



ECONOMIC RESEARCH
FEDERAL RESERVE BANK OF ST. LOUIS
WORKING PAPER SERIES

Economic Surveillance using Corporate Text

Authors	Tarek Hassan Stephen Hollander, Aakash Kalyani, Markus Schwedeler, Ahmed Tahoun, and Laurence van Lent
Working Paper Number	2024-022A
Creation Date	September 2024
Citable Link	https://doi.org/10.20955/wp.2024.022
Suggested Citation	Hollander, T.H.S., Kalyani, A., Schwedeler, M., Tahoun, A., van Lent, L., 2024; Economic Surveillance using Corporate Text, Federal Reserve Bank of St. Louis Working Paper 2024-022. URL https://doi.org/10.20955/wp.2024.022

Federal Reserve Bank of St. Louis, Research Division, P.O. Box 442, St. Louis, MO 63166

The views expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Federal Reserve System, the Board of Governors, or the regional Federal Reserve Banks. Federal Reserve Bank of St. Louis Working Papers are preliminary materials circulated to stimulate discussion and critical comment.

Economic Surveillance using Corporate Text

Tarek A. Hassan, Stephan Hollander, Aakash Kalyani, Laurence van Lent, Markus Schwedeler, and Ahmed Tahoun

Abstract

This article applies simple methods from computational linguistics to analyze unstructured corporate texts for economic surveillance. We apply text-as-data approaches to earnings conference call transcripts, patent texts, and job postings to uncover unique insights into how markets and firms respond to economic shocks, such as a nuclear disaster or a geopolitical event—insights that often elude traditional data sources. This method enhances our ability to extract actionable intelligence from textual data, thereby aiding policy-making and strategic corporate decisions. By integrating computational linguistics into the analysis of economic shocks, our study opens new possibilities for real-time economic surveillance and offers a more nuanced understanding of firm-level reactions in volatile economic environments.

■ *Tarek A. Hassan is Professor of Economics at Boston University, Boston, Massachusetts. Stephan Hollander is Professor of Financial Accounting at Tilburg University, Tilburg, the Netherlands. Aakash Kalyani is Research Economist at the Federal Reserve Bank of St. Louis, St. Louis, Missouri. Laurence van Lent is Professor of Accounting and Economics, Frankfurt School of Finance and Management, Frankfurt, Germany. Markus Schwedeler is PhD student in Economics at Boston University, Boston, Massachusetts. Ahmed Tahoun is Professor of Accounting at London Business School, London, the United Kingdom.*

■ *This article reflects the views of the authors and not necessarily those of the Federal Reserve Bank of St. Louis or the Federal Reserve System. We thank Jiarui Wang for providing excellent research assistance. We appreciate helpful suggestions from Eric Hurst (editor). Tahoun's research is supported by grants from the Institute for New Economic Thinking. Van Lent gratefully acknowledges funding from the Deutsche Forschungsgemeinschaft Project ID 403041268-TRR 266.*

Tracking the expected, ongoing, and realized impacts of economic shocks is a fundamental challenge in economics, crucial for formulating effective economic policy responses and enhancing our understanding of the functioning of the economy. Comprehending how events such as natural disasters, wars, or political crises affect businesses and employment is essential for developing policies to mitigate adverse events and manage systemic risks. Similarly, discerning why firms respond differently to proposed reforms or technological changes can inform more effective policy design and facilitate smoother technological transitions. At a broader level, tracing the propagation of supply, demand, and other economic shocks across firms, sectors, and countries is key to testing and refining our understanding of how the economy functions.

Traditionally, economic surveillance has relied on data from surveys or accounting sources, such as firms' profit, investment, or employment figures. We argue that systematically analyzing earnings conference call transcripts, job postings, patents, and other texts produced as part of routine business operations represents a sea change in our ability to track and understand economic shocks. These texts vastly expand the available data and contain information about decision-makers' expectations, perceived risks, costs, and opportunities that are difficult or impossible to extract from conventional sources. Given the richness of this information, we anticipate that future economic surveillance will rely as heavily on the analysis of these texts as it does on survey data and accounting numbers today.

