Does finance benefit society? A language embedding approach

Manish Jha, Hongyi Liu, and Asaf Manela

Washington University in St. Louis

October 2020

Abstract

We measure popular sentiment toward finance using a computational linguistics approach applied to millions of books published in eight countries over hundreds of years. We document persistent differences in finance sentiment across countries despite ample time-series variation. Finance sentiment declines after epidemics and earthquakes, but rises following droughts, floods, and landslides. These heterogeneous effects of natural disasters suggest finance sentiment responds differently to the realization of insured versus uninsured risks. Using local projections, we find that positive shocks to finance sentiment have positive and persistent effects on economic growth. Our estimates predict a contraction in finance sentiment due to the COVID-19 pandemic that will exacerbate its long-term economic damage.

Keywords: sentiment, natural disasters, text analysis, word embedding, COVID-19

We thank Yves-Paul Auffray, Oscar Iglesias, Jay Kloecker, David Lindequist, Rodrigo Moser, Andrea Paloschi, Rita Podledneva, Tatiana Vdovina, Juan Ignacio Vizcaino, and seminar participants at the Virtual Finance Seminar and Washington University in St. Louis for helpful comments. E-mail: mjha@wustl.edu, hongyi.liu@wustl.edu, or amanela@wustl.edu. Computations used the Washington University Center for High Performance Computing, which is partially funded by NIH grant S10 OD018091.

``As finance academics, we should care deeply about the way the financial industry is perceived by society. Not so much because this affects our own reputation, but because there might be some truth in all these criticisms, truths we cannot see because we are too embedded in our own world. And even if we thought there were no truth, we should care about the effects that this reputation has in shaping regulation and government intervention in the financial industry. Last but not least, we should care because the positive role that finance can play in society depends on the public's perception of our industry.''

(Zingales, 2015, AFA Presidential Address)

1 Introduction

Positive popular sentiment toward finance can spread its benefits widely, while suspicion toward financial services can restrict credit, risk-sharing, and competition (Zingales, 2012, 2015). Survey evidence reveals that trust in bankers fell sharply following the 2007–2008 financial crisis (Sapienza and Zingales, 2012), that such public perceptions often diverge from those of economists (Sapienza and Zingales, 2013), and that low trust can hinder insurance market efficiency (Gennaioli, Porta, Lopez-de-Silanes, and Shleifer, 2020). The relatively short time series of survey data restricts our understanding of how finance sentiment changes over time and differs across countries, how it responds to rare disasters like the currently spreading pandemic, and how such changes relate to economic and financial outcomes. While we cannot survey those who lived through the devastating wars and natural disasters of the 20th century, books allow us to travel through time and across borders, and to study public perceptions about the benefits of finance to society.

We measure popular sentiment toward finance in an annual panel covering eight large economies from 1870 to 2009 using a computational linguistics approach applied to the text of millions of books. Our finance sentiment index relies on a recently developed language model (BERT, Devlin, Chang, Lee, and Toutanova, 2018) to measure whether references to finance are, on average, semantically closer to positive versus negative words. BERT and its offsprings have shattered records on multiple natural language processing tasks, surpassing human ability on many. We use BERT to embed sentences into a relatively low dimensional numerical vector. Following Kozlowski, Taddy, and Evans (2019), we measure the angle between the embedding of sentences mentioning ``finance'' and the ``positive'' minus ``negative'' dimension. This approach goes beyond the dictionary or bag-of-words approach to sentiment analysis (Zhou, 2018) by capturing not only whether a book is positive or negative, but also the degree to which the *context* of the word ``finance'' is positive. By aggregating this positivity angle for all finance-mentioning sentences in each language and in each year, we construct a novel finance sentiment panel.

We find highly persistent differences in finance sentiment across languages. Generally, books written in languages of more capitalist countries tend to discuss finance in a more positive context. Russian finance sentiment is lowest by far throughout our long sample, followed by German, Italian, Chinese, French, and Spanish, with British and American English at the top. Despite considerable within-country variation, this ordering persists throughout our long sample, with the exception of British English finance sentiment, which is slightly higher than American English sentiment until 1912, and slightly lower thereafter. Interestingly, Chinese finance sentiment is about as positive as the French one, though more volatile, temporarily plummeting in 1971 when the People's Republic of China (PRC) is admitted into the United Nations (UN) then rising by a similar amount the following year when US President Nixon visits the PRC, and the Shanghai Communiqué (1972) is issued. Other significant changes in sentiment coincide with major historical events, like wars and revolutions. These major events could affect finance sentiment, but they could also be affected by it or jointly determined by other socio-economic changes. To better understand how the finance sentiment evolves, we analyze how it responds to plausibly exogenous natural disasters (Baker, Bloom, and Terry, 2020).

We document that finance sentiment declines by about 1 percent, one year after a country suffers a severe natural disaster. This average treatment effect, however, hides ample heterogeneity across disaster types. In particular, epidemics and earthquakes reduce finance sentiment by about 4 percent, while droughts, floods, and landslides increase it by 3, 2 and 5 percent, respectively. The effects of extreme temperature, storms, and smog are statistically no different from zero. These results hold controlling for wars, fatalities, and for country and year fixed effects. Thus, our panel allows us to overcome a common concern about cross-country comparisons that other sources of heterogeneity may be omitted (Guiso, Sapienza, and Zingales, 2004). The inclusion of year fixed effects also means these estimates are not driven by a single common shock such as the 1918 flu pandemic.

What explains these disparate effects? Disaster insurance data, available over the later part of our sample, suggests that epidemic and earthquake risks are largely uninsured by insurance companies, while extreme temperatures, floods, wildfires, and storms are relatively well covered by insurance. Thus, one potential answer, is that finance facilitates risk sharing of some types of risks through insurance, securitization or derivatives, but financial contracts and intermediaries are often designed to prevent ex-post renegotiation (Diamond and Rajan, 2001; Agarwal, Amromin, Ben-David, Chomsisengphet, Piskorski, and Seru, 2017). When insured disasters hit, economic costs are shared broadly across households and generations. But as the COVID-19 pandemic illustrates, when uninsured disasters strike (Walsh, 2020), their damage can be concentrated in parts of the population (Mongey, Pilossoph, and Weinberg, 2020), destroy fragile businesses (Chetty, Friedman, Hendren, and Stepner, 2020), and generate resentment against financial intermediaries (Scism, 2020). Another explanation is that insurance claim disputes affect finance sentiment. Gennaioli, Porta, Lopez-de-Silanes, and Shleifer (2020) show that insurance claims are frequently disputed and result in rejections or lower payments. Sentiment toward insurers may worsen if households learn they are uninsured only after disaster strikes.

The above findings beg the question, does finance sentiment matter? To answer it, we estimate impulse responses to finance sentiment shocks on GDP growth using local projections (Jordà, 2005). We study GDP growth, an imperfect measure of economic well being, simply because it is available for all countries in our panel. We find that a 1 percent improvement in finance sentiment leads to a gradual and persistent increase in GDP growth of about 20 basis points in each of the ten years following the shock. For a subset of countries that excludes China and Russia, we can include credit growth in the local projections. We find that some but not all of the positive effects of finance sentiment on GDP can be attributed to its positive effect on credit growth.

What do these results imply for the currently spreading COVID-19 pandemic? Assuming its death toll is no greater than the 1918 flu pandemic, our estimates predict a 4 percent contraction in finance sentiment after one year. Such a shock would exacerbate the direct negative effect of the health crisis on economic growth and further reduce cumulative GDP and credit growth by 4 and 5 percentage points, respectively, over the next 5 years.

Our paper relates to recent work on the measurement of public attitude toward the financial sector. Stulz and Williamson (2003) find that a country's language and religion predict its creditor rights. Guiso, Sapienza, and Zingales (2008) find that a general lack of trust reduces stock market participation. Giannetti and Wang (2016) document that after the revelation of corporate fraud in a state, household participation, and trust in the stock market decreases. D'Acunto, Prokopczuk, and Weber (2019) find that present-day demand for finance is lower in German counties where historical antisemitism (and therefore distrust in finance) was higher. Gurun, Stoffman, and Yonker (2018) find that communities indirectly exposed to a Ponzi scheme withdraw assets from investment advisers. Levine, Lin, and Xie (2019) link the African slave trade to household demand and trust of financial services. We contribute to this work by providing a novel measure of sentiment toward finance that spans over a century and several large economies, and documenting how finance sentiment is shaped by natural disasters.

Natural disasters and their effects on the economy are of great interest since the onset of the COVID-19 pandemic. Eisensee and Stromberg (2007) study disaster relief and news coverage. Baker, Bloom, and Terry (2020) use natural disasters as instruments for stock market uncertainty. Jordà, Singh, and Taylor (2020) document persistent declines in real rates of return and increases in wages after pandemics. Closely related is Aksoy, Eichengreen, and Saka (2020), who find that epidemic exposure in an individual's impressionable years negatively affects their confidence in political institutions and leaders.

A broader related literature considers the measurement of culture and its effects on economic outcomes (Guiso, Sapienza, and Zingales, 2006). Cultural differences can persist for generations (Spolaore and Wacziarg, 2013). Changes in culture, ideas, and in particular language, have been tied to the dramatic enrichment the world experienced starting in the 19th century (Mokyr, 2016; McCloskey, 2016). It remains unclear, however, exactly why cultural changes occur (Guiso, Sapienza, and Zingales, 2015). Our results about natural disasters provide one plausibly exogenous cause for such cultural changes.

A recent increase in the availability of textual data has prompted great interest in its use for analysis of culture in particular (Michel, Shen, Aiden, Veres, Gray, Pickett, Hoiberg, Clancy, Norvig, Orwant, et al., 2011), and in economics and finance more broadly (Gentzkow, Kelly, and Taddy, 2019; Loughran and McDonald, 2020). While this literature has yet to study sentiment toward finance or any particular sector, textual analysis has been used to analyze partisanship (Gentzkow and Shapiro, 2010; Luo, Manconi, and Massa, 2020; Goldman, Gupta, and Israelsen, 2020; Engelberg, Henriksson, Manela, and Williams, 2019), product markets (Hoberg and Phillips, 2016), central bank communication (Hansen, McMahon, and Prat, 2018; Cieslak and Vissing-Jorgensen, forthcoming), corporate culture (Grennan, 2019), asset market sentiment (Antweiler and Frank, 2004; Tetlock, 2007; García, 2013; Soo, 2018; Ke, Kelly, and Xiu, 2019), employee expectations (Sheng, 2019), financial constraints (Bodnaruk, Loughran, and McDonald, 2015), subjective wellbeing (Hills, Proto, Sgroi, and Seresinhe, 2019), uncertainty (Baker, Bloom, and Davis, 2016; Manela and Moreira, 2017; Goetzman, Kim, and Shiller, 2017; Hassan, Hollander, van Lent, and Tahoun, 2017; Boudoukh, Feldman, Kogan, and Richardson, 2018), and emerging risks (Hanley and Hoberg, 2019; Bybee, Kelly, Manela, and Xiu, 2019).

While early work relied on simple word counts (the bag-of-words approach), recent work starting with Mikolov, Chen, Corrado, and Dean (2013) shows that using neural networks to embed words in vector spaces improves learning algorithms' performance in natural language processing tasks. Kozlowski, Taddy, and Evans (2019) demonstrate that such word embeddings produce richer insights into cultural associations and categories than prior methods. Our work builds and improves on their methodology by using a pre-trained language model designed to capture context (BERT), both to embed sentences mentioning our object of interest (finance) and to define the dimension on which we project these embeddings (positive − negative). This ``transfer learning'' approach lowers both estimation error and computation costs.

We proceed as follows. Section 2 describes our text-based finance sentiment measure and how it evolves over time and across countries. Section 3 studies how natural disasters affect finance sentiment. Section 4 analyzes how finance sentiment relates to economic growth. Section 5 concludes. Additional results are provided in an online appendix.

2 Text-based sentiment toward finance

In this section, we describe our text data and how we measure a language's sentiment toward finance across time. For each language and year, we start with a sample of finance-mentioning sentences published in the language and year. Next, we measure the degree to which each sentence places finance in a positive context. We then aggregate these scores to an average finance sentiment that reflects the mean sentiment toward finance of books written in the language in that year.

We assume throughout that the choice of words used by book authors, magazine publishers, and journalists whose work is archived in libraries, reflects the sentiment of the average denizen of that language during the time, or at least that of an influential literary elite. For example, in our dataset, the sentence "correcting corruption or financial malpractice" appears first in 1951 and then appears every year after 1959. The sentence was part of the 1959 US labor management reform legislation hearings, when correcting corruption or financial malpractice became an allowable purpose for establishing a trusteeship by labor unions. Hence, ``financial malpractice'' is more frequently used in subsequent legal documents and books. The context for the word ``financial'' here is clearly negative. In this particular case, we assume that the labor unions in particular, and the US English-speaking public in general, are more likely to associate finance with malpractice around that time.

2.1 Data

Our text data includes five-word sentences (5-grams) containing the word "finance" across eight languages, during 1870–2009, extracted from the 2012 edition of the Google Books Ngram Corpus (Michel, Shen, Aiden, Veres, Gray, Pickett, Hoiberg, Clancy, Norvig, Orwant, et al., 2011; Lin, Michel, Aiden Lieberman, Orwant, Brockman, and Petrov, 2012). The corpus consists of words and phrases and their annual usage frequency from 1500 to 2009. The data originates from Google scanning over 8 million books or 6% of all books ever published in American English, British English, Simplified Chinese, French, German, Italian, Russian, and Spanish.[1]

Although the original data provides lower complexity n-grams counts as well, we focus on 5-grams because for sentiment analysis, especially with BERT, a word's context is essential. We start our study in 1870 (Google corpora is available from 1500) because

[1] We reluctantly omit Hebrew because its word for finance (Mimun) without niqqud is also the name of Maimonides---a famous Jewish philosopher (Rabbi Moshe ben Maimon).

from that year, we have more confidence in the accuracy of our macro data. Moreover, the number of sentences becomes sparser as we go back in time, and there are fewer mentions of finance before 1870, which increases the measurement error of our sentiment index.

Table 1: Finance mentions across languages

Language	Finance word stem	Unique sentences	Total sentences
American English	financ	220k	79m
British English	financ	48k	15m
Simplified Chinese	金融, 金 融, 金_融	196k	305m
French	financ	100k	43m
German	finanz	28k	7m
Italian	finanz	23k	9m
Russian	финан	187k	250m
Spanish	finan	89k	33m

Note: We report the number of mentions of the word finance, translated and stemmed, in a five-word sequence (5-gram) for each language in the Google Book Ngram Corpus. Our dataset covers the period 1870–2009.

We preprocess the Book Corpus by stripping case, symbols, double spaces, part of speech tags, and positional tags. Next, we extract all sentences mentioning the stem of the word for finance. The finance stem word is different across languages, as listed in Table 1. We use the word stem "financ" for English to include sentences that contain either "finance" or "financial." Similarly, for other languages, we use a word stem common to the different verb and noun forms of "finance." For example, for Simplified Chinese we use "金融" (financial) but also include base words where there is space and underscore between 金 (gold) and 融 (melt). The filtering yields a set of unique sentences mentioning finance for each language. In our data set, American English has the highest number of unique sentences that mention "financ", followed by Simplified Chinese and Russian. Although the simplified Chinese is issued starting from the 1950s, the Google ngram data for the Chinese has been translated to the simplified version throughout the whole dataset.

