total word count is 79,922, the equivalent of two fiction novels filled only with investor beliefs.[12] Each response varies somewhat in length with a standard deviation of 39 words, and some responses are very lengthy, with a maximum of 474 words.

To illustrate the type of responses, their elicitation of investor beliefs, and their relation to narrative economics, we present here part of selected quotes:

> *"[...] I believe that it is due to fear. People are afraid of the unknown and there is so much we don't know about this virus. [...] No one knows*

[12]The minimum length of a novel is 40,000, according to the Science Fiction and Fantasy Writers of America which is a nonprofit organization of professional science fiction and fantasy writers, cf. https://web.archive.org/web/20090319043837/http://sfwa.org/awards/faq.htm.

> *what is going to happen here although it does seem that conditions will worsen. [...] There seem to be conflicting information coming from our leaders. President Trump is downplaying the risks of this virus while his administration are saying different things that sometimes sound contradictory. This confuses us and makes us resistant to investing in the markets. [...] This mistrust makes us feel less confident and more financially insecure which causes us to pull back from spending freely. This negatively impacts the markets.*" ID#4, 29 February, 2020

> "*[...] Now that we know how infectious this virus is, much of the world is quarantined thus, travel is slowed, production is slowed, consumption is slowed and supply chains interrupted. This may very well lead to the recession we have been due for for a while.*" ID#314, 14 March, 2020

The (part of the) first response strongly indicates a situation of fear and uncertainty among investors. This is mainly driven by a lack of information and mistrust in that provided by public authorities. The (part of the) second response demonstrates a degree of resolution of uncertainty, going from fear and lack of (reliable) information to acknowledging a present economic recession. This change happens within a span of two weeks, indicating the necessity in obtaining real-time and frequent responses to capture those rapid shifts in beliefs. Importantly, the first response reveals a clear link to narrative economics. The narrative that information quality is uncertain causes individuals to be resistant to investing in the financial markets and to pull back from spending freely. The narrative directly influences the individual's investment-consumption decisions, as argued by Shiller (2019). Lastly, part of another response reads:

> "*[...] most people have been brainwashed into thinking they need to stop doing everything and go live in a bunker.*" ID#457, 27 March, 2020

This response links to a narrative of widespread lockdown and that individuals should isolate themselves. The respondent indicates that this narrative has spread, affecting the majority of people, just as proposed by Shiller (2019).

III. Narrative retrieval via topic modelling

In this section, we outline our approach for retrieving narratives from the survey data via textual analysis. We then quantify these via time series that each represents the prevalence of a given narrative at every time point t. Our exposition is mainly intuitive, influenced by the introduction in Blei (2012). We refer to the Appendix for a detailed explanation of data preparation, cleaning, model representation, and estimation.

A. Topic modeling and interpretation

In order to extract the most salient semantic themes of the high-dimensional corpus of survey responses, we follow the latent Dirichlet allocation (LDA) topic modeling approach of Blei et al. (2003). LDA is an unsupervised learning algorithm that seeks a tractable thematic summary of the survey responses into N_n so-called topics, with subscript n abbreviating "narratives". To achieve this, LDA imposes a factor structure on the text corpus with a resemblance to conventional factor modeling like Principal Component Analysis (PCA). The idea is to represent the entire text corpus by a set of topics, each of which is a grouping of terms. This is similar to a common approach to textual analysis in economic research which relies on dictionary methods in which the researcher pre-defines a set of terms of interest and then computes their counts across documents. However, the important advantage of LDA over dictionary-based methods is that it determines, in a data-driven manner, which terms are the most important for discriminating between themes in the survey responses rather than imposing this on the data. This appealing feature has motivated some very recent uses in economics and finance, e.g. Hansen et al. (2018), Larsen and Thorsrud (2019), Adämmer and Schüssler (2020), and Bybee et al. (2020).[13]

Crucially, the output is highly interpretable. To understand this, let $\boldsymbol{W}$ be a $S \times V$ matrix capturing the "bag-of-words" representation of the survey responses. Row indices correspond to the list of different survey responses and column indices to the vocabulary of V many different unique terms in the total text corpus. The $w_{s,v}$ element of $\boldsymbol{W}$ represents the number of times term v appears in survey response s. Naturally, $\boldsymbol{W}$ is high-dimensional – in our case, it is $1,812 \times 5,144$ after cleaning (e.g. removing stopwords and numbers), including bi-grams of terms, and applying TF-IDF weights, cf. Appendix A for additional details.[14] To summarize the thematic content of $\boldsymbol{W}$, LDA assumes a distribution for the V-dimensional term count vector of the s'th survey response, where the expected term counts of given response are summarized compactly as $\boldsymbol{\Phi}'\boldsymbol{\theta}_s$. Here, $\boldsymbol{\Phi} = [\boldsymbol{\phi}_1, \ldots, \boldsymbol{\phi}_{N_n}]'$ are the text corpus-wide topics, functioning as common factors, and $\boldsymbol{\theta}_s$ is the survey response-specific allocation towards each of the topics, functioning as factor loadings. Specifically, the n'th topic, $\boldsymbol{\phi}_n$, is a probability distribution over terms such that $\phi_{n,v} \geq 0$ for all v and $\sum_v \phi_{n,v} = 1$. This makes the output highly interpretable, because the set of terms that take a high probability in $\boldsymbol{\phi}_n$, which we refer to as *key terms*, convey the thematic content of the topic. This is as an essential input into labeling the topics.

[13] Hassan et al. (2019) use a different tool from computational linguistics to measure political risk faced by individual U.S. firms.

[14] We remove three responses, compared to the previous section, since they have no terms left after removing stopwords and low TF-IDF ranked terms.

The value of N_n, i.e. the number of topics, is taken to be much smaller than V, the vocabulary, to enforce dimension reduction and facilitate our interpretation of the prevailing narratives in the corpus of survey responses. That is, it retrieves salient COVID-19 narratives. Moreover, $\boldsymbol{\theta}_s = [\theta_{s,1}, \ldots, \theta_{s,N_n}]'$ is also a probability vector, which specifies how the topics are prevailing within the s'th survey response. In that way, LDA views each survey response as a mixture of topics with $\theta_{s,n}$ representing the prevalence of topic n in survey response s. This facilitates further interpretation, since inspection of $\boldsymbol{\theta}_s$ allows the researcher to associate certain survey responses with specific topics and enhance their understanding of the thematic content of retrieved narratives. We also use these as input in interpreting the topics and provide examples for the reader in Appendix B.

As informally put by Blei (2012), LDA balances the trade-off between assigning high probabilities to as few terms as possible within a given topic n and allocating terms in the survey response s to as few topics as possible. Both objectives enhance interpretability. This goal is achieved by identifying a small number of clusters of terms that frequently appear together across the corpus of survey responses and assign these as topics. The estimation of $\boldsymbol{\Phi}$ and $\boldsymbol{\theta}_s$ for each survey is conducted via Bayesian methods and the Gibbs sampler from Griffiths and Steyvers (2004), cf. Appendix A. The approach can be understood intuitively from the algorithm's "writing" of text. First, for each survey, it draws randomly from the list of topics with the sampling probabilities determined by $\boldsymbol{\theta}_s$. Suppose it draws the n'th topic. Second, to draw the first term of the survey, it then draws randomly from $\boldsymbol{\phi}_n$. The process is repeated as many times as there are terms in the survey. This procedure is conducted for each of the survey responses in the corpus, and the Gibbs sampling estimation obtains those $\boldsymbol{\Phi}$ and $\boldsymbol{\theta}_s$ that best simulates the writing of our entire text corpus, with the dimensionality reduction and interpretation goals of LDA in mind. This setting, therefore, naturally leads to intuitive metrics we may use for representation of the output that facilitates interpretation.

A.1. Topic number selection

A commonly used data-driven approach to select the number of narratives to extract is via Bayes Factors (BF). The higher the BF, the better the statistical fit of the text corpus obtained from the model. Since the null model is the same across all models (that only differ in the number of topics N_n), it is sufficient to consider the pattern of maximized log-likelihood values as it is the key constituent of the posterior model probability. This is depicted in Figure 3. Increasing the number of topics naturally improves the model's ability to describe the pattern in the survey

Figure 3: Log-likelihood versus number of topics

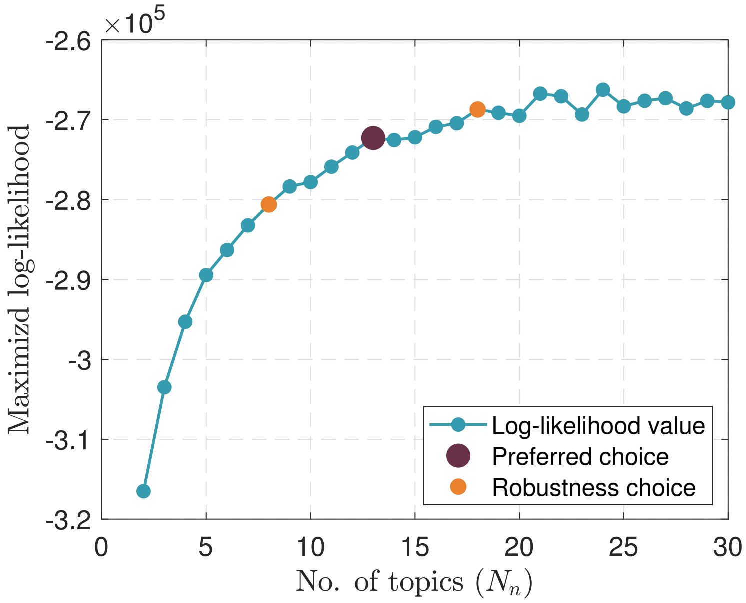

This figure depicts the maximized log-likelihood à la Griffiths and Steyvers (2004) for models with topics N_n ranging from 2 to 30 in increments of one. We mark our preferred choice at $N_n = 13$ with a large circle in purple, and our two choices for robustness check ($N_n = 8$ and $N_n = 18$) with a smaller orange circle.

responses initially, yet the improvement in fit levels off quite quickly after about ten topics. Nevertheless, the choice of the number of topics involves a trade-off between interpretability of the output of the model (lower N_n) and statistical goodness-of-fit (higher N_n). Moreover, picking too few topics will mix relatively distinct narratives into one broad topic and clutter the picture, whereas too many topics will be highly specific to particular limited sets of words or phrases. Following Hansen et al. (2018), we overcome this challenge by favoring interpretability in the sense that we choose a value for N_n that is marginally smaller than the value that optimizes the BF.[15] We therefore settle at $N_n = 13$ topics. At this point, there is little improvement in statistical goodness-of-fit by increasing N_n and interpretability of the topics remains. We find that our main conclusions are robust for $N_n = 8$ and $N_n = 18$, yet both choices lead to less interpretable outcomes.

A.2. Narrative retrieval and their labels

We represent a narrative by a topic from the output of the LDA, that is, each of the topics represents a prevailing story explained by the survey respondents. We,

[15] Blei (2012) argues that interpretability is indeed a legitimate reason to choose a number of topics different from the one that optimizes statistical goodness-of-fit. Generally, a lower number of topics favors interpretability.

therefore, refer to each topic as a narrative and use the wording interchangeably throughout. We manually assign a label to each of the thirteen retrieved narratives based on our reading of each narrative's key terms and the survey responses that load the highest on that topic. Although somewhat subjective, these labels are intuitive and serve as a shorthand for referring to narratives throughout the paper.

In Table 1, we report the top ten key terms for each narrative along with their assigned label. We report the stemmed terms used in the estimation of the topics, cf. Appendix A. Following, e.g., Bybee et al. (2020) we base the labelling on re-weighted term proportions

$$\tilde{\phi}_{n,v} = \frac{\hat{\phi}_{n,v}}{\sum_{n=1}^{N_n} \hat{\phi}_{n,v}}, \tag{1}$$

where hats indicate estimated values. The re-weighting reduces the importance of terms that are common across many topics and emphasizes those that have an uncommonly large proportion in topic n. This allows us to best identify the unique semantic of each topic by sorting the elements in $\tilde{\phi}_{n,v}$. The estimated topics are highly interpretable and can easily be linked to distinct and salient narratives about financial markets and the economy during various stages of the COVID-19 pandemic.[16] We identify:

1. a *consumer confidence* narrative concerning individuals confidence and lack thereof,
2. a *personal spending* narrative about individuals' personal spending and the stimulus check provided to individuals and households,[17]
3. a *stock market crash* narrative reflecting the pronounced stock market crash as the initial response to the severe outbreak of COVID-19,
4. a *monetary policy intervention* about the interest rate cuts by the Federal Reserve at unscheduled meetings,[18]
5. a *supply disruption* narrative concerning the impact of the pandemic on the (foreign) supply chain of goods,

[16] In contrast, a choice of, say, $N_n = 8$ mixes several narratives rendering interpretability unclear. An estimated topic, for instance, include elements from the *consumer confidence*, *fiscal policy intervention*, and *stay at home*, as evident from the following 10 key terms: *distanc*, *consum spend*, *consum*, *economi*, *initi*, *predict*, *govern*, *packag*, *consum confid*, *stabil*.

[17] The federal stimulus bills passed by Congress included a one-time payment of 1,200 dollars per qualifying adult and 500 dollars per child. Additionally, they also include expanded unemployment benefits for a limited period of time.

[18] There were two cuts to the Fed Funds rate announced by the Fed in unscheduled meetings during our sample period. The first one on March 3, 2020 lowered the rate by 50 basis points while the second one on March 15, 2020 lowered it further by 100 basis points.

Table 1: Narrative labels and top ten key terms

This table shows the top ten key terms for each labelled narrative based on the estimated $\mathbf{\Phi}$, based on their re-weighted values $\tilde{\phi}_{n,v}$ in (1). The topics are not ordered in any way. The ordering of the key terms, however, is so that the first term has the highest probability and the last the tenth largest probability associated with the n'th topic. We bold terms that we primarily associate with the labelling of each of the topics.

No.	Label	Key terms	Abbreviation
1.	**Consumer confidence**	**confid**, **lack**, **consum confid**, **consum**, reduc, **investor confid**, restrict, **lack confid**, set, **faith**	CCF
2.	**Personal spending**	**peopl spend**, peopl aren, **spend**, **money peopl**, **stimulus check**, aren, **money spend**, **peopl money**, don money, **check peopl**	PSP
3.	**Stock market crash**	**stock**, **stock market**, price, **drop**, **market crash**, **crash**, **short**, dollar, short term, trillion dollar	SMC
4.	**Monetary policy intervention**	week, **cut**, diseas, **rate**, **feder**, **reserv**, **percent**, anticip, crisi, correct	MPI
5.	**Supply disruption**	**industri**, **china**, **travel**, **suppli**, **airlin**, affect, **suppli chain**, **chain**, **manufactur**, intern	SPL
6.	**Business closure**	**busi close**, **busi**, **busi shut**, slowli, **close**, **busi force**, **essenti busi**, **shut busi**, busi busi, employ	BUC
7.	**Job loss**	**job**, **lost**, **job loss**, **peopl lose**, **lost job**, cure, miss, overreact, fear miss, miss ralli	JBL
8.	**Infection worry**	peopl buy, **sick**, item, **leav**, **hous**, **afraid**, **peopl don**, **worri**, **catch**, feel	INW
9.	**Financial market impact**	main, **financi market**, **coronavirus negat**, **effect financi**, **negat effect**, reason, main reason, **effect market**, **market experienc**, **experienc recoveri**	FMI
10.	**Fiscal policy intervention**	**pay**, economi, **packag**, **individu**, **economi peopl**, unemployment rate, rent, **bill**, **bank**, peopl	FPI
11.	**Investor fear**	**invest**, start, **hit**, **peopl invest**, **peopl scare**, **investor**, market return, **scare**, **peopl fear**, prepar	IVF
12.	**Stay at home**	**home**, mean, due, **stay at home**, **peopl stay**, **stay**, **social**, **prevent**, **distanc**, **isol**	SAH
13.	**COVID-19 status**	**panic**, **vaccin**, **media**, **corona**, **corona virus**, **news**, **hope**, incit, volatil, market react	C19

6. a *business closure* narrative about the lockdown of businesses, particularly non-essential retail establishments, as a result of government imposed restrictions,
7. a *job loss* narrative that focuses on individuals' worsened and generally uncertain employment situation,
8. an *infection worry* narrative about anxiety of catching the COVID-19 virus, making them afraid of leaving their houses,
9. a *financial market impact* narrative that addresses the impact of COVID-19 on financial markets in general,
10. a *fiscal policy intervention* narrative that relates to the stimulus packages

provided by the US government,[19]

11. an *investor fear* narrative representing individuals' general state of mind in this period,
12. a *stay at home* narrative about the widespread effects of the lockdown that forced individuals to spent more time at home and enforce social distancing, and
13. a *COVID-19 status* narrative concerning the reported status of coronavirus in the media, resulting panic, and the hope for a vaccine.

We provide elaborating details about the context for each of these narratives in Appendix B, including the survey responses that has the purest degree and is most representative of each of the topics.

Unlike conventional PCA, the topics do not have a natural ordering. Hence, as common when interpreting the output of LDA, we use cluster methods to examine the hierarchical structure of the topics. Based on the semantic distance between the topics, as measured by their $\boldsymbol{\phi}_n$ estimates, we estimate a dendrogram with a recursive agglomeration (Murtagh and Legendre, 2014). Furthermore, to get a better sense of the distance between the clusters, we reduce their dimension by PCA and plot them relative to the first two principal components. The resulting Figure 4 reveals a flat hierarchical topic structure judging on the distance among the clusters with none of the topics forming a distinctive cluster. In fact, it indicates that our identified topics belong to the same broad COVID-19 cluster, with no apparent hierarchy among them. This appears natural as all narratives concern the first wave of COVID-19 so that they, to some degree, are related as they address the same overall environment, yet represent heterogeneous aspects within this.

Evidently, the retrieved narratives from the survey responses capture numerous prevailing stories during the COVID-19 crisis. Some are closely related to the outbreak of the pandemic and therefore linked to the beginning of the period, such as *stay at home*, *supply disruption*, and *stock market crash*, and others are more closely linked to the later stage of the crisis as, for instance, *fiscal policy intervention*, and some likely follow in part the spreading of COVID-19 like *infection worry* and *COVID-19 status*. The following section, therefore, translates those labeled narratives into time series that measure their prevalence through time, providing a

[19] The US Federal Government has enacted several stimulus packages in response to the COVID-19 pandemic. The two largest ones are the Coronavirus Aid, Relief, and Economic Security Act (CARES) and the Health and Economic Recovery Omnibus Emergency Solutions Act or Heroes Act (HEROES), passed on March 27 and May 15 of 2020, respectively. Together, they amount to more 5 trillion dollars in additional federal spending.

Figure 4: Hierarchical taxonomy and cluster of narratives

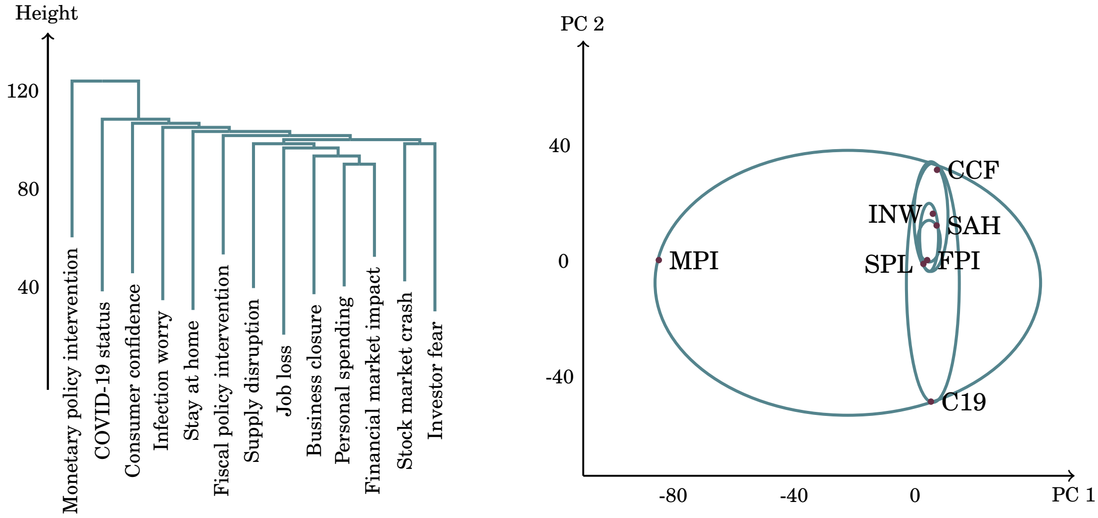

This figure depicts the hierarchical agglomerative clustering dendrogram (left plot) and nested cluster diagram (right plot) based on the estimated $\mathbf{\Phi}$, that is, the semantic distance (measured by the Euclidean norm) among the $N_n = 13$ topics. The nested cluster diagram relies on a PCA and shows the clusters relative to their first two principal components (PCs). For expositional reasons, the nested cluster diagram displays only the most outer clusters and the labels (with abbreviations from Table 1) of the first seven narratives in hierarchical ordering.

further interpretation as well as facilitates quantitative analysis of their interplay with the real economy and financial markets.

A.3. Quantifying time-varying prevalence of narratives

The survey-specific estimates $\boldsymbol{\theta}_s$ describes the proportion of each topic in survey response s. Since each survey is associated with a timestamp, we may within each day aggregate all topic proportions per survey to measure the daily prevalence of each individual narrative. We then normalize within each day, so that each narrative has an associated share of the total daily narrative prevalence and they sum to unity. Once put together over calendar time, this forms a daily time series for each narrative that measures its prevalence over time. Their time series dynamics facilitate further interpretation, yet it also enables an understanding of the propagation of each narrative during the (first wave) COVID-19 period. These are plotted in Figure 5. To further enhance the signal and emphasize trends, we also depict the 5-day moving average of the raw series. Some of the narratives are prevalent and relatively stable across the whole sample period. For example, the *investor fear* and *COVID-19 status* narratives have a notable increase in mid-March when the national state of emergency was announced and a wave of lockdown and stay-at-home orders were

Figure 5: Narrative prevalence over time

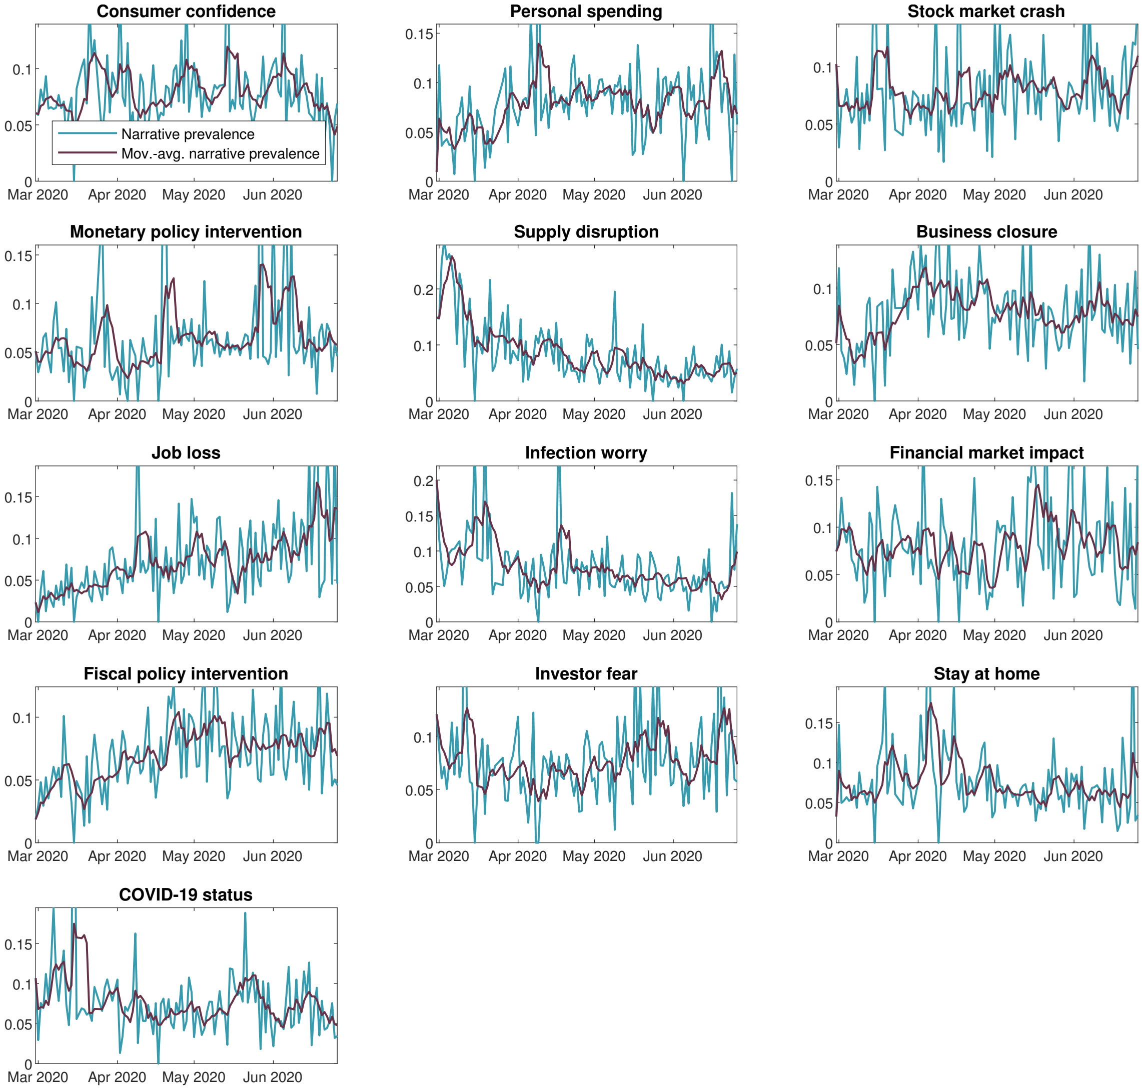

This figure depicts the time series of each narrative's prevalence during the February to June, 2020, period. The values on the second axis measure the share of total prevalence at a given point in day, summing to unity across all narratives each day. The blue (purple) line measures the raw (five-day moving average) time series based on estimated $\boldsymbol{\theta}_s$ for all surveys.

imposed across the US territory.[20] However, these narratives are otherwise stable, reflecting the general apprehensive state of mind of investors during the early stages of the pandemic and the general attention of the population towards its development and course. Other narratives display a more pronounced episodic and trending behavior. For example, the *supply disruption* narrative has a much higher prevalence in the beginning of our sample when the worry for the majority of the population centered around the potential shortage of goods and supply chain disruptions in

[20]Lockdown orders across the US mainland varied, but they span between March 18 when California issued the first state-wide stay at home order to April 7 when South Carolina and Missouri issued the decree as the last two states.