To illustrate our approach, consider Russia's large-scale invasion of Ukraine on February 24, 2022. Within hours of this escalation, Geox S.p.A., an Italian shoe manufacturer, held its quarterly earnings call. The call focused on the unfolding crisis, with analysts and investors inquiring about contingency plans, revenue dependencies, and strategic options in light of potential global embargoes, disrupted supply chains, and interrupted interbank communications. Geox was not unique; management of numerous corporations engaged in similar discussions during earnings calls in the following days, anticipating impacts ranging from high natural gas prices to potential asset seizures and increased defense budgets. As

political leaders and policymakers were still assessing the situation, decision-makers within firms and capital market participants were already discussing the potential effects of the war on their business activities and operations.

Systematic analysis of earnings-call discussions and other corporate textual sources enhances our ability to comprehend the impact and transmission of shocks in four distinct ways. It enables near real-time measurement of risks, costs, and opportunities firms associate with specific shocks, policies, regulations, and crises, eliminating the need for costly surveys. This approach traces shock-induced risk transmission and potential contagion across firms, sectors, and regions without imposing structural assumptions or relying on input-output tables. It also facilitates measurement of past expectations related to historical events, such as the 2003 SARS outbreak ([Hassan et al., 2023](#)) or the 2017 Tax Cuts and Jobs Act ([Gallemore et al., 2024](#)), overcoming post-event surveys' limitations in reconstructing pre-event expectations. Finally, it enables cardinal comparison of different risks, revealing, for example, that US-based firms in 2023 devoted twice as much attention to inflation-related risks as to those stemming from geopolitics or trade wars.

In the subsequent sections, we present and discuss simple computational linguistics techniques that enable these capabilities, primarily using transcripts of firms' quarterly earnings conference calls. Many scholars, central banks, think tanks, and international organizations are currently integrating these techniques into their economic surveillance activities.¹ The article concludes by presenting applications that extend these techniques to other corporate text sources, including patents and job postings. We demonstrate how to use these sources to trace the origin and spread of technological innovations across geography, industries, firms, and occupations.

¹Recent applications include [Correa et al. \(2023\)](#) on trade uncertainty and bank lending, [Albrizio et al. \(2023\)](#) on cross-country firm expectations, [Gosselin and Taskin \(2023\)](#) on economic slack, [Alfaro and Chor \(2023\)](#) on supply chain reallocation, and [Ma and Zimmermann \(2023\)](#) on monetary policy and innovation. The IMF's World Economic Outlook (October 2022, April 2023, October 2023) analyzes supply chains, inflation expectations, and other macroeconomic risks ([Staff, 2022](#); [Ahn et al., 2023](#); [Albrizio et al., 2023](#)). The European Central Bank tracks corporate sentiment ([Andersson et al., 2023](#)), and the European Bank for Reconstruction and Development's 2022 Transition Report examines supply chain resilience post Covid-19 ([Faella et al., 2022](#)).

A unifying principle in all of these applications is the effectiveness of simplicity and clarity over complexity. Our approaches are typically straightforward and easily interpretable, avoiding intricate linguistic and econometric models. We contend that the primary frontier in economic text analysis lies not in the development of advanced techniques, but in making vast, untapped datasets available for research ([Bae et al., 2023](#)).

We begin by introducing earnings conference call transcripts as a data source, followed by an outline of the building blocks underlying our approach and the basic text processing prerequisites, before going into the details of various economic surveillance options.

1. Earnings Conference Calls: A Marketplace of Information

Until about 30 years ago, publicly listed firms primarily communicated with shareholders through accounting disclosures. However, financial statements and regulatory filings provide an incomplete representation of economic reality, mainly recording past events due to accounting standards that permit only gradual reflection of news in earnings. Moreover, these numbers offer limited context about how or why shocks impact or are expected to impact firms.