2.2 Methods

We measure finance sentiment across languages at an annual frequency. We employ a three-step process to measure the finance sentiment (i) embed each sentence in our corpus into a 768-dimensional vector space (ii) measure the cosine similarity of this sentence embedding with respect to positive minus negative embedding (iii) average the cosine similarity of all finance mentioning sentences in each year, weighted by their frequency. We next describe how we calculate the sentence embedding, the positive minus negative embedding, and their cosine similarity.

2.2.1 BERT

Recent work in natural language processing (NLP) has been increasingly successful in capturing the complexity of language by considering words in sequence rather than in isolation. One of the ways this is accomplished is by representing words as embeddings. Word embeddings are high-dimensional vector-space models of text in which each unique word in the corpus is represented as a vector in a shared vector space (Mikolov, Chen, Corrado, and Dean, 2013). The vector for each word is based on the context the word shares with other words in the sentence. The classic flavors of word embeddings, such as Word2Vec (Mikolov, Chen, Corrado, and Dean, 2015), GloVe (Pennington, Socher, and Manning, 2014), and FastText (Bojanowski, Grave, Joulin, and Mikolov, 2016; Joulin, Grave, Bojanowski, and Mikolov, 2016) rely on the Distributional Hypothesis (Harris, 1954) to capture relationships in the embedding space. The hypothesis states that words that occur in the same contexts tend to have similar meanings, with the underlying idea that "a word is characterized by the company it keeps" popularized by Firth (1957). However, there are certain downsides with these flavors. First, the traditional methods assign embeddings from the ground up; this is an issue for our data set in earlier years, where the number of words in corpus is less than a million (Altszyler, Sigman, Ribeiro, and Slezak, 2017).

Second, while these embedding methods work well for word-level embeddings, they are poor sentence encoders, which extend the word embedding approach to sentences. (Cer, Yang, Kong, Hua, Limtiaco, John, Constant, Guajardo-Cespedes, Yuan, Tar, Sung, Strope, and Kurzweil, 2018; Perone, Silveira, and Paula, 2018). Thus, we move away from the traditional shallow neural network methods.

We employ a deep neural network-based natural language processing method, Bidirectional Encoder Representations from Transformers (BERT) developed by Devlin, Chang, Lee, and Toutanova (2018). BERT produces meaningful results even with smaller training data and can provide context for words in sentences. The key advantage of this method over classic word vector models is transfer learning – where a model developed for a task is reused as the starting point for a model on a second task. BERT's neural network is pre-trained on 800 million BooksCorpus and 2,500 million Wikipedia words. Thus, the model knows which words have a similar meaning, based on pre-training. BERT is a state of the art NLP model, and Google applies it to both rankings and featured snippets in search. BERT is expected to improve around 10% US English search queries currently, and Google is bringing it to other languages soon.[2] BERT surpassed human performance on the reading comprehension questions provided by the Stanford Question Answering Dataset (SQuAD).

While a full treatment of BERT is beyond our scope, we wish to provide an intuitive understanding of this method and the structure that it implicitly imposes on the data. BERT uses Transformers (Vaswani, Shazeer, Parmar, Uszkoreit, Jones, Gomez, Kaiser, and Polosukhin, 2017), a mechanism that learns contextual relations between words in a text. The model processes each word in relation to all other words in a sentence, rather than one-by-one in order. BERT is also bidirectional, which allows the model to learn the context of a word based on all its surroundings, as opposed to a directional model, which reads the text sequentially. To train the model from unlabeled text from BooksCorpus and Wikipedia

[2]https://www.blog.google/products/search/search-language-understanding-bert/

text, BERT employs two strategies: (i) Masked Language Modeling – where 15% of the input words are masked out and then predicted (ii) Next Sentence Prediction – predict if Sentence B is the actual sentence that proceeds Sentence A. Solving the above two problems using its large corpora, BERT is able to place words in the embedding space. Google shares two versions of the pre-trained model: Base (12-layer, 768-hidden features) and Large (24-layer, 1024-hidden features). Both models are available in a cased and uncased variant. We use the base uncased model for English, and Chinese since the extra efficiency we get from the large and cased model is not significant enough to spend more time and resources on them. For French, German, Italian, Russian, and Spanish, we use the cased multilingual model as recommended by Google Research. Thus, we use BERT Base for American and British English, BERT Base Chinese for Simplified Chinese and BERT Base Multilingual Cased for French, German, Italian, Russian and Spanish.[3]

The following features of BERT make it especially useful for our purposes: First, BERT comes pre-trained, so it works well out of the box. A pre-trained model is important to us, especially in earlier years of our sample, where the Google Books corpus is considerably smaller. Second, it offers contextualized embedding. For example, the word "bank" has a different meaning in the following two sentences "In a crisis, we could bank on financing from the government", and "Government's financing for the bank is in crisis." In the first, it means "to rely upon," while in the second, it refers to a financial institution. The context changes how the author feels about the situation. BERT could distinguish the connotation difference between the two sentences, resulting in different embeddings. By contrast, in classical word vector models, where each word has a unique embedding. Third, to reduce the number of unique words that feature in the model, BERT breaks each word into smaller subwords or tokens. For example, “wonderful“ is tokenized to “won #der ##ful,“ where # denotes subwords. The dimensionality reduction is especially important in the multilingual model. Finally, BERT is designed to encode entire sentences, up to 512 subwords.

[3] The pre-trained models are available at https://github.com/google-research/bert

The tokenization process adds [CLS], which stands for "classification" at the beginning of each sentence. The embedding for [CLS] is used as the embedding for the entire sentence that follows it.

2.2.2 Cosine Similarity

A major advantage of word embeddings is that they allow language features (such as words, sentences, etc.) to be treated like vector spaces with intuitive mathematical properties. A common example from Mikolov, Yih, and Zweig (2013) is king – man + woman ~ queen. That is, subtracting the male gender vector and adding the female gender vector to the king vector corresponds to a vector that is close to the queen vector. Thus the word queen could be seen as starting at the word king and then moving in the feminine gender direction. Similarly, we could think of dictator + positive - negative ~ king; here, positive minus negative represent a displacement in the positive direction. Thus, if we start from the dictator vector and move a step in the positive direction, we get the king vector. Other word pairs also correspond to the positive dimension, such as (benefit − damage), (good − bad), (good − corrupt), and (help − hurt).

To define our positive minus negative dimension, we average the sentence embedding differences across sentences containing ``finance'' or ``financial'' together with the above words, similar to Kozlowski, Taddy, and Evans (2019). The list of sentences for English (both American and British) are shown in Table 2. The corresponding sentence pairs for other languages are included in Appendix A.1. We focus on the broader notion of ``finance'', as opposed to more specific financial activities or players (e.g. ``bank'', ``lender'', etc.), because this sentiment measure speaks directly to our question of interest, attitude toward finance. More specific related words would be close in vector space to ``finance'' because they are frequently mentioned together, so we expect them to generate similar sentiment estimates, but each brings along its own identification issues. For example, ``bank turmoil'' can often refer to the financial institution but also to the contested West Bank territory.

Table 2: Positive $-$ negative defining sentences for English

Positive sentences	Negative sentences
financial services benefit society	financial services damage society
finance is good for society	finance is bad for society
finance professionals are mostly good people	finance professionals are mostly corrupt people
finance positively impacts our world	finance negatively impacts our world
financial system helps the economy	financial system hurts the economy

Note: To define the positive minus negative dimension, we average the embeddings of positive sentences less that of their negative counterparts.

To measure sentiment toward finance, for each finance-mentioning sentence j in language i with embedding s_{ji}, we calculate the orthogonal projection of the sentence vector onto the language-specific positivity embedding p_i using cosine similarity:

$$a_{ji} = \frac{s_{ji} \cdot p_i}{|s_{ji}||p_i|} = \frac{\sum_d s_{jid} p_{id}}{\sqrt{\sum_d s_{jid}^2}\sqrt{\sum_d p_{id}^2}}, \tag{1}$$

where d enumerates the elements of s_{ji} and p_i, both 768-dimension vectors. By construction, the cosine similarity in Equation 1 of two positive vectors is bounded between -1 and +1, with zero indicating a neutral sentence. A more negative cosine similarity indicates that the sentence has a more negative sentiment, while a more positive cosine similarity indicates a more positive sentence.

Figure 1 illustrates this method in a two-dimensional space. The five positive (and negative sentences), from Table 2 bunch together in the embedding space, as similar sentences keep similar companies Firth (1957). We take the vector difference between positive and negative sentence embeddings to define our positivity dimension. Next, we project finance sentences onto the positive minus negative dimension. Sentences tend to be close to the dimension, which is closer to their connotation. For example, a sentence such as "financial sector supports economic development" lies closer to the positive sentences, at a smaller angle with the positive dimension.

Cosine similarity measures the position between -1 and 1 where the shadow of a given sentence vector falls. If the sentence has a positive connotation, such as the one in our example, we will have a smaller angle between the sentence vector and the positive dimension. A smaller angle is associated with a higher cosine similarity. On the other hand, for a negative sentence such as "financial malpractices stunted our growth" would be closer to the negative dimension, or $\theta_{ij} > 90°$. Thus the cosine similarity for a sentence with negative connotation is negative. Similarly, a neutral sentence such as "finance lessons from the pandemic" would be equidistant from both positive and negative dimensions ($\theta_{ij} \approx 90°$), and thus a cosine similarity close to zero. Table 3 lists the sentences with the most positive and most negative finance sentiment for American English. Appendix A.2 provides similar lists for all languages.

Table 3: Sentences assigned the most positive and negative finance sentiment for English

Positive sentiment sentences	Negative sentiment sentences
financial support of the science	turmoil in the financial markets
financial management of the school	instability in the financial markets
financial support of the research	lack of money to finance
financial management of the business	a financial panic
financial support of this project	the financial panic
financial management initiative	financial panic in the united
financial support of the work	international financial instability
understanding of the financial system	lack of funds to finance
finance for small and medium	my finances falling short
finance graduate school of	the financial deficit

Note: A sentence is assigned positive or negative finance sentiment, based on its projection onto the finance positivity dimension (cosine similarity). Sentences at the top are the most positive or negative in their respective column, and the absolute value of finance sentiment decreases down each list.

Finally, we calculate an annual finance sentiment for each language i based on the cosine similarity of all finance-mentioning sentences that occurs in that language in each

Figure 1: Conceptual diagram of finance sentiment measurement

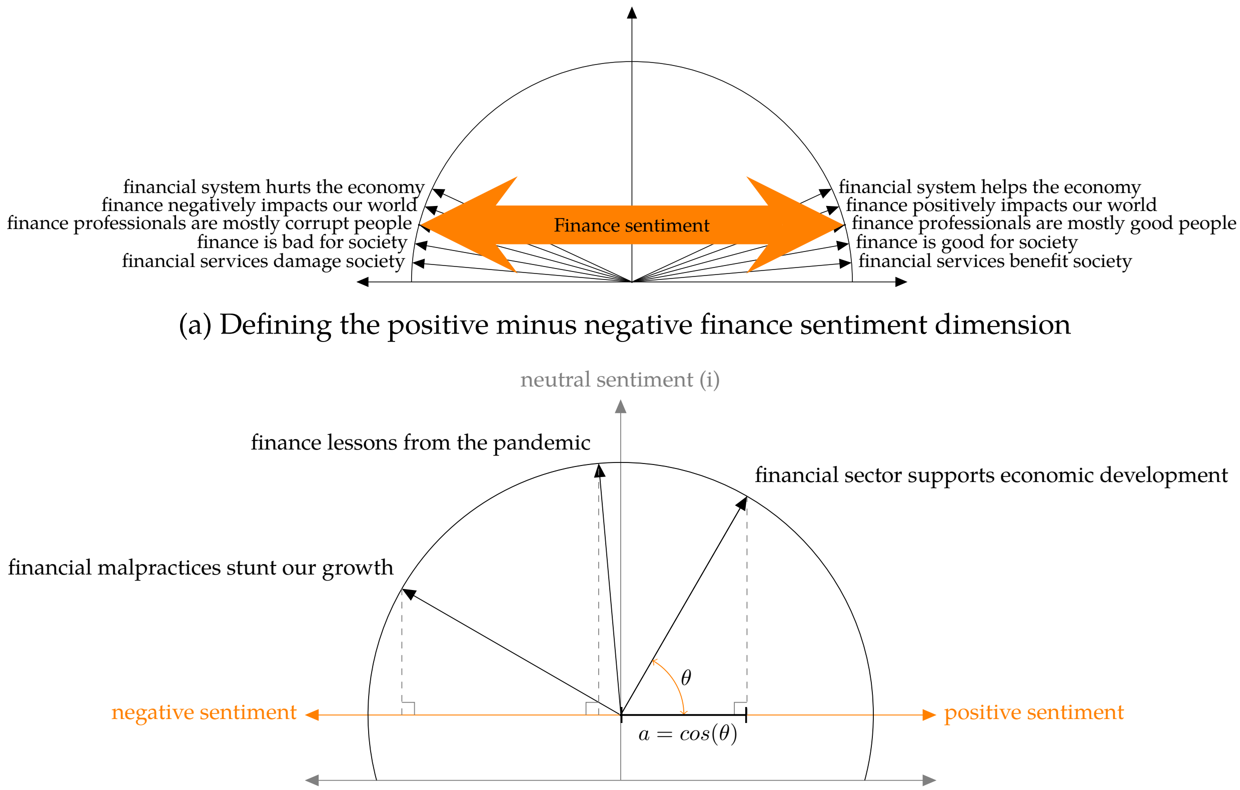

(a) Defining the positive minus negative finance sentiment dimension

(b) Projection of sentences onto positive minus negative sentiment dimension

Note: Panel (a) shows a conceptual diagram of how similar sentences aggregate to a two-dimensional embedding space. We take the vector difference between positive and negative embeddings to define the finance sentiment dimension. Panel (b) illustrates the classification of three example sentences by projecting them onto this dimension. For one of the sentences, we illustrate cosine similarity, defined as the cosine of the angle between two vectors. Sentences that are close in terms of meaning have a smaller angle between them in this vector space, thus higher cosine similarity. Positive finance sentences have a smaller angle to the positive dimension and a larger positive projection on the finance sentiment dimension.

year t, weighted by the number of times the sentence occurred that year,

$$f_{it} = \sum_j a_{ji} \times \frac{c_{jit}}{\sum_k c_{kit}}. \tag{2}$$

The frequency weighted f_{it} varies over time only because of changes in sentence occurrence c_{jit}, because the sentiment of particular sentences a_{ji} in each language i, does not vary over time. This is an important distinction from the approach of Kozlowski, Taddy, and Evans (2019), who train a language model for each language in each year, and then measure the orthogonal projections based on these year-specific models. While their approach may be more robust when very large amounts of data are available throughout the sample, our approach is more efficient and avoids issues with measurement error in small samples, which are particularly acute early in the Google Books corpus. Computationally, our approach is considerably cheaper, because training neural networks like the one behind BERT is still fairly expensive.

The cost of this reduction in measurement error and computation cost is that we implicitly assume that the language model is constant over time and that only the frequency of language use varies over time. This is obviously not exactly right. Languages evolve. We lack the data to measure the extent to which such changes in language matter for our conclusions. Encouragingly, evidence from US newspapers over a similar period suggests that changes to the English language do not affect much the ability to predict with text (Manela and Moreira, 2017).

We calculate finance sentiment for every year from 1870 to 2009 for American English, British English, French, German, Italian, Russian, and Spanish. We have a shorter 95-year sample for Chinese because prior to 1922, its corpus is highly sparse and most years feature no mentions of finance. These languages can all be traced to a major geographical area, centered around a distinct country, throughout most of our sample. For example, the concentration of Russian speakers is highest in Russia. Therefore, in what follows, we refer to the finance sentiment of these languages and countries interchangeably, but note that this requires a modest leap of faith. We expect it to introduce more error into our measurement toward the end of our sample, when Spanish books, for example, may be published in Latin American countries whose economic condition is no longer highly correlated with that of Spain.