China as well as on the effects of travel restrictions. These worries faded away as infections started increasing in the US and local restrictions were imposed. This latter development is captured by the increased prevalence of the *business closure* narrative which increases sharply after compulsory closure of certain businesses were introduced across the US territory in the second half of March and beginning of April. Interestingly, the *job loss* narrative has an increasing prevalence thoroughout the sample period, reflecting the increasing worry about the labor market.[21]

B. Implications for people's mobility

There is ample evidence on large changes in mobility behavior during the pandemic. Some studies indicate that the decrease in mobility is largely voluntary and driven by information on the effects of social distancing during a pandemic (Maloney and Taskin, 2020; Gupta et al., 2020) while other studies point toward the effect of mandatory stay-at-home orders (Abouk and Heydari, 2020; Andersen, 2020). The effect of narratives on mobility has, to our knowledge, not been analyzed. Although we are careful to claim a causal link from narratives to mobility, since the effects can be bi-directional, some of the narratives have clear implications for human mobility.[22] To test this hypothesis, we make use of Google Mobility data.[23] This novel data set includes daily time series of visits and length of stay (mobility) at different places compared to an average baseline value in percent. The baseline is the median value for the corresponding day of the week during the pre-COVID-19 five-week period of 3 January to 6 February, 2020. The source of the data is anonymized mobile phone data from Google Maps. We are interested in the following three hypotheses:

1. the *stay at home* narrative relates positively to mobility at people's own residence and negatively to mobility at their workplace;

2. the *business closure* narrative relates positively to mobility at people's own residence and negatively to mobility at their workplace;

3. the *supply disruption* narrative relates positively to mobility at retail and groceries.

[21]Although the unemployment rate reached a peak in April 2020, initial unemployment claims have remained elevated above one million during our sample period. For comparison, the peak during the Great Recession in 2008-2009 was about 900,000.

[22]As an additional robustness, we also test the intertemporal causality of the relation by running at VAR(1) model with narratives and Google Mobility data. The general pattern that emerges is that narratives are Granger causal towards Google Mobility data but not the other way around.

[23]The data is available at https://www.google.com/covid19/mobility/.

Table 2: Implications for people's mobility

This table reports the estimates of ϕ_1, associated t-statistics and R^2 from the following regression $G_{it} = \phi_0 + \phi_1 N_{jt} + \varepsilon_{ijt}$ where G_{it} denotes Google Mobility data for the i'th group and N_{jt} the j'th narrative series corresponding to each of the three hypothesis listed in Subsection B. The t-statistics are computed using heteroskedasticity- and autocorrelation-robust Newey and West (1987) standard errors.

Google mobility	*Stay at home*			*Business closure*			*Supply disruption*		
	$\hat{\phi}_1$	t-stat.	R^2	$\hat{\phi}_1$	t-stat.	R^2	$\hat{\phi}_1$	t-stat.	R^2
Residential	2.75	4.92	21.04%	4.35	6.54	52.69%			
Workplace	-5.63	-3.86	14.07%	-10.95	-5.70	53.23%			
Retail and recreation							7.71	2.53	22.56%
Groceries and pharmacy stores							3.18	2.26	10.40%

The Google Mobility data enables a test of these hypotheses as it is grouped into *residential*, *workplace*, *groceries and pharmacy stores*, and *retail and recreational* mobility.[24] We regress each of these four groups of Google Mobility data onto each of the three narrative time series involved in the hypotheses above. For the i'th Google Mobility group and the j'th narrative time series the regression reads

$$G_{it} = \phi_0 + \phi_1 N_{jt} + \varepsilon_{ijt}. \tag{2}$$

To make the estimates comparable, the narrative series have been standardized to have a mean value of zero and a standard deviation of one. This implies that one standard deviation increase in the narrative time series results in ϕ_1 percentage increase in mobility with respect to the pre-COVID-19 baseline. The results are collected in Table 2.

We confirm each of the three hypotheses. The *stay at home* narrative is significantly and positively (negatively) related to residential (workplace) mobility at a 1% significance level. The estimate of ϕ_1 suggest that a standard deviation higher prevalence of the *stay at home* narrative corresponds to 2.75% higher (5.63% lower) presence at people's homes (workplace). The *business closure* narrative stands further out. A standard deviation increased talk about business closures corresponds to a 4.35% higher (10.95% lower) presence at people's homes (workplace), both effects being significant at the 1% level. The R^2 is notably high at about 53% in both cases, implying that more than half of humans' mobility at their homes and workplaces in this period is related to changes in the prevalence of the *business closure* narrative. Finally, the *supply disruption* narrative is positively (significant at the 5% level) related to both the presence at retail, recreation, groceries, and pharmacies consistent

[24] We de-seasonalize each series using day of the week dummies.

with the hypothesis that spreading stories about goods shortage lead to individuals running to stores. The estimates of ϕ_1 suggest that a standard deviation increase in the prevalences of stories on *supply disruption* corresponds to a 7.71% (3.18%) higher presence at retail and recreation (groceries and pharmacies).[25] Altogether, these findings document that the narrative time series are strongly related to individuals' decisions on to their mobility and, at a broader level, verify that they capture relevant dynamics for individuals' behavior. We study the economic implications of this in the following section.

IV. Real economic and financial implications

In this section we analyze the real economic and financial implications by means of two main analyses.

The first analysis quantifies the interaction of narratives with macro-finance variables using a variance decomposition of directional impulse responses in a Vector AutoRegressive (VAR) system that contains several macro-financial variables along with the time series of narrative prevalence.[26] Since the relationship between the narratives and macro-financial variables is hypothesized to be bi-directional (Shiller, 2019, 2020), simple one-sided regression will face endogeneity issues. Moreover, as argued earlier, there is no structural theory guiding us in the form of the relationship nor which macro-financial variables that should enter this analysis, making us at risk of an omitted variable bias. This motivates an agnostic approach. For these reasons, we consider a large VAR system that contains numerous macro-financial variables and allows each variable to have a bi-directional role. This analysis thus measures the amount of unexpected fluctuations of each variable or group (e.g. one or more macro-finance variables) due to shocks elsewhere in the VAR (e.g. due to one or more narratives). Diebold and Yilmaz (2014) show that such measures, referred to as connectedness, are sophisticated networks. We use this appealing intuition further below to represent and analyse our findings as networks.

The second analysis examines asset pricing implications by estimating the risk premia for each of the narratives (and the macro-finance variables). This analysis is motivated by the Intertemporal Capital Asset Pricing Model (ICAPM) of Mer-

[25] In unreported results, we find that other narratives have much weaker link to mobility and this link is often not statistically significant. For example, *infection worry* is positively related to residential mobility, possibly implying that people stay more at home when the fear of contagion increases, but the slope is not statistically significant.

[26] Variance decompositions of reduced form shocks in a VAR have a long tradition in macroeconomics and finance, for example for characterizing the dynamics of stock or bond returns (Campbell, 1991; Campbell and Ammer, 1993) or the interaction between the macreoeconomy and the yield curve (Diebold et al., 2006).

ton (1973) which predicts that innovations to state variables that affect investors' marginal utility of consumption constitute systematic risk factors that demand a risk premia. Since the hypothesis of narrative economics is that narratives influence individuals' decision making and, thus, their marginal utility, this second analysis quantifies to what extend and how this manifests in asset prices and expected returns. We use the innovations from the above-mentioned VAR system as state variables.

A. VAR system and data

To adequately capture the complex and multi-faceted aspects of the macro-financial part of the VAR system, we rely on a large number of daily indicators. In that way, we are agnostic in not imposing a specific structure on any potential relationship between narratives and the real economy or financial markets, letting data decide on the appropriate structure. Moreover, a high-dimensional system reduces concerns of omitted variable bias (Christiano et al., 1999; Giannone and Reichlin, 2006).

Table 3 lists our $N_{mf} = 17$ daily macro-financial variables, with mf indicating macro-finance variables, along with any transformation made to obtain stationarity, their definition, reference if applicable, data source, and sector of the economy they belong to. We refer to the table for a full overview. We include variables representing the real economy (e.g., the ADS index of Aruoba et al. (2009), inflation expectations, or a real estate investment trust index), equity markets (e.g., S&P500 returns and VIX), bond markets (e.g., term spread and default spread), commodities (gold returns and oil returns), and equity risk premia (size and value premium). We also include two news media variables: the economic policy uncertainty index (Baker et al., 2016) and a recently developed infectious disease tracker of the equity market volatility (Baker et al., 2020). We include the two latter for two reasons. First, it is natural to conjecture that narratives are shaped by what is reported in the media or, conversely, that individuals' demand for news shapes what is reported in the media.[27] Second, news media play a fundamental and important role in information dissemination and track the economy well (Bybee et al., 2020). Hence, we include these news-based variables to control for the role of news media in the economy.

We gather all $N = N_{mf} + N_n = 30$ variables into a vector

$$\boldsymbol{X}_t = [\boldsymbol{X}_t^{mf}, \boldsymbol{X}_t^{n}]' = [RET_t, \ldots, IDT_t, CCF_t, \ldots, C19_t]', \tag{3}$$

[27]As noted in Mullainathan and Shleifer (2005), the news presented in the media is an equilibrium outcome driven by consumer preferences, news production technologies, and the industry competition. In that way, news likely mirror aspects that are important to both news consumers and producers.

Table 3: Description of macro-finance variables

This table lists the 17 daily macro-finance variables along with their transformations to obtain stationarity, their definition, reference if applicable, data source, and sector of the economy they belong to. For z_{it} denoting the original time series, the transformation methods used to turn z_{it} stationary denoted by (1)-(3) corresponds to (1) $x_{it} = z_{it}$, (2) $x_{it} = \Delta z_{it}$, and (3) $x_{it} = \Delta \ln(z_{it})$. The transformations are consistent with Andreou et al. (2013).

No.	Abbreviation	Definition	Sector	Source	Trans. method	Reference
1.	RET	S&P 500 stock price index price	Equity	FRED St. Louis	3	
2.	VLM	S&P 500 stock price index volume	Equity	Yahoo Finance	3	
3.	VIX	CBOE equity market volatility index	Equity	FRED St. Louis	1	
4.	SMB	Small-minus-Big (size premium, Fama and French)	Equity	Kenneth R. French's website	1	Fama and French (1993)
5.	HML	High-minus-Low (value premium, Fama and French)	Equity	Kenneth R. French's website	1	Fama and French (1993)
6.	TMS	Term spread: yield difference between 10-year and 3-month Treasuries	Bond market	FRED St. Louis	1	
7.	TED	TED spread: yield difference between 3-month LIBOR and 3-month Treasury Bill	Bond market	FRED St. Louis	1	
8.	DEF	Default spread: yield difference between Baa- and Aaa-rated corporate bonds	Bond market	FRED St. Louis and FRED St. Louis	1	
9.	FFR	Federal funds rate (% p.a.)	Bond market	FRED St. Louis	2	
10.	DOL	Trade weighted U.S. dollar index: broad, goods and services	Real economy	FRED St. Louis	3	
11.	ADS	Daily Aruoba-Diebold-Scotti business conditions index	Real economy	Philadelphia Fed	1	Aruoba et al. (2009)
12.	INF	Expected inflation: rate at which 5-year Treasury Note and TIPS achieve same yield	Real economy	FRED St. Louis	1	
13.	RIT	Wilshire U.S. real estate investment trust total market index	Real economy	FRED St. Louis	3	
14.	GLD	Gold fixing price (3pm London fixing time)	Commodities	FRED St. Louis	3	
15.	OIL	Crude oil prices: West Texas Intermediate (WTI)	Commodities	FRED St. Louis	3	
16.	EPU	News-based economic policy uncertainty index	News	Economic Policy Uncertainty website	1	Baker et al. (2016)
17.	IDT	News-based infectious disease equity market volatility tracker	News	Economic Policy Uncertainty website	1	Baker et al. (2020)

where $\boldsymbol{X}_t^{mf}$ and $\boldsymbol{X}_t^{n}$ denote the group of macro-financial and narrative time series, respectively. A p-lag VAR, VAR(p), is then given by

$$\boldsymbol{X}_t = \boldsymbol{\Psi}_0 + \sum_{i=1}^{p} \boldsymbol{\Psi}_i \boldsymbol{X}_{t-i} + \boldsymbol{\varepsilon}_t, \tag{4}$$

with $\boldsymbol{\varepsilon}_t \sim (\mathbf{0}, \boldsymbol{\Omega})$ and where $\boldsymbol{\Psi}_i$ is a coefficient matrix. This system forms the basis of the following analyses.

B. Elastic Net for high-dimensionality and interpretability

For any value of the lag-order p conventional least-squares estimation of the high-dimensional macro-financial VAR system is intractable, causing a curse of dimensionality. For this reason, we resort to a machine learning technique known as Elastic Net (ENet) (Zou and Hastie, 2005). This is based on a penalized regression which alleviates the curse of dimensionality by augmenting the typical least-squares objective function with ℓ_1 and ℓ_2 penalty terms. The former allows for variable selection by setting coefficients possibly to zero, and the latter enforces shrinkage of coefficients. Since we are interested in every potential node of the network, regularization is preferred over factor-augmenting the VAR system to handle the data-rich environment as in, e.g. Bernanke and Boivin (2003). Moreover, it is important to allow narratives to interact with one another. The procedure will only leave in those potentially Granger causal relations by setting irrelevant coefficients equal to zero. The ENet thus retains economic interpretation and does not overfit. The estimates of the VAR(p) system are defined as per the following regularized multivariate least-squares problem

$$\underset{\boldsymbol{\Psi}}{\arg\min} \sum_{t=1}^{T} ||\boldsymbol{X}_t - \boldsymbol{\Psi}_0 - \sum_{i=1}^{p} \boldsymbol{\Psi}_i \boldsymbol{X}_{t-i} + \boldsymbol{\varepsilon}_t||_F^2 + \lambda \mathcal{P}_\alpha(\boldsymbol{\Psi}), \tag{5}$$

where $||\boldsymbol{\Psi}||_F^2 = (\sum_{i,j} \Psi_{i,j}^2)^{0.5}$ is the Frobenius norm, and the penalty function

$$\mathcal{P}_\alpha(\boldsymbol{\Psi}) = 0.5(1-\alpha)||\boldsymbol{\Psi}||_2^2 + \alpha||\boldsymbol{\Psi}||_1, \tag{6}$$

with $||\boldsymbol{\Psi}||_k$ is the ℓ_k norm, and $\alpha \in [0,1]$ is a blending parameter for the ℓ_1 and ℓ_2 components of the penalty term. If $\alpha = 1$, the penalty term reduces to that of the LASSO. The LASSO effectively selects relevant variables yet tends to arbitrarily select one variable from a group of highly correlated variables. The ENet is a refinement that mitigates this potential issue by adding the ℓ_2 (ridge) component. We tune λ and α in a data-driven manner via cross-validation.

It is important to note that the regularized VAR system does not necessarily impose regularization on a resulting macro-financial-narrative network since the variance decomposition a nonlinear transformation of the VAR coefficients and is, therefore, generally not necessarily sparse (Demirer et al., 2018). By encouraging sparsity via variable selection, we also benefit from reducing the probability of spuriously-inflated connectedness for large N that otherwise may occur simply by enabling more and more possible connections.

B.1. Regularized macro-financial-narrative VAR system

Since some of the included variables may exhibit non-negligible persistence, we estimate a VAR(5) system. Hence, we include a generous amount of lags in the model and let the ENet decide which ones are important. This ensures that the shocks of the system are not driven by unaccounted for serial correlation in the variables. Figure 6 depicts the selected variables in the VAR system from the ENet estimation divided into five matrices associated with each of the lags. We again emphasize that the essential interest is in the network connectedness, which is related to but does not correspond to these VAR estimates. Nevertheless, several important patterns emerge. First, there is a clear indication that our system captures the required persistence of the included variables, as the ENet selects diagonal terms (lagged dependent variables) for the first and partly for the second lag. Higher-order lags are not selected. Secondly, the structure is sparse.[28] There exist multiple selected relations at all lags, with lower lags tending to exhibit the most selected variables. Yet, the majority of coefficients are set identically to zero. This underpins the success of the ENet to select only those relationships that are potentially Granger causal and discard the remainder. Lastly, we see clear indications as to important relationships between narratives ($\boldsymbol{X}_t^n$) and macro-financial variables ($\boldsymbol{X}_t^{mf}$) exist. For instance, the upper right block for the first lag that measures the effect of narratives on next-day macro-finance has multiple selected coefficients, e.g., *financial market impact*, *stay at home*, and *COVID-19 status* on S&P500 and REITS returns or *infection worry* on expected inflation rates. On the other hand, the Fed funds rate and the TED spread affect next-period *investor fear*, as an example. *Supply disruption* is particularly effectual at the second lag, driving VIX, S&P500, REITS, and High-minus-Low returns. We take these findings as motivation for further analysis of the network connectedness below.

[28]This is also consistent with the fact that the tuned α lies at the unity boundary with full weight on the ℓ_1 term. That is, the cross-validation favours LASSO and an emphasis on variable selection rather than on coefficient shrinkage.

Figure 6: Selected variables from Elastic Net estimation

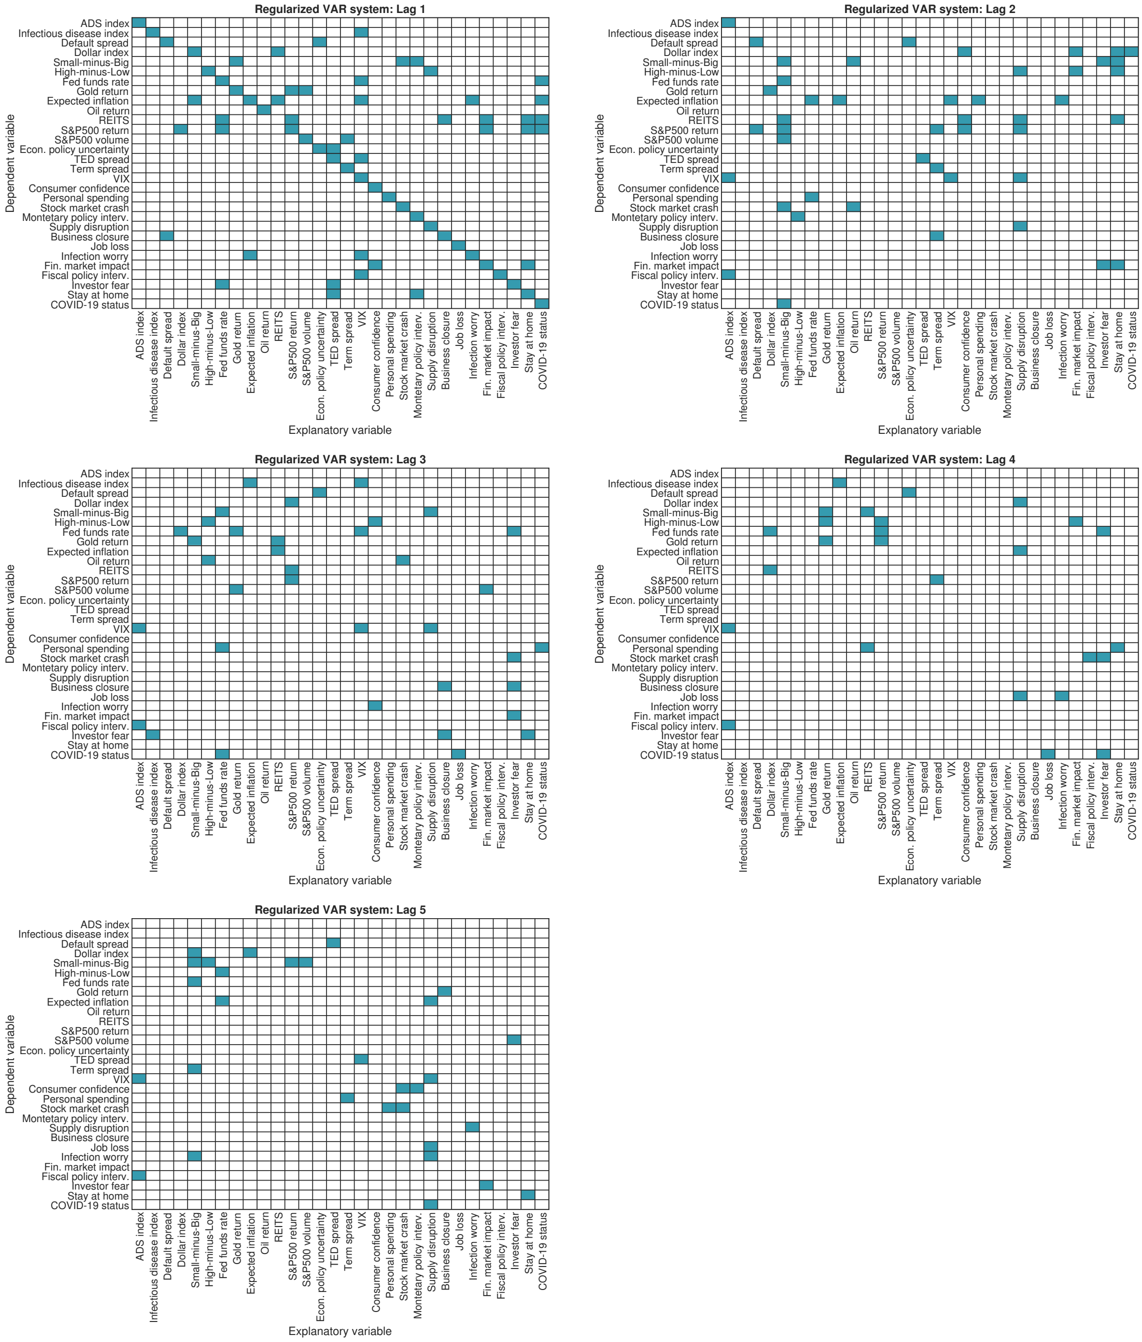

This figure depicts the selected variables by the ENet estimation. Each matrix corresponds to each of the 1-5 lags in the VAR system. The y-axis shows the dependent variable for each equation in the system, and the x-axis the lagged variables. A full blue square indicates selection, whereas an empty square indicates the ENet discarded the variable.