In recent decades, technological advancements and regulatory changes such as Regulation Fair Disclosure in the US have elevated earnings conference calls as outlets for corporate information. These calls serve as platforms for market participants to gain additional insights into companies' outlooks, encompassing past achievements and future potential. Typically held quarterly following earnings releases, these calls have become a primary means for global publicly listed firms to communicate with shareholders and stakeholders.

Earnings conference calls, usually broadcasted live over the Internet, differ significantly from written regulatory filings such as SEC Forms 10-K and 10-Q. They allow financial analysts to probe management with questions, forcing anticipation and response to possibly uncomfortable inquiries. Being largely unscripted and conducted live, they often lead to more candid and open discussions. Management is less likely to use legal language and boilerplate

jargon, making information more accessible and comprehensible (Bae et al., 2023; Frankel et al., 1999).² Additionally, international listed firms either conduct their earnings calls in English or provide translations, enabling rapid assembly of directly comparable datasets across more than 12,500 firms in nearly 90 countries.

In summary, for essentially each set of accounting numbers available through databases such as Standard & Poor’s Compustat, earnings call transcripts offer substantial complementary textual narrative. Released simultaneously, this narrative provides valuable context about firms’ current and future situations. Consequently, any economic analysis relying on firms’ accounting data could benefit from leveraging this additional information.

Earnings call transcripts are widely accessible through various channels. These include firms’ investor relations websites, financial news platforms such as Yahoo! Finance, investor community websites like Seeking Alpha, transcription services such as EarningsCast, the SEC’s EDGAR database, and data providers like London Stock Exchange Group (LSEG, formerly Refinitiv, Eikon). This wide availability ensures that researchers and policymakers have easy access to this rich source of corporate information.

2. Decoding Shocks: Exposure, Risk, and Sentiment

Exposure

Our basic premise is that the time people devote to discussing a particular topic in their communication reflects its importance. This intuition extends to firms and their communication practices, including their regular earnings conference calls. These calls typically last between 45 minutes and an hour (Matsumoto et al., 2011), requiring participants to judiciously select the most significant topics for discussion within this limited time frame.

Based on this intuition, we informally refer to the number of sentences participants

²Under the “safe harbor” protection, US securities law shields managers making forward-looking statements in earnings calls. This enables more candid disclosures about ongoing and future issues not yet reflected in accounting numbers, economic shocks influencing current or future results, and the company’s challenges, risks, or opportunities.

dedicate to a given topic in firm i 's quarter t earnings conference call as the firm's "exposure" to that topic:

$$(1) \quad TopicExposure_{i,t} = \frac{\text{Count of sentences on a given topic}_{i,t}}{\text{Total number of sentences}_{i,t}}$$

This simple metric allows us to quantify a firm's exposure to specific issues. For instance, by counting the sentences discussing matters related to Russia's invasion of Ukraine as a percentage of the total sentences in the earnings conference call, we can gauge which firms are exposed to the Russian-Ukrainian conflict and the relative importance of this exposure compared to other topics.³

For certain research questions, this off-the-shelf measure of topic-specific exposure suffices (Hassan et al., 2023). However, our approach is adaptable and can be extended along various dimensions. For example, by analyzing the context in which a topic is mentioned, we can infer *how* a given shock is expected to affect the firm.

Risk

For some applications, we might be interested in the level of risk or uncertainty a firm encounters due to a shock, specifically focusing on the shock's second-moment impact. Our text-based approach offers a means to measure the ebbs and flows in the level of risk or uncertainty faced by individual firms and, when aggregated, by the economy as a whole.