Another caveat to our finance sentiment index is that it is based on published books, which may not represent the average citizen, especially early in our sample, when large parts of the world were illiterate. As a result, we may miss marginalized and under-represented groups of the population. Nonetheless, this ``literary elite'' has historically commanded a disproportionately large share of wealth, power, and exerted considerable influence on the opinions of the rest of society.

2.3 Finance sentiment over time and across countries

Table 4 describes finance sentiment over our sample. We see that sentiment toward finance in languages spoken in more capitalist countries tends to be above that of communist countries. In our sample, American English, on average, is at the very top followed closely by British English. The next set of languages that follow are Spanish, French, Chinese, and Italian. German and Russian are the two languages with an overall negative connotation for finance. The standard deviation is highest for Russian, followed by Spanish and French.

Figure 2 plots the finance sentiment and shows some salient features. American English has the most positive sentiment towards finance after 1912; before that, it was slightly below British English. The trend across the language is of an improving finance sentiment across time, with a slight dip at the very end in 2007–2008. A possible reason for that could be the great recession, whose impact could be felt across the globe. We see a 5% drop in US finance sentiment in 1874, a year after the Panic of 1873, which triggered economic depression in Europe and North America. A similar decline in finance sentiment is in 1896, after the country's gold reserves had dwindled and saved by JP Morgan's, and the Rothschild's gold loan. We see an increase in sentiments in 1885 and 1887, after labor union strikes, which eventually led to the eight-hour workday.[4]

Languages do not seem to cross each other, apart from Chinese, which exhibits significant changes in finance sentiment over time. This volatility is in line with historical events.

[4]See Wikipedia for historical events mentioned in this section.

Table 4: Finance sentiment and other summary statistics

Country (language)	Variable, %	Mean	Std. Dev.	Obs.
China	Finance sentiment index	7.5	0.5	95
(Chinese)	Finance sentiment growth	0.1	7.9	88
	GDP growth	3.2	7.1	119
France	Finance sentiment index	7.6	0.7	140
(French)	Finance sentiment growth	0.2	1.3	139
	GDP growth	1.9	6.4	139
	Credit growth	4.9	12.9	101
Germany	Finance sentiment index	-4.9	0.5	140
(German)	Finance sentiment growth	0	4.4	139
	GDP growth	2.1	8.1	139
	Credit growth	8.9	17.8	129
Italy	Finance sentiment index	4.4	0.4	140
(Italian)	Finance sentiment growth	0.4	5.2	139
	GDP growth	2	4.7	139
	Credit growth	6.1	13.9	139
Russia	Finance sentiment index	-11.9	0.8	140
(Russian)	Finance sentiment growth	0.1	1.5	139
	GDP growth	2	8.5	139
Spain	Finance sentiment index	8.3	0.7	140
(Spanish)	Finance sentiment growth	0.2	3.4	139
	GDP growth	2.1	5	139
	Credit growth	7.4	11.1	98
UK	Finance sentiment index	14.2	0.3	140
(British English)	Finance sentiment growth	0	1.5	139
	GDP growth	1.5	2.9	139
	Credit growth	4	8.2	129
US	Finance sentiment index	14.5	0.6	140
(American English)	Finance sentiment growth	0.1	1.3	139
	GDP growth	2.1	5	139
	Credit growth	4.5	6.7	129
Total	Finance sentiment index	4.8	8.7	1075
	Finance sentiment growth	0.2	3.7	1061
	GDP growth	2.1	6.2	1092
	Credit growth	5.9	12.5	725

Note: The sample spans from 1870 to 2009 for 8 country-language pairs. The corpus of sentences for each language is gathered from the Google Book Ngram Corpus. The connotation for each finance-mentioning sentence is measured based on its cosine similarity with respect to the positive minus negative vector. Finance sentiment for each year is the weighted average of the cosine similarity, weighted by the frequency of sentences in the language corpus for the year. GDP and credit data are from Jordà, Schularick, and Taylor (2017) and Barro and Ursua (2010) when available.

Figure 2: Sentiment toward finance

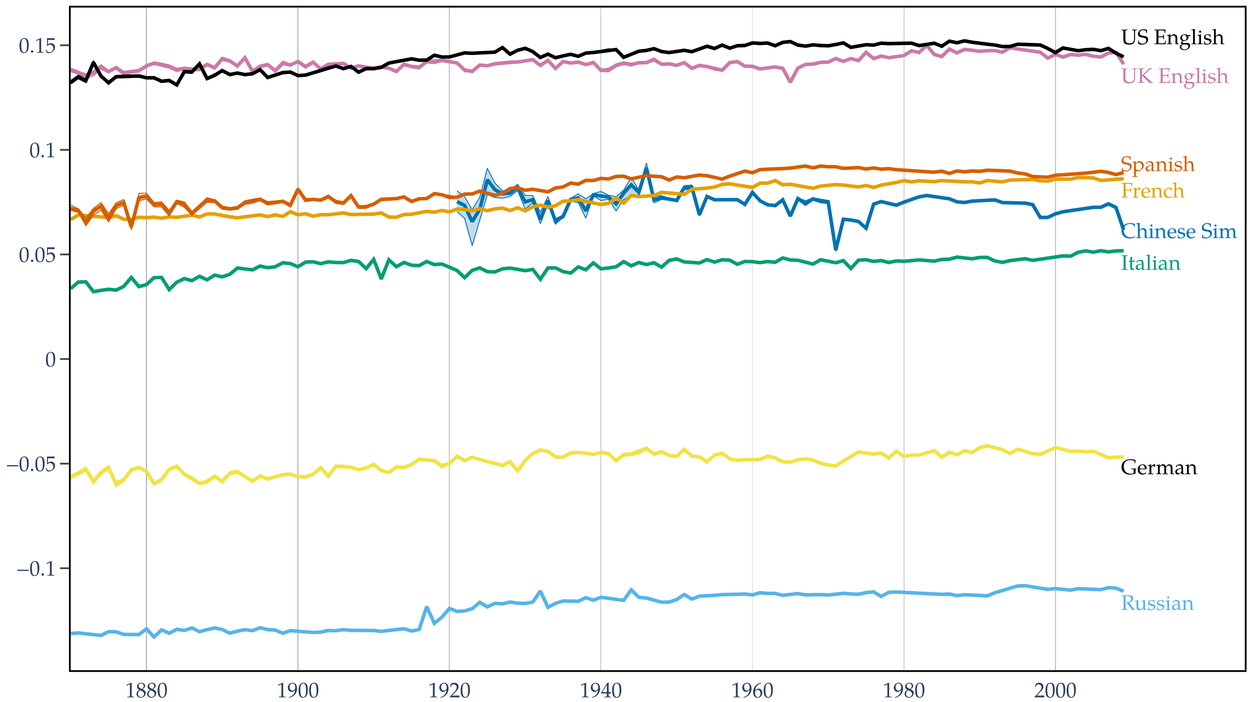

Note: Finance sentiment is based on the annual average projection of finance-mentioning sentences' embeddings onto the positive minus negative finance sentiment dimension. Sentences are from the Google Books Ngram corpus and embedded using BERT. Bands represent 95 percent confidence intervals produced by subsampling.

Chinese finance sentiment plunges in 1971 by 31%, just after Mao suggested the end of the Cultural Revolution. However, there is an uptick of 28% in 1972, one year after the United Nations recognized the People's Republic of China as "the only legitimate representative of China", followed by a visit from US President Nixon. We also see an 18% increase in 1976, a year after the constitution of the People's Republic of China was formalized.

The three Romance languages in our sample, French, Italian, and Spanish, have similar attitude towards finance. Spanish has the most favorable view, followed closely by French. We see higher volatility and an uptick in Spanish finance sentiments at the start of the 1874 Bourbon Restoration, which restored the monarchy. French finance sentiment is more volatile during World War II, with the most significant drop of 3% in 1943 when the French surrendered to Germany. The highest surge in French finance sentiment is in 1944, the year Paris was liberated. Sentiment dips and recovers for Italian in 1911–1912 at the start of the Italo-Turkish war, then rises in 1933 by 14%, when Fascist membership becomes compulsory for University teachers, prompting more favorable and nationalistic literature.

The two languages in which we find a negative finance sentiment are German and Russian. Similar to Italian, we see finance sentiment becoming more positive as the Nazi party gains power in Germany. Finance sentiment increases by 9% in 1930, the year the Nazi party gained its first minister. For Russia, we see a permanent increase in finance sentiment in 1917, coinciding with the Russian revolution. We also see a permanent increase at the beginning of 1990s after the collapse of the USSR, as Russian speaking countries adopt a more capitalist system. The largest drop for Russian finance sentiment is in 1933, a year after the Soviet famine of 1932–1933.

Given the positive trend apparent from Figure 2, our analysis below focuses on finance sentiment growth $\Delta f_{i,t}$, which characterizes the relative change of finance sentiment towards either positive or negative change direction given the absolute value of the previous year's sentiment for country i and year t:

$$\Delta f_{it} = \frac{f_{i,t} - f_{i,t-1}}{|f_{i,t-1}|} \times 100. \tag{3}$$

Figure 3 plots finance sentiment growth for each country in our panel. It shows more clearly that China exhibits the greatest volatility (7.9), followed by Italy (5.2) and Germany (4.4) (see Table 4 for summary statistics). We can also see that sharp changes in finance sentiment growth tend to partially reverse within a year. We formally investigate this pattern in Section 4.

A potential concern may be that it is not just finance-specific sentiment that is changing over time and across countries, but rather sentiment more generally. To measure general sentiment for a language, we use the sentiment associated with the fairly generic word "January" across time, following Gentzkow, Glaeser, and Goldin (2004), who use it to deflate for changes in newspaper reporting volume over time. Appendix Figure 7, shows this

Figure 3: Finance Sentiment Growth

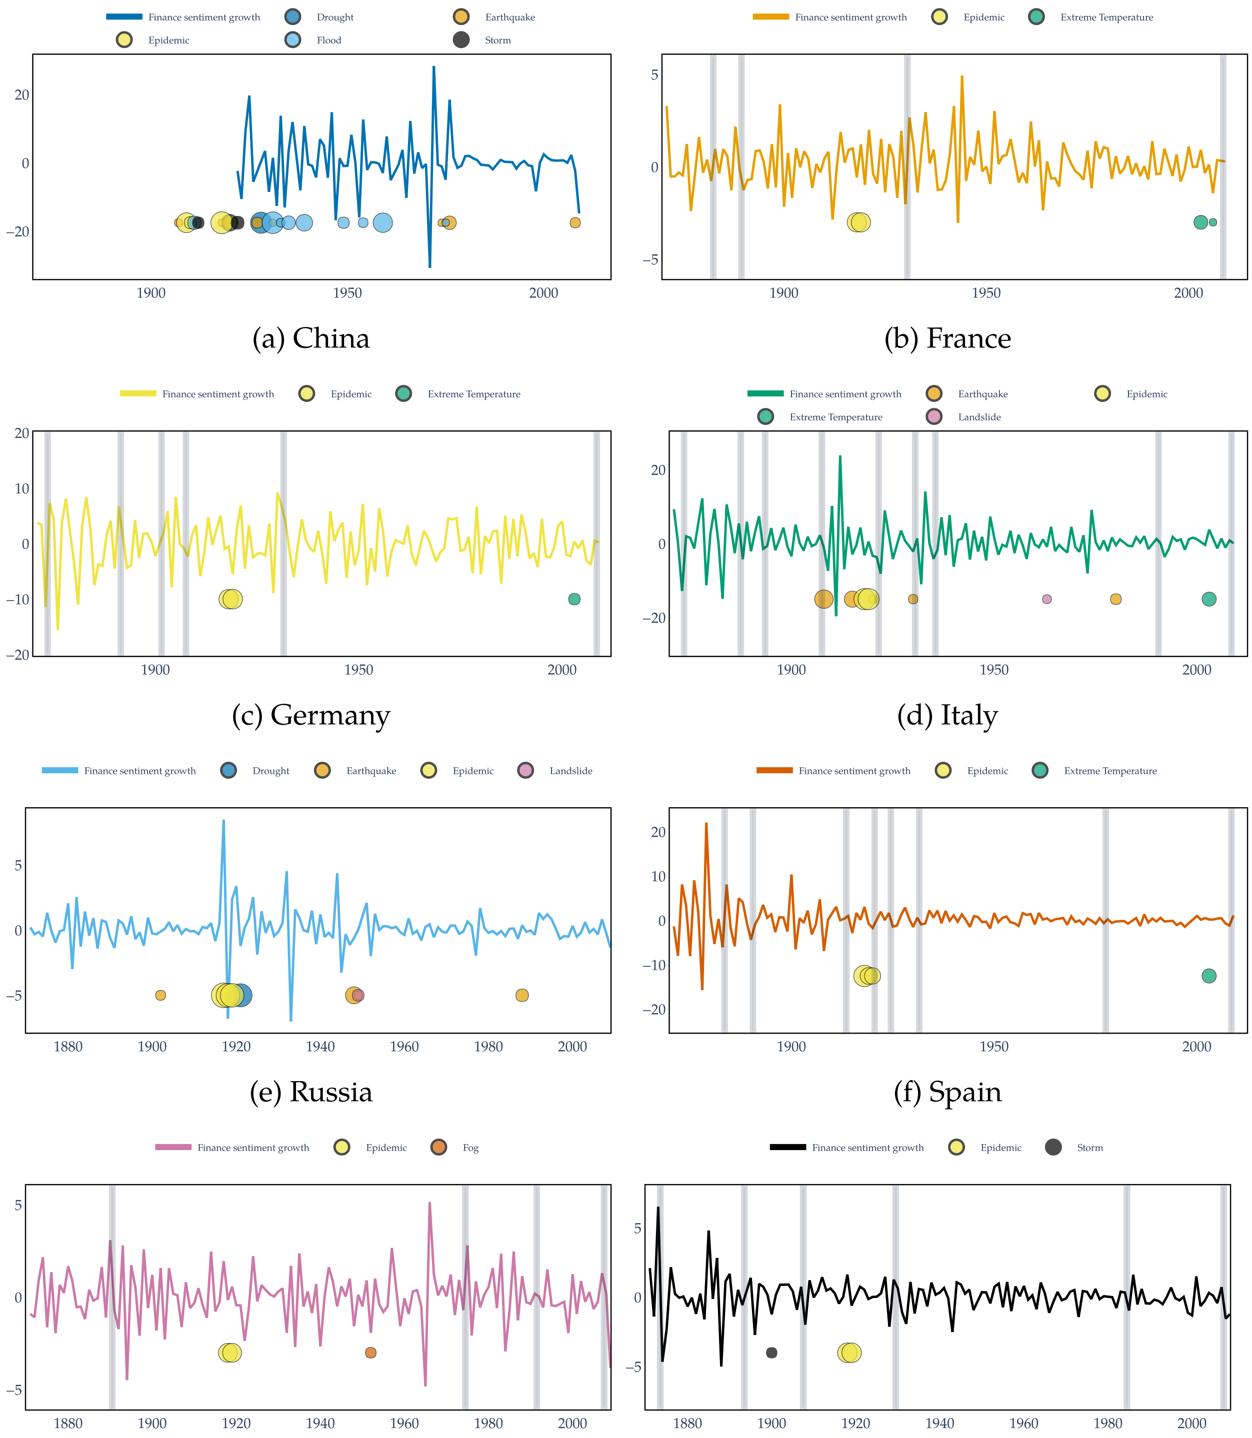

(a) China

(b) France

(c) Germany

(d) Italy

(e) Russia

(f) Spain

(g) UK

(h) US

Note: The figure shows finance sentiment growth during the period 1870–2009 for each country. Shades indicate financial crises, as defined by Jordà, Schularick, and Taylor (2017) (not available for China and Russia). Severe natural disasters are indicated by circles whose size is proportional to log deaths.

general sentiment. The general sentiment is quite flat across years, without an upward or a downward trend, for each language. The flat general sentiment suggests that languages have not changed drastically across years. Furthermore, it points to changes in peoples' perception of finance as the underlying cause of the upward trend apparent in Figure 2.