C. Narratives' macro-financial connectedness

In this section, we present our quantitative results for the interaction of narratives with macro-finance variables. We first examine this directional impact among the macro-finance group ($\boldsymbol{X}_t^{mf}$) and narratives group ($\boldsymbol{X}_t^{n}$) as a whole. We then represent the results as an elaborate network graph, which reveals the links among individual variables in the macro-finance-narratives system at the most granular level. Appendix G report results at an intermediate level where the top transmitters and receivers of shocks are outlined.

To understand the methodology, which has been popularized in a series of papers that includes Diebold and Yilmaz (2009, 2012, 2014), the ij'th, $i, j = 1, \ldots N$, h-step variance decomposition component measuring the fraction of variable i's unexpected fluctuation due to shocks in variable j is denoted by $v^h_{i,j}$. This element measures the effect to variable i cumulated over the time period t to $t+h$ due to a shock to variable j at time t. Let $GIRF^h_{i,j}$ denote the generalized impulse response function of variable i over h periods due to a shock to variable j, we may write

$$v^h_{i,j} = \frac{\sum_{q=0}^{h-1}(GIRF^q_{i,j})^2}{\sum_{j=1}^{N}\sum_{q=0}^{h-1}(GIRF^q_{i,j})^2} \in [0,1], \quad \sum_{j=1}^{N} v^h_{i,j} = 1, \tag{7}$$

where the numerator is the effect of the j'th shock, and the denominator measures the aggregate (over j) of all the shocks in the VAR system.[29] As such, the element $v^h_{i,j}$ captures the response of a given variable, for instance, S&P500 returns over some future horizon h due to a shock to another variable in the system, say the narrative *investor fear*. In other words, $v^h_{i,j}$ captures how much of the unexpected variation in variable i over the future horizon h is due to shocks to variable j, with the total effect across j summing to unity. A value of, say $v^h_{i,j} = 5\%$ would mean 5% of the unexpected fluctuations over h periods of variable i is due to shocks in variable j at time t. Importantly, these quantities are directional, that is, the effect from i to j is not necessarily equivalent to the effect from j to i. Since no theory nor previous empirical evidence is available to guide us in any credible structural identification restrictions, we remain objective in not imposing any restriction to a dominating direction. We let data inform us about the direction of transmission.[30] The estimation of the variance decomposition relies on the ENet estimates of the VAR system, see Appendix C for further details. Diebold and Yilmaz (2014) show that these generalized variance decompositions, referred to as connectedness, are

[29] We base our analysis on the generalized VAR framework of Koop et al. (1996) and Pesaran and Shin (1998), following Diebold and Yilmaz (2012, 2014), that are invariant to the ordering of the VAR system (as opposed to, e.g., Cholesky factorization). The identification does not orthogonalize shocks, leaving them correlated, relying on a large data-driven identification scheme accounting for their correlation.

[30] Rambachan and Shephard (2019) show that under appropriate assumptions, generalized impulse response functions has a direct causal interpretation, which, as per (7) renders the generalized variance decomposition causal. We do not check nor claim that the assumptions are appropriate in our context and prefer being less strict in the interpretation, yet remain positive on the possibility of directional causality.

Table 4: Schematic connectedness table of network

This table depicts the generalized variance decomposition schematic based on $\boldsymbol{V}^h$, including the directional from- and to-connectedness in the rows and columns, respectively. The lowest right element measures the total connectedness in the system. The block shaded with blue (orange) indicates the quantities relating the individual connectedness from narratives (macro-finance) to macro-finance (narratives). Numbers indicate row and column indices.

		Macro-finance			Narratives			
		1	$\cdots$	17	18	$\cdots$	30	
		RET	$\cdots$	IDT	CCF	$\cdots$	C19	From others
1	RET	$v^h_{RET,RET}$	$\cdots$	$v^h_{RET,IDT}$	$v^h_{RET,CCF}$	$\cdots$	$v^h_{RET,C19}$	$\sum_j^N v^h_{RET,j}/N, j \neq \text{RET}$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$
17	IDT	$v^h_{IDT,RET}$	$\cdots$	$v^h_{IDT,IDT}$	$v^h_{IDT,CCF}$	$\cdots$	$v^h_{IDT,C19}$	$\sum_j^N v^h_{IDT,j}/N, j \neq \text{IDT}$
18	CCF	$v^h_{CCF,RET}$	$\cdots$	$v^h_{CCF,IDT}$	$v^h_{CCF,CCF}$	$\cdots$	$v^h_{CCF,C19}$	$\sum_j^N v^h_{CCF,j}/N, j \neq \text{CCF}$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$
30	C19	$v^h_{C19,RET}$	$\cdots$	$v^h_{C19,IDT}$	$v^h_{C19,CCF}$	$\cdots$	$v^h_{C19,C19}$	$\sum_j^N v^h_{C19,j}/N, j \neq \text{C19}$
	To others	$\sum_i^N v^h_{i,RET}/N$	$\cdots$	$\sum_i^N v^h_{i,IDT}/N$	$\sum_i^N v^h_{i,CCF}/N$	$\cdots$	$\sum_i^N v^h_{i,C19}/N$	$\sum_{i,j=1}^N v^h_{i,j}/N^2$
		$i \neq \text{RET}$	$\cdots$	$i \neq \text{IDT}$	$i \neq \text{CCF}$	$\cdots$	$i \neq \text{C19}$	$i \neq j$

sophisticated networks where, for instance, the directed connectedness from j to i is the directed edge between the two nodes (i and j). We use this appealing intuition further below to represent and analyse our findings.[31]

For a fixed h, we may gather the variance decomposition matrix $\boldsymbol{V}^h$, collecting $v^h_{i,j}$ for all combinations of i and j. The elements are represented schematically in a connectedness table illustrated in Table 4. The table further adds a "To" and "From" row and column, respectively, which provide the average off-diagonal elements of $\boldsymbol{V}^h$ in either row or column direction. In general, one can sum across elements of $\boldsymbol{V}^h$ to construct any measure of connectedness of interest across, for example, certain groups (like macro-finance versus narratives). We elaborate on the various ways of aggregating connectedness in the technical Appendix C.

Since the connectedness measures are some nonlinear transformation of the VAR coefficient estimates, its asymptotic distribution cannot easily be obtained. This bears a resemblance to making inference on impulse response functions that can, similar to connectedness, be obtained via a moving-average representation of the associated VAR model, cf. Lütkepohl (2000). Inspired by this resemblance and the

[31] It is important to note that alternative frameworks that attempt to characterize connectedness directly from the fitted (sparse) VAR model, e.g., Bonaldi et al. (2015), are likely incomplete as they only allow for connectedness through cross-*lag* linkages and ignores the relationships from the disturbance covariance matrix. Our approach utilizes a sparse VAR model that incorporates those variables that are likely to be Granger causal and incorporate effects from the contemporaneous disturbance covariance matrix.

typical approach in that literature, we rely on bootstrapped confidence intervals on estimated connectedness measures for inference. The procedure is tailored to our context using a residuals-based parametric bootstrap which is outlined in Appendix D.

Table 5 presents results that tests whether narratives as a group are driving unexpected fluctuations in macro-financial variables and vice versa. We consider the horizons $h = 1, 2, 3, 4, 5, 6, 7, 14, 30$, spanning daily, intra-weekly, weekly, and monthly cumulative effects. For the narratives-to-macro-finance direction,

$$C^h_{n \to mf} = \sum_{\substack{i \in mf \\ j \in n}} \tilde{v}^h_{i,j} \Big/ \sum_{i \in mf} \sum_{j=1}^{N} \tilde{v}^h_{i,j} \tag{8}$$

measures the share of total unexpected fluctuations of the macro-finance group which is due to the narratives group. This amounts to summing over all elements in the blue shared area in Table 4 and dividing by the sum of the entire area between rows 1 to 17 and columns 1 to 30. The reverse direction, $C^h_{mf \to n}$, from macro-finance to narratives is defined analogously by summing over the orange shaded area and dividing by the sum over the area between rows 18 to 30 and columns 1 to 30. Since it is natural that a large part of shocks from i transmits to itself, we also report those quantities where we mask the diagonal of $\boldsymbol{V}^h$, i.e., remove "own" shares. To measure the net effect, i.e. the strength of either direction denoted by C^h_{n-mf}, we subtract the total connectedness from macro-finance to narratives (sum over orange shaded area) from total connectedness from narratives to macro-finance (sum over blue shaded area) and divide by the total connectedness (sum over all elements in $\boldsymbol{V}^h$). We also here report those quantities where we mask the diagonal of $\boldsymbol{V}^h$ to illustrate the net effect as a fraction of total non-own shares.

The transmission from narratives to macro-finance amounts to 12% at the daily horizon. That is, out of total unexpected fluctuations in the macro-finance group, about 12% is directly attributable to shocks to narratives. This amounts to 28% out of total received variation from other variables except for the variables themselves which is directly attributable to the narratives group. These numbers are statistically (at the 1% level) and economically significant. The effects cumulate over time, rendering the one-week effect 15%, two-week effect 17%, and monthly effect 20%. Again, all the effects are statistically (at the 1% level) and economically significant. The share of non-own unexpected fluctuations amounts to between 28–32%. Switching the direction, we observe a similar pattern. The transmission of shocks from the macro-finance group to narratives increases from 17% to 32% over the daily to the monthly horizon, with all being statistically significant at the

Table 5: Group-wise connectedness: narratives to and from macro-finance
This table reports the cumulative transmission of shocks (connectedness) from the narratives group ($\boldsymbol{X}_t^n$) to the macro-financial group ($\boldsymbol{X}_t^{mf}$), and vice versa, for the horizons $h = 1,2,3,4,5,6,7,14,30$. The numbers represent shares of total forecast error variance, cf. (8), or when the diagonal (own share) is masked. The latter is reported in parenthesis and is not indicated by statistical significance as it is identical to the original number, $C^h_{mf\to n}$ or $C^h_{n\to mf}$. Statistical significance is based on bootstrapped confidence intervals outlined in Appendix D: asterisks "***", "**", and "*" indicates significance on the 1% level, 5% level, and 10% level, respectively.

	Effect horizon in days (h)								
Direction of connectedness	1	2	3	4	5	6	7	14	30
Narratives to macro-finance: $C^h_{n\to mf}$	12.3***	13.1***	13.8***	13.9***	14.5***	14.7***	15.0***	17.1***	20.2***
(Share of non-own)	(28.4)	(29.3)	(29.4)	(28.1)	(28.1)	(27.8)	(27.7)	(29.1)	(31.9)
Macro-finance to narratives: $C^h_{mf\to n}$	17.2***	18.7***	19.9***	21.1***	22.3***	23.3***	24.2***	28.0***	31.9***
(Share of non-own)	(47.5)	(49.4)	(50.2)	(50.9)	(51.2)	(51.3)	(51.2)	(51.2)	(52.4)
Net effect: C^h_{n-mf}	-0.5	-0.7	-0.8	-1.3	-1.5	-1.8*	-2.0*	-2.4*	-2.4
(Share of non-own)	(-1.2)	(-1.7)	(-1.8)	(-2.8)	(-3.1)	(-3.6)	(-3.9)	(-4.3)	(-3.8)

1% level as well. The share of non-own connectedness is about 48–52%. The net effect, that is, the difference between what narratives drive of macro-finance and what macro-finance drive of narratives is economically small between −0.5% and −2.4%. They are statistically significant (at the 10% level) for about one to two week horizons, yet insignificant otherwise.

It stands very clear that narratives are connected to macro-finance as they drive their fluctuations and, at the same time, are also shaped by their shocks. In other words, narratives are firmly integrated within the real economy and financial markets and play a significant bi-directional role for economic fluctuations. The network graphs presented below illustratively underpin this conclusion further. These effects exist even at the daily horizon, yet almost double over a month. We include results in Appendix F using eight and eighteen topics/narratives from the LDA estimation as robustness. Qualitative conclusions are identical. For instance, the share of unexpected fluctuations of macro-financial variables driven by narratives are 19% and 22% for eight and eighteen topics, respectively, compared to 20% for our main results. On the shorter horizons, the shares are moderately smaller for eight topics and moderately larger for eighteen topics, though all are statistically (at 1% level) and economically significant at all horizons. Results for the net effect are also similar, with the macro-finance to narratives connection sometimes being statistically significant, yet the difference is economically moderate.

C.1. The narrative-enhanced macro-financial network

Our macro-finance-narrative system consists of $N = 30$ variables, thus the are potentially $N^2 = 900$ pairwise directed connections. We utilize that the connectedness measures have a direct mapping to a network representation and, thus, visualize the entire macro-finance-narrative system in network graphs. This presents the relationships between all variables at the most granular level possible. The network representation is based on the connectedness schematic in Table 4, which is reported in Appendix E with estimated values at the daily, weekly, and monthly horizon. We use values from these tables in our commenting below.

The visualization is conducted via the open-source Gephi software, as in Demirer et al. (2018), and we generally follow the authors' approach in constructing the graphs.[32] The networks are characterized by four devices; node size, node color, node location, and edge sizes (there are two edges per node as the network is directed). The nodes represent each of the variables in the network, and their size indicates their total to-connectedness. The larger the node, the more important the variable is for driving fluctuations elsewhere in the network. The edges are directed and measure the to- and from-connectedness for each pair of variables (nodes). Their thickness depends on the size of the connectedness. The node color or grouping is based on the community detection algorithm of Blondel et al. (2008) and assigns variables to groups where they are more densely connected than the rest of the network. This coloring is also used for the edges, indicating the direction of the connectedness. The node location is based on the ForceAtlas2 algorithm (Jacomy et al., 2014), which determines a so-called force-directed layout (Battista et al., 1999). The idea is that nodes are charged with a repulsive force that drives them apart and edges are the attractive force between the nodes that connect, much alike how similar poles of magnets repel each other (representing nodes) and springs attract each other (representing the edges between nodes). The equilibrium is a steady state where these forces balance. The resulting location is visually intuitive, as connected nodes will find themselves close in the network graph (and likely of similar color/group). Therefore, the denser the network, the more connected it is in total. We focus on the daily ($h = 1$), weekly ($h = 7$), and monthly ($h = 30$) horizons, generating Figures 7–9. Each figure contains all non-zero edges (left panel) and the same network that masks those of small magnitude to emphasize individual relations for interpretation. We highlight a number of important results that can be derived from the network graphs. First, looking across horizons, the network becomes denser.

[32]The software is available at https://gephi.org.

Figure 7: Network graph of macro-financial-narrative system: daily effects ($h = 1$)

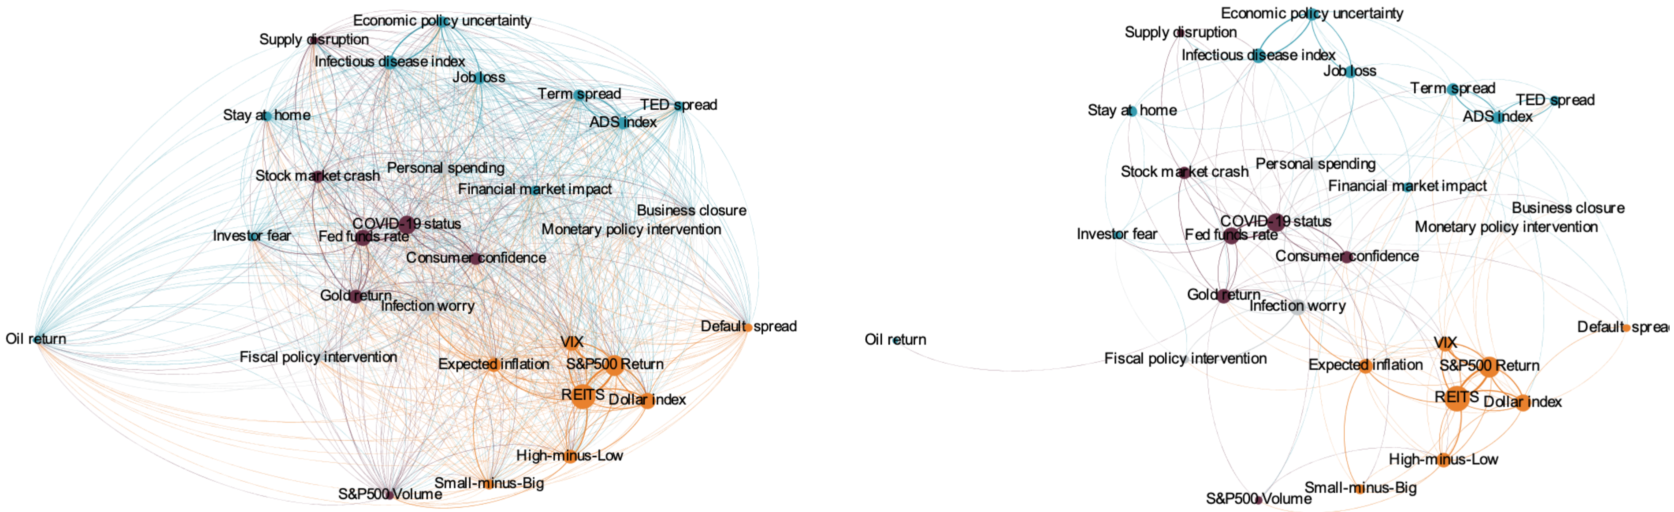

This figure depicts the estimated network graph associated with the macro-finance-narrative system comprised by $[\boldsymbol{X}_t^{mf}, \boldsymbol{X}_t^{n}]$. The horizon is one day ($h = 1$). Each node (indicated by a circle) represents one variable, and the size of the node depends on the size of to-connectedness in the network. The edges (indicated by lines) represent the directional connectedness among the variables, and their thickness depends on the size of the connectedness. The coloring of the nodes into groups is based on their connectedness so that nodes within the same color indicate larger interconnectedness. The distance among nodes is based on the ForceAtlas2 algorithm (Jacomy et al., 2014) as implemented in Gephi and represents a steady-state equilibrium where repelling and attracting forces balance. That is, nodes isolated (like oil returns) are generally less connected to the other variables in the network. The left plot shows all the network edges, whereas the right plot masks those of small magnitude to emphasize individual relations for interpretation.

Figure 8: Network graph of macro-financial-narrative system: weekly effects ($h = 7$)

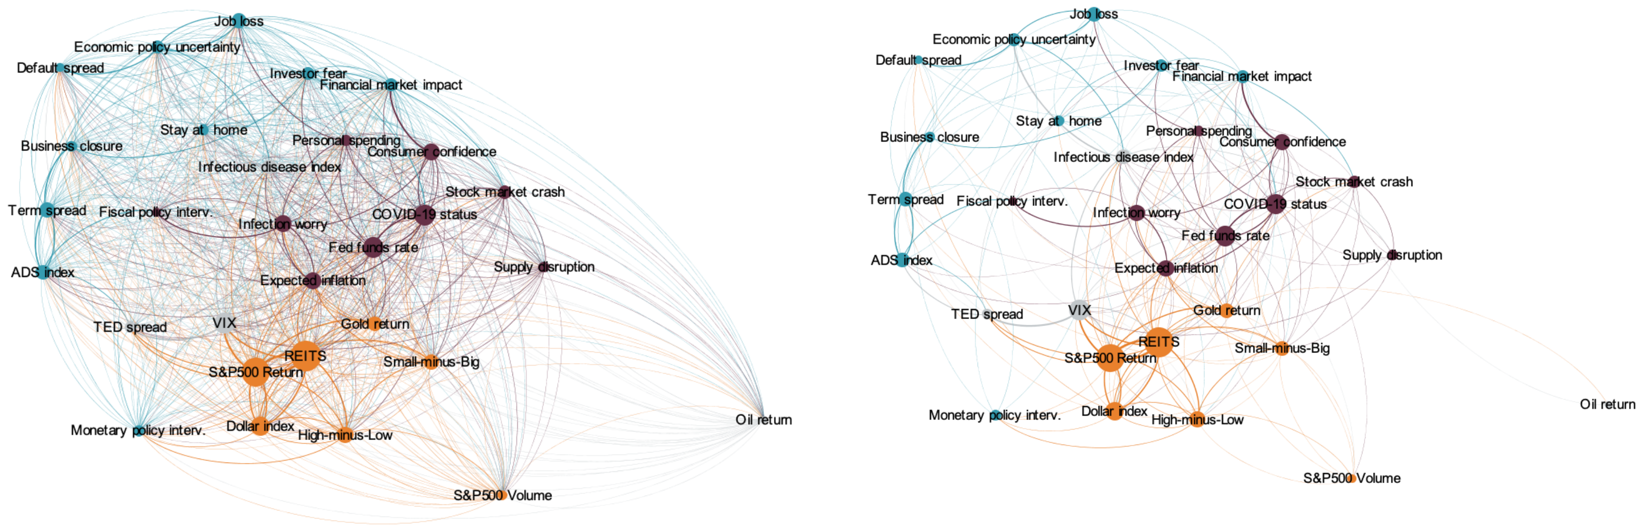

This figure depicts the estimated network graph associated with the macro-finance-narrative system comprised by $[\boldsymbol{X}_t^{mf}, \boldsymbol{X}_t^{n}]$. The horizon is one day ($h = 7$). Each node (indicated by a circle) represents one variable, and the size of the node depends on the size of to-connectedness in the network. The edges (indicated by lines) represent the directional connectedness among the variables, and their thickness depends on the size of the connectedness. The coloring of the nodes into groups is based on their connectedness so that nodes within the same color indicate larger interconnectedness. The distance among nodes is based on the ForceAtlas2 algorithm (Jacomy et al., 2014) as implemented in Gephi and represents a steady-state equilibrium where repelling and attracting forces balance. That is, nodes isolated (like oil returns) are generally less connected to the other variables in the network. The left plot shows all the network edges, whereas the right plot masks those of small magnitude to emphasize individual relations for interpretation.