Figure 1, Panel A illustrates this concept with a time series plot of *OverallRisk* faced by US-listed firms. We measure this as the average percentage of sentences in these firms'

³Our earlier work (e.g., Hassan et al., 2019, 2023, 2024) includes extensive audits with human coders evaluating the odds of type 1 (i.e., false-positive) and type 2 (i.e., false-negative) classification errors. These audits support the validity and reliability of our classification algorithms.

earnings call transcripts that mention “risk,” “uncertainty,” or their synonyms:⁴

$$(2) \text{ OverallRisk}_{i,t} = \frac{\text{Count of sentences mentioning “risk,” “uncertainty,” or synonym}_{i,t}}{\text{Total number of sentences}_{i,t}}$$

During our sample period, discussions of risk and uncertainty on earnings conference calls show distinct trends. The early years (2002-2007) reflect the “great moderation” (Benati and Surico, 2009; Galí and Gambetti, 2009), with fewer than 1.6% of sentences on average mentioning risk or uncertainty. This changes dramatically in 2008 Q4 when the Global Financial Crisis leads to a peak of 2.2%. Additional peaks occur during the European Sovereign Debt Crisis of 2011-2012 and the Covid-19 pandemic, reaching up to 2.8%.

At this aggregate level, this basic tally of risk and uncertainty discussions during earnings conference calls by publicly listed international firms shows a strong correlation with the Chicago Board Options Exchange (CBOE) volatility index (VIX): $\rho = 0.635$ (not tabulated). Compared to VIX and other aggregate indicators, a big advantage of our approach is the ease of disaggregation. Starting with generic unconditional risk discussed in the earnings call, we can easily obtain disaggregated metrics down to the firm, sector, and regional level (independently of whether the firm has options traded on its stocks), and also by topic.

Figure 1, Panel B illustrates such a topic decomposition by enumerating only those sentences that feature a discussion of uncertainty and risk that also contain the word “Brexit.” This decomposition forms the basis for our paper investigating the global impact of Brexit uncertainty, Hassan et al. (2024), in which we analyze the effects of uncertainty related to the UK’s departure from the European Union (EU) on listed firms from 82 different countries.⁵

The time-series pattern shows distinct peaks in Brexit-related concern expressed in earnings call discussions immediately after the UK’s vote to leave the EU in 2016. Although

⁴We use the Oxford English Dictionary to obtain all single-word synonyms of “risk,” “risky,” “uncertain,” and “uncertainty” (excluding “question” and “questions”). In total we have 131 synonyms, the top 10 of which (when sorted by frequency) are: “variable,” “possibility,” “pending,” “chance,” “doubt,” “prospect,” “exposed,” “variability,” “likelihood,” and “threat.” This is as in Hassan et al. (2019).

⁵Another recent example of *TopicRisk* is a measure of trade policy uncertainty (Hassan et al., 2019; Correa et al., 2023).

average *BrexitRisk* subsides in 2017, it increases sharply in the second half of 2018 when the details of Theresa May’s agreement with the EU became clear, along with the difficulties in obtaining parliamentary approval. In 2019, the prospect of a no-deal Brexit under Boris Johnson kept uncertainty high until Brexit finally took effect on January 31, 2020. Excerpts from top-scoring earnings-call transcripts explicitly highlight these concerns and enable us to pinpoint these periods of high uncertainty to specific Brexit-related events.

More generally, we refer to the number of sentences in firm i ’s quarter t earnings call that mention a given topic (e.g., Brexit, Russia’s invasion of Ukraine, artificial intelligence) in conjunction with a synonym of risk or uncertainty, as *TopicRisk*. Specifically:

$$(3) \text{ TopicRisk}_{i,t} = \frac{\text{Count of sentences on a given topic and a risk/uncertainty synonym}_{i,t}}{\text{Total number of sentences}_{i,t}}$$

This approach to identifying topic-specific risk discussions can provide powerful insights into the uncertainty generated by specific shocks and events. Moreover the time series of topic-specific risk are “self-labeling,” allowing us to better understand each peak by systematically reading the excerpts driving it. Building on this conditional analysis of text, we can further distinguish whether a given shock represents “good news” or “bad news” for a firm, capturing its first-moment impact.