3 Natural disasters affect finance sentiment

We next study how natural disasters and wars affect sentiment toward finance. We first introduce our natural disasters sample and define severe natural disasters. In the next subsection, we present our empirical model consisting of heterogeneity across disaster types and discuss our results.

3.1 Disaster data

Natural disasters data are from the Emergency Events Database (EM-DAT) maintained by the Centre for Research on the Epidemiology of Disasters (CRED). EM-DAT records an event as a disaster if it kills 10 or more people, if it affects 100 or more people, or if there is a formal declaration of a state of emergency or an appeal for international assistance.

To match with our text data, we extract an 8-country subsample of data from EM-DAT dataset, including mainland China, France, Germany, Italy, Russia, Spain, UK and US. Following Eisensee and Stromberg (2007), we focus on natural disasters and omit complex disasters (e.g., famine) and technological disasters (e.g. coal mine collapse), which are likely human-made. We manually add the death tolls caused by the 1918 Flu Pandemic when missing.[5]

[5] The estimated death number from the 1918 Flu Pandemic for each country comes from the academic article and/or the CDC (Center of Disease Control) website or the Spanish flu feature news. Due to the difficulty of accurately determining the real death number caused by the Spanish flu during the period of 1918 to 1919, we distribute the estimated death number equally to each year for those countries where EM-DAT does not specify the death toll.

Table 5: Natural Disasters Summary Statistics

Disaster Group	Type	Obs.	Severe	Mean Killed	Damage, $M	Insured, %	Pub. Lag
Biological	Epidemic	46	19	378133			0.58
Climatological	Drought	20	3	783922	1830		0.00
	Wildfire	53	0	41	504	37.22	
Geophysical	Earthquake	150	18	7534	1744	21.23	0.28
	Volcano	5	0	206	431		
	Mass move.	8	0	79			
Hydrological	Flood	189	9	38949	859	42.97	0.00
	Landslide	66	2	321	224		3.50
Meteorological	Storm	217	3	951	1132	101.20	0.00
	Extreme Temp.	70	5	1068	2233	36.26	0.00
	Fog (Smog)	1	1	4000			0.00
All				35175	1116	83	0.38

Note: Natural disasters by group and type that affect countries in our sample, 1900--2009. For each disaster type, we report the number of disasters, the number of severe disasters, the mean number of people killed, the mean damage (in current USD millions), and the mean percent of damage that is insured. Severe disasters are those that killed at least 20 people per million population. Damage, when available, is the total estimated value of damages and economic losses directly or indirectly related to the disaster in USD millions. Insured losses, when available, are the percent of total damage covered by insurance companies, which sometimes exceed the damage. Publication Lag is the number of years from severe disaster occurrence to its first mention in the corpus.

The EM-DAT dataset classifies disasters by (sub) group and type. Our sample includes 11 distinct natural disaster types, belonging to 5 broader groups. Some countries encounter more than one natural disaster in the same year while other countries experience none. We thus sum the death toll by disaster type within the same year for each country.

Table 5 reports summary statistics for the matched sample of disasters, which includes 825 naturals disasters from 1900 to 2009, 733 of which caused fatalities. While some of these events are clearly salient disasters, some are of a more local nature, and unlikely to change popular sentiment. We therefore, classify a disaster as severe if it kills at least 20per million population, and focus on severe disasters for the most part. As we show below, the exact choice of cutoff is not as important as having a cutoff. The cutoff filters out disasters that affected many people but killed few.

Our analysis thus focuses on 60 severe disasters, of which 32% are epidemics, 30% are earthquakes, and 15% are floods. The table shows that droughts and epidemics were most lethal, killing on average 783 and 378 thousand people. EM-DAT includes an estimate of total damage in current US dollars for some disasters. For a subset of disasters with damage information, EM-DAT reports how much of that damage is covered by insurance companies. Insurance is available for wildfires, earthquakes, floods, storms, and extreme temperature disasters, with sample coverage increasing over time. Floods and storms are well covered by insurance, with storms having higher than 100% coverage. However, earthquakes and especially epidemics are rarely covered by insurance. Since 2017, Munich Re, a large reinsurance company, tried to start underwriting business insurance for epidemics. This effort was mostly unsuccessful until the COVID-19 pandemic hit in late 2019 (Walsh, 2020).

The last column of Table 5 shows that the average severe disaster first appears in the Google Books corpus in the same year that it occurs. This short publication lag is partly the result of books' prominence over most of our sample period as a source of timely information. For example, after a severe hurricane hit the city of Galveston, Texas on September 8, 1900, a book describing the disaster was published in the same year as a fundraising device for the area's devastated public schools (Ousley, 1900). Another reason for this modest average lag is that the Google Books corpus includes many library-stored serial publications. For example, the 1930 Irpinia Earthquake is first mentioned in an information bulletin of the Italian National Research Council (Consiglio nazionale delle ricerche, 1930). Based on this average lag, below we regress finance sentiment growth on one-year lagged disaster indicators.

To investigate the effects of war, we rely on war death tolls from http://necrometrics.com, and concentrate on severe ones as well, which killed a similar fraction of a country's population.

Figure 4 visualizes the distribution of disasters across time and countries. It provides a simple explanation for the lack of insurance coverage for epidemics --- when the 1918 Flu

Figure 4: Severe natural disasters and mortality rates

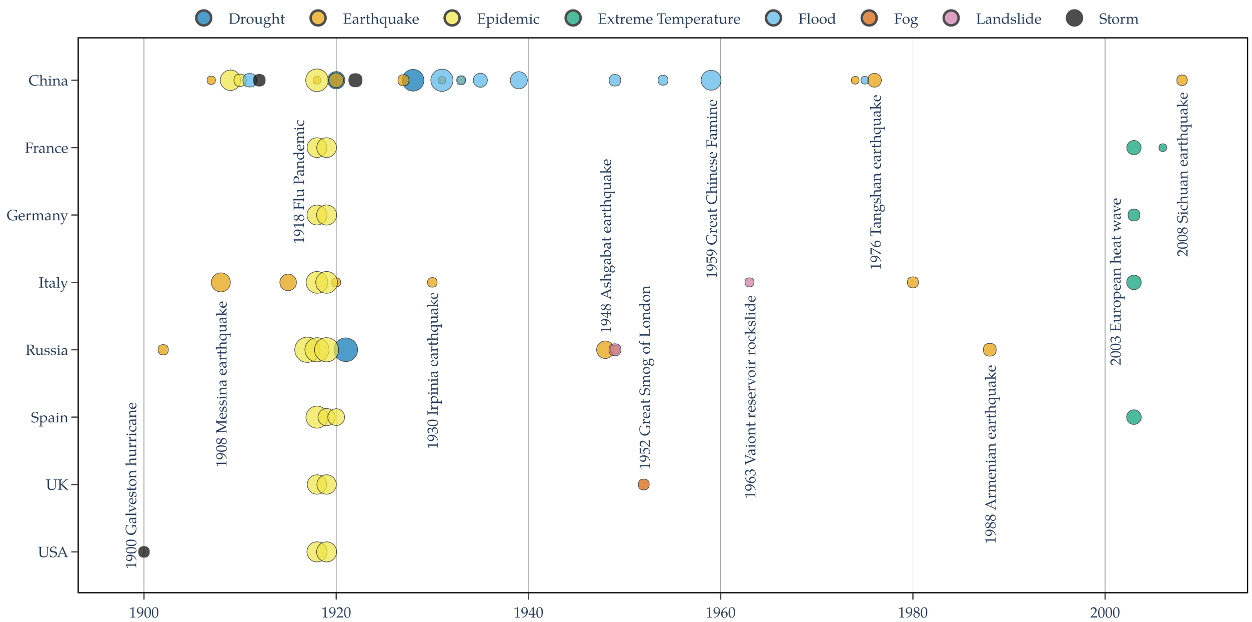

Note: Circles indicate severe natural disasters, with size proportional to the logarithm of the number of death. We define severe disasters, as disasters with death above 20 per million population.

hit, it spread across the globe within a year and affected all major countries in our sample. Such systematic sources of risk may not provide enough opportunities for risk sharing, and therefore feature high insurance premia and low take up. The figure also shows that severe disasters occur more frequently in developing countries such as China and Russia. UK and US, on the other hand, experienced the 1918 flu and one additional disaster. The dearth of disasters for the more developed economies means that estimates of the effects of natural disasters largely originate in developing counties. Because of the clustering of disasters in time and across countries shown by the figure, we include country and year fixed effects in our analysis below.

3.2 Results

Table 6, Column (1) shows that the mean severe natural disaster hitting a country i in year t decreases next year finance sentiment growth $\Delta f_{i,t+1}$ by about one percentage point. Column (2) considers war as another source of variation in finance sentiment, but shows that wars do not have a material effect on finance sentiment. Unlike natural disasters, war and its timing, in particular, is endogenous, as it is under the control of the aggressor and partially under the control of the retaliating country, and therefore likely shaped by other economic and political considerations. It is also possible that our war severity indicator is based on realized casualties, but a country's citizens respond more to war news and to expectations of damage (Verdickt, forthcoming). One may expect the number of fatalities caused by a disaster to be more important than the mere occurrence of a natural disaster. However, Column (3) shows that controlling for the number of people killed hardly changes the natural disaster indicator's coefficient or the R-squared. These panel regressions control for country and year fixed effects. Thus unobserved sources of heterogeneity across countries or time, do not confound this result.

The average treatment effect of a severe natural disaster, however, masks considerable heterogeneity. In Column (4), we replace the single disaster indicator, with type-specific disaster indicators that turn on if a disaster of the listed type hits a country in year t. We find that droughts, floods and landslides tend to increase future finance sentiment, while epidemics and earthquakes decrease it significantly. Storms and fog (smog) disasters have large economic effects as well, but cannot be statistically distinguished from zero. Column (5) shows that this result is robust to controlling for disaster fatalities.

The differential effect of a low insurance disaster is considerably negative. As Column (6) shows, the considerable heterogeneity in effects across disaster types can be explained by the variation in insurance coverage mentioned above. To investigate this hypothesis a bit more formally, we define an indicator for low insurance taking the value of 1, if insurance companies covers less than a third of the damage caused by the disaster. Because data on insurance is sparse and concentrated in the latter part of our sample, we impute missing insured percentages for each disaster type, assuming no coverage for missing droughts, epidemics, landslides, volcano eruptions and fogs.[6]

[6]Our conclusions are robust to allowing the cutoff to increase or decrease by 20 percent.

Table 6: Natural disasters affect financial sentiment

	Finance sentiment growth$_{t+1}$					
	(1)	(2)	(3)	(4)	(5)	(6)
Natural Disaster$_t$	-0.88**	-0.88**	-0.89**			2.01**
	(0.32)	(0.33)	(0.33)			(0.70)
War$_t$		0.10	0.08			
		(0.40)	(0.42)			
Natural Disaster$_t$ × Low Insured$_t$						-4.44**
						(1.70)
logKilled$_t$			0.10		0.12	
			(0.09)		(0.09)	
Drought$_t$				3.27*	3.60*	
				(1.39)	(1.55)	
Earthquake$_t$				-4.57**	-4.64**	
				(1.88)	(1.92)	
Epidemic$_t$				-4.13**	-4.16**	
				(1.64)	(1.69)	
Extremetemp$_t$				-0.07	-0.05	
				(0.35)	(0.37)	
Flood$_t$				2.39**	2.42***	
				(0.68)	(0.68)	
Landslide$_t$				5.20***	5.41***	
				(1.08)	(1.26)	
Storm$_t$				-5.87	-5.93	
				(4.90)	(5.19)	
Fog$_t$				3.31	3.37	
				(2.57)	(2.50)	
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.13	0.13	0.13	0.16	0.17	0.14
Obs	851	851	851	851	851	851

Note: The dependent variable is finance sentiment percent growth from year t to $t+1$. Natural Disaster and War indicate that a country suffers a severe natural disaster or war in year t, killing at least 20 people per million population. Type-specific indicators for severe disasters are similarly defined. Low insured indicates that the no more than a third of the damage caused by the disaster is covered by insurance. $logKilled$ is the logarithm of the number of deaths plus 1. Standard errors clustered by country are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

A potential explanation is the dual roles of finance. Finance facilitates risk sharing through insurance, securitization or derivatives, so that when insured disasters hit, their economic costs are shared broadly, across households and generations. But financial contracts and intermediaries are often designed to prevent ex-post renegotiation (Diamond and Rajan, 2001; Agarwal, Amromin, Ben-David, Chomsisengphet, Piskorski, and Seru, 2017). As a result, damage caused by uninsured disasters can be concentrated in parts of the population and generate resentment against financial intermediaries.

A related explanation is that sentiment toward financial intermediaries, and insurers in particular, may worsen if households and businesses learn they are uninsured only after the fact. Consistent with this channel, Gennaioli, Porta, Lopez-de-Silanes, and Shleifer (2020) document that insurance claims are frequently disputed, and in countries where this is the norm, insurance policies are more expensive and purchased less.

While the effect heterogeneity between insured and uninsured disasters is intriguing, we note that unlike the disasters themselves, which are plausibly exogenous, unobservable omitted variables may confound the differential effect of insurance. Identifying the exact mechanism is as usual harder than identifying the reduced form effect without an instrument for insurance coverage. For example, insurance markets may be more sophisticated in developed economies or in recent periods due to other technological changes we do not observe.

4 Sentiment and economic growth

We next analyze how finance sentiment growth affects macroeconomic activity. We describe our macroeconomic data, discuss our model specification, and the empirical results.

4.1 Macroeconomic data

The economic and credit data that we use are from the macrohistory dataset compiled by Jordà, Schularick, and Taylor (2017). The macrohistory dataset covers annual data for 17 advanced countries from 1870 to 2016. To merge consistently with our text-based finance sentiment index, we only utilize 6 of them: France, Germany, Italy, Spain, UK, and US, spanning from 1870 to 2009. Together, these 6 countries make up more than 40% of the world economy throughout our sample period. This dataset lacks the GDP and population of China and Russia, which we supplement from the Barro-Ursua Macroeconomic Data (Barro and Ursua, 2010). We incorporate credit growth as one of key control variables in our model as credit plays an important role in the macroeconomy and financial development. Following Schularick and Taylor (2012), we use total loans to non-financial private sector as credit proxy.

Table 4 summarizes these macroeconomic variables. GDP growth is highest for China at 3.2%, and all other economies hover around 2%. Germany and Spain exhibit the highest average credit growth.

4.2 Impulse response estimates by local projection

To analyze whether shocks to finance sentiment growth affect economic and credit growth, we estimate cumulative impulse response functions via local projections Jordà (2005):

$$\Delta_h y_{i,t+h} = \alpha_i^h + \sum_{k=1}^{3} \beta_k^h \Delta f_{i,t-k} + \sum_{k=1}^{3} \gamma_k^h X_{i,t-k} + \epsilon_{i,t+h}, \quad h = 0, ..., H, \tag{4}$$

where i and t represent the country and year, respectively. α_i^h are country fixed effects, $\Delta_h y_{i,t+h} = y_{i,t+h} - y_{i,t-1}$ indicate the h-year cumulative growth of interest, e.g. GDP growth rate and credit growth, with $h = 0, ..., H$. f_i is finance sentiment index, X_i is vector of control variables, and $\epsilon_{i,t+h}$ are disturbance terms.