Figure 9: Network graph of macro-financial-narrative system: monthly effects ($h = 30$)

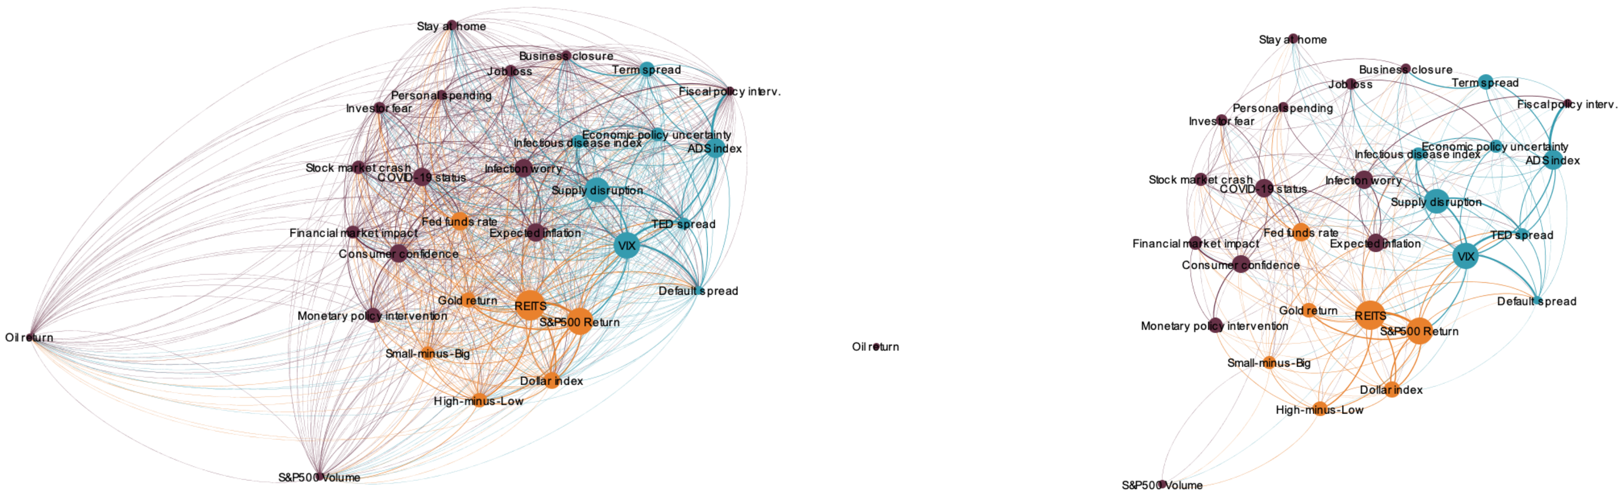

This figure depicts the estimated network graph associated with the macro-finance-narrative system comprised by $[\boldsymbol{X}_t^{mf}, \boldsymbol{X}_t^{n}]$. The horizon is one day ($h = 30$). Each node (indicated by a circle) represents one variable, and the size of the node depends on the size of to-connectedness in the network. The edges (indicated by lines) represent the directional connectedness among the variables, and their thickness depends on the size of the connectedness. The coloring of the nodes into groups is based on their connectedness so that nodes within the same color indicate larger interconnectedness. The distance among nodes is based on the ForceAtlas2 algorithm (Jacomy et al., 2014) as implemented in Gephi and represents a steady-state equilibrium where repelling and attracting forces balance. That is, nodes isolated (like oil returns) are generally less connected to the other variables in the network. The left plot shows all the network edges, whereas the right plot masks those of small magnitude to emphasize individual relations for interpretation.

This results from the evident connectedness cumulating over horizons so that the variables' fluctuations become to a larger extent dependent upon one another. The overall connectedness in the network is, nevertheless, statistically significant at all horizons, cf. Appendix E.

Second, the daily and weekly network graphs contain four communities or groups (colors) of variables, whereas the monthly network graphs contains three. While the overlapping groups roughly represent the same variables across horizons, some variables do switch groups. For instance, gold returns belong to the narratives-dominated purple group at the daily horizon but switch to the macro-finance-dominated orange group at the weekly and monthly horizons. This suggests gold returns are connected to narratives at the short horizons possibly via its role as a safe haven, yet turn to more macro-economic drivers for longer horizons. Two (purple and blue) out of the three dominating groups contain a mix of macro-financial variables and narratives, while the orange group mainly captures macro-financial variables. It is important to stress that it does not mean this group is isolated from the other variables in the network. It reflects the intuitive fact that those variables are, on average, more densely connected to one another than to the rest of the network. For example, the strongest link in the monthly network is from the ADS index (blue group) towards the narrative *fiscal policy intervention* (purple group), that is, a connectedness across groups. Another example is the bi-directional link between the Fed funds rate (orange group) and *COVID-19 status* (purple group). The purple group, which tend to be at the center of the network, is a mix of the Fed funds rate, expected inflation, and gold returns at different horizons and narratives on *infection worry*, *COVID-19 status*, *consumer confidence*, and *stock market crash*. The orange group, usually located to the bottom of the network, is mainly a macro-financial group consisting of the S&P500 returns, risk premia, the dollar index, and REITS. The blue group, located mainly at the top-right of the network, contains the ADS index, the default and term spreads, the news-based variables, and narratives like *supply disruption* (at $h = 30$), *job loss*, *stay at home*, *business closure*, and *investor fear*. The gray group, mainly located to the middle and right of the network, is small and exists only for the daily and weekly horizon with a varying composition. There are generally strong links between all of the groups.

Third, commenting on individual links (including across horizons), we observe a number of interesting patterns. The *supply disruption* narrative plays a significant role in the network at the monthly horizon, driving a sizeable 15.0% of the unexpected fluctuations of the VIX and 14.2% of the TED spread. The reverse direction is negligible, with the VIX and the TED spread driving 1.0% and 0.4%, respectively, of *supply disruption*. This is firm evidence of individuals' fear of supply shortage in connection with the world lockdown and the emergency orders enforced were driving a substantial part of the variations in the VIX and the TED spread; variables that are commonly referred to as reflecting investor fear, and that exhibited peaks at the Great Recession levels during March, 2020, the time at which the *supply disruption* shortage prevailed. Our findings strongly support this notion. These effects accumulate over horizons, as *supply disruption* moves from the periphery of the network to the center going from $h = 1$ to $h = 30$.

There is also a strong one-sided connection between the ADS index and the *fiscal policy intervention* narrative at the weekly and monthly horizon. A notable 20.9% of the unexpected variation in narratives on fiscal policy and stimulus packages are driven solely by unexpected changes to the ADS index at the monthly horizon. On the other hand, a mere 0.7% of unexpected fluctuations in the ADS is attributable to *fiscal policy intervention*. That is, individuals' stories on government actions as response to mitigate the economic impact of COVID-19 is largely inferred by their observations of the current economic environment.

The Fed funds rate is always at the center of the network. It is strongly shaped by the *COVID-19 status* narrative with 14.4% (10.3% at weekly horizon and 9.7% at the monthly horizon) of its unexpected fluctuations attributable to *COVID-19 status* shocks. The opposite direction is equally strong at 13.4% (11.7% at weekly horizon and 9.7% at monthly horizon), that is, unexpected changes in the Fed funds rate shapes how people talk about the societal impact of COVID-19. The effects decrease somewhat over horizons, which may partly be explained by the fact that the Fed funds rate only exhibits notable variation at the beginning of the sample in March, where the Federal Reserve lowered the rate twice, yet the stories and talk about monetary policy intervention did not stop. At the monthly horizon, *monetary policy intervention* drives 9.9% of the *consumer confidence* (and not the opposite direction, 1.1%) which in turn drives 13.3% of individuals' *financial market impact* narrative (and not the opposite direction, 1.5%).

We also observe that the news-based variables are generally less connected with macro-finance variables and narratives compared to the connectedness between macro-finance variables and narratives. Moreover, the main connection goes from narratives towards news as an indication that narratives influence reported news. This supports the equilibrium predictions of Mullainathan and Shleifer (2005) that news content is (partly) driven by consumer preferences. Oil returns are located far from the dense part of the network, suggesting that the extreme events with negative oil prices during this period largely were an isolated event having no significant impact on the remaining of the network. It is also evident that the network plot visualizes aspects from Figures A.2–A.3 in Appendix G above that show how oil returns are generally the weakest transmitter and receiver of shocks in the system.

Many more interesting connections exist. In the interest of space, we refrain from commenting further on those but urge the reader to explore the network graphs further in connection with the connectedness tables in Appendix E. On this basis, we firmly conclude that there are statistical and economically significant individual links between macro-financial variables and narratives and, judging by the visualization of the network, find clear evidence that narratives are firmly integrated within the real economy and financial markets. That is, there is significant quantitative evidence in favor of the narrative basis of economic fluctuations as hypothesized by Shiller (2017, 2019, 2020).

D. Asset pricing implications

In this section we estimate risk premia for each of the variables included in the network with an emphasis on the narratives. We use the recently developed three-pass methodology of Giglio and Xiu (2021) which recovers the risk premium for a given observable (traded or nontraded) factor, while accounting for the omission of relevant risk factors and measurement error. Robustness to omitted variables is crucial given the plethora of potential factors available (Harvey et al., 2016) and the associated uncertainty of the true asset pricing model. To understand the methodology, consider the case where excess returns $\boldsymbol{r}_t$, of dimension q, are governed by the linear factor model dynamics, see e.g. Ross (1976),

$$\boldsymbol{r}_t = \boldsymbol{\beta}'\boldsymbol{\gamma} + \boldsymbol{\beta}'\boldsymbol{f}_t + \boldsymbol{\varepsilon}_t, \tag{9}$$

where $\boldsymbol{\gamma}$ is a p-vector of risk premia for the unobservable fundamental factors, $\boldsymbol{f}_t$ is the p-vector of those mean-zero fundamental factor innovations, and $\boldsymbol{\beta}$ is the $p \times q$ matrix of fundamental factor exposures such that $\mathbb{E}[\boldsymbol{r}_t] = \boldsymbol{\beta}\boldsymbol{\gamma}$. The interest is in estimating the risk premium for a given observable factor g_t (e.g. a narrative or a macro-finance state variable), where the law of motion reads

$$g_t = \delta + \boldsymbol{\eta}'\boldsymbol{f}_t + \nu_t. \tag{10}$$

The first step of the three-pass methodology is to apply principal component analysis (PCA) to a large set of de-meaned test asset returns to estimate the unobservable fundamental factors $\hat{\boldsymbol{f}}_t$. The second pass estimates the risk premia of those fundamental factors by a standard two-pass Fama-MacBeth procedure; run a time series regression of r_{it} on $\hat{\boldsymbol{f}}_t$ for each i of the set of test assets to obtain $\hat{\boldsymbol{\beta}}_i$ and run a cross-sectional regressions of $\bar{\boldsymbol{r}}$, denoting mean returns, onto $\boldsymbol{\beta}$, which yields $\hat{\boldsymbol{\gamma}}$. The third and last pass runs a time series regression of the observable risk factor g_t onto the fundamental factors $\hat{\boldsymbol{f}}_t$ as per equation (10), which yields coefficients $\hat{\boldsymbol{\eta}}$. The product of the second and third pass estimates provides the estimate of the risk premia of the candidate risk factor via $\hat{\gamma}_g = \hat{\boldsymbol{\eta}}'\hat{\boldsymbol{\gamma}}$. If g_t is strongly related to the pervasive fundamental factors, $\boldsymbol{f}_t$, one obtains $\hat{\boldsymbol{\eta}} \neq 0$ and the risk premia is expected to be non-zero. A byproduct of the third pass regression is a measure of factor strength captured by the coefficient of determination, $\mathrm{R}^2_{\mathrm{g}}$. This measures how related the observable candidate risk factor is to the fundamental factors. A weak factor would have a value $\mathrm{R}^2_{\mathrm{g}}$ close to zero. Giglio and Xiu (2021) show that $\hat{\gamma}_g$ is a consistent estimator for the true risk premia and provides the relevant asymptotic distribution for inference.

We gather a large set of 198 (value-weighted) portfolios with daily observations as tests assets from Kenneth French's Data Library: 25 portfolios sorted on size and book-to-market ratio, 25 portfolios sorted on size and momentum, 25 portfolio sorted on size and investment, 25 portfolios sorted on size and operating profitability, 25 portfolios sorted on size and short-term reversal, 25 portfolios sorted on size and long-term reversal, and 48 portfolios sorted on industry classification.[33]

To ensure that the principal components in the first step adequately span the space of test assets and estimate the fundamental factors, we use the first 12 principal components. The choice is guided by a plot of the eigenvalues and the cross-sectional explanatory power of the test assets.[34] They explain a total of 98% of the time series variation of the test asset returns and a 56.5% of the cross-sectional variation in their average returns. This number compares to Giglio and Xiu (2021) who settle at a value of 59%. As candidate risk factors we use the innovations to each of the variables that enter the VAR system from Subsection B.1 in accordance with the ICAPM.[35] These innovations are normalized by their standard deviation for comparability. Table 6 reports their estimated risk premia, associated statistical significance, and the strength of each factor as measured by $\mathrm{R}^2_{\mathrm{g}}$.

We identify multiple significant risk premia, two of which are associated with the narratives. The negative risk premium on *supply disruption* suggests that an

[33]These constitute all portfolios available in the library, cf. http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

[34]Additional principal components have very small eigenvalues and contributes marginally to the explanation of the cross-sectional variation of expected test asset returns. Further details can be found in Appendix H.

[35]Results are qualitatively similar if innovations to an AR(1) or first-differences are used as state variables.

Table 6: Risk premia and factor strength

This table reports three-pass regression results for standardized innovations to state variables includes in the VAR system that is comprised by 17 macro-finance variables and the time series prevalence of 13 narratives. The risk premia estimates, $\hat{\gamma}_g$ are obtained via the methodology of Giglio and Xiu (2021), using 12 principal components extracted from de-meaned returns of a panel of 198 equity portfolios as test assets. The R^2_g measure the strength of each state variable from a regression onto the 12 principal components. Statistical inferences are based on t-statistics that use heteroskedasticity- and autocorrelation-robust Newey and West (1987) standard errors and the asymptotic theory in Giglio and Xiu (2021): asterisks "***", "**", and "*" indicates significance on the 1% level, 5% level, and 10% level, respectively.

Macro-finance			Narratives		
	$\hat{\gamma}_g$	R^2_g		$\hat{\gamma}_g$	R^2_g
ADS index	-3.94	41.38%	*Consumer confidence*	3.67	14.91%
Infectious disease index	6.02*	12.01%	*Personal spending*	-0.18	16.93%
Default spread	-12.24**	27.31%	*Stock market crash*	0.97	15.35%
Dollar index	9.92	24.89%	*Monetary policy interv.*	-8.44**	18.34%
Small-minus-Big	1.27	64.85%	*Supply disruption*	-7.02**	21.77%
High-minus-Low	-2.57	63.23%	*Business closure*	-1.87	13.25%
Fed funds rate	9.70*	36.68%	*Job loss*	-2.00	11.23%
Gold returns	-0.72	17.17%	*Infection worry*	4.49	15.44%
Inflation expectations	-7.76**	13.99%	*Financial market impact*	3.42	21.65%
Oil returns	4.30	8.28%	*Fiscal policy interv.*	3.00	17.62%
REITS	-9.71	49.70%	*Investor fear*	-0.12	12.71%
S&P500 returns	-12.09	60.55%	*Stay at home*	0.77	17.84%
S&P500 volume	-1.58	7.09%	*COVID-19 status*	0.61	14.56%
Economic policy uncertainty	11.32**	25.21%			
TED spread	3.49	34.21%			
Term spread	5.13	30.80%			
VIX	5.69	60.74%			

increase in the prevalence of the narrative is associated with an increase in marginal utility of consumption, that is, an increase in the stories on *supply disruption* are perceived as a bad state by investors. It also represents the strongest factor with the highest R^2_g among all narratives, further noting that it is dominant in the network analysis in the previous section. The *monetary policy intervention* narrative also carries a significant risk premium and relatively high R^2_g, suggesting that stories spreading on the interest rate cuts by the Federal Reserve are related to periods of high marginal utility states. This makes intuitive sense, since investors are more likely to talk about rate cuts when the economic situation deteriorates. The Fed funds rate is also statistically significant, yet with an opposite sign.[36] The opposite sign of the risk premia is explained by the fact that an increase in the prevalence of the *monetary policy intervention* narrative on rate cuts by the Fed is negatively correlated to actual rate cuts (i.e., decreases in the Fed funds rate). The sign of the Fed funds rate risk premium is also consistent with the findings of Maio and Santa-Clara (2012), who note that the risk-free rate predicts positive future returns. In the present case, the Fed funds rate is selected by the Elastic net estimator in the equation for S&P500 returns at a one-day lag with a positive coefficient, indicating that it predicts future returns positively. This is consistent with a positive risk premia in the ICAPM. Four additional macro-finance variables have risk premia that are statistically significant at conventional significance levels; news variables (infectious disease index and economic policy uncertainty), the default spread, and inflation expectations. These findings altogether support the role of narratives in the financial markets as they carry significant risk premia, and it underscores the conclusions from the previous network analysis.

V. Concluding remarks

In contrast to other social sciences, economics has been relatively reluctant to analyze the effects of popular narratives on its field of study. The advent of narrative economics and the integration of techniques to gather and analyze text to explain economic behavior into the discipline is set to change this. However, many questions remain to be answered. This paper is our attempt to shed light on some of the most important and general questions in the field, namely, to answer whether popular narratives have an effect on the economy and financial markets. We find strong quantitative support for the role of narrative economics. Since our approach does not impose any structure, we can answer these questions in a general sense, but we think our findings open the door to structural approaches that can narrow the specificity of the questions asked and establish a more transparent chain of causality.

This paper also shows how services such as Amazon's MTurk allow economists to survey investors directly using open-ended questionnaires. Although economists are not alien to using surveys, the tradition in the field has been to use categorical or ordinal questions that can easily be quantified. The incorporation of text mining techniques into the economics discipline, such as LDA, provides us with much larger flexibility to deal with complex text data and quantify it. We hope that the approach

[36] The fact that short-term interest rates carry a significant risk premium is supportive of Merton (1973), who suggest this as natural candidate state variable.

taken in this paper can open the door to future research that directly aims at obtaining popular narratives from respondents.

Although the COVID-19 pandemic is still ongoing and has had tremendous negative effects on millions of people's lives, it has also provided fertile ground to look at the impact of economic narratives. This context is interesting since, as noted by Cinelli et al. (2020), social media has had an enormous effect on the transmission of information and misinformation. Economists like to treat individuals as rational agents, but it is naïve to think that misinformation spreading through social media and the internet has no effect on economic behavior. We think that this particular field opens very interesting avenues of future research.

References

ABOUK, R. AND B. HEYDARI (2020): "The Immediate Effect of COVID-19 Policies on Social-Distancing Behavior in the United States," *Public Health Reports*, 0033354920976575.

ADÄMMER, P. AND R. A. SCHÜSSLER (2020): "Forecasting the equity premium: mind the news!" *Review of Finance*.

ANDERSEN, M. (2020): "Early evidence on social distancing in response to COVID-19 in the United States," *Available at SSRN 3569368*.

ANDREOU, E., E. GHYSELS, AND A. KOURTELLOS (2013): "Should macroeconomic forecasters use daily financial data and how?" *Journal of Business & Economic Statistics*, 31, 240–251.

ARUOBA, S. B., F. X. DIEBOLD, AND C. SCOTTI (2009): "Real-time measurement of business conditions," *Journal of Business & Economic Statistics*, 27, 417–427.

ASUNCION, A., M. WELLING, P. SMYTH, AND Y. W. TEH (2009): "On smoothing and inference for topic models," *Proceedings of the twenty-fth conference on uncertainty in articialintelligence*, 27–34.

BAKER, S. R., N. BLOOM, AND S. J. DAVIS (2016): "Measuring economic policy uncertainty," *The Quarterly Journal of Economics*, 131, 1593–1636.

BAKER, S. R., N. BLOOM, S. J. DAVIS, K. KOST, M. SAMMON, AND T. VIRATYOSIN (2020): "The unprecedented stock market reaction to COVID-19," *The Review of Asset Pricing Studies*, 10, 742–758.

BATTISTA, G. D., P. EADES, R. TAMASSIA, AND I. G. TOLLIS (1999): *Graph drawing: algorithms for the visualization of graphs*, Upper Saddle River: Prentice Hall.

BELLSTAM, G., S. BHAGAT, AND J. A. COOKSON (2020): "A text-based analysis of corporate innovation," *Management Science*, Forthcoming.

BERNANKE, B. S. AND J. BOIVIN (2003): "Monetary policy in a data-rich environment," *Journal of Monetary Economics*, 50, 525–546.

BHOOTRA, A. AND J. HUR (2013): "The timing of 52-week high price and momentum," *Journal of Banking & Finance*, 37, 3773–3782.

BIETTI, L. M., O. TILSTON, AND A. BANGERTER (2019): "Storytelling as adaptive collective sensemaking," *Topics in cognitive science*, 11, 710–732.

BLEI, D. (2012): "Probabilistic topic models," *Communications of the ACM*, 55, 77–84.

BLEI, D. M. AND J. D. LAFFERTY (2009): "Topic models," *Text mining: classification, clustering, and applications*, 10, 34.

BLEI, D. M., A. Y. NG, AND M. I. JORDAN (2003): "Latent dirichlet allocation," *Journal of Machine Learning Research*, 3, 993–1022.

BLONDEL, V. D., J.-L. GUILLAUME, R. LAMBIOTTE, AND E. LEFEBVRE (2008): "Fast unfolding of communities in large networks," *Journal of Statistical Mechanics: Theory and Experiment*, 2008, P10008.

BONALDI, P., A. HORTAÇSU, AND J. KASTL (2015): "An empirical analysis of funding costs spillovers in the euro-zone with application to systemic risk," Tech. rep., National Bureau of Economic Research.

BROCK, T. C. AND J. L. BALLOUN (1967): "Behavioral receptivity to dissonant information." *Journal of Personality and Social Psychology*, 6, 413.

BYBEE, L., B. T. KELLY, A. MANELA, AND D. XIU (2020): "The structure of economic news," Tech. rep., National Bureau of Economic Research.

CAMPBELL, J. Y. (1991): "A variance decomposition for stock returns," *The Economic Journal*, 101, 157–179.

CAMPBELL, J. Y. AND J. AMMER (1993): "What moves the stock and bond markets? A variance decomposition for long-term asset returns," *The Journal of Finance*, 48, 3–37.

CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (1999): "Monetary policy shocks: What have we learned and to what end?" *Handbook of Macroeconomics*, 1, 65–148.

CINELLI, M., W. QUATTROCIOCCHI, A. GALEAZZI, C. M. VALENSISE, E. BRUGNOLI, A. L. SCHMIDT, P. ZOLA, F. ZOLLO, AND A. SCALA (2020): "The covid-19 social media infodemic," *Nature Scientific Reports*.

COOKSON, J. A., J. ENGELBERG, AND W. MULLINS (2020): "Echo Chambers," *Available at SSRN*.

DEMIRER, M., F. X. DIEBOLD, L. LIU, AND K. YILMAZ (2018): "Estimating global bank network connectedness," *Journal of Applied Econometrics*, 33, 1–15.

DIEBOLD, F. X., G. D. RUDEBUSCH, AND S. B. ARUOBA (2006): "The macroeconomy and the yield curve: a dynamic latent factor approach," *Journal of Econometrics*, 131, 309–338.

DIEBOLD, F. X. AND K. YILMAZ (2009): "Measuring financial asset return and volatility spillovers, with application to global equity markets," *The Economic Journal*, 119, 158–171.

——— (2012): "Better to give than to receive: Predictive directional measurement of volatility spillovers," *International Journal of Forecasting*, 28, 57–66.

——— (2014): "On the network topology of variance decompositions: Measuring the connectedness of financial firms," *Journal of Econometrics*, 182, 119–134.

FAMA, E. F. AND K. R. FRENCH (1993): "Common risk factors in the returns on stocks and bonds," *Journal of Financial Economics*, 33, 3–56.

GENTZKOW, M., B. KELLY, AND M. TADDY (2019): "Text as data," *Journal of Economic Literature*, 57, 535–74.

GENTZKOW, M. AND J. M. SHAPIRO (2011): "Ideological segregation online and offline," *The Quarterly Journal of Economics*, 126, 1799–1839.

GIANNONE, D. AND L. REICHLIN (2006): "Does information help recovering structural shocks from past observations?" *Journal of the European Economic Association*, 4, 455–465.

GIGLIO, S. AND D. XIU (2021): "Asset pricing with omitted factors," *Journal of Political Economy*, Forthcoming.

GOETZMANN, W. N., D. KIM, AND R. J. SHILLER (2017): "Crash beliefs from investor surveys," Tech. rep., National Bureau of Economic Research.

GRIFFITHS, T. L. AND M. STEYVERS (2004): "Finding scientific topics," *Proceedings of the National Academy of Sciences*, 101, 5228–5235.

GRUBER, M. J. (1996): "Another puzzle: the growth in actively managed mutual funds," *The Journal of Finance*, 51, 783–810.

GUPTA, S., T. D. NGUYEN, F. L. ROJAS, S. RAMAN, B. LEE, A. BENTO, K. I. SIMON, AND C. WING (2020): "Tracking public and private responses to the COVID-19 epidemic: evidence from state and local government actions," Tech. rep., National Bureau of Economic Research.

HANSEN, S., M. MCMAHON, AND A. PRAT (2018): "Transparency and deliberation within the FOMC: a computational linguistics approach," *The Quarterly Journal of Economics*, 133, 801–870.