Sentiment

Motivated by [Loughran and McDonald \(2011\)](#), who developed a sentiment lexicon tailored to financial text, we repurposed this lexicon-based sentiment analysis to understand the directional impact (i.e., positive versus negative) of specific shocks at the firm level.⁶

Following our approach to measuring topic-specific risk (see expression 3), we capture textual sentiment for topic-related sentences, denoted *TopicSentiment*, as:

⁶[Loughran and McDonald \(2011\)](#) recognized that words like “liabilities” carry different connotations in business communication compared to everyday language.

$$(4) \quad TopicSentiment_{i,t} = \frac{\text{Count of positive minus negative sentences on a given topic}_{i,t}}{\text{Total number of sentences}_{i,t}}$$

A sentence is flagged positive if it contains at least one positive tone word (e.g., “good,” “strong,” “great”) and no negative tone words, and negative if it has at least one negative tone word (e.g., “loss,” “decline,” “difficult”) and no positive tone words.

Figure 2 demonstrates how $TopicSentiment_{it}$ can assess the directional impact of shocks or events on various firms, sectors, or countries.^{7,8} Panel A compares the average sentiment firms associate with different shocks: Russia’s invasion of Ukraine in 2022, the 2020 outbreak of Covid-19, the 2016 prospect of Brexit, and major artificial intelligence (AI) developments (e.g., breakthroughs in Generative Pre-trained Transformer (GPT) models, including ChatGPT) in 2023.⁹

AI-related discussions are generally more positive, with 42.1% of all AI sentences containing positive sentiment words and only 10.2% negative, resulting in a net sentiment of +0.32. This contrasts with the significantly negative average sentiment for Brexit (2016), Covid-19 (2020), and Russia (2022, postinvasion). Whiskers indicate standard errors, estimated in panel regressions.

Despite these average sentiments, substantial heterogeneity emerges when decomposed by firm, country, and sector. Panel B of Figure 2 shows highly positive AI-related discussions for Alphabet, NVIDIA, and Microsoft, companies heavily involved in AI development. AI discussions are quite negative for firms concerned about intellectual property rights, such as News Corp, Universal Music Group, and Warner Music Group.

Panels C and D of Figure 2 demonstrate the usefulness of aggregated textual analysis

⁷Executives have gradually adopted more optimistic vocabulary in earnings calls over time (Cao et al., 2023), resulting in a noticeable trend in *OverallSentiment* (not shown for brevity). In most applications, this inflation of positive language can be easily controlled by including time fixed effects.

⁸Another potential concern is the use of negation, such as “not good” or “not terrible” (Loughran and McDonald, 2011). However, we have found this to be rare, with negligible effects once measures are aggregated to the transcript (firm-quarter) level.

⁹Taskin and Ruch (2023) measure demand- and supply-related sentiment in earning calls.

by focusing on Russia (2022) and Covid-19 (2020) sentiment at country (Panel C) and industry sector (Panel D) levels, respectively.¹⁰ While firms generally associate negative sentiment with Russia-related geopolitical developments and economic uncertainties, this effect is strongest for firms in Finland, Germany, and Denmark, with concerns revolving around regional security, energy markets and supplies, and economic ties. Targeted reading of the underlying transcripts reveals that Japanese and South Korean firms are concerned about the impact on trade access to, from, and through Russian territory.

Covid-19, another negative shock, shows varied impacts across industries even in its early stages. Panel D reveals anticipated benefits in sectors like Video Tape Rental and Offices and Clinics of Doctors and Medicine. This effect may be due to an increase in the demand for streaming and health care services. Contrasting with these benefits are expected economic headwinds in Life Insurance and Motion Picture Theaters sectors, with Covid mortality-related payouts and restrictions on public gatherings taking their tolls.

Lexicon Construction

Each of our measures of topic exposure, risk, and sentiment requires a systematic procedure for flagging sentences that mention a given topic. For this task, researchers have three basic options.