This model estimates the response of $\Delta_h y_{i,t+h}$ from a shock in Δf_i. To capture the direct link of such shocks to economic growth, we control for the first 3 lags of credit growth. Similarly, the first 3 lags of economic growth are control variables when credit growth is our target variable.

The results in Figure 5 include all countries in the sample, and therefore focus on GDP and finance sentiment alone (no credit data is available for China or Russia). The bottom left panel shows that a 1 percentage point increase in finance sentiment growth increases

Figure 5: Impulse response of GDP growth to a finance sentiment growth shock

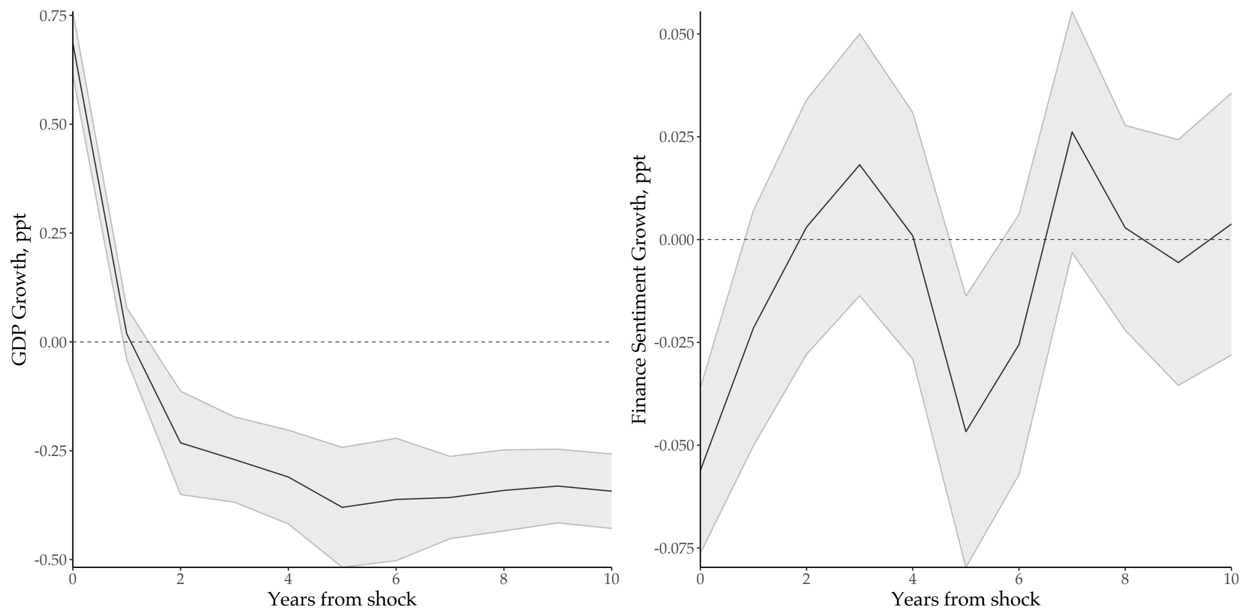

(a) GDP growth shock

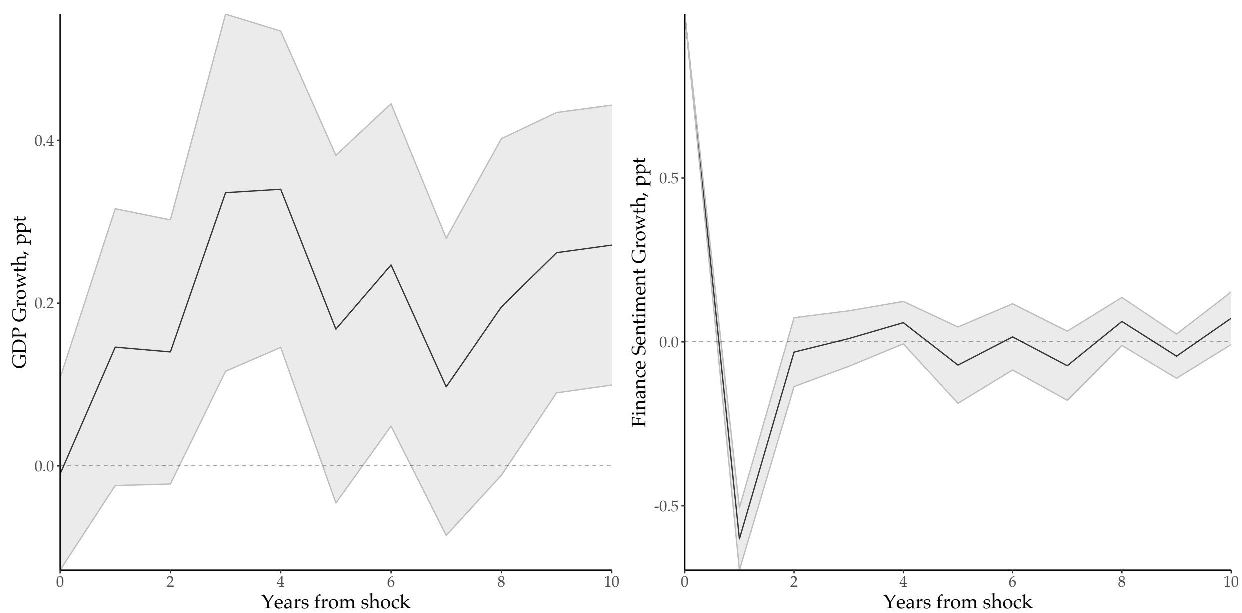

(b) Finance sentiment growth shock

Note: Impulse responses estimated via local projections indicate the change of the cumulative response to a unit shock. Bands are 90% confidence intervals based on Driscoll and Kraay nonparametric robust standard errors.

GDP growth by about 0.3 percentage points 4 years out, though it has no contemporaneous effect. This effect is quite large compared with the mean annual GDP growth of 2.1 percent.

The top right panel shows increases in GDP tend to coincide with declines in finance sentiment growth. While this latter effect is statistically different from zero, its economic magnitude is quite modest.

Interestingly, the bottom right panel reveals that finance sentiment growth tends to oscillate after shocks. This is somewhat surprising, as we expected it to gradually mean-revert like GDP growth does on the top right panel. These oscillations could be the result of book writers and publishers attempting to continuously innovate with contrarian books.

It is likely, however, that finance sentiment affects economic growth not directly, but indirectly, by changing the demand for financial services. It may also affect the supply for financial services by changing how the sector is regulated. Both mechanisms should manifest as changes in the quantity of credit. To investigate this channel, we focus next on the subsample of advanced economies for which we have credit data.

Figure 6 depicts the response of economic growth, finance sentiment growth, and credit growth to shocks by the same three variables. Regardless of the oscillating impact of finance sentiment growth, the cumulative response to the shock in finance sentiment growth is positive for both economic growth and credit growth after a year. A one percentage point increase in finance sentiment growth is associated with a 0.4 percentage point increase in credit growth. The addition of credit growth also reduces somewhat the impulse response of GDP to finance sentiment. It seems, therefore, that some but not all of the effect of finance sentiment on GDP is through credit growth.

What do these estimates imply about the COVID-19 pandemic? From Figure 6, the average effect over the 5 years following the pandemic is a 0.2 percentage point reduction in annual GDP growth, and a 0.25 percentage point reduction in credit growth. Table

Figure 6: Impulse response of GDP growth and credit growth to finance sentiment growth shocks (without China and Russia)

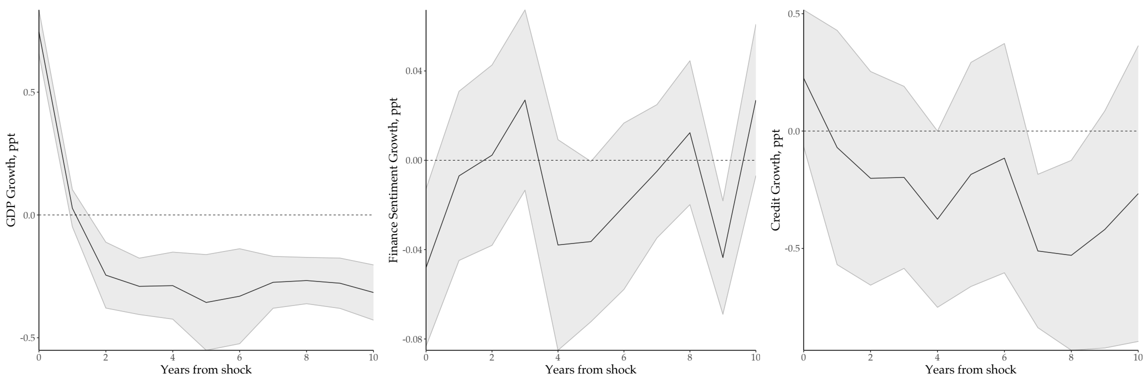

(a) GDP growth shock

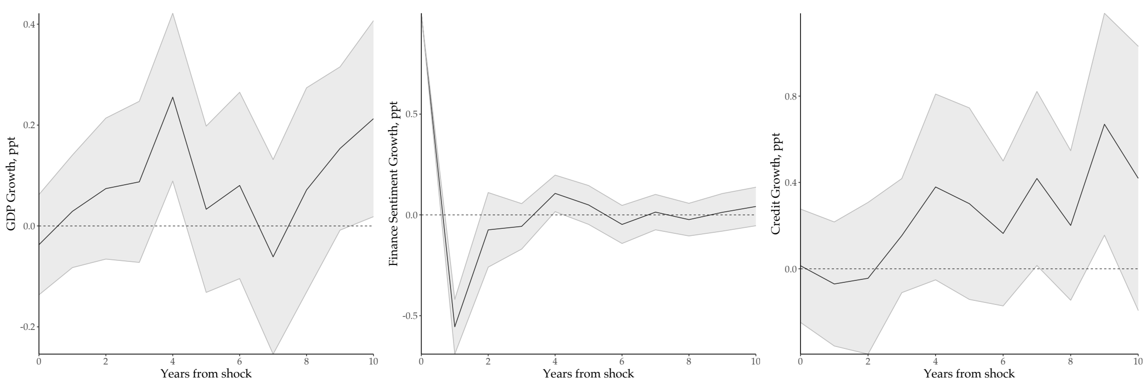

(b) Finance sentiment growth shock

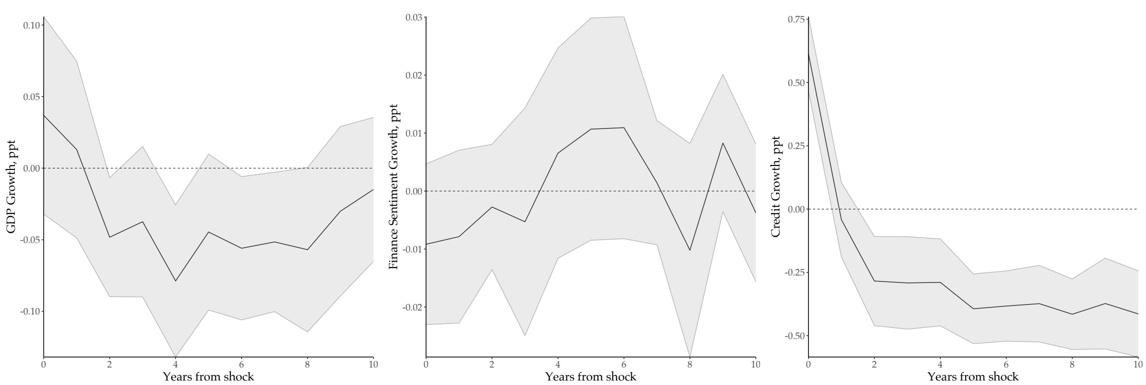

(c) Credit growth shock

Note: Impulse responses estimated via local projections indicate the change of the cumulative response to a unit shock. Bands are 90% confidence intervals based on Driscoll and Kraay nonparametric robust standard errors.

6 shows that the effect of a severe epidemic is a 4 percentage point reduction in finance sentiment growth. Therefore, the cumulative effect of such a decline in finance sentiment on GDP growth over the subsequent 5 years, is about 4 percentage points ($(1.002 \times 4)^5 - 1$). The corresponding effect on credit growth is about 5 percentage points ($(1.0025 \times 4)^5 - 1$).

5 Conclusion and implications for COVID-19

We measure popular sentiment toward finance using a computational linguistics approach applied to millions of books published in eight countries over hundreds of years, and document several new facts.

Finance sentiment differences across countries mostly persist throughout our long sample, with the exception of China, which exhibits great volatility and a level of finance sentiment about as positive as that of Italy and France. Finance sentiment responds negatively to uninsured natural disasters and positively to insured ones. Epidemics and earthquakes, in particular, reduce finance sentiment by about 4 percent within a year. In the VAR sense, shocks to finance sentiment positively affect long term economic and credit growth.

Our estimates imply that beyond its health crisis, the COVID-19 pandemic may reduce GDP growth by 4 percentage points and reduce credit growth by 5 percentage points over the next five years by worsening attitudes toward finance. This back of the envelope calculation assumes, of course, that the COVID-19 pandemic affects finance sentiment like previous severe epidemics. Governments and central bank interventions will hopefully alleviate the pandemic's physical and financial damage and reduce any damage to public sentiment toward finance.

References

Agarwal, Sumit, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, Tomasz Piskorski, and Amit Seru, 2017, Policy Intervention in Debt Renegotiation: Evidence from the Home Affordable Modification Program, *Journal of Political Economy* 125, 654--712.

Aksoy, Cevat Giray, Barry Eichengreen, and Orkun Saka, 2020, The Political Scar of Epidemics, Working Paper 27401 National Bureau of Economic Research.

Altszyler, Edgar, Mariano Sigman, Sidarta Ribeiro, and Diego Fernández Slezak, 2017, Comparative study of LSA vs Word2vec embeddings in small corpora: A case study in dreams database, *Consciousness and Cognition* 56, 178--187.

Antweiler, Werner, and Murray Z. Frank, 2004, Is all that talk just noise? The information content of Internet stock message boards, *Journal of Finance* 59, 1259--1293.

Baker, Scott, Nicholas Bloom, and Stephen Terry, 2020, Using Disasters to Estimate the Impact of Uncertainty, Discussion Paper w27167 National Bureau of Economic Research Cambridge, MA.

Baker, Scott R, Nicholas Bloom, and Steven J Davis, 2016, Measuring economic policy uncertainty, *Quarterly Journal of Economics* 131, 1593--1636.

Barro, Robert J, and Jose F Ursua, 2010, Barro-ursua macroeconomic data, .

Bodnaruk, Andriy, Tim Loughran, and Bill McDonald, 2015, Using 10-K Text to Gauge Financial Constraints, *Journal of Financial and Quantitative Analysis* 50, 623--646.

Bojanowski, Piotr, Edouard Grave, Armand Joulin, and Tomas Mikolov, 2016, Enriching Word Vectors with Subword Information, .

Boudoukh, Jacob, Ronen Feldman, Shimon Kogan, and Matthew Richardson, 2018, Information, Trading, and Volatility: Evidence from Firm-Specific News, *Review of Financial Studies*.