HARVEY, C. R., Y. LIU, AND H. ZHU (2016): "... and the cross-section of expected returns," *The Review of Financial Studies*, 29, 5–68.

HASSAN, T. A., S. HOLLANDER, L. VAN LENT, AND A. TAHOUN (2019): "Firm-level political risk: Measurement and effects," *The Quarterly Journal of Economics*, 134, 2135–2202.

HOBERG, G. AND C. LEWIS (2017): "Do fraudulent firms produce abnormal disclosure?" *Journal of Corporate Finance*, 43, 58–85.

HORNIK, K. AND B. GRÜN (2011): "topicmodels: An R package for fitting topic models," *Journal of Statistical Software*, 40, 1–30.

HUANG, A. H., R. LEHAVY, A. Y. ZANG, AND R. ZHENG (2018): "Analyst information discovery and interpretation roles: A topic modeling approach," *Management Science*, 64, 2833–2855.

JACOMY, M., T. VENTURINI, S. HEYMANN, AND M. BASTIAN (2014): "ForceAtlas2, a continuous graph layout algorithm for handy network visualization designed for the Gephi software," *PlOS ONE*, 9, e98679.

KOGAN, S., T. J. MOSKOWITZ, AND M. NIESSNER (2019): "Fake news: Evidence from financial markets," *Available at SSRN 3237763*.

KOOP, G., M. H. PESARAN, AND S. M. POTTER (1996): "Impulse response analysis in nonlinear multivariate models," *Journal of Econometrics*, 74, 119–147.

LARSEN, V. H. AND L. A. THORSRUD (2019): "The value of news for economic developments," *Journal of Econometrics*, 210, 203–218.

LOWRY, M., R. MICHAELY, AND E. VOLKOVA (2020): "Information revealed through the regulatory process: Interactions between the SEC and companies ahead of their IPO," *The Review of Financial Studies*, 33, 5510–5554.

LOWRY, P. B., J. D'ARCY, B. HAMMER, AND G. D. MOODY (2016): ""Cargo cult" science in traditional organization and information systems survey research: A case for using nontraditional methods of data collection, including Mechanical Turk and online panels," *The Journal of Strategic Information Systems*, 25, 232–240.

LÜTKEPOHL, H. (2000): "Bootstrapping impulse responses in VAR analyses," in *COMPSTAT*, Springer, 109–119.

MAIO, P. AND P. SANTA-CLARA (2012): "Multifactor models and their consistency with the ICAPM," *Journal of Financial Economics*, 106, 586–613.

MALONEY, W. AND T. TASKIN (2020): *Determinants of social distancing and economic activity during COVID-19: A global view*, The World Bank.

MERTON, R. C. (1973): "An intertemporal capital asset pricing model," *Econometrica*, 867–887.

MULLAINATHAN, S. AND A. SHLEIFER (2005): "The market for news," *American Economic Review*, 95, 1031–1053.

MURDOCK, B. B. J. (1962): "The serial position effect of free recall." *Journal of Experimental Psychology*, 64, 482.

MURTAGH, F. AND P. LEGENDRE (2014): "Ward's hierarchical agglomerative clustering method: which algorithms implement Ward's criterion?" *Journal of Classification*, 31, 274–295.

NEWEY, W. K. AND K. D. WEST (1987): "A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix." *Econometrica*, 55, 703–708.

NOFSINGER, J. R. AND A. VARMA (2013): "Availability, recency, and sophistication in the repurchasing behavior of retail investors," *Journal of Banking & Finance*, 37, 2572–2585.

PAPARODITIS, E. (1996): "Bootstrapping autoregressive and moving average parameter estimates of infinite order vector autoregressive processes," *Journal of Multivariate Analysis*, 57, 277–296.

PESARAN, H. AND Y. SHIN (1998): "Generalized impulse response analysis in linear multivariate models," *Economics Letters*, 58, 17–29.

PORTER, M. F. (1980): "An algorithm for suffix stripping," *Program*, 14, 130–137.

RAMBACHAN, A. AND N. SHEPHARD (2019): "Econometric analysis of potential outcomes time series: instruments, shocks, linearity and the causal response function," *arXiv preprint arXiv:1903.01637*.

ROSS, S. A. (1976): “The Arbitrage Theory of capital asset pricing,” *Journal of Economic Theory*, 13, 341–360.

SCHANK, R. C. AND R. P. ABELSON (1977): *Scripts, plans, goals and understanding, an inquiry into human knowledge*, Hillsdale NJ: Lawrence Erlbaum Associates.

SHILLER, R. J. (2017): “Narrative economics,” *American Economic Review*, 107, 967–1004.

——— (2019): *Narrative economics: How stories go viral and drive major economic events*, Princeton University Press.

——— (2020): “Popular economic narratives advancing the longest US expansion 2009-2019,” *Journal of Policy Modeling*.

SMITH, D., P. SCHLAEPFER, K. MAJOR, M. DYBLE, A. E. PAGE, J. THOMPSON, N. CHAUDHARY, G. D. SALALI, R. MACE, L. ASTETE, ET AL. (2017): “Cooperation and the evolution of hunter-gatherer storytelling,” *Nature communications*, 8, 1–9.

TUBBS, R. M., W. F. J. MESSIER, AND W. R. KNECHEL (1990): “Recency effects in the auditor’s belief-revision process,” *Accounting Review*, 452–460.

ZAK, P. J. (2015): “Why inspiring stories make us react: The neuroscience of narrative,” in *Cerebrum: the Dana forum on brain science*, Dana Foundation, vol. 2015.

ZOU, H. AND T. HASTIE (2005): “Regularization and variable selection via the elastic net,” *Journal of the Royal Statistical Society: Series B*, 67, 301–320.

Supplementary Appendix for

Tell me a story: Quantifying economic narratives and their role during COVID-19

A. Data cleaning and topic modelling

This section outlines our data processing steps and presents the LDA model representation and its estimation. The exposition is inspired by Blei (2012), Hansen et al. (2018), and Bybee et al. (2020).

A.1. Preparing the data set

The initial data set consists of 2,076 survey responses over the period February 29 to June 26. We conduct the following data processing steps:

1. Remove invalid responses based on manual reading, e.g. those that clearly misunderstood or did not reply to the questions (41 in total).

2. Remove responses with less than five words (93 in total).

3. Remove duplicated responses (128 in total).

4. Remove days with no responses (four in total)

This leaves 1,815 responses with a total of 79,922 words. We then conduct the following preparation steps:

1. Set non-alphabetical characters to an empty string and set all terms to lowercase letters.

2. Tokenize each of the responses, that is, split the responses into individual words.

3. Remove stopwords based on the snowball stopword list used in Hansen et al. (2018), cf. `http://snowball.tartarus.org/algorithms/english/stop.txt`. We also remove those stopwords that appear in lexicons SMART and onix, cf. `https://www.lextek.com/manuals/onix/stopwords2.html` and `http://www.lextek.com/manuals/onix/stopwords2.html`, respectively.

4. Convert all terms to their linguistic roots using the Porter algorithm for suffix stripping (Porter, 1980) which is the standard stemming tool for English language text (Gentzkow et al., 2019).

5. Remove words with less than three letter.

6. From the resulting unigrams, generate bigrams from all pairs of unigrams.

This leaves a data set of 21,858 unique terms. Finally, we follow the recommendations by Blei and Lafferty (2009) and Gentzkow et al. (2019) and rank those terms using the Term Frequency-Inverse Document Frequency (TF-IDF) metric. The objective is to remove those terms that are rare and very common. Very common words often ental stopwords, that has already been removed, yet also includes conjunctions like "and" or forms of the verb "to be", for instance. Those words are important for the grammatical structure of the sentences, yet convey little meaning on their own. For instance, the frequency of the word "the" carry no semantic information. Excluding very rare terms remove unnecessary "outliers" from the corpus of text and facilitate easier identification of common, important themes via LDA estimation. A useful approach is the TF-IDF. It is defined for the vth term

$$\text{TF-IDF}_v = \text{TF}_v \cdot \text{IDF}_v, \tag{A.1}$$

with

$$\text{TF}_v = 1 + \log(w_v), \tag{A.2}$$

where w_v is the total count of word v in the corpus of text, and

$$\text{IDF}_v = \log\left(\frac{S}{S_v}\right), \tag{A.3}$$

where S is the total number of survey responses, and S_v is the number of survey responses containing the vth term. As such, very rare words will have low TF-IDF scores because TF_v is low, whereas very common words that appear in most or all documents will have low TF-IDF scores because IDF_v is low. It is important to note that this approach is desirable compared to just removing words that occur frequently, because TF-IDF keeps in those that occur frequently in some responses, yet not so much in others. Sorting all terms based on their TF-IDF, one has to decide upon a threshold or cutoff rank below which terms are eliminated. Figure A.1 depicts the resulting TF-IDF weights and ranks for all 21,858 terms, as well as our choice of threshold. We decide upon a threshold at rank 5,144, dropping all terms with this rank or lower. This choice is at the point just before last large plateau at which terms have identical TF-IDF weights and their ranking is random. This approach is guided by Hansen et al. (2018). Since removal of stopwords and these low TF-IDF ranked terms renders two responses empty of terms, we end up with a total of 1812 survey responses.

Finally, we represent the text corpus in a "bag of words". Let $\boldsymbol{W}$ be a $S \times V$ matrix that

Figure A.1: Ranking of terms with TF-IDF and threshold

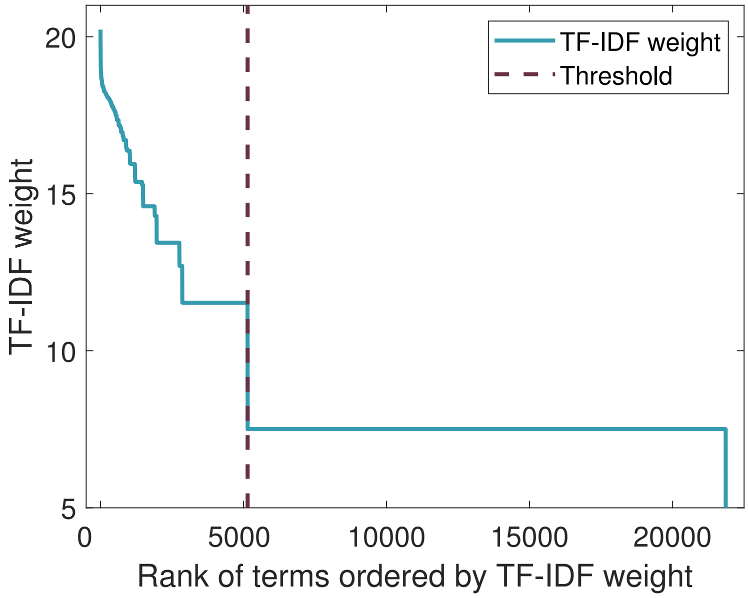

This figure depicts the TF-IDF weight for each of the 21,858 terms (blue line). The terms are ranked according to their TF-IDF weight. The vertical dashed purple line indicates the threshold or cutoff for which we eliminate terms ranked lower.

captures this representation of the survey responses, where row indices correspond to the list of different survey responses and column indices to the vocabulary of V many different unique terms in the total text corpus. The $w_{s,v}$ element of $\boldsymbol{W}$ represents the number of times term v appears in survey response s. Naturally, $\boldsymbol{W}$ is high-dimensional – in our case, it is $1,812 \times 5,144$.

A.2. Model representation

LDA is a generative probabilistic Bayesian factor model for discrete data. It assumes that the V-dimensional term counts vector for each survey response, $\boldsymbol{w}_s$ follows the multinomial distribution

$$\boldsymbol{w}_s \sim M(\boldsymbol{\Phi}\boldsymbol{\theta}_s, C_s), \tag{A.4}$$

where $M(\cdot,\cdot)$ denotes the multinomial distribution. Here C_s is the total number of terms in article s, governing the number of trials in $M(\cdot,\cdot)$, $\boldsymbol{\Phi} = [\boldsymbol{\phi}_1, \ldots, \boldsymbol{\phi}_{N_n}]'$ are the text corpus-wide topics and $\boldsymbol{\theta}_s$ is the survey response-specific allocation towards each of the topics. Specifically, the n'th topic, $\boldsymbol{\phi}_n$, is a probability distribution over terms such that $\phi_{n,v} \geq 0$ for all v and $\sum_v \phi_{n,v} = 1$. The value of N_n, i.e. the number of topics, is taken to be much smaller than V, the vocabulary, to enforce dimension reduction and facilitate our interpretation of the prevailing narratives in the corpus of survey responses. Moreover, $\boldsymbol{\theta}_s = [\theta_{s,1},\dots,\theta_{s,N_n}]'$ is also a probability vector, which specifies how the topics are prevailing within the s'th survey response. In that way, LDA views each survey response as a mixture of topics with $\theta_{s,n}$ representing the prevalence of topic n in survey response s.

A.3. Estimation

Given the probability density function of the multinomial distribution used in (A.4), several approaches to estimation exists including maximum likelihood estimation and Bayesian approaches. Asuncion et al. (2009) conclude that there is no substantial differences in their ability to obtain a proper empirical fit, yet maximum likelihood becomes computationally infeasible for high-dimensional text corpus. For these reasons, and following Hansen et al. (2018) and Bybee et al. (2020), we resort to the popular Bayesian approach of Griffiths and Steyvers (2004) using the Gibbs sampler for estimation. The estimation relies on the generative nature of the LDA model, that is, a set of sampling rules that simulates the writing of the text corpus. First, for each survey, it draws randomly from the list of topics with the sampling probabilities determined by $\boldsymbol{\theta}_s$ using a multinomial distribution. Suppose it draws the n'th topic. Second, to draw the first term of the survey, it then draws randomly from $\boldsymbol{\phi}_n$. The process is repeated as many times as there are terms in the survey, that is, C_s many times. This procedure is conducted for each of the survey responses in the corpus, and the Gibbs sampling estimation obtains those $\boldsymbol{\Phi}$ and $\boldsymbol{\theta}_s$ that best simulates the writing of our entire text corpus, with the dimensionality reduction and interpretation goals of LDA in mind. The generative procedure is initiated with a Dirichlet priors for $\boldsymbol{\phi}_n$ and $\boldsymbol{\theta}_s$, i.e. $\boldsymbol{\phi}_n \sim \text{Dir}(\beta)$ and $\boldsymbol{\theta}_s \sim \text{Dir}(\alpha)$, respectively. We use a high number of Gibbs samplers, specifically 100,000, to ensure proper convergence over the entire corpus. Our implementation is done in R using the *topicmodels* package (Hornik and Grün, 2011), and hyperparameters are set to $\beta = 0.1$ and $\alpha = 50/N_n$ like Griffiths and Steyvers (2004) and Hansen et al. (2018), which is also the default values in the *topicmodels* package. Our program is available upon request.

B. Narrative elaboration and associated survey response

In this appendix, we present a short description of the labeled narratives extracted from the survey data using the LDA methodology described in Section III. Although the labels are necessarily subjective, we show that they are readily intuitive. For each narrative, we also show the survey response with the highest loading on that particular narrative. These examples should not be taken as a prototypical example of a narrative, but simply as an illustration of an example response. We present the raw responses before cleaning.

Consumer confidence (CCF): This narrative captures the degree of optimism and pessimism that consumers feel about the overall state of the economy and their financial situation as a result of the COVID-19 pandemic. It includes key terms like *confidence* and *consumer confidence* that relate directly to its name, but also terms like *lack* and *reduced* which reflect the negative effect of the pandemic on consumer (and investor) confidence.

> "*It has had a negative effect because it has driven down demand and consumer confidence. When demand is undercut, it does not matter how much supply you have it will have an impact on confidence. It has also caused volatility in the markets, because we are unsure when people can return to work, increasing demand back to pre-pandemic levels. The government injecting liquidity into the market to shore up the market against volatility has created temporary stability. However, this stability will evaporate if we have a second wave in the pandemic or if there is no real national response to help alleviate people of the pandemic.*"

Personal spending (PSP): This narrative captures the sharp decrease in spending as a result of job losses, government imposed restrictions and increased uncertainty. The narrative also includes *stimulus check* as part of its key terms, hence, federal government support and its effects on spending are also part of the narrative.

> "*Businesses aren't open so people aren't supporting them. People are out of work and don't have money to spend. This is slowly turning around came as stimulus checks and unemployment checks came and people have money to spend.* "

Stock market crash (SMC): This narrative focuses on the negative effect on financial markets from the pandemic and has a particular focus on the crash that started on February 20 and resulted in a 34% drop in US stock market in less than a month. Since our survey collection started on February 29, the equity market crash plays a large role. Key terms include *stock*, *drop* and *market crash*.

> *"When the US economy shuts down, the supply of dollars to the rest of the world seizes up. This is further being exacerbated by the crash in oil prices caused by the Russia and KSA OPEC dispute, the lower the price of oil: the less dollars flow out. The effects of this dollar shortage in the offshore and eurodollar markets are causing the dollar to rise. This can be seen in the crash in commodity prices and the unwinding of the yen carry trade along with the rapid depreciation of emerging market currencies compared to the dollar. As the dollar rises, dollar-priced assets fall. The rapid drop in securities and commodities that have been driven up by leverage is causing a cascade of margin calls that have forced selling at the market bid price, further lowering prices and causing more forced selling and hence a crash in the stock market. To avoid margin calls, funds are selling whatever they can to put dollars into their accounts, hoarding even more dollars that have nowhere to go."*

Monetary policy intervention (MPI): This narrative relates to policy interventions by the Fed. Although similar to the FPI narrative, one key difference is that this narrative includes *federal*, *reserve*, *rate* and *cut* as a key terms, implying a focus on monetary policy interventions, i.e. rate cuts by the Fed. There were two cuts to the Fed Funds rate announced by the Fed in unscheduled meetings during our sample period. The first one on March 3, 2020 lowered the rate by 50 basis points while the second one on March 15, 2020 lowered it further by 100 basis points, effectively bringing the rate to zero.

> *"No one is working. Unemployment is at record high. The world is shutdown. The stimulus packages were passed and the Federal reserve has lowered interest rates by a lot, this has given the investors more confidence in the market."*

Supply disruption (SPL): This narrative focuses on the effect of the pandemic on Chinese manufacturing and supply chains. This focus is clearly shown by the following key terms *industrial*, *China*, *supply chain*, *manufacturing*. This narrative also captures restrictions on travel which is shown by key terms such as *travel* and *airlines*.

> *"I believe the Coronavirus is affecting the markets because of how much we trade with China, companies import a lot of goods from China. If there are factories being shut down or less material available, the prices will go up, thus leading to less sales because the demand does not instantly change. Ultimately leading to decreased revenue for the corporations."*

Business closure (BUC): This narrative relates to the business closure orders imposed as part of the state of emergency proclamations during the pandemic. Most of these orders were at the state level between March 22 and April 7 2020, depending on the state. The focus on business closure is evident from key terms like *business close*, *business shut* and *close*.

> "*With the ongoing spread of COVID-19, and most if not all business are being shut down for the time being. A lot of small and family-owned businesses will not be able to survive being shut down as long as it will be required during this outbreak. With a lot of businesses going under, that hurts the overall economy a lot.*"

Job loss (JBL): This narrative is related to the economic fallout from the pandemic as people lost their jobs. The unemployment rate peaked at 14.7% in April 2020, from a historically low level of 3.5% two months before. Thus, one would naturally expect that a narrative about job losses would be a dominant topic during the sample period.

> "*Fear of how badly job loss and company loss of income will affect profitability of the market and lower stock prices. Job losses affect people's ability to invest. Companies will be unable to see their products. Service industries have less customers.*"

Infection worry (INW): This narrative relates to the general anxiety of the population about catching the virus and to preventive ways to avoid infection. Key terms include words that are semantically related to disease sich as *sick* or *catch*, but also to preventing measures such as *leave*, *house* and *afraid*.

> "*I believe that it is due to fear. People are afraid of the unknown and there is so much we don't know about this virus at yet as it is so new. It also can be deadly which frightens people. When people are afraid, this is reflected in the stock market because they seek safe havens and don't want to invest in stocks which are seem as risky. Another reason is uncertainty. It is a black swan event. No one knows what is going to happen here although it does seem that conditions will worsen. People are trying to keep safe and not get sick by avoiding crowds and situations like travelling to minimize their risk of catching the virus. They are worried and don't spend as much which is not good for the financial markets. A third reason is there seem to be conflicting information coming from our leaders. President Trump is downplaying the risks of this virus while his administration are saying different things that sometimes sound contradictory. This confuses us and makes us resistant to investing in the markets. A fourth reason is we don't trust President Trump and some of our political leaders. This mistrust makes us feel less confident and more financially insecure which causes us to pull back from spending freely. This negatively impacts the markets.*"

Financial market impact (FMI): This narrative relates to the general volatility in stock markets during the sample period. It includes both negatively loaded terms like *negative effect* and *coronavirus negative* and positive terms like *experienced recovery*. This is not surprising given that the US stock market experienced its fastest fall since the fall of 1929 in March 2020, followed by a quick recovery that started the following month.

> *"The spread of the corona virus has had a negative effect on the financial markets because of how many people are out of work. Less people working means the economy is going to suffer tremendously and this is why we've seen the financial markets drop significantly. The hope of a new vaccine by the Massachusetts based biotech company called Moderna is why the financial markets have rebounded since march 23rd. The hope that their vaccine will cure corona virus is what this hope is hinging upon."*

Fiscal policy intervention (FPI): This narrative relates to policy interventions by the federal government as a reaction to the pandemic. The US government passed two large aid packages during our sample period. The Coronavirus Aid, Relief, and Economic Security Act (CARES) and the Health and Economic Recovery Omnibus Emergency Solutions Act (HEROES), passed on March 27 and May 15, respectively. These packages included direct economic assistance for American workers, families and small business. Key terms include *package*, *bill* and *pay*.

> *"Too many people were scared, so they stopped purchasing and stopped investing. People have started purchasing again. The stimulus money that people received have given people confidence and businesses are reopening. People felt that the efforts to minimize the spread of the disease have worked and are getting back out into the economy again. "*

Investor fear (IVF): This narrative captures the general anxiety and fear from investors that result from uncertainty about the ultimate outcome and course of the pandemic. It includes key terms like *fear*, *investor* and *market return*.

> *"The biggest factor for the effects on the financial markets, in my opinion, is just fear of the unknown. People who actually invest are unsure of what is actually going to happen in the near future, so they pull/sell their investments, even taking a large loss, so they can have immediate cash available. The government, local and federal, has basically brought everything to a halt by shutting everything down, which just gives more reason for fear to take over, and more people panic sell. When people get scared, they no longer pay attention to logic. Most people alive in America today have never been through, or experienced, an actual pandemic or large catastrophe. I think this also adds to the fear factor. The market*

has also been up for so long, this gives people more inclination to sell as they are still ahead on their investments. A lot of people are also nearing retirement age and need to conserve their money, so adding to the fear of the unknown, I am sure a lot of people sold investments to save their capital. People are starting to realize that the world is actually not ending and getting back into the market. People, like myself, start seeing great deals on stock and decide to invest large amounts of money as the returns are much greater at these stock prices. Things are slowly starting to open and attempting to return to normal again. I am also sure there is a lot of FOMO (fear of missing out) in the stock markets."

Stay at home (SAH): This narrative is related to the stay at home orders imposed by the majority of US states in March and April. The first state in the main territory to impose it was California (March 19 2020) while the last state to impose it was South Carolina (April 7 2020). Regulations across states varied, but most of the orders ordered closures of schools, daycare centers, bars and restaurants and non-essential retail.

"People are unable to work due to stay at home orders. The stay at home orders prevent people from making money. Not having money leads to businesses getting hurt because people can't purchase anything."

COVID-19 status (C19): This narrative is related to the status of the COVID-19 pandemic, its coverage by the media and the possibility of a resolution due to a vaccine. Key terms are relatively diverse. For example some relate directly to the corona virus: *corona*, *corona virus* while other relate to media coverage of the pandemic: *media*, *news* and yet others relate to a vaccine and the possibility of a resolution to the pandemic: *vaccine* and *hope*.