Keyword-based approach. The most straightforward approach is keyword-based, especially when the topic of interest pertains to a specific, well-defined event or policy.¹¹ To digress briefly, this process usually begins with a set of annotated sentences obtained from human readers, flagging those that discuss the topic of interest. Next, we select keywords or phrases that closely mimic this human classification within the sample of previously annotated text. These keywords or phrases are frequent in sentences discussing our topic and

¹⁰To measure *RussiaExposure*, we search the text of earnings calls for a list of about a hundred frequently used uni- and bigrams related to Russia, identified through word embeddings trained on our repository of earnings-call transcripts.

¹¹For example, in Hassan et al. (2024) the single keyword “Brexid” sufficed to accurately identify conversations on that topic.

infrequent in those that do not. To check their effectiveness out-of-sample, we sample new sentences from our corpus and assess the rate of false positives and false negatives. This process may prompt changes to our set of keywords, followed by further sampling and iteration until we identify a set that performs well on our corpus.¹² To illustrate, to identify supply chain-related discussions in Hassan et al. (2023), we started with “supplier” and “supply” and, thus, iteratively built a comprehensive set of keywords. While flexible, this approach can be quite labor intensive.

An alternative approach is semi-supervised, which does not necessitate beginning with an annotated sample of sentences. In Kalyani et al. (2023), for example, to identify job postings for research and development positions, we started with a set of phrases (“research,” “and develop,”...), and used a large language model (LLM) to identify other phrases commonly used in similar context – effectively using the LLM like a custom-trained thesaurus. We, then, iteratively refined our list by sampling excerpts to check for false positives and adding suggested phrases that met a threshold of true positives.

Similarly, in Sautner et al. (2023), we employed a machine learning approach to identify climate change discussions in earnings-call transcripts, again starting with a small set of climate-related words, and iteratively expanding this set to obtain a comprehensive library of relevant phrases. In addition to its capability to capture context-specific language and track the evolution of climate change discourse, researcher involvement in this approach is confined to selecting initial seed words.

Training libraries. If the topic of interest is more abstract or general, a keyword-based approach can often be unwieldy. In such instances, replacing human judgement with training libraries proves to be more straightforward. Here, the simplest approach to consider is the numerical statistic used in information retrieval and text mining: term frequency-inverse

¹²This systematic approach requires that the topic has existed long enough for researchers to annotate a sufficient number of sentences. For Covid-19 (Hassan et al., 2023), no such training data were available at the outbreak’s onset. Instead, we relied on a list of Covid-19 synonyms maintained by the World Health Organization, supplemented by cross-referencing with contemporaneous newspaper articles and online resources.

document frequency, or $tf \times idf$ in short (Song and Wu, 2008; Manning et al., 2008). We apply this approach in Hassan et al. (2019) to differentiate between political and non-political risk discussions, and in Hassan et al. (2023) to assess the risks firms associate with various countries. This approach requires a large corpus of text representative of discussions on the topic in question, such as textbooks from relevant fields (e.g., in Hassan et al. (2019), one on politics and another on non-politics). From these sources, we extract and tally the occurrences of bigrams (two-word combinations), identifying those that are frequent in one dataset but not in the other. This typically results in a list of thousands of topic-related bigrams, which are assigned weights proportional to their frequency in the training library. Compared to a keyword-based approach, reliant on a handful of manually selected keywords, a training library-based approach generally captures a much broader range of language used to discuss a topic. From our experience, this breadth allows for the detection of more subtle and indirect references to the topic.

Large language models. Recently, researchers have started to use LLMs either extrapolate from an initial annotated set (similar to Sautner et al. 2023) or, in the extreme, to feed an entire earnings-call transcript into ChatGPT and let it classify what topic a given text fragment (e.g., sentence, paragraph, section) is about (Jha et al., 2024). While promising, it remains unclear whether fully automated approaches outperform the lower-tech methods outlined above (e.g., Vafa et al., 2024).

In sum, different applications warrant slightly different methods for picking keywords. Our view is to err on the side of simplicity and use relatively straightforward approaches whenever feasible.