Bybee, Leland, Bryan T. Kelly, Asaf Manela, and Dacheng Xiu, 2019, The Structure of Economic News, *SSRN Electronic Journal*.

Cer, Daniel, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil, 2018, Universal Sentence Encoder, *arXiv:1803.11175 [cs]*.

Chen, Yanqing, and Steven Skiena, 2014, Building Sentiment Lexicons for All Major Languages, in *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)* pp. 383--389 Baltimore, Maryland. Association for Computational Linguistics.

Chetty, Raj, John N. Friedman, Nathaniel Hendren, and Michael Stepner, 2020, Real-Time Economics: A New Platform to Track the Impacts of COVID-19 on People, Businesses, and Communities Using Private Sector Data, Discussion paper Mimeo.

Cieslak, Anna, and Annette Vissing-Jorgensen, forthcoming, The Economics of the Fed Put, *Review of Financial Studies*.

Consiglio nazionale delle ricerche, 1930, *Bollettino d'informazioni* (Consiglio nazionale delle ricerche).

D'Acunto, Francesco, Marcel Prokopczuk, and Michael Weber, 2019, Historical Antisemitism, Ethnic Specialization, and Financial Development, *Review of Economic Studies* 86, 1170--1206.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, 2018, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, *arXiv:1810.04805 [cs]*.

Diamond, Douglas W., and Raghuram G. Rajan, 2001, Liquidity Risk, Liquidity Creation, and Financial Fragility: A Theory of Banking, *Journal of Political Economy* 109, pp. 287--327.

Eisensee, T., and D. Stromberg, 2007, News Droughts, News Floods, and U. S. Disaster Relief, *Quarterly Journal of Economics* 122, 693--728.

Engelberg, Joseph, Matthew Henriksson, Asaf Manela, and Jared Williams, 2019, The Partisanship of Financial Regulators, Working Paper.

Firth, JR, 1957, Applications of General Linguistics - Firth - 1957 - Transactions of the Philological Society - Wiley Online Library, .

García, Diego, 2013, Sentiment during recessions, *Journal of Finance* 68, 1267--1300.

Gennaioli, Nicola, Rafael La Porta, Florencio Lopez-de-Silanes, and Andrei Shleifer, 2020, Trust and Insurance Contracts, Discussion Paper w27189 National Bureau of Economic Research Cambridge, MA.

Gentzkow, Matthew, Edward L Glaeser, and Claudia Goldin, 2004, The Rise of the Fourth Estate: How Newspapers Became Informative and Why It Mattered, Working Paper 10791 National Bureau of Economic Research.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy, 2019, Text as Data, *Journal of Economic Literature* 57, 535--574.

Gentzkow, Matthew, and Jesse M. Shapiro, 2010, What Drives Media Slant? Evidence From U.S. Daily Newspapers, *Econometrica* 78, 35--71.

Giannetti, Mariassunta, and Tracy Yue Wang, 2016, Corporate Scandals and Household Stock Market Participation, *Journal of Finance* 71, 2591--2636.

Goetzman, W., D. Kim, and R. J. Shiller, 2017, Affect, Media and Earthquakes: Determinants of Crash Beliefs from Investor Surveys, Discussion paper Yale University Working Paper.

Goldman, Eitan, Nandini Gupta, and Ryan D. Israelsen, 2020, Political Polarization in Financial News, SSRN Scholarly Paper ID 3537841 Social Science Research Network Rochester, NY.

Grennan, Jillian, 2019, A Corporate Culture Channel: How Increased Shareholder Governance Reduces Firm Value, Working Paper ID 2345384 Rochester, NY.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales, 2004, Does Local Financial Development Matter?, *Quarterly Journal of Economics* 119, 929--969.

———, 2006, Does Culture Affect Economic Outcomes?, *Journal of Economic Perspectives* 20, 23--48.

———, 2008, Trusting the stock market, *Journal of Finance* 63, 2557--2600.

———, 2015, Corporate Culture, Societal Culture, and Institutions, *American Economic Review* 105, 336--339.

Gurun, Umit G., Noah Stoffman, and Scott E. Yonker, 2018, Trust Busting: The Effect of Fraud on Investor Behavior, *Review of Financial Studies* 31, 1341--1376.

Hanley, Kathleen Weiss, and Gerard Hoberg, 2019, Dynamic Interpretation of Emerging Risks in the Financial Sector, *Review of Financial Studies*.

Hansen, Stephen, Michael McMahon, and Andrea Prat, 2018, Transparency and Deliberation Within the FOMC: A Computational Linguistics Approach, *Quarterly Journal of Economics* 133, 801--870.

Harris, Zellig S., 1954, Distributional Structure, *WORD* 10, 146--162.

Hassan, Tarek A., Stephan Hollander, Laurence van Lent, and Ahmed Tahoun, 2017, Firm-Level Political Risk: Measurement and Effects, SSRN Scholarly Paper ID 2838644 Social Science Research Network Rochester, NY.

Hills, Thomas T., Eugenio Proto, Daniel Sgroi, and Chanuki Illushka Seresinhe, 2019, Historical analysis of national subjective wellbeing using millions of digitized books, *Nature Human Behaviour* 3, 1271--1275.

Hoberg, Gerard, and Gordon Phillips, 2016, Text-Based Network Industries and Endogenous Product Differentiation, *Journal of Political Economy* 124, 1423--1465.

Jordà, Òscar, 2005, Estimation and Inference of Impulse Responses by Local Projections, *American Economic Review* 95, 161--182.

———, Moritz Schularick, and Alan M. Taylor, 2017, Macrofinancial History and the New Business Cycle Facts, *NBER Macroeconomics Annual* 31, 213--263.

Jordà, Òscar, Sanjay R Singh, and Alan M Taylor, 2020, Longer-run Economic Consequences of Pandemics, Working Paper 26934 National Bureau of Economic Research.

Joulin, Armand, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov, 2016, Bag of Tricks for Efficient Text Classification, *arXiv e-prints* 1607, arXiv:1607.01759.

Ke, Zheng Tracy, Bryan T. Kelly, and Dacheng Xiu, 2019, Predicting Returns with Text Data, Working Paper.

Kozlowski, Austin C., Matt Taddy, and James A. Evans, 2019, The Geometry of Culture: Analyzing the Meanings of Class through Word Embeddings:, *American Sociological Review*.

Levine, Ross, Chen Lin, and Wensi Xie, 2019, The African Slave Trade and Modern Household Finance, SSRN Scholarly Paper ID 3031310 Social Science Research Network Rochester, NY.

Lin, Yuri, Jean-Baptiste Michel, Erez Aiden Lieberman, Jon Orwant, Will Brockman, and Slav Petrov, 2012, Syntactic Annotations for the Google Books NGram Corpus, in *Proceedings of the ACL 2012 System Demonstrations* pp. 169--174 Jeju Island, Korea. Association for Computational Linguistics.

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

Loughran, Tim, and Bill McDonald, 2020, Textual Analysis in Finance, SSRN Scholarly Paper ID 3470272 Social Science Research Network Rochester, NY.

Luo, Mancy, Alberto Manconi, and Massimo Massa, 2020, Blinded by Perception? The Stock Market's Reaction to Politically Aligned Media, SSRN Scholarly Paper ID 2879939 Social Science Research Network Rochester, NY.

Manela, Asaf, and Alan Moreira, 2017, News implied volatility and disaster concerns, *Journal of Financial Economics* 123, 137--162.

McCloskey, Deirdre Nansen, 2016, *Bourgeois Equality: How Ideas, Not Capital or Institutions, Enriched the World* (University of Chicago Press).

Michel, Jean-Baptiste, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K Gray, Joseph P Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, et al., 2011, Quantitative analysis of culture using millions of digitized books, *Science* 331, 176--182.

Mikolov, Tomas, Kai Chen, Greg Corrado, and Jeffrey Dean, 2013, Efficient Estimation of Word Representations in Vector Space, *arXiv:1301.3781 [cs]*.

Mikolov, Tomas, Kai Chen, Gregory S. Corrado, and Jeffrey A. Dean, 2015, Computing numeric representations of words in a high-dimensional space, .

Mikolov, Tomas, Wen-tau Yih, and Geoffrey Zweig, 2013, Linguistic Regularities in Continuous Space Word Representations, in *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* pp. 746--751 Atlanta, Georgia. Association for Computational Linguistics.

Mokyr, Joel, 2016, *A Culture of Growth: The Origins of the Modern Economy* (Princeton University Press).

Mongey, Simon, Laura Pilossoph, and Alex Weinberg, 2020, Which Workers Bear the Burden of Social Distancing Policies?, Working Paper 27085 National Bureau of Economic Research.

Ousley, Clarence, 1900, *Galveston in Nineteen Hundred: The Authorized and Official Record of the Proud City of the Southwest as It Was Before and After the Hurricane of September 8, and a Logical Forecast of Its Future* (Brookhaven Press).

Pennington, Jeffrey, Richard Socher, and Christopher Manning, 2014, GloVe: Global Vectors for Word Representation, in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* pp. 1532--1543 Doha, Qatar. Association for Computational Linguistics.

Perone, Christian S., Roberto Silveira, and Thomas S. Paula, 2018, Evaluation of sentence embeddings in downstream and linguistic probing tasks, *arXiv:1806.06259 [cs]*.

Sapienza, Paola, and Luigi Zingales, 2012, A Trust Crisis, *International Review of Finance* 12, 123--131.

——— , 2013, Economic Experts versus Average Americans, *American Economic Review* 103, 636--642.

Schularick, Moritz, and Alan M Taylor, 2012, Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870–2008, *American Economic Review* 102, 1029--1061.

Scism, Leslie, 2020, Companies Hit by Covid-19 Want Insurance Payouts. Insurers Say No., *Wall Street Journal*.

Sheng, Jinfei, 2019, Asset Pricing in the Information Age: Employee Expectations and Stock Returns, SSRN Scholarly Paper ID 3321275 Social Science Research Network Rochester, NY.

Soo, Cindy K., 2018, Quantifying Sentiment with News Media across Local Housing Markets, *Review of Financial Studies* 31, 3689--3719.

Spolaore, Enrico, and Romain Wacziarg, 2013, How Deep Are the Roots of Economic Development?, *Journal of Economic Literature* 51, 325--369.

Stulz, René M., and Rohan Williamson, 2003, Culture, openness, and finance, *Journal of Financial Economics* 70, 313--349.

Tetlock, Paul C., 2007, Giving Content to Investor Sentiment: The Role of Media in the Stock Market, *Journal of Finance* 62, 1139--1168.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin, 2017, Attention Is All You Need, *arXiv:1706.03762 [cs]*.

Verdickt, Gertjan, forthcoming, The Effect of War Risk on Managerial and Investor Behavior: Evidence from the Brussels Stock Exchange in the Pre-1914 Era, *Journal of Economic History*.

Walsh, Mary Williams, 2020, Coronavirus Will Cost Businesses Billions. Insurance May Not Help., *The New York Times*.

Zhou, Guofu, 2018, Measuring Investor Sentiment, *Annual Review of Financial Economics* 10, 239--259.

Zingales, Luigi, 2012, *A Capitalism for the People: Recapturing the Lost Genius of American Prosperity* (Basic Books).

——— , 2015, Presidential Address: Does Finance Benefit Society?, *Journal of Finance* 70, 1327--1363.

A Online Appendix

A.1 Positive and negative sentences used to define the positivity dimension across languages

Table 7: Positive and negative sentences

Positive sentences	Negative sentences
financial services benefit society	financial services damage society
finance is good for society	finance is bad for society
finance professionals are mostly good people	finance professionals are mostly corrupt people
finance positively impacts our world	finance negatively impacts our world
financial system helps the economy	financial system hurts the economy

(a) English

金融服务有益社会	金融服务损害社会
金融对社会好	金融对社会不好
财务专业人员大多很好	财务专业人员大多邪恶
金融对世界产生积极影响	金融对世界产生消极影响
金融系统帮助经济	金融系统有害金融

(b) Chinese

les services financiers profitent à la société	les services financiers nuisent à la société
la finance est bonne pour la société	la finance est mauvaise pour la société
les professionnels de la finance sont surtout bons	les professionnels de la finance sont surtout mauvais
la finance a un impact positif sur notre monde	la finance a un impact négatif notre monde
le système financier aide l'économie	le système financier nuit à l'économie

(c) French

Table 7: Positive and negative sentences, continued.

Positive sentences	Negative sentences
Finanzdienstleistungen kommen der Gesellschaft zugute	Finanzdienstleistungen schaden der Gesellschaft
Finanzen sind gut für die Gesellschaft	Finanzen sind schlecht für die Gesellschaft
Finanzprofis sind meistens gut	Finanzprofis sind meistens böse
Finanzen wirken sich positiv auf unsere Welt aus	Finanzen wirken sich negativ aus unsere Welt
Finanzsystem hilft der Wirtschaft	Finanzsystem schadet der Wirtschaft

(d) German

i servizi finanziari avvantaggiano la società	i servizi finanziari danneggiano la società
la finanza fa bene alla società	la finanza fa male alla società
i professionisti della finanza sono per lo più buoni	i professionisti della finanza sono principalmente cattivi
la finanza ha un impatto positivo sul nostro mondo	la finanza ha un impatto negativo il nostro mondo
il sistema finanziario aiuta l'economia	il sistema finanziario danneggia l'economia

(e) Italian

общество оказывает финансовую помощ	общество наносит ущерб финансовым услугам
финансы полезны для общества	финансы вредны для общества
профессионалы в области финансов в основном хороши	профессионалы в области финансов в основном злые
финансы положительно влияют на наш мир	финансы негативно влияют наш мир
финансовая система помогает экономике	финансовая система наносит ущерб экономике

(f) Russian

los servicios financieros benefician a la sociedad	los servicios financieros perjudican a la sociedad
los profesionales financieros son en su mayoría buenos	los profesionales financieros son en su mayoría malos
las finanzas impactan positivamente en nuestro mundo	las finanzas impactan negativamente nuestro mundo
el sistema financiero ayuda a la economía	el sistema financiero perjudica a la economía

(g) Spanish

Note: In line with Kozlowski, Taddy, and Evans (2019), we start with five pairs of words for the positive minus negative dimension for English (both American and British). The word pair includes: (positive − negative), (benefit − damage), (good − bad), (good − corrupt), and (help − hurt). We then create positive and negative sentences which discuss finance, using these words. For other languages, we translate these sentences with the help of native speakers.