"The main reason that the corona virus has had a negative impact on the financial markets is panic. The media, politicians, and people in general have incited panic. This has led to panic selling, panic buying, panic reactions. This is why we're seeing the financial markets tank and represent a bear market. The stock market has experienced a large recovery since March 23rd is hope. The hope being that a vaccine will be discovered soon, false narratives creating false hope from the news media and politicians. When the reality of a vaccine being a year plus away is realized by the masses, the financial markets will tank once again and we will face one of the worst recessions that we've ever seen as a race."

C. On connectedness measurement

A pairwise directional connectedness from i to j and from i to j are defined as

$$C^h_{i\to j} = v^h_{j,i} \quad \text{and} \quad C^h_{i\leftarrow j} = v^h_{i,j}, \tag{C.5}$$

respectively, noting again that, in general, $C^h_{i\to j} \neq C^h_{i\leftarrow j}$. Our interest is in those pairwise connectedness to and from ("→" and "←", respectively) macro-financial variables and narratives. The definition of *net* pairwise directional connectedness naturally becomes

$$C^h_{i,j} = C^h_{i\to j} - C^h_{i\leftarrow j} = v^h_{j,i} - v^h_{i,j} \tag{C.6}$$

as the difference between the share of shocks transmitted from variable i to j and how much variable i receives from variable j. This allows for quantification of the strength of the bi-directional links between two variables, informing about which direction dominates. If $C^h_{i,j} > 0$ variable i dominates and vice versa.

Extending these notions further, the average of the off-diagonal row elements defined as the "From others" column in the schematic table, captures the amount of connectedness of variable i due to all the other variables in the network. For instance, the first row of from-connectedness measures the fraction of forecast error variances of S&P500 returns due to shocks in all other variables. Analogously, the "To others" column captures the average off-diagonal elements in each column. Continuing the example, this measures the average fraction of forecast error variance of all other variables in the network due to the transmission of shocks to S&P500 returns. That is, from-connectedness measures how much variable i receives from other variables in the system and to-connectedness how much it transmits to the other variables. We define variable i's total directional connectedness to others as

$$C^h_{i\to\bullet} = \frac{1}{N}\sum_{j=1, j\neq i}^{N} v^h_{j,i} \tag{C.7}$$

and its total directional connectedness from others as

$$C^h_{i\leftarrow\bullet} = \frac{1}{N}\sum_{i=1, i\neq j}^{N} v^h_{i,j}. \tag{C.8}$$

The difference is referred to as net total directional connectedness, given by $C^h_i = C^h_{i\to\bullet} - C^h_{i\leftarrow\bullet}$ with a similar intuition as the previous variable-specific net connectedness measure. Naturally, we may also be interested in the average fraction to and from narratives vis-á-vis macro-financial variables. This is achieved by averaging over the last N_n rows (from direction) or columns (to direction), referred to as total *narrative* directional connectedness. Moreover, a net total narrative direction connectedness is obtained as their difference, measuring in total how much of the fluctuations in macro-financial variables is due to shocks in narratives versus how much of the fluctuations in narratives is due to shocks in macro-financial variables.

Lastly, the average across all off-diagonal elements in $\boldsymbol{V}^h$ measures total connectedness in the network as per

$$C^h = \frac{1}{N^2} \sum_{i,j=1, i \neq j} v_{i,j}^h. \tag{C.9}$$

Averaging over the off-diagonal elements in the $N_n \times N_n$ ($N_{mf} \times N_{mf}$) lower right (upper left) block of $\boldsymbol{V}^h$ provides a total connectedness measure among the narratives (macro-financial variables) themselves.

C.1. Estimates of connectedness

The p-lag VAR, VAR(p),

$$\boldsymbol{X}_t = \sum_{i=1}^{\cdot} ..inkp \boldsymbol{\Psi}_i \boldsymbol{X}_{t-i} + \boldsymbol{\varepsilon}_t, \tag{C.10}$$

$\boldsymbol{\varepsilon}_t \sim (\mathbf{0}, \boldsymbol{\Omega})$, obeys the moving-average representation $\boldsymbol{X}_t = \sum_{i=0}^{\infty} \boldsymbol{\Theta}_i \boldsymbol{\varepsilon}_t$ where $\boldsymbol{\Theta}_i$ are $N \times N$ coefficient matrices that satisfy the recursion $\boldsymbol{\Theta}_i = \boldsymbol{\Psi}_1 \boldsymbol{\Theta}_{i-1} + \boldsymbol{\Psi}_2 \boldsymbol{\Theta}_{i-2} + \cdots + \boldsymbol{\Psi}_p \boldsymbol{\Theta}_{i-p}$ where $\boldsymbol{\Psi}_0$ is the $N \times N$ identity matrix, and $\boldsymbol{\Psi}_i$ is filled with zeros for $i < 0$. In general, the computation of variance decompositions requires orthogonal innovations, whereas those from the VAR system are generally not. A solution to this conundrum is via identification schemes. We base our analysis on the generalized VAR framework of Koop et al. (1996) and Pesaran and Shin (1998), following Diebold and Yilmaz (2012, 2014), that are invariant to the ordering of the VAR system (as opposed to, e.g., Cholesky factorization). The identification does not orthogonalize shocks, leaving them correlated, relying on a large data-driven identification scheme accounting for their correlation. To obtain estimates of connectedness, i.e. elements of $\boldsymbol{V}^h$, in practice, we estimate the VAR(p) system and translate its coefficient estimates into those compatible with the moving-average representation above. The following expressions provide the relevant estimate of the i, j-th variance decomposition

$$\hat{v}_{i,j}^{h} = \hat{\sigma}_{j,j}^{-1} \sum_{q=0}^{h-1} \left(e_i' \hat{\mathbf{\Theta}}_h \hat{\mathbf{\Omega}} e_i\right)^2 \Big/ \sum_{q=0}^{h-1} \left(e_i' \hat{\mathbf{\Theta}}_h \hat{\mathbf{\Omega}} \hat{\mathbf{\Theta}}_h' e_i\right)^2, \tag{C.11}$$

where hats denote estimates of population counterparts, $\sigma_{j,j}$ is the standard deviation of the error term for the j-th equation in the VAR system, and e_i is a selection vector with unity for the i-th element and zero otherwise. Since the identification scheme is not orthogonalizing shocks, we normalize each entry of $\boldsymbol{V}^h$ row-wise as

$$\tilde{v}_{i,j}^{h} = \hat{v}_{i,j}^{h} \Big/ \sum_{j=1}^{N} \hat{v}_{i,j}^{h}, \tag{C.12}$$

such that $\sum_{j=1}^{N} \tilde{v}_{i,j}^{h} = 1$ and $\sum_{i,j=1}^{N} \tilde{v}_{i,j}^{h} = N$.

D. Bootstrap inference on network connectedness

We implement the following residuals-based parametric bootstrap procedure for obtaining confidence intervals on estimated connectedness measures, facilitating inference. It is performed for each horizon h, yet we drop this superscript for notational simplicity. Denote any connectedness measure by C for simplicity, its estimate by $\hat{C}$, and its bootstrap estimate by $\hat{C}^b$.

1. Collect residuals, $\hat{\boldsymbol{\varepsilon}}_t$, $t = p,\dots,T$, from the VAR(p) system estimated using the Elastic Net on actual data $\boldsymbol{X}_t$.

2. Construct $b = 1,\dots,B$ bootstrap series of residuals (using re-sampling with replacement) of $T-p$ length and re-centering to obtain $\tilde{\boldsymbol{\varepsilon}}_t^b = \hat{\boldsymbol{\varepsilon}}_t^b - (T-p)^{-1}\sum_{t=p}^{T}\hat{\boldsymbol{\varepsilon}}_t^b$.

3. Construct for each b a bootstrapped multivariate time series recursively via

$$\boldsymbol{X}_t^b = \sum_{i=1}^{p}\hat{\boldsymbol{\Psi}}_i\boldsymbol{X}_{t-i}^b + \tilde{\boldsymbol{\varepsilon}}_t^b. \tag{D.13}$$

   Obtain VAR coefficient estimates $\hat{\boldsymbol{\Psi}}_i^b$ from a restricted VAR system, which sets coefficient estimates equal to zero at relevant places whenever the initial Elastic Net estimation does so. This delivers unbiased bootstrap coefficient estimates (with respect to the Elastic Net estimation) and accurate coverage of the confidence intervals constructed in a following step. We have checked this in a simulation experiment with a two-variable VAR(1) model and results are available upon request.

4. Following the principles of constructing connectedness measures of Section IV, generate B many bootstrap estimates of variance decompositions and associated connectedness measure $\hat{C}^b$. From these any bootstrap total and directional connectedness measure can be constructed, cf. Section IV.

5. Using their resulting bootstrap empirical distribution establish (1-α)% confidence intervals as $[2\hat{C} - Q_{1-\alpha/2}^b, 2\hat{C} - Q_{\alpha/2}^b]$ where $Q_{\alpha/2}^b$ and $Q_{1-\alpha/2}^b$ are the $\alpha/2$ and $1-\alpha/2$ percentiles of the bootstrap empirical distributions. This way of constructing the confidence is required, since $\hat{C}^b \geq 0$ for all b, rendering significance testing unreliable based on $[Q_{\alpha/2}^b, Q_{1-\alpha/2}^b]$ as it will never include zero. The simple transformation is appropriate under asymptotic validity of

the bootstrap (e.g. Theorem 3.2 in Paparoditis (1996)) since

$$
\begin{aligned}
1-\alpha &= \mathbb{P}^b[Q^b_{\alpha/2} < \hat{C}^b < Q^b_{1-\alpha/2}],\\
&= \mathbb{P}^b[Q^b_{\alpha/2} - \hat{C} < \hat{C}^b - \hat{C} < Q^b_{1-\alpha/2} - \hat{C}],\\
&\approx \mathbb{P}^b[Q^b_{\alpha/2} - \hat{C} < \hat{C} - C < Q^b_{1-\alpha/2} - \hat{C}],\\
&= \mathbb{P}^b[2\hat{C} - Q^b_{1-\alpha/2} < C < 2\hat{C} - Q^b_{\alpha/2}],
\end{aligned}
\tag{D.14}
$$

where $\mathbb{P}^b[\cdot]$ denotes the bootstrap probability. We have examined the empirical coverage of this procedure in a two-variable VAR(1) system, delivering coverages close to the true $1-\alpha$ for a sample size equal to the one applied in the paper. Results are available upon request.

Note that this approach takes no stance on the asymptotic, nor finite-sample distribution. The method relies only on the bootstrap empirical distribution for inference for a given choice of significance level α. We use $B = 9999$ in the implementation.

E. Connectedness tables

This section reports the connectedness tables for the daily ($h = 1$), weekly ($h = 7$), and monthly ($h = 30$) horizons in Tables A.1–A.3. The structure follows that of the schematic in Table 4.

Table A.1: Connectedness table: daily horizon ($h = 1$)

This table reports the estimated connectedness over a daily horizon, $h = 1$, (in percentages), as per the schematic in Table 4. The abbreviations are from Tables 1–3. The block shaded with blue (orange) indicates the quantities relating the individual connectedness from narratives (macro-finance) to macro-finance (narratives). Values in bold indicate statistical significance at a 10% level using bootstrapped confidence intervals outlined in Appendix D.

	Macro-finance																	Narratives													
	ADS	IDT	DEF	DOL	SMB	HML	FFR	GLD	INF	OIL	RIT	RET	VLM	EPU	TED	TMS	VIX	CCF	PSP	SMC	MPI	SPL	BUC	JBL	INW	FMI	FPI	IVF	SAH	C19	From others
ADS	**57.2**	0.3	0.0	1.4	1.2	0.2	0.8	1.0	**4.3**	0.2	0.2	1.1	0.1	2.8	**8.7**	**10.2**	0.2	0.1	0.5	0.0	1.9	0.0	**2.0**	0.3	**2.8**	0.6	0.7	0.7	0.1	0.6	**1.4**
IDT	0.3	**56.4**	0.2	0.3	0.1	0.6	2.4	0.1	0.1	0.3	1.2	0.8	1.8	**12.8**	0.9	0.2	1.3	3.5	0.0	1.4	1.1	**4.6**	0.2	0.4	0.1	0.4	0.0	0.8	3.4	**4.1**	**1.5**
DEF	0.1	0.2	**74.4**	**5.4**	0.5	0.1	0.6	**2.8**	1.4	0.0	1.6	1.4	0.0	0.8	0.1	3.0	0.2	0.0	0.2	0.0	0.9	0.0	0.3	1.0	1.4	2.7	0.1	0.0	0.1	0.7	0.9
DOL	1.1	0.1	3.2	**44.5**	1.1	**6.2**	0.1	1.1	**3.7**	0.1	**13.0**	**12.3**	0.9	0.2	0.1	2.2	**4.3**	0.4	0.0	0.0	2.1	0.7	0.3	1.0	0.2	0.2	0.1	0.1	0.6	0.1	**1.8**
SMB	1.6	0.1	0.5	0.3	**72.2**	**5.1**	0.2	0.8	0.8	0.4	**4.8**	0.6	1.3	0.1	0.3	0.5	0.0	2.5	0.5	0.9	0.8	0.2	0.2	0.1	0.0	0.9	1.1	1.4	0.2	1.7	**0.9**
HML	0.2	0.4	0.1	**7.0**	3.6	**52.1**	0.1	0.2	2.0	0.1	**12.4**	**7.2**	3.2	0.2	0.1	1.0	2.7	2.9	0.0	0.2	0.6	0.6	0.1	0.6	0.2	0.2	0.0	0.0	0.8	1.1	**1.6**
FFR	0.7	1.8	0.4	0.1	0.1	0.2	**52.0**	**6.6**	0.8	0.0	1.6	0.2	0.4	0.1	0.0	2.7	3.9	1.1	0.5	**4.6**	0.4	0.6	0.1	0.0	2.5	1.0	0.6	1.6	0.8	**14.4**	**1.6**
GLD	0.9	0.1	2.0	1.3	0.6	0.3	**7.0**	**53.6**	2.9	1.4	0.9	0.5	3.3	0.0	0.3	0.0	3.0	0.6	1.0	**6.1**	0.0	0.4	0.0	0.0	**6.5**	2.8	0.1	0.1	3.0	1.4	**1.5**
INF	2.4	0.2	0.9	**3.6**	**8.2**	**3.3**	1.6	**2.7**	**37.0**	0.5	**8.5**	**5.1**	0.6	0.3	0.4	0.5	2.1	1.0	1.3	0.2	1.1	0.2	0.5	0.1	**8.5**	2.3	1.2	0.3	0.3	**5.2**	**2.1**
OIL	0.3	0.5	0.0	0.2	0.5	0.1	0.0	2.4	1.7	**88.2**	0.0	0.0	0.2	0.0	1.7	0.1	0.0	0.4	0.0	0.4	0.0	0.1	0.0	0.2	0.0	0.2	0.6	0.5	1.0	0.8	0.4
RIT	0.1	0.3	0.7	**9.0**	2.0	**7.4**	2.9	1.0	**3.3**	0.0	**30.5**	**22.0**	0.2	0.2	1.2	0.0	**22.0**	1.5	0.6	0.6	0.9	0.0	0.7	0.5	0.9	1.8	0.3	0.0	0.8	1.8	**2.3**
RET	0.6	0.1	0.8	**12.3**	0.2	**4.9**	1.8	0.3	**3.0**	0.0	**21.5**	**32.5**	0.0	0.1	0.7	0.0	**11.8**	0.4	0.3	0.4	1.3	0.5	0.8	1.0	0.9	0.9	0.5	0.0	0.7	1.7	**2.3**
VLM	0.1	2.6	0.0	1.6	1.4	**4.9**	0.6	0.1	0.3	0.2	0.3	0.1	**78.5**	0.1	0.1	0.2	0.5	1.8	0.0	0.5	1.4	0.1	0.3	0.2	0.5	0.0	0.1	0.7	0.0	2.8	**0.7**
EPU	3.3	**14.0**	0.0	0.2	0.0	0.2	0.1	0.0	0.3	0.0	0.3	0.3	0.0	**60.5**	0.1	2.3	0.3	0.2	0.1	0.3	0.8	0.6	0.1	**9.3**	2.2	1.7	1.5	0.2	0.6	0.2	**1.3**
TED	**9.7**	1.2	0.0	0.1	0.3	0.2	0.0	0.2	0.5	1.3	3.6	2.5	0.1	0.3	**66.4**	3.0	2.7	0.3	0.4	0.0	0.3	0.2	1.1	1.1	0.4	**3.3**	0.3	0.0	0.3	0.2	**1.1**
TMS	**11.1**	0.2	2.3	3.1	0.4	1.2	3.2	0.0	1.0	0.1	0.0	0.0	0.0	2.2	2.9	**62.1**	0.0	0.1	0.0	0.7	1.1	1.6	0.1	2.7	0.0	0.1	0.8	1.9	1.0	0.1	**1.3**
VIX	0.2	0.0	0.1	**4.3**	0.0	2.4	3.6	2.6	3.0	0.0	**13.0**	**17.5**	0.3	0.2	0.5	0.0	**45.7**	0.7	0.1	0.9	0.1	0.3	1.5	0.3	1.9	0.2	0.1	0.2	0.3	0.0	**1.8**
CCF	0.1	3.5	0.0	0.4	2.2	3.4	1.3	0.6	0.0	0.3	**2.8**	0.8	1.4	0.2	0.2	0.1	1.0	**62.2**	1.4	1.5	0.0	1.2	2.0	1.1	0.5	0.9	1.1	0.3	0.2	**9.2**	1.3
PSP	0.6	0.0	0.2	0.0	0.5	0.0	0.6	1.2	0.2	0.0	1.2	0.6	0.0	0.1	0.4	0.0	0.1	1.5	**63.6**	2.5	1.3	0.0	1.5	**7.0**	**8.7**	1.3	2.0	0.2	0.7	**3.9**	**1.2**
SMC	0.0	1.3	0.0	0.1	0.0	0.2	**5.3**	**6.8**	0.1	0.3	0.7	0.5	0.4	0.3	0.0	0.7	1.1	1.4	2.4	**60.0**	0.3	**4.7**	0.1	2.1	2.0	0.9	0.7	1.0	1.1	**5.4**	**1.3**
MPI	2.2	1.5	0.7	3.1	0.4	0.8	0.5	0.0	0.9	0.0	2.0	2.7	1.2	0.9	0.3	1.2	0.2	0.0	1.4	0.3	**68.3**	0.7	3.8	2.4	1.7	0.1	0.1	0.0	1.8	0.8	**1.1**
SPL	0.0	**5.6**	0.0	1.1	0.2	0.8	0.9	0.5	0.0	0.1	0.1	1.3	0.1	0.8	0.2	1.9	0.6	1.4	0.1	**5.7**	0.8	**72.0**	0.5	0.5	1.4	0.0	0.0	0.0	0.0	3.4	**0.9**
BUC	2.5	0.4	0.4	0.4	0.2	0.2	0.1	0.0	0.9	0.0	1.6	1.6	0.3	0.2	1.0	0.2	2.3	2.3	1.6	0.1	3.9	0.5	**72.0**	0.4	1.9	0.0	0.3	2.9	0.2	1.8	**0.9**
JBL	0.3	0.4	0.5	1.3	0.1	0.6	0.0	0.0	0.1	0.1	1.0	1.9	0.2	**8.9**	1.1	2.5	0.4	1.0	**6.4**	2.0	2.0	0.4	0.3	**58.9**	0.8	**4.2**	0.2	1.0	2.8	0.3	**1.4**
INW	2.4	0.1	0.8	0.3	0.0	0.2	2.5	**6.1**	**9.0**	0.0	1.6	1.8	0.3	1.8	0.2	0.0	2.4	0.4	**6.4**	1.6	1.2	0.9	1.2	0.6	**48.1**	2.5	**6.6**	0.0	0.0	1.0	**1.7**
FMI	0.7	0.8	2.0	0.4	1.1	0.5	1.4	3.3	2.2	0.1	1.7	0.6	0.0	1.6	**3.1**	0.1	0.4	3.5	1.5	1.1	0.2	0.1	0.1	**4.5**	3.3	**62.3**	0.1	0.9	0.2	2.2	**1.3**
FPI	0.9	0.1	0.1	0.1	1.0	0.0	0.9	0.2	2.6	0.5	0.4	0.6	0.1	1.8	0.3	0.9	0.6	1.1	2.1	0.9	0.1	0.0	0.2	0.2	**9.9**	0.0	**69.1**	2.6	0.9	1.7	**1.0**
IVF	1.1	1.4	0.0	0.1	1.3	0.1	**5.2**	0.4	0.9	0.4	0.1	0.0	0.7	0.3	0.1	2.6	0.2	0.2	0.3	1.7	0.0	0.1	2.7	1.1	0.2	0.8	2.9	**71.6**	2.5	0.8	**0.9**
SAH	0.1	**4.2**	0.1	1.0	0.2	1.1	1.1	**4.0**	0.5	0.9	0.1	0.0	0.0	0.7	0.3	1.0	0.4	0.2	0.8	1.3	2.3	0.0	0.1	**3.4**	0.0	0.5	1.0	2.7	**70.6**	1.1	**1.0**
C19	0.5	**3.7**	0.5	0.0	1.1	1.0	**13.4**	1.1	1.6	0.4	0.7	0.0	1.8	0.1	0.2	0.0	0.0	**7.4**	**3.0**	**4.5**	0.6	2.4	1.3	0.3	1.1	1.2	1.3	0.1	0.8	**49.8**	1.7
To others	**1.5**	**1.5**	0.5	**2.0**	**0.9**	**1.5**	**1.9**	**1.5**	**1.6**	0.3	**3.2**	**2.8**	0.6	2.8	0.9	1.2	**1.7**	**1.3**	**1.1**	**1.3**	0.9	0.7	0.7	**1.4**	**2.0**	**1.1**	**0.8**	0.7	0.8	**2.3**	1.3

Table A.2: Connectedness table: weekly horizon ($h = 7$)

This table reports the estimated connectedness over a weekly horizon, $h = 7$, (in percentages), as per the schematic in Table 4. The abbreviations are from Tables 1–3. The block shaded with blue (orange) indicates the quantities relating the individual connectedness from narratives (macro-finance) to macro-finance (narratives). Values in bold indicate statistical significance at a 10% level using bootstrapped confidence intervals outlined in Appendix D.