3. Cardinal Risk Comparison

Our approach enables a cardinal comparison of different risks, which is valuable in many contexts. The axes in Figure 1, Panels A and B, are directly comparable, both measured as

percentages of total sentences in earnings calls. For instance, if in 2016 Q3, listed firms globally averaged 6.85 sentences discussing risks, with 0.098 related to Brexit, we can conclude that Brexit contributed 1.43 percent to overall risk discussions at that time.

This cardinal decomposition, which assigns a share of overall risk to specific sources, is new in the literature.¹³ It quantifies the time devoted to discussing risk related to particular topics during earnings calls. Panel C of Figure 1 illustrates this approach, showing the share of risk-related sentences attributable to Covid-19, inflation, supply chain disruptions, Brexit, trade policy, Russia, and AI. It reveals the evolving prominence of different risk sources over time. Trade policy risk peaked during the Trump years at about 0.98 percent of risk mentions in 2018 Q3, comparable to Brexit’s impact. The Covid-19 crisis followed, accounting for 9.88 percent of all risk sentences in 2020 Q2. Later, in 2021, Covid-19’s direct impact gave way to uncertainty about inflation and supply chain challenges. Globally, these dwarf uncertainty related to Russia’s invasion of Ukraine, though this pattern reverses for German firms, where 4.02% of risk sentences in 2022 Q1 relate to Russia. AI discussions remain a minor risk driver for the average firm in our sample.

4. First- vs. Second-Moment Impacts

The ability to use text to break down shock exposure represents notable progress. From an econometric perspective, it is important to consider both the uncertainty (second-moment) effect and the first-moment impact of a shock (Berger et al., 2019). As Hassan et al. (2019) highlight, a key difficulty in isolating the second-moment effect is the potential strong correlation between innovations in shock variance and innovations in their conditional mean. For instance, consider a music company that becomes aware of advances in AI models capable of imitating artists’ voices. These advances represent both bad news (a potential drop in the net present value of its earnings due to competition from AI-produced content) and increased

¹³The closest analog is attributing excess returns to risk factors in empirical asset pricing. However, these methods generally do not allow an assessment of which shock is responsible for what share of the overall uncertainty faced by a given firm at a particular time, even when combined with an event study.

uncertainty (the potential for legal action or changes in legislation to protect its copyright).

If we observe a drop in the music company’s hiring or planned investments, it raises a crucial question: Is this decline due to the anticipated negative impact of the new technology on its business (first-moment impact), or is it due to uncertainty about future legislation or court decisions? Disentangling these two factors is key for policymakers deciding whether to pass new legislation or not.

To illustrate these ideas, Table 1 presents a compilation of excerpts related to AI from earnings-call transcripts, organized into Panels A through C depending on whether management expresses uncertainty, negative or positive sentiment, respectively, toward AI. The excerpts highlight the concerns of Spotify and News Corp, prominent players in the media and entertainment industry, about uncertainty about who will own the copyright and intellectual property of AI’s creations. In terms of negative sentiment, Chegg Inc., an educational support services company, experienced a significant drop of more than 40% in its shares after management announced during the earnings call for 2023 Q1 that ChatGPT has a detrimental impact on its business (Min, 2023). Lastly, the management of Alphabet and Microsoft expresses their anticipation of benefiting from AI, such as exploring new business opportunities.

Table 2 presents ordinary least squares regressions of the quarterly capital investment rate of firms against risk and sentiment measures for a sample of publicly listed US companies. In the first column, we report the association between the investment rate and our text-based measure of unconditional risk: a higher overall risk is negatively correlated with investment (coeff.=−0.039, s.e.= 0.006). Faced with a spike in uncertainty, firms prefer to ‘wait and see’ before committing to any new hiring or investment projects (e.g., Pindyck (1991); Dixit and Pindyck (1994)). Moving on to column 2, in which we examine specifically AI sentiment and AI risk. Breaking down the AI shock into risk and sentiment components, we see that, on the margin, additional risk arising from AI also has a significantly negative effect on investment (coeff.=−0.333, s.e.=0.125). At the same time, we see firms with positive AI

sentiment significantly expanding their investment (coeff.=0.073, s.e.=0.041), while we do not see a significant association between negative AI sentiment and investment (though the coefficient has the intuitively plausible negative sign).