A.2 Top ten worst to best ngrams sorted by finance sentiment

Table 8: Worst to best sentence sorted by finance sentiment

American English	British English
turmoil in the financial markets	turmoil in the financial markets
finances become disordered the	instability in the financial markets
financial panic swept the country	lack of money to finance
turmoil in financial markets	a financial panic
financial panic swept the nation	the financial panic
instability in the financial markets	financial panic in the united
financial panic in the country	international financial instability
severe financial setbacks	lack of funds to finance
a major financial panic	my finances falling short
world wide financial panic	the financial deficit
⋮	⋮
knowledge of the financial structure	finance graduate school of
financial support of the field	finance for small and medium
financial support of the course	understanding of the financial system
financial support of the science	financial support of the work
financial support of the graduate	financial management initiative
business and financial experience	financial support of this project
financial management of the organization	financial management of the business
financial support of the center	financial support of the research
finance in the graduate school	financial management of the school
the goal of financial management	financial support of the science

(a) English

Table 8: Worst to best sentence sorted by finance sentiment, continued

Chinese	English Translation
严重 扰乱 了 金融 秩序	Seriously disturbed the financial order
扰乱 了 国家 金融 秩序	Disrupt the national financial order
严重 扰乱 了 金融	Seriously disrupting the financial
扰乱 了 正常 的 金融	Disrupt the normal financial
扰乱 了 金融 秩 序	Disrupt the financial order of rank
扰乱 了 金融 秩序	Disrupt the financial order
扰乱 了 金融 市场	Disrupt the financial markets
干扰 了 金融 秩序	Disturb financial order
既 不 利于 金融	Not only is not conducive to financial
扰乱 了 金融 序	Disrupt the financial order
⋮	⋮
经济 发展 提供 金融	Economic development has provided financial
农村 发展 提供 金融	Rural Development provides financial
金融 推动 发展	Promote the development of financial
金融 服务 促进 农村	Promotion of rural financial services
金融 务 促进	Promote financial affairs
服务 促进 金融	Promoting financial services
金融 立足	Financial foothold
服务 农村 金融	Financial services in rural areas
金融 服务 社会	Financial services community
服务 规范 发展 金融	Regulate the development of financial services

(b) Chinese

Table 8: Worst to best sentence sorted by finance sentiment, continued

French	English Translation
ny a pas de finances	ny no finances
ministre des finances rené pleven	Finance Minister Rene pleven
the financial revolution in england	the financial revolution in england
état des finances était déplorable	financial condition was deplorable
bérenger finances et absolutisme	bérenger Finance and absolutism
finances est rejeté	Finance is rejected
mauvais état des finances royales	poor state of the royal finances
état des finances na pas	financial condition didnt
finances étaient en mauvais état	finances were in bad condition
finances na pu être déposé	na been filed Finance
⋮	⋮
encourager et à soutenir financièrement	encourage and support financially
mobiliser les ressources financières et	mobilize financial resources and
la gestion financière en	financial management
assistance financière et technique avec	financial and technical assistance with
à la coopération financière avec	financial cooperation with
réaliser la solidarité financière des	achieve financial solidarity
assurer la gestion financière et	the financial management and
organiser et de financer les	organize and finance
de promouvoir et de financer	promote and finance
à promouvoir et à financer	to promote and finance

(c) French

Table 8: Worst to best sentence sorted by finance sentiment, continued

German	English Translation
christian watrin bochum finanzpolitik	christian watrin Bochum financial policy
finanzmarkt kapitalismus	financial market capitalism
renzsch wolfgang finanzverfassung	renzsch wolfgang financial constitution
imperialismus staatsfinanzen rüstung	imperialismus government finances armor
gemeindefinanzgesetz vom dezember	community financial law from december
neoabsolutismus staatsfinanzen und politik	Absolutism government finances and politics
r a finanztheorie	r a financial theory
r a musgrave finanztheorie	r a musgrave finance theory
schmölders finanzpolitik berlin	Schmölders financial policy berlin
mayer geschichte der finanzwirtschaft	mayer history of finance economy
⋮	⋮
ist zuständig für die finanzielle	is responsible for the financial
finanziell und organisatorisch zu unterstützen	financial and organizational support
hilfe bei der finanzierung der	help with the financing of
finanzen die zur durchführung und	finance the implementation and to
hilfe bei der finanzierung von	help with the financing of
finanzieren mit	fund with
finanzierung erfolgt durch beiträge der	financed through contributions of the
unternehmen damit derartige finanzierungen	company so that such financing
sorgt für die finanzierung	provides for the financing
finanzen und mit zustimmung des	Finance and with the approval of

(d) German

Table 8: Worst to best sentence sorted by finance sentiment, continued

Italian	English Translation
dimissioni del ministro delle finanze	resignation of Finance Minister
il finanziamento è stato concesso	The loan was granted
le finanze sono condannate dai	finances are condemned by
scioglimento del contratto di finanziamento	termination of the loan agreement
ministro delle finanze è autorizzato	Minister of Finance is authorized
grave crisi finanziaria	serious financial crisis
il ministro delle finanze dichiarava	Finance Minister declared
le finanze saranno emanate	finances will be issued
la finanza sabauda allaprirsi	finance Savoy allaprirsi
l esercizio finanziario ha inizio	the financial year
⋮	⋮
disponibilità di risorse finanziarie che	availability of financial resources
di gestire le risorse finanziarie	to manage the financial resources
a soddisfare le esigenze finanziarie	to meet the financial needs
effettuare la gestione finanziaria di	make the financial management of
gestione delle risorse finanziarie e	management of financial resources and
relazioni economiche e finanziarie con	economic and financial relations with
coordinamento della finanza regionale con	coordination of regional finance with
assistere tecnicamente e finanziariamente i	assist technically and financially
gestione delle risorse finanziarie idonee	management of the financial resources
relazioni commerciali e finanziarie con	commercial and financial relations with

(e) Italian

Table 8: Worst to best sentence sorted by finance sentiment, continued

Russian	English Translation
обращение финансы кредит	recourse finance loan
плутократия бароны финансового	plutocracy financial barons
буржуазии финансовый срыв	Financial breakdown of the bourgeoisie
протекционизм господство финансистов	Protectionism domination of financiers
финансов кредита социализме	Finance socialism loan
империализм финансовый капитализм	financial capitalism, imperialism
обращение кредит финансы	recourse loan finance
финансовое банкротство	financial bankruptcy
империализма колониальное финансовое порабощение	colonial imperialism financial enslavement
страшных финансовых грозных	terrible financial formidable
⋮	⋮
оказывает и финансовую поддержку	and providing financial support
оказывает колхозам финансовую помощь	It provides financial assistance to collective farms
финансовая деятельность колхоза осуществляется на	financial activities carried out on a collective farm
обеспечивается финансирование мероприятий	provided funding
оказывает финансовую и помощь	and provides financial assistance
оказывает финансовую и поддержку	It provides financial support and
оказывает финансовую и политическую поддержку	It provides financial and political support
оказывает большую финансовую помощь	providing more financial aid
оказывает значительную финансовую помощь	providing substantial financial assistance
оказывает финансовую и техническую помощь	It provides financial and technical assistance

(f) Russian

Table 8: Worst to best sentence sorted by finance sentiment, continued

Spanish	English Translation
la finanza no era	finance was not
la situación financiera no era	the financial situation was not
el capital financiero se sentirá	financial capital will feel
el capital financiero no es	financial capital is not
el mercado financiero no es	the financial market is not
la especulación financiera domine su	financial speculation dominates its
el sistema financiero se vio	the financial system was
una desgraciada situación financiera pudiese	an unfortunate financial situation could
el déficit se financió	The deficit was financed
su situación financiera no era	its financial situation was not
⋮	⋮
actividades financieras y de servicios	financial activities and services
asesoría técnica y apoyo financiero	technical advice and financial support
financiamiento de las diversas actividades	financing various activities
apoyo técnico y financiero internacional	international technical and financial support
apoyo financiero y asistencia técnica	financial support and technical assistance
apoyo financiero a las actividades	financial support to activities
apoyo financiero para las actividades	Financial support for activities
asistencia técnica y recursos financieros	technical assistance and financial resources
financiamiento de las actividades culturales	financing of cultural activities
asistencia técnica y de financiamiento	technical assistance and financing

(g) Spanish

Note: The sentences are sorted from worst to best in terms of their cosine similarity with positive minus negative vector. The English translation is provided using Google Translate.

A.3 General sentiment for each language, not just sentences mentioning finance

In Figure 7 we repeat our exercise, but instead of finance-mentioning sentences, we measure general sentiment by focusing on January-mentioning sentences. We find no time trend in general sentiment, and a ranking across languages that is quite different from that of the finance sentiment shown in Figure 2.

Figure 7: General sentiment

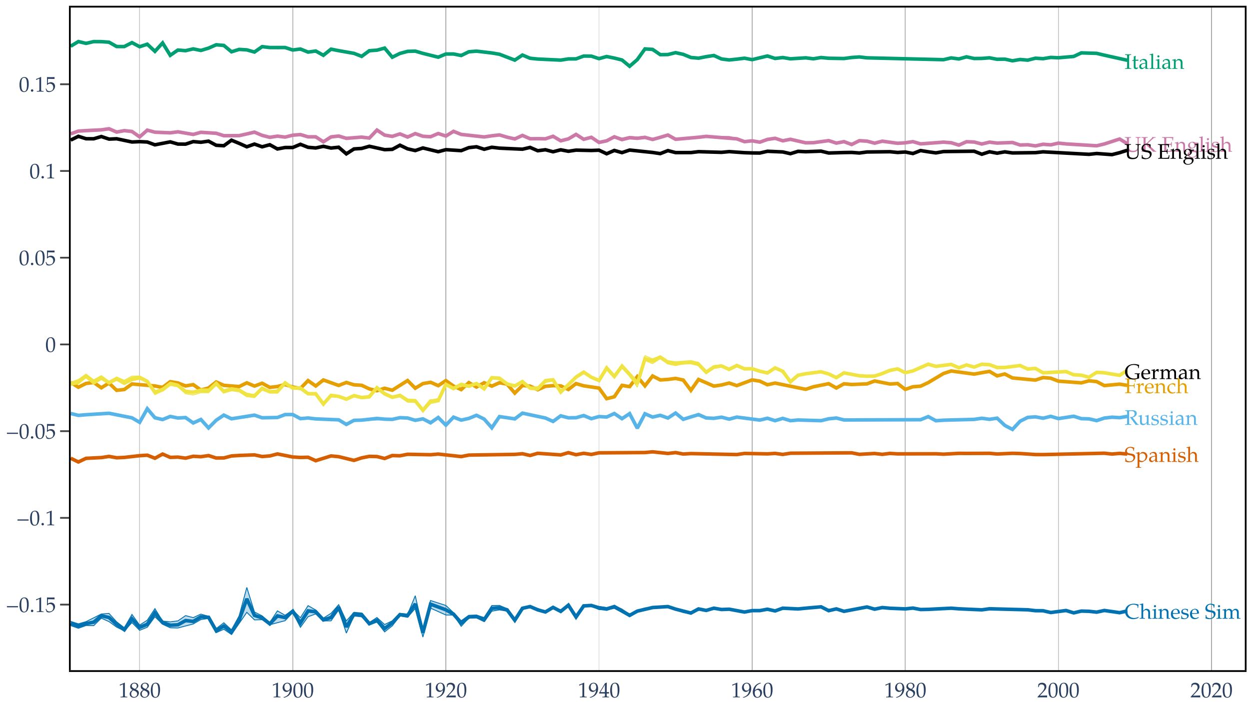

Note: General sentiment is based on the annual average projection of January-mentioning sentences' embeddings onto the positive minus negative January sentiment dimension. To define the positivity dimension, we average the difference in embedding for following tuples (and their translations to each language): [("january is good for society","january is bad for society"), ("january is mostly good","january is corrupt"), ("january positively impacts our world","january negatively impacts our world"), ("january helps the economy","january hurts the economy"), ("january benefits society","january damages society")]. Sentences are from the Google Books Ngram corpus and embedded using BERT. Bands represent 95 percent confidence intervals produced by subsampling.

A.4 Comparison with alternative text-based approaches

A.4.1 Dictionary-based approach

A considerably simpler and popular method than ours, counts positive versus negative words to measure sentiment (Zhou, 2018). One limitation of this approach is that it often misses the context and subtleties of language, which humans would quickly discern from reading words in sequence. In fact a major engineering feat of BERT is that its underlying neural network pays attention to longer sequences of words (Vaswani, Shazeer, Parmar, Uszkoreit, Jones, Gomez, Kaiser, and Polosukhin, 2017). However, the dictionary-based approach may be a reasonable alternative due to its simplicity.

Figure 8: Sentiment toward finance using an alternative dictionary-based approach

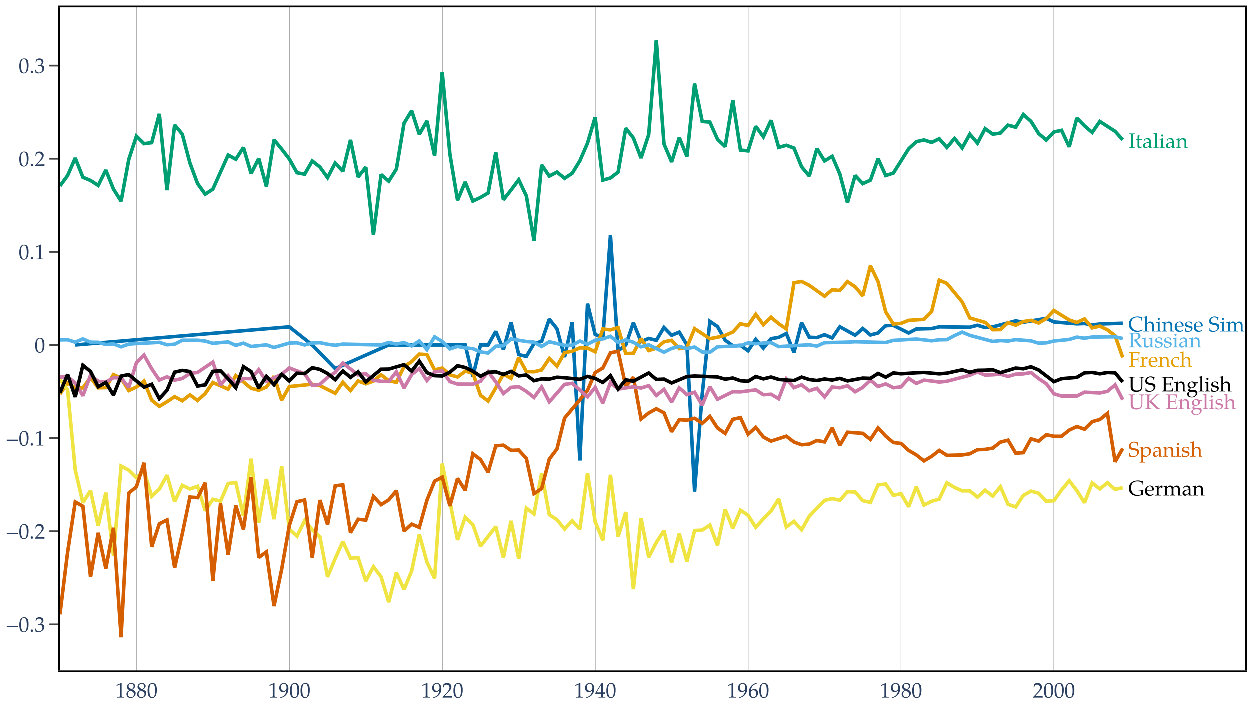

Note: Dictionary-based finance sentiment is based on the annual average sentiment of finance-mentioning sentences. Sentiment for a sentence is net positive words in a sentence normalized by total positive and negative words in the sentence. Sentences are from the Google Books Ngram corpus and the positive and negative words are from Loughran and McDonald (2020) and Chen and Skiena (2014).

We use a list of positive and negative word for each language. For English we use the Loughran and McDonald (2011) dictionary. For all other language we rely on Chen and Skiena (2014). The sentiment for each sentence is the number of net positive words in a sentence normalized by total positive and negative words in the sentence. We aggregate the sentence sentiments, weighted by their frequency in a year, to get sentiment for the year. Based on the dictionary approach, we get a more volatile score, illustrated in Figure 8, and a non-significant relationships with disasters, reported in Table 9.