	Macro-finance																	Narratives													
	ADS	IDT	DEF	DOL	SMB	HML	FFR	GLD	INF	OIL	RIT	RET	VLM	EPU	TED	TMS	VIX	CCF	PSP	SMC	MPI	SPL	BUC	JBL	INW	FMI	FPI	IVF	SAH	C19	From others
ADS	**57.2**	0.3	0.0	1.4	1.2	0.2	0.8	1.0	**4.3**	0.2	0.2	1.1	0.1	2.8	**8.7**	**10.2**	0.2	0.1	0.5	0.0	1.9	0.0	**2.0**	0.3	**2.8**	0.6	0.7	0.7	0.1	0.6	**1.4**
IDT	0.6	**33.8**	0.3	**2.0**	1.0	**1.8**	**3.9**	1.6	**6.2**	0.2	**6.6**	**5.6**	1.1	**7.4**	0.6	0.2	**9.1**	2.8	0.4	1.1	0.8	2.8	0.1	0.2	**2.5**	1.0	0.3	0.5	2.0	**3.7**	**2.2**
DEF	1.9	1.4	**57.5**	**4.4**	0.4	0.1	0.4	2.2	1.2	0.0	1.4	1.1	0.0	9.4	3.0	4.0	0.2	0.0	0.2	0.1	1.1	0.1	0.3	**3.4**	1.6	3.4	0.3	0.1	0.1	0.7	**1.4**
DOL	1.0	0.2	2.8	**38.8**	2.2	**5.7**	0.7	1.3	**3.5**	0.1	**12.2**	**11.6**	0.9	0.2	0.1	1.9	**4.2**	2.6	0.1	0.3	1.8	0.9	0.3	1.1	0.4	1.8	0.2	0.3	1.2	1.8	**2.0**
SMB	1.0	1.1	0.4	**3.1**	**37.5**	**7.3**	**4.4**	1.7	1.0	1.9	**7.3**	**5.4**	2.7	0.0	0.3	0.4	2.1	2.8	0.3	**4.2**	1.3	**3.9**	0.3	0.2	0.5	1.0	0.9	2.2	0.3	4.3	**2.1**
HML	0.3	0.7	0.1	**6.7**	3.0	**43.9**	1.3	1.6	1.7	0.1	**11.6**	**8.5**	2.6	0.2	0.5	0.9	**3.0**	3.1	0.1	0.3	0.7	1.4	0.1	0.6	0.4	2.4	0.0	0.1	2.5	1.3	**1.9**
FFR	0.6	1.3	0.7	3.7	3.0	1.2	**35.8**	**6.2**	1.4	0.1	**5.9**	**6.0**	0.4	0.1	0.1	2.0	**5.7**	1.3	0.4	3.3	0.5	0.7	0.1	0.2	2.1	1.5	0.7	3.5	1.1	**10.3**	**2.1**
GLD	0.7	0.1	1.5	**6.1**	1.6	**1.9**	**5.6**	**33.9**	**2.2**	1.0	**7.3**	**11.1**	2.5	0.0	0.6	0.1	**5.1**	0.5	0.7	**4.1**	0.5	0.3	1.4	0.3	**4.0**	2.1	0.2	0.3	2.2	1.9	**2.2**
INF	1.5	0.3	0.5	2.4	**7.9**	**2.3**	**6.8**	**5.0**	**21.1**	0.3	**9.0**	**4.8**	0.5	0.4	0.4	0.2	3.4	**3.0**	3.0	0.6	0.6	0.7	0.4	0.1	**9.7**	**3.2**	0.9	0.8	1.2	**9.2**	**2.6**
OIL	0.2	0.5	0.0	0.4	0.5	1.1	0.1	2.2	1.8	**85.6**	0.3	0.2	0.2	0.0	1.7	0.2	0.1	0.4	0.0	1.0	0.0	0.1	0.0	0.3	0.0	0.2	0.7	0.5	0.9	0.8	0.5
RIT	0.2	0.3	0.7	**8.3**	2.6	**6.7**	2.9	1.5	**3.1**	0.0	**27.3**	**20.2**	0.3	0.3	1.1	0.0	**8.1**	**4.0**	0.5	0.7	0.8	0.1	0.8	0.9	1.0	**3.1**	0.4	0.1	1.5	2.6	**2.4**
RET	0.5	0.1	1.0	**12.1**	0.9	**5.0**	1.8	0.5	**3.0**	0.0	**21.0**	**31.1**	0.1	0.2	0.6	0.1	**11.1**	1.2	0.2	0.4	1.3	0.5	0.9	1.0	0.9	1.4	0.5	0.1	0.7	1.9	**2.3**
VLM	0.2	2.5	0.2	1.5	3.3	**4.5**	1.2	2.3	0.5	0.3	0.7	0.2	**70.0**	0.1	0.1	0.2	0.6	1.8	0.1	1.0	1.3	0.2	0.3	0.2	0.9	1.2	0.1	1.4	0.2	**2.8**	**1.0**
EPU	**4.5**	**13.0**	0.0	0.2	0.0	0.2	0.1	0.0	0.2	0.0	0.2	0.2	0.0	**58.9**	1.0	2.9	0.2	0.1	0.1	0.3	0.9	0.6	0.1	**9.8**	1.9	2.2	1.4	0.3	0.5	0.2	**1.4**
TED	5.7	0.9	0.0	0.8	0.2	0.9	0.7	0.4	1.6	0.8	**8.1**	**8.1**	0.1	0.5	**46.6**	1.7	**15.9**	0.5	0.2	0.1	0.1	0.2	1.7	0.5	1.0	1.8	0.3	0.1	0.4	0.1	**1.8**
TMS	**10.1**	0.2	2.4	3.1	3.1	1.7	3.2	0.1	1.2	0.1	0.3	0.1	0.1	2.2	2.6	**59.5**	0.0	0.1	0.0	0.9	1.1	1.5	0.1	2.5	0.0	0.1	0.7	1.6	1.0	0.1	**1.4**
VIX	1.8	0.4	0.1	**3.8**	0.1	2.0	3.1	2.3	**3.3**	0.0	**12.2**	**15.7**	0.3	0.7	0.5	0.8	**42.3**	0.5	0.1	0.6	0.4	4.3	1.5	0.6	1.6	0.3	0.1	0.2	0.3	0.2	**1.9**
CCF	0.2	3.1	0.0	0.6	2.2	**3.4**	1.2	0.6	0.1	0.3	**2.8**	0.9	1.4	0.3	0.2	0.2	0.9	**58.8**	1.9	2.4	2.0	1.2	2.2	1.7	0.4	0.8	1.2	0.3	0.3	**8.5**	**1.4**
PSP	1.0	0.1	0.2	0.7	1.1	0.4	**6.9**	**3.5**	1.3	0.1	**4.5**	2.2	0.2	0.4	0.3	1.0	1.3	1.3	**43.3**	**3.1**	0.8	0.1	1.0	**4.7**	**7.2**	1.4	1.4	0.4	3.3	**6.8**	**1.9**
SMC	0.4	1.4	0.0	0.1	3.2	0.7	**5.3**	**5.3**	0.3	2.2	1.0	0.6	0.6	0.3	0.1	1.3	0.9	1.9	3.1	**46.6**	0.3	**3.9**	0.1	2.0	2.0	1.1	2.0	6.5	1.4	**5.5**	**1.8**
MPI	2.1	1.1	0.7	**4.3**	0.9	**6.0**	0.4	0.1	1.4	0.0	**3.9**	**3.8**	1.8	1.0	0.2	1.4	0.5	0.4	1.3	0.4	**57.4**	0.6	3.4	2.5	1.2	0.4	0.1	0.0	1.6	1.1	**1.4**
SPL	0.1	**5.8**	0.0	1.0	0.1	0.9	0.8	0.5	0.3	0.1	0.1	1.1	0.1	1.0	0.3	1.9	0.5	1.3	0.4	**5.3**	1.0	**70.9**	0.5	0.5	2.0	0.1	0.3	0.0	0.0	3.2	**1.0**
BUC	4.4	0.7	2.7	0.8	0.2	0.3	0.2	0.0	1.0	0.1	1.0	1.0	0.3	0.8	1.6	**10.3**	1.5	1.7	1.1	0.8	2.5	0.6	**55.4**	0.5	1.4	0.3	0.3	**7.0**	0.2	1.2	**1.5**
JBL	0.3	0.4	0.5	1.4	0.1	0.7	0.0	0.1	0.2	0.1	0.9	1.9	0.2	**8.2**	1.1	2.3	0.4	1.1	**7.1**	1.9	1.8	1.1	0.3	**57.7**	2.0	**3.9**	0.4	0.9	**2.7**	0.3	**1.4**
INW	1.9	0.2	0.7	0.6	3.3	0.8	2.8	**5.6**	**8.0**	0.0	**3.5**	2.4	0.3	1.6	0.3	0.0	3.2	7.3	**4.6**	1.4	1.0	0.9	0.8	0.4	**37.9**	2.6	**4.7**	0.2	0.1	2.6	**2.1**
FMI	0.8	**3.0**	1.3	0.8	1.1	1.4	2.9	2.4	1.6	0.2	2.2	0.8	0.6	1.3	1.8	0.4	0.5	**14.1**	1.2	1.6	0.6	0.4	1.5	**3.8**	2.8	**40.0**	0.1	4.8	1.4	**4.7**	**2.0**
FPI	7.0	0.1	0.1	0.4	1.1	0.1	1.5	0.6	**3.9**	0.4	0.5	0.9	0.1	2.3	0.5	2.2	2.9	0.8	1.8	0.8	0.2	0.1	0.2	0.1	**9.9**	0.1	**56.9**	2.2	0.7	1.5	**1.4**
IVF	0.9	4.1	0.1	0.7	1.3	0.5	**6.4**	0.7	1.0	0.4	1.0	0.7	0.8	0.2	0.1	2.4	0.7	0.9	0.3	1.8	0.5	0.3	3.9	0.8	0.4	2.5	2.4	**57.7**	5.2	1.4	**1.4**
SAH	0.4	**3.7**	0.1	1.0	0.2	1.2	1.1	**3.7**	0.7	0.9	1.3	1.1	0.0	0.7	4.7	0.9	1.8	0.3	0.7	1.2	2.9	0.0	0.4	**3.1**	0.2	0.7	0.9	2.4	62.5	1.0	**1.2**
C19	0.5	**3.1**	0.4	0.2	**4.2**	1.1	**11.7**	1.2	1.4	0.4	0.9	0.2	1.6	0.7	0.3	0.5	0.1	**6.3**	**3.1**	**3.6**	0.6	2.3	1.4	4.2	1.0	1.1	1.0	5.6	0.7	**40.7**	**2.0**
To others	1.7	**1.7**	0.6	**2.4**	**1.7**	**2.0**	**2.6**	**1.8**	**1.9**	0.4	**4.3**	**3.9**	0.7	**1.4**	1.1	**1.7**	**2.8**	**2.1**	**1.1**	**1.4**	1.0	1.0	0.9	**1.6**	**2.1**	**1.4**	0.8	**1.4**	1.1	**2.7**	**1.7**

Table A.3: Connectedness table: monthly horizon ($h = 30$)

This table reports the estimated connectedness over a monthly horizon, $h = 30$, (in percentages), as per the schematic in Table 4. The abbreviations are from Tables 1–3. The block shaded with blue (orange) indicates the quantities relating the individual connectedness from narratives (macro-finance) to macro-finance (narratives). Values in bold indicate statistical significance at a 10% level using bootstrapped confidence intervals outlined in Appendix D.

	Macro-finance																	Narratives													
	ADS	IDT	DEF	DOL	SMB	HML	FFR	GLD	INF	OIL	RIT	RET	VLM	EPU	TED	TMS	VIX	CCF	PSP	SMC	MPI	SPL	BUC	JBL	INW	FMI	FPI	IVF	SAH	C19	From others
ADS	**57.2**	0.3	0.0	1.4	1.2	0.2	0.8	1.0	**4.3**	0.2	0.2	1.1	0.1	2.8	**8.7**	**10.2**	0.2	0.1	0.5	0.0	1.9	0.0	**2.0**	0.3	**2.8**	0.6	0.7	0.7	0.1	0.6	**1.4**
IDT	5.5	**18.8**	0.3	**1.4**	1.1	**1.1**	**3.5**	2.3	**6.8**	0.1	**6.4**	**4.7**	0.6	**5.3**	0.4	1.9	**8.5**	2.3	1.3	1.3	1.5	12.7	0.1	0.2	**5.4**	1.3	0.7	0.6	1.3	**2.5**	**2.7**
DEF	4.6	1.4	**19.8**	**2.5**	0.2	0.5	1.1	1.7	3.3	0.2	6.6	7.0	0.1	7.1	8.1	2.9	16.2	0.1	0.2	0.2	1.0	7.4	0.7	2.2	2.0	**1.9**	0.4	0.2	0.1	0.4	**2.7**
DOL	1.0	0.2	2.6	**36.8**	2.2	**5.5**	1.0	1.5	**3.5**	0.1	**11.8**	**11.1**	0.9	0.2	0.1	1.9	**4.2**	2.7	0.1	0.7	2.2	1.5	0.4	1.1	0.6	1.8	0.3	0.3	1.5	2.0	**2.1**
SMB	0.9	1.1	0.4	**3.1**	**34.6**	**6.9**	**4.3**	2.1	1.0	1.9	**7.0**	**5.3**	2.5	0.1	0.4	0.4	2.2	3.5	0.8	**4.1**	1.4	**3.9**	0.5	0.3	0.6	1.8	1.1	2.6	1.0	**4.2**	**2.2**
HML	0.4	0.7	0.1	**6.6**	3.0	**41.3**	1.5	1.6	1.8	0.1	**11.4**	**8.5**	2.5	0.3	0.5	1.0	**2.9**	3.6	0.2	0.4	1.4	1.5	0.3	0.8	0.7	2.4	0.1	0.3	2.9	1.4	**2.0**
FFR	0.7	1.6	0.7	4.0	3.0	1.6	**32.7**	**5.7**	1.4	0.1	**6.1**	**6.3**	0.5	0.1	0.2	1.9	**5.6**	1.6	0.5	3.5	0.7	2.0	0.4	0.2	2.0	1.7	0.7	3.5	1.3	**9.7**	**2.2**
GLD	0.7	0.2	1.5	**6.3**	1.6	**2.4**	**5.4**	**31.8**	**2.2**	1.0	**7.8**	**11.7**	2.4	0.1	0.6	0.2	**5.2**	0.8	0.7	**4.0**	0.6	0.5	1.6	0.4	**3.8**	2.0	0.3	0.4	2.1	1.9	**2.3**
INF	2.0	1.3	0.5	2.0	**6.9**	**1.9**	**5.4**	**4.2**	**17.0**	0.3	**7.5**	**4.0**	0.4	0.9	0.4	0.9	2.9	3.1	3.2	1.7	1.1	8.7	0.4	0.2	**9.3**	**2.8**	1.0	1.7	1.0	**7.4**	**2.8**
OIL	0.3	0.5	0.0	0.4	0.5	1.1	0.2	2.2	1.8	**84.7**	0.4	0.2	0.2	0.0	1.6	0.3	0.1	0.4	0.1	1.0	0.0	0.1	0.0	0.3	0.1	0.2	0.7	0.6	1.0	0.8	0.5
RIT	0.3	0.4	0.6	**8.0**	2.6	**6.5**	2.9	1.5	**3.1**	0.0	**26.1**	**19.4**	0.3	0.3	1.1	0.1	**7.8**	**4.4**	0.5	0.9	1.5	0.5	0.9	0.9	1.2	**3.1**	0.4	0.2	1.9	2.7	**2.5**
RET	0.6	0.1	1.0	**11.9**	0.9	**4.9**	1.8	0.5	3.0	0.0	**20.6**	**30.6**	0.1	0.2	0.6	0.1	**10.9**	1.3	0.3	0.5	1.5	0.8	0.9	1.0	1.0	1.4	0.5	0.1	0.8	1.9	**2.3**
VLM	0.2	2.5	0.2	1.6	3.2	**4.5**	1.4	2.3	0.6	0.4	1.0	0.5	**67.5**	0.1	0.1	0.2	0.8	2.2	0.1	1.0	1.5	0.3	0.5	0.3	1.0	1.4	0.1	1.4	0.2	**2.9**	**1.1**
EPU	**4.9**	**8.9**	0.0	0.6	0.0	0.3	0.6	0.6	1.8	0.1	3.4	3.6	0.1	**41.3**	3.4	2.6	8.8	0.2	0.1	0.3	0.9	4.7	0.3	**7.1**	2.0	1.8	1.0	0.3	0.3	0.2	**2.0**
TED	10.1	2.1	0.1	1.4	0.1	0.7	1.5	1.5	4.9	0.3	**7.9**	**8.1**	0.1	2.6	**11.8**	3.2	**18.9**	0.5	0.5	0.4	1.6	**14.2**	0.8	0.6	4.1	0.8	0.8	0.1	0.2	0.2	**2.9**
TMS	**8.2**	1.0	2.1	2.8	3.4	2.2	2.9	0.3	1.1	0.2	1.0	0.7	0.4	1.9	2.2	**50.2**	0.5	1.5	0.4	3.5	1.2	4.4	0.4	2.1	0.2	0.6	0.7	1.5	1.0	1.5	**1.7**
VIX	12.0	1.7	0.2	**2.2**	0.1	1.0	2.0	2.0	**4.9**	0.1	**7.5**	**8.9**	0.2	2.6	0.5	3.6	**23.2**	0.9	0.6	0.8	2.0	**15.0**	0.7	0.7	4.7	0.7	0.8	0.1	0.1	0.3	**2.6**
CCF	1.4	2.2	0.2	1.8	1.8	**3.9**	2.2	1.0	1.1	0.2	**3.8**	1.8	1.6	0.8	0.2	1.1	1.1	**40.0**	1.8	3.2	**9.9**	1.1	2.7	2.6	1.2	1.5	0.9	1.1	0.8	**7.1**	**2.0**
PSP	1.4	0.2	0.4	1.0	2.0	0.7	**6.6**	**3.6**	1.4	0.1	**4.7**	2.6	0.2	0.5	0.9	3.0	2.3	1.9	**35.4**	**3.5**	0.7	1.2	1.0	**4.0**	**6.0**	2.2	1.2	0.6	4.3	**6.3**	**2.2**
SMC	1.2	1.8	0.1	1.0	3.1	1.5	**6.5**	**4.9**	1.2	1.5	3.2	2.0	0.7	0.3	0.1	1.2	2.0	3.2	3.2	**33.4**	0.8	**3.9**	0.5	1.6	2.8	2.5	2.0	5.9	1.5	**6.3**	**2.2**
MPI	2.2	0.9	0.7	**4.1**	1.1	**5.8**	0.8	0.3	1.5	0.0	**4.0**	**3.6**	1.8	1.1	0.3	1.6	0.5	1.6	1.3	0.4	**53.7**	0.8	3.4	2.8	1.4	1.1	0.1	0.1	1.6	1.5	**1.5**
SPL	1.1	**4.7**	0.2	0.7	0.9	0.7	0.8	1.1	2.3	0.1	1.3	1.1	0.1	2.0	0.4	1.8	1.0	3.1	1.0	**4.4**	2.4	**57.5**	0.4	0.5	6.8	0.6	0.9	0.2	0.0	2.2	**1.4**
BUC	3.8	2.1	2.0	1.1	0.3	1.0	1.4	0.6	1.4	0.3	4.1	3.5	0.4	1.5	2.5	**11.2**	5.6	3.6	0.7	2.2	2.4	1.5	**35.3**	0.6	1.6	0.9	0.2	**5.4**	0.3	2.3	**2.2**
JBL	0.7	0.9	0.4	1.1	0.4	0.6	0.1	0.4	1.2	0.1	1.2	1.6	0.2	**7.0**	1.1	2.1	0.6	1.8	**6.5**	2.0	2.3	8.5	0.3	**46.7**	4.7	**3.2**	0.7	0.9	**2.2**	0.3	**1.8**
INW	2.9	0.7	0.6	0.7	3.2	0.9	2.5	**4.5**	**6.9**	0.0	**3.6**	2.2	0.3	2.1	0.3	1.0	2.8	7.6	**3.5**	3.0	3.4	7.3	0.8	0.8	**29.0**	2.4	**3.7**	0.6	0.2	2.3	**2.4**
FMI	1.8	**2.1**	0.8	2.0	1.1	2.5	3.8	2.0	2.8	0.1	5.2	3.2	0.9	1.3	1.0	0.7	2.9	**13.3**	1.0	1.0	5.7	0.9	3.4	**3.6**	3.3	**23.4**	0.1	4.3	1.3	**4.7**	**2.6**
FPI	**20.9**	0.5	0.1	0.8	1.0	0.2	1.4	1.0	**4.8**	0.3	0.9	1.5	0.1	3.0	1.9	5.0	3.3	0.6	1.4	0.5	1.1	2.2	0.4	0.3	**8.1**	0.4	**35.6**	1.5	0.4	1.0	**2.1**
IVF	1.5	3.2	0.1	1.2	1.3	1.0	**6.0**	1.5	2.7	0.3	4.3	2.9	0.7	0.5	0.3	1.8	3.8	4.9	0.2	1.4	2.3	1.7	4.2	1.0	2.4	4.2	1.6	**36.7**	3.5	2.7	**2.1**
SAH	2.5	**3.4**	0.1	1.1	0.2	1.0	1.4	**3.5**	1.7	0.8	2.5	2.3	0.1	1.2	3.6	1.5	5.3	0.4	0.7	1.2	2.7	3.5	0.5	**2.6**	1.1	0.7	0.9	2.0	**50.8**	0.8	**1.6**
C19	0.4	**4.0**	0.3	1.1	**3.4**	1.8	**9.7**	1.0	1.1	0.3	1.4	1.0	1.6	0.7	0.4	0.7	0.5	**5.9**	**3.4**	**3.4**	0.8	5.9	2.0	5.8	1.1	1.8	0.9	7.1	0.7	**31.8**	**2.3**
To others	**3.1**	**1.7**	0.6	**2.5**	**1.7**	**2.1**	**2.6**	**1.9**	**2.5**	0.3	**4.9**	**4.3**	0.7	**1.5**	1.4	**2.1**	**4.2**	**2.6**	**1.2**	**1.7**	1.9	3.9	1.0	**1.5**	**2.7**	**1.6**	0.8	**1.5**	1.1	**2.6**	2.1

F. Robustness of main findings to number of topics

This section presents robustness of the connectedness between narratives and macro-finance variables as a function of the number of topics selected from the LDA estimation. It addresses the question whether narratives as a group drives unexpected fluctuations of macro-finance variables and vice versa, though we stress that interpretability of the used topic model is notably reduced, cf. Section III. Table A.4 contains results for $N_n = 8$, and Table A.5 contains results for $N_n = 18$.

Table A.4: Narratives to and from macro-finance, $N_n = 8$

This table reports the cumulative transmission of shocks (connectedness) from the narratives group ($\boldsymbol{X}_t^n$) to the macro-financial group ($\boldsymbol{X}_t^{mf}$), and vice versa, for the horizons $h = 1,2,3,4,5,6,7,14,30$ and a choice of topics equal to $N_n = 8$. The numbers represent shares of total forecast error variance, cf. (8), or when the diagonal (own share) is masked. The latter is reported in parenthesis and is not indicated by statistical significance as it is identical to the original number, $C^h_{mf\to n}$ or $C^h_{n\to mf}$. Statistical significance is based on bootstrapped confidence intervals outlined in Appendix D: asterisks "***", "**", and "*" indicates significance on the 1% level, 5% level, and 10% level, respectively.

	Effect horizon in days (h)								
Direction of connectedness	1	2	3	4	5	6	7	14	30
Narratives to macro-finance: $C^h_{n\to mf}$	7.7***	8.3***	9.2***	9.6***	10.3***	10.8***	11.3***	14.5***	19.3***
(Share of non-own)	(18.6)	(19.2)	(20.2)	(19.8)	(20.3)	(20.6)	(21.2)	(24.8)	(30.7)
Macro-finance to narratives: $C^h_{mf\to n}$	17.6***	19.2***	21.1***	22.1***	22.8***	23.6***	24.3***	28.4***	31.0***
(Share of non-own)	(55.4)	(57.7)	(60.2)	(47.2)	(61.0)	(61.8)	(61.9)	(63.4)	(61.8)
Net effect: C^h_{n-mf}	-0.4	-0.5	-0.5	-0.5	-0.3	-0.2	-0.1	0.8	3.2
(Share of non-own)	(-0.9)	(-1.2)	(-1.2)	(-0.7)	(-0.5)	(-0.1)	(1.4)	(-2.8)	(5.4)

Table A.5: Narratives to and from macro-finance, $N_n = 18$

This table reports the cumulative transmission of shocks (connectedness) from the narratives group ($\boldsymbol{X}_t^n$) to the macro-financial group ($\boldsymbol{X}_t^{mf}$), and vice versa, for the horizons $h = 1, 2, 3, 4, 5, 6, 7, 14, 30$ and a choice of topics equal to $N_n = 18$. The numbers represent shares of total forecast error variance, cf. (8), or when the diagonal (own share) is masked. The latter is reported in parenthesis and is not indicated by statistical significance as it is identical to the original number, $C^h_{mf\to n}$ or $C^h_{n\to mf}$. Statistical significance is based on bootstrapped confidence intervals outlined in Appendix D: asterisks "***", "**", and "*" indicates significance on the 1% level, 5% level, and 10% level, respectively.