In combination with our summary statistics from Figure 2, we have now gleaned some insight into firms' attitudes toward AI and how recent advances in AI are shaping the investment activities of listed firms.

Column 3 expands on our previous work on Brexit ([Hassan et al., 2024](#)). The Brexit referendum took place in 2016. The UK and the EU ratified the withdrawal agreement at the end of 2019, and the UK officially left the European Union in 2020. In our paper, we analyze the effects of Brexit-related uncertainty up to the ratification of the exit agreement. We document significant impacts of this uncertainty on investment and employment growth in firms outside the UK. For these firms, Brexit acted as an uncertainty shock until the terms of Brexit became clear. During the pre-ratification period, we observe a large negative coefficient on Brexit Risk (coeff.=-0.298, s.e.=0.136) but no significant effects of Brexit sentiment, indicating that Brexit mainly delayed firm investment.

With the benefit of hindsight, we now see that after the implementation of hard Brexit, the effect of Brexit uncertainty turns insignificant. Instead, Brexit develops into a negative shock for affected firms: the coefficient on negative Brexit sentiment is negative and significant (coeff.=-0.256, s.e.=0.154), highlighting the adverse impact on firms following this major political event.

5. Shock-Induced Risk Transmission and Contagion

Another promising application of our methodology is to measure the risk that international firms associate with a specific country ([Hassan et al., 2023](#)). This approach makes visible how a crisis in one country affects risk perceptions of firms in other countries.

In Panel A of Figure 4, we provide an example of shock-induced risk transmission. The Figure illustrates how sudden geopolitical events, such as the invasion of Ukraine by Russia

on February 24, 2022, can ripple through the global economy, showing the varying effect on firms in different countries.

The figure plots the relationship between the average discussions of Russia risk before and after the invasion. Significant increases in risk discussions postinvasion, as indicated by the position of data points on the y-axis compared to their preinvasion levels on the x-axis, highlight the invasion’s impact on firms’ perceptions of risk associated with Russia.

The purple line represents the predicted value of postinvasion discussions based on preinvasion levels. Points above this line, such as the one representing Austria, indicate countries where the increase in Russia-related risk discussions exceeds expected values. Conversely, points below the purple line, such as Finland, represent countries where the increase is less than expected.

This representation allows us to quickly characterize the ripple effect that a crisis has on international firms. The 45-degree line represents the typical transmission of risk from Russia to firms in other countries during normal times. The farther a destination country is above this line, the more concerned its firms are about risks from Russia during the crisis. In our previous work, we refer to the distance between the 45-degree line and the median destination country as the crisis’ “global impact.” The slope of the regression line indicates the degree of “bilateral transmission,” showing how much more concerned traditionally affected countries are. Finally, the R^2 of the regression line measures the “regularity” of transmission, indicating how closely the crisis transmission follows the usual pattern during noncrisis periods ([Hassan et al., 2023](#)).

Overall, in the case of the Russian invasion, we observe a global impact of 1.51 (the median country’s distance to the 45-degree line), a high degree of bilateral transmission (the regression line’s steep slope of 6.97), and a high level of regularity (the fitted line’s R^2 of 73%). In other words, the surge in concerns about risks from Russia was significant but concentrated mainly in European firms with historically high economic interactions with Russia, with little unexpected transmission beyond this group.

Contrast these metrics with the Global Financial Crisis (GFC) originating in the United States in 2008 Q1-Q3. The GFC had a greater global impact (2.27), with all countries far above the 45-degree line, regardless of their historical exposure to the US. However, it exhibited more moderate bilateral transmission (the fitted line is close to one). Furthermore, the GFC stands out for its high degree of irregular transmission (0.55—much lower than most other crises in our sample), consistent with its global transmission pattern. Thus, the GFC was a more global crisis, affecting firms in many countries in unpredictable ways, whereas the crisis