A.4.2 Alternative language embedding-based approaches

As mentioned, our language embedding approach builds on Kozlowski, Taddy, and Evans (2019), but differs in an important way. Kozlowski, Taddy, and Evans (2019) fit a word embedding model (e.g. word2vec, glove) to each decade of sentences. They then measure the

Table 9: Dictionary-based approach: Natural disaster effects on financial sentiment

	Finance sentiment growth$_{t+1}$					
	(1)	(2)	(3)	(4)	(5)	(6)
Natural Disaster$_t$	593.65	596.80	-139.39	595.86		
	(569.80)	(572.67)	(111.36)	(575.24)		
War$_t$		-52.24		-47.27		
		(104.78)		(98.86)		
Natural Disaster$_t$ × Low Insured$_t$			1094.32			
			(614.91)			
logKilled$_t$				-11.27		-14.58
				(14.72)		(19.50)
Drought$_t$					150.53	115.07
					(253.08)	(280.46)
Earthquake$_t$					1426.29	1435.65
					(1213.20)	(1236.52)
Epidemic$_t$					-71.98	-67.91
					(218.51)	(207.31)
Extremetemp$_t$					48.80	45.61
					(60.12)	(55.22)
Flood$_t$					-479.58	-483.40
					(260.25)	(265.39)
Landslide$_t$					-790.37	-816.91
					(1148.38)	(1188.96)
Storm$_t$					54.44	9.49
					(33.94)	(42.08)
Fog$_t$					-111.12	-130.96
					(145.37)	(135.32)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.17	0.17	0.18	0.17	0.20	0.20
Obs	851	851	851	851	851	851

cosine similarity once for each phrase of interest. The variation in their measures of culture come from variation in term frequencies but also from estimation error that generates variation in these fitted language models. By contrast, we use a pretrained language model (BERT), measure cosine similarity once for each phrase of interest, and then average these cosine similarities for each year (and language). Variation in culture in our approach is due only to term frequencies, as language model error is held fixed over time.

We attempt to apply the Kozlowski, Taddy, and Evans (2019) method by fitting three word embedding models, word2vec, glove, and fasttext, to every language-year in our panel. We find, however, that the finance sentiment series this approach generates are

Table 10: Word2Vec as an alternative model: Natural disaster effects on financial sentiment

	Finance sentiment growth$_{t+1}$					
	(1)	(2)	(3)	(4)	(5)	(6)
Natural Disaster$_t$	30.06	35.08	-1.07	35.02		
	(41.27)	(43.93)	(76.60)	(44.05)		
War$_t$		-93.86		-92.07		
		(61.62)		(61.79)		
Natural Disaster$_t$ × Low Insured$_t$			46.44			
			(96.54)			
logKilled$_t$				-7.41		-7.39
				(4.62)		(4.64)
Drought$_t$					120.48**	98.49
					(44.95)	(53.27)
Earthquake$_t$					-14.06	-9.57
					(52.08)	(52.22)
Epidemic$_t$					126.78	128.75
					(164.78)	(159.69)
Extremetemp$_t$					-82.12	-83.56
					(278.56)	(279.97)
Flood$_t$					153.86*	149.83*
					(75.71)	(74.62)
Landslide$_t$					173.61	160.23
					(135.12)	(137.62)
Storm$_t$					-52.45	-48.14
					(109.40)	(106.87)
Fog$_t$					28.81	25.49
					(92.94)	(89.12)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.17	0.17	0.17	0.17	0.17	0.17
Obs	833	833	833	833	833	833

highly noisy. In Tables 10, 11, and 12, we report the natural disaster regression estimates using these alternative text-based measures. The tables show that, as one may expect, the noisier measures generate considerable parameter uncertainty.

A.5 Severe disaster cutoff

Figures 9 and 10 shows our estimates of the effect of natural disasters are mostly robust to varying the cutoff for the fraction of the population killed by the disaster. Lower cutoffs include more benign natural disasters, while higher cutoffs concentrate the treatment effect estimates on fewer but more fatal disasters. As a result, the point estimates for more fatal

Table 11: GloVe as an alternative model: Natural disaster effects on financial sentiment

	Finance sentiment growth$_{t+1}$					
	(1)	(2)	(3)	(4)	(5)	(6)
Natural Disaster$_t$	600.28	596.84	1441.70	596.19		
	(472.00)	(473.03)	(1614.60)	(467.86)		
War$_t$		111.32		115.63		
		(151.28)		(151.56)		
Natural Disaster$_t$ × Low Insured$_t$			-1302.80			
			(1616.99)			
logKilled$_t$				-16.00		-14.97
				(17.79)		(18.47)
Drought$_t$					321.47*	260.52
					(147.45)	(158.39)
Earthquake$_t$					71.10	78.30
					(108.60)	(114.66)
Epidemic$_t$					83.13	61.36
					(118.79)	(141.98)
Extremetemp$_t$					4343.79	4341.38
					(4766.25)	(4764.87)
Flood$_t$					33.42	25.34
					(141.04)	(144.75)
Landslide$_t$					-124.93***	-151.91***
					(34.28)	(41.61)
Storm$_t$					-303.68	-292.19
					(257.38)	(298.14)
Fog$_t$					284.91	278.55
					(386.15)	(387.92)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.13	0.13	0.14	0.13	0.15	0.15
Obs	808	808	808	808	808	808

disasters are generally larger in magnitude and feature greater parameter uncertainty.

Figure 9: Robustness to the severe disaster cutoff for natural disaster groups

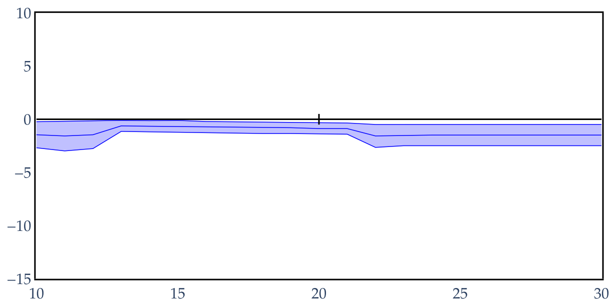

(a) Natural Disaster (any)

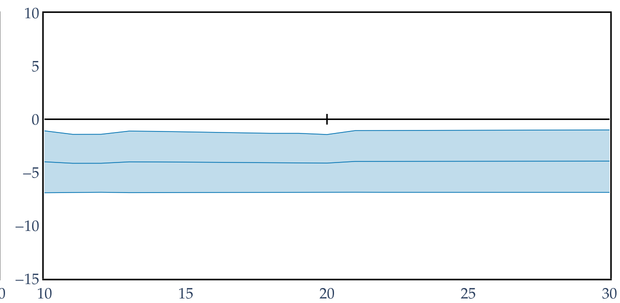

(b) Biological

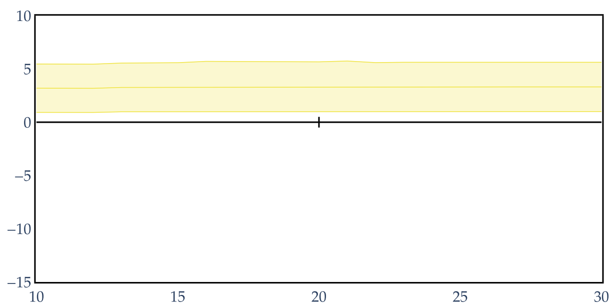

(c) Climatological

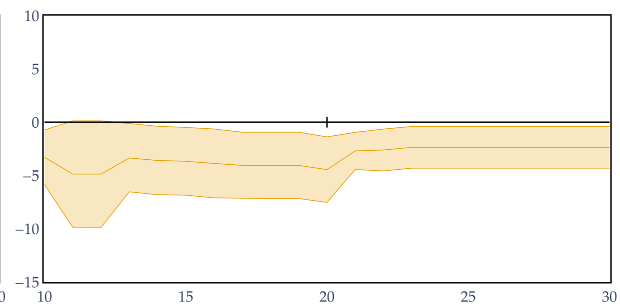

(d) Geophysical

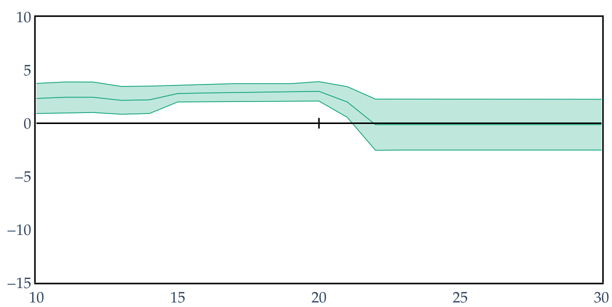

(e) Hydrological

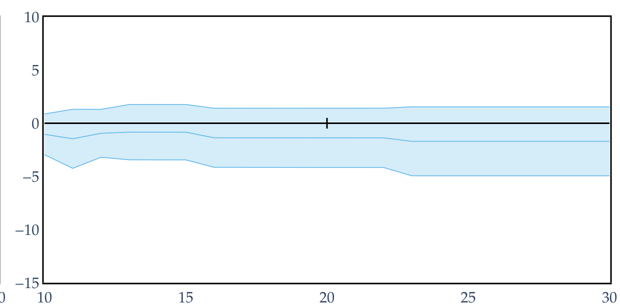

(f) Meteorological

Note: The figure shows how the estimated treatment effects of severe natural disaster groups change as we vary the minimum number of deaths per million for a disaster to be considered severe, thus filtering out less devastating disasters. Bands indicate 90% confidence intervals.

Table 12: FastText as an alternative model: Natural disaster effects on financial sentiment

	Finance sentiment growth$_{t+1}$					
	(1)	(2)	(3)	(4)	(5)	(6)
Natural Disaster$_t$	-34.41	-39.17	-103.05	-39.26		
	(81.37)	(78.75)	(72.71)	(79.01)		
War$_t$		88.20		90.97		
		(176.68)		(177.71)		
Natural Disaster$_t$ $\times$ Low Insured$_t$			102.32			
			(168.59)			
logKilled$_t$				-11.09		-10.89
				(12.51)		(12.69)
Drought$_t$					25.67	-6.74
					(199.08)	(216.77)
Earthquake$_t$					-13.88	-7.27
					(111.39)	(114.81)
Epidemic$_t$					176.59	179.55
					(139.59)	(147.85)
Extremetemp$_t$					-28.82	-30.92
					(115.34)	(116.23)
Flood$_t$					-260.87***	-266.80***
					(49.75)	(52.43)
Landslide$_t$					206.24*	186.51*
					(93.04)	(86.99)
Storm$_t$					204.43	210.79
					(187.13)	(211.38)
Fog$_t$					-6.41	-11.32
					(75.17)	(78.11)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.14	0.14	0.14	0.14	0.14	0.14
Obs	840	840	840	840	840	840

Figure 10: Robustness to the severe disaster cutoff for natural disaster types

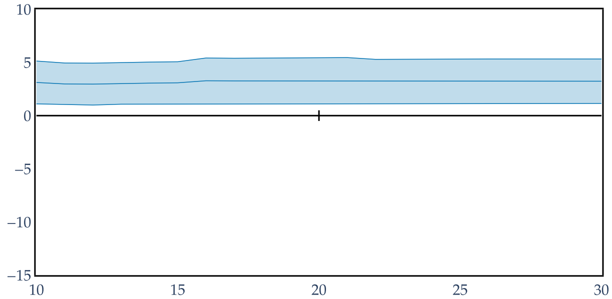

(a) Drought

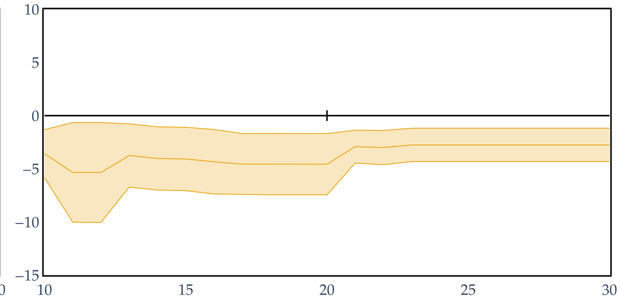

(b) Earthquake

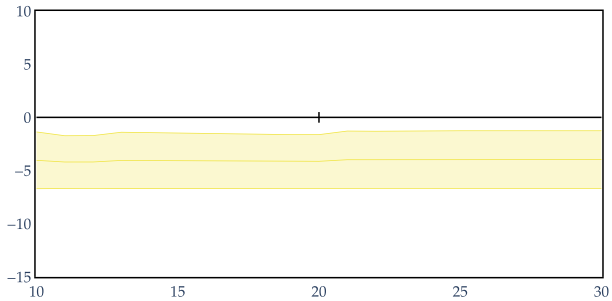

(c) Epidemic

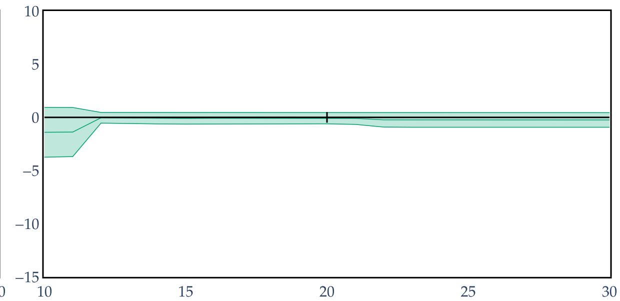

(d) Extreme Temperature

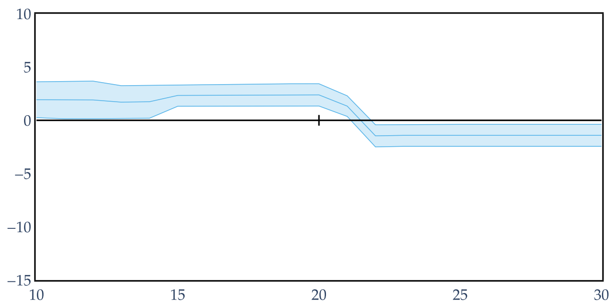

(e) Flood

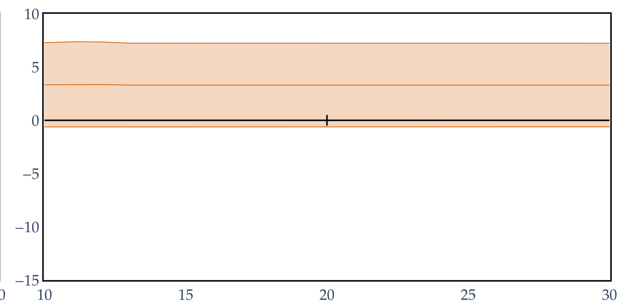

(f) Fog

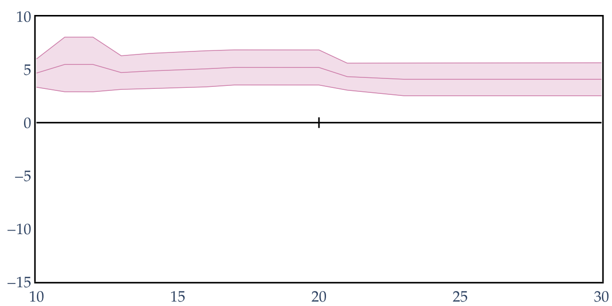

(g) Landslide

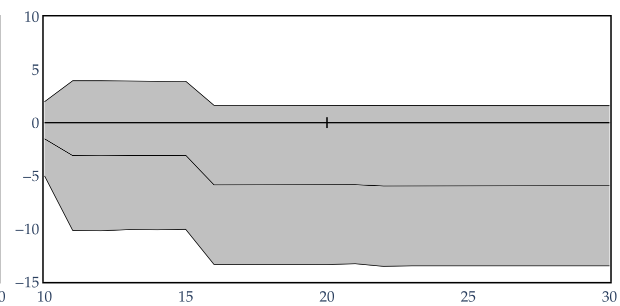

(h) Storm

Note: The figure shows how the estimated treatment effects of severe natural disaster types change as we vary the minimum number of deaths per million for a disaster to be considered severe, thus filtering out less devastating disasters. Bands indicate 90% confidence intervals.