	Effect horizon in days (h)								
Direction of connectedness	1	2	3	4	5	6	7	14	30
Narratives to macro-finance: $C^h_{n\to mf}$	13.4***	14.1***	15.0***	15.3***	15.9***	16.3***	16.7***	19.2***	21.5***
(Share of non-own)	(30.5)	(31.1)	(31.7)	(30.7)	(30.7)	(30.6)	(30.6)	(32.2)	(33.8)
Macro-finance to narratives: $C^h_{mf\to n}$	15.4***	16.8***	17.7***	18.6***	19.2***	19.5***	19.8***	21.2***	22.9***
(Share of non-own)	(46.7)	(47.8)	(47.5)	(47.2)	(46.0)	(44.5)	(43.3)	(40.9)	(41.4)
Net effect: C^h_{n-mf}	-1.4***	-1.8***	-1.8***	-2.2***	-2.1***	-2.1**	-2.1**	-1.6	-1.3
(Share of non-own)	(-3.7)	(-4.4)	(-4.3)	(-4.9)	(-4.5)	(-4.3)	(-4.1)	(-2.8)	(-2.2)

Figure A.2: Individual variable to-connectedness: monthly horizon

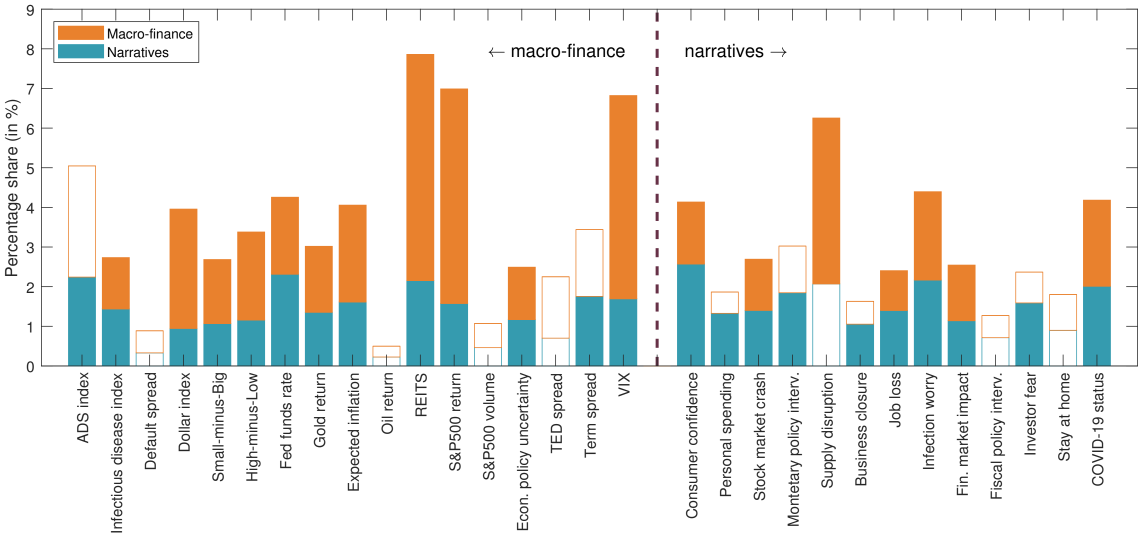

This figure depicts the amount of to-connectedness from the ith variable as a share of the total to-connectedness across all variables in the network. Orange bars indicate the amount transmitted towards macro-finance variables for each variable and blue bars indicate that towards narratives. Shaded bars indicate significance at the 10% level, and hollow bars indicate insignificance at the same level. Statistical significance is based on bootstrapped confidence intervals outlined in Appendix D. The purple dashed line separates the macro-finance and narratives groups.

G. Transmitters and receivers of shocks

To get a deeper understanding of the main drivers of the macro-finance-narratives network, we now examine the total to- and from-connectedness for each individual variable. We further decompose these quantities into those driving (other) macro-finance variables or narratives to understand whether a given variable's role is mainly because it drives the macro-finance or narrative side, or both. We focus on the monthly horizon and report plots for the daily and weekly horizon further below.

G.1. Top transmitters

Figure A.2 depicts the share of total to-connectedness attributed to each of the variables in the system. That is, the share of unexpected fluctuation of other variables in the network attributable to the i-th variable. The bars sum to unity. It also includes a decomposition into whether this transmission is towards (other) macro-financial variables or narratives. The larger the values, the more important the variables are for driving fluctuations of the network's other variables. The ten most important variables in the network are the REITS (8%) and S&P500 returns (7%), VIX (7%), *supply disruption* (6%), ADS index (5%), *infection worry* (4%), Fed funds rate (4%), *COVID-19 status* (4%), *consumer confidence* (4%), and expected inflation (4%). The least important variables are oil returns, the default spread, and S&P500 volume. Generally, the amount of transmissions from narratives to macro-finance is of similar magnitude as that from macro-finance to narratives.

The majority of the unexpected fluctuations due to the *supply narrative* runs towards the macro-finance variables. This link is also statistically significant. On the other hand, most important macro-finance variables predominantly drive unexpected fluctuations of other macro-finance variables with a smaller role for narratives, except for the ADS index, which has a significant influence on narratives and not on macro-finance variables. Stories about the COVID-19 disease captured by *COVID-19 status* and *infection worry* have an equal and significant influence on both other narratives and the macro-finance variables, whereas *consumer confidence* mainly influence the other narratives. The most important driver of narratives within the macro-finance group is the Fed funds rate, indicating that individuals' views are driven by the Fed's monetary policy decisions.

G.2. Top receivers

In a similar style as above, Figure A.3 depicts the share of total from-connectedness received by each of the variables in the system. The bars sum to unity. It also includes a decomposition into whether the receiving shocks are due to (other) macro-financial variables or narratives. The larger the values, the larger the share of a variable's fluctuations is due to other variables in the network. The most important recipient variable is the TED spread. Out of total from-connectedness, it receives about 5%. It is followed expected inflation (4.4%), the infectious disease index (4.4%), the default spread (4.3%), VIX (4.1%), *financial market impact* (4.1%), REITS (4.0%), *infection worry* (3.8%), S&P500 returns (3.7%), and *COVID-19 status* (3.7%). These effects are statistically significant at conventional significance levels. The least receiving variable is the oil return, which, holding together with its low transmission, indicates it has little role in the network as a whole. Generally, the fluctuations of the macro-financial variables are driven more by other elements of the network compared the narratives. The largest macro-finance recipient of narrative shocks is expected inflation. This is followed by the infectious disease index, indicating that narratives are significantly connected with the news media and, in particularly, influence what is reported. Stories about economic policy intervention captured by *monetary policy intervention* and *fiscal policy intervention* are relatively more susceptible to macro-finance shocks. That is, individuals' talk about economic policy intervention is largely driven by the development of the economy surrounding them.

Figure A.3: Individual variable from-connectedness: monthly horizon

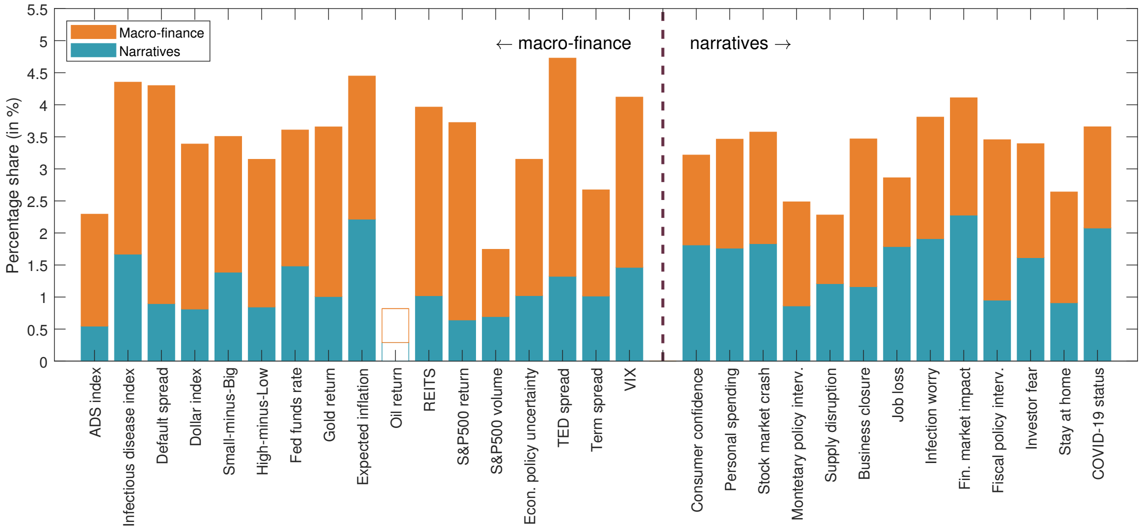

This figure depicts the amount of from-connectedness for the ith variable as a share of the total from-connectedness across all variables in the network. Orange bars indicates the amount from macro-finance variables for each variable and blue bars indicate that from narratives. Shaded bars indicate significance at the 10% level, and hollow bars indicate insignificance at the same level. Statistical significance is based on bootstrapped confidence intervals outlined in Appendix D. The purple dashed line separates the macro-finance and narratives groups.

G.3. Daily and weekly horizons

Figures A.4 and A.5 depict the decompositions of to-connectedness for the daily and weekly horizon, respectively. Similarly, Figures A.6 and A.7 depict the decompositions of from-connectedness for the daily and weekly horizon, respectively.

Figure A.4: Individual variable to-connectedness: daily horizon

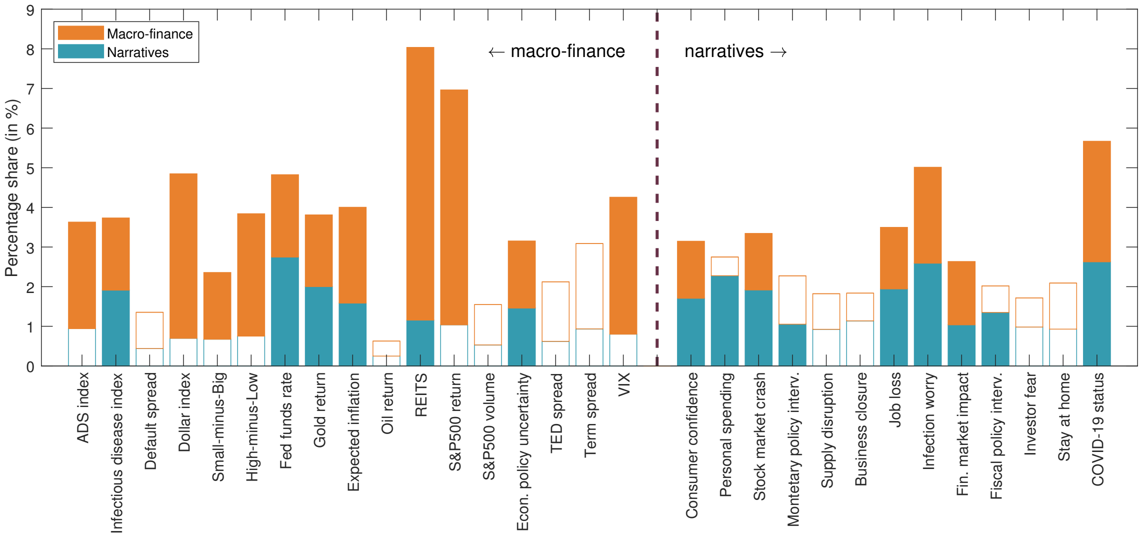

This figure depicts the amount of to-connectedness from the ith variable as a share of the total to-connectedness across all variables in the network. Orange bars indicate the amount transmitted towards macro-finance variables for each variable and blue bars indicate that towards narratives. Shaded bars indicate significance at the 10% level, and hollow bars indicate insignificance at the same level. Statistical significance is based on bootstrapped confidence intervals outlined in Appendix D. The purple dashed line separates the macro-finance and narratives groups.

Figure A.5: Individual variable to-connectedness: weekly horizon

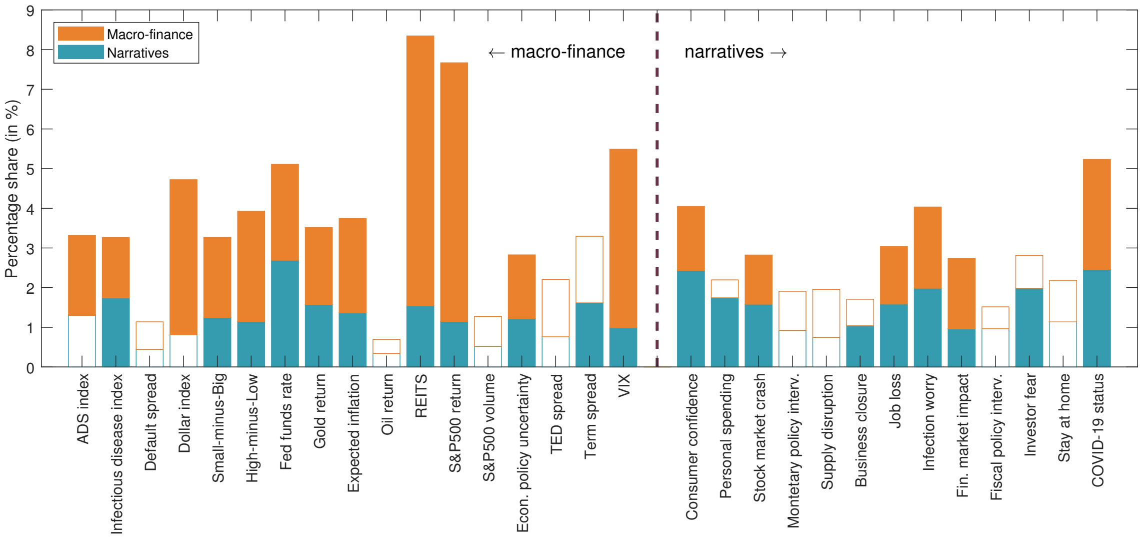

This figure depicts the amount of to-connectedness from the ith variable as a share of the total to-connectedness across all variables in the network. Orange bars indicate the amount transmitted towards macro-finance variables for each variable and blue bars indicate that towards narratives. Shaded bars indicate significance at the 10% level, and hollow bars indicate insignificance at the same level. Statistical significance is based on bootstrapped confidence intervals outlined in Appendix D. The purple dashed line separates the macro-finance and narratives groups.

Figure A.6: Individual variable from-connectedness: daily horizon

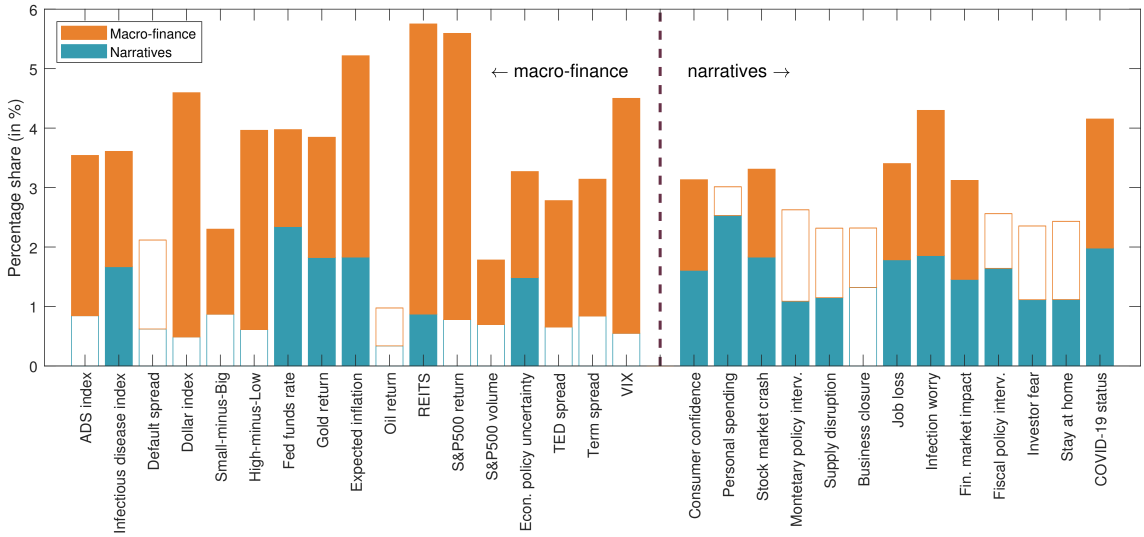

This figure depicts the amount of from-connectedness for the ith variable as a share of the total from-connectedness across all variables in the network. Orange bars indicates the amount from macro-finance variables for each variable and blue bars indicate that from narratives. Shaded bars indicate significance at the 10% level, and hollow bars indicate insignificance at the same level. Statistical significance is based on bootstrapped confidence intervals outlined in Appendix D. The purple dashed line separates the macro-finance and narratives groups.

Figure A.7: Individual variable from-connectedness: weekly horizon

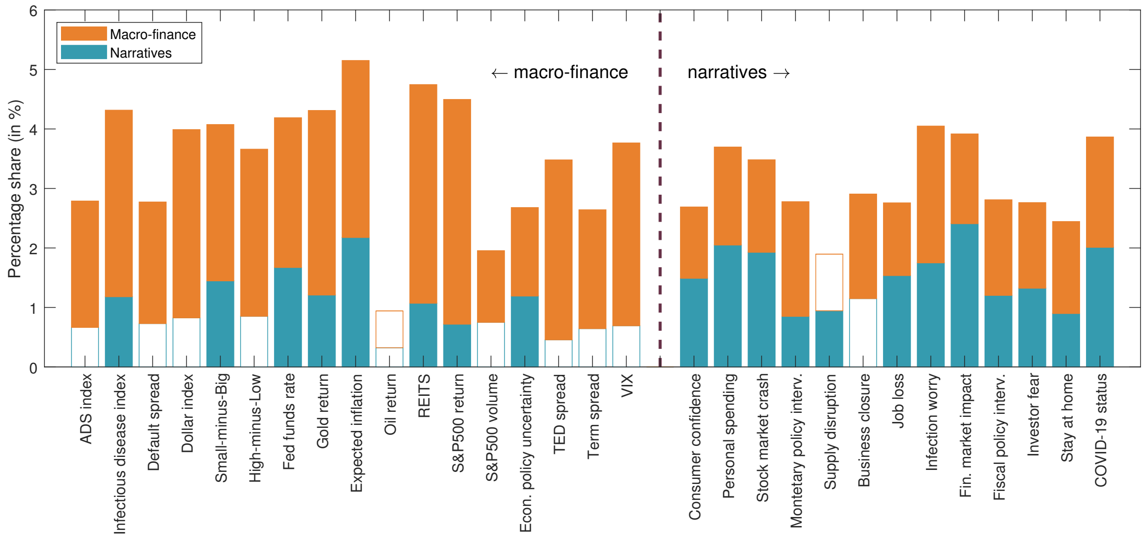

This figure depicts the amount of from-connectedness for the ith variable as a share of the total from-connectedness across all variables in the network. Orange bars indicates the amount from macro-finance variables for each variable and blue bars indicate that from narratives. Shaded bars indicate significance at the 10% level, and hollow bars indicate insignificance at the same level. Statistical significance is based on bootstrapped confidence intervals outlined in Appendix D. The purple dashed line separates the macro-finance and narratives groups.

H. Principal components for estimating risk premia

Like Giglio and Xiu (2021) we use a plot of the first twenty eigenvalues and the degree of cross-sectional explanatory power of the test assets to determine the dimension of the latent fundamental factor model. Figure A.8 depicts the first twenty eigenvalues of the covariance matrix of de-meaned returns from the 198 equity portfolios used as test assets. As is typical for large panels, the first eigenvalue tends to be much larger than the others, which is also why we zoom in on the eigenvalues five through 20 in the right panel of the figure. An important feature of the Giglio and Xiu (2021) is that the estimate of the risk premia γ_g is only consistent as long as the chosen dimensions is at least as large as true dimension. In other words, it is important to select a dimension that is not too small. With this in mind, we note that the eigenvalues are quite steadily decreasing with no notable kinks. The marginal contribution from each eigenvalue, however, tends to level off after about 12 principal components. Indeed, the time series variance explained by the first 12 principal components is 98%.

Figure A.8: First 20 eigenvalues of the covariance matrix of test assets

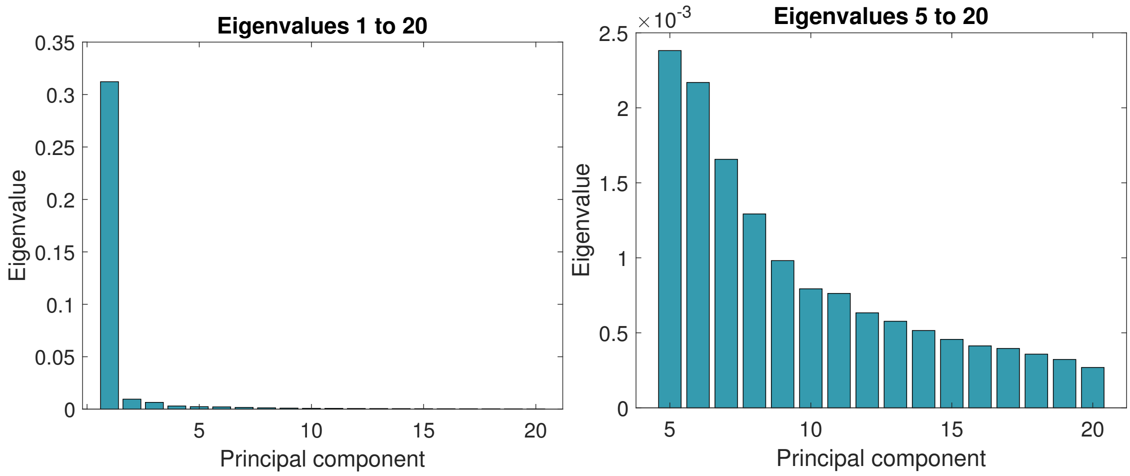

This figure depicts the first 20 eigenvalues of the covariance matrix of de-meaned returns from 198 equity portfolios as test assets. The right panel masks the first 4 eigenvalues and zoom in to those from five through 20.

To further guide us in choosing the dimension of the latent factor model, we depict in Figure A.9 the pricing ability of the average test asset returns as a function of number of principal components included. It is evident that at least ten factors are needed, though with the requirement that the dimension is not selected too small for consistency of $\hat{\gamma}_g$ we settle at 12 principal components. This explains about 56.5% of the cross-sectional variation in average returns. The contribution after the 12'th component again seem marginal. It is important to note that an explanatory of this

Figure A.9: Explained cross-sectional variation of test assets

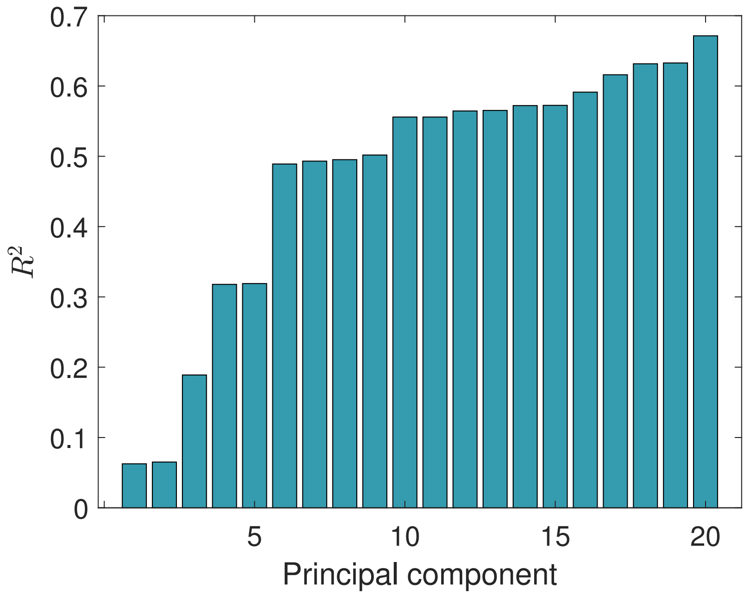

This figure depicts the cross-sectional R^2 from the second pass of the three-pass methodology of Giglio and Xiu (2021) that relates the a given number of principal components to the average returns of the test assets.

magnitude suggests a reasonable approximation of the latent factor model as it is close to numbers obtained by the three-factor Fama and French (1993) model only on the cross-section of 25 portfolios sorted by size and book-to-market (Giglio and Xiu, 2021), yet on a vastly greater cross-section, which is about eight times larger.

We note that our main findings presented in the paper are qualitatively similar when the dimension equals 10 and 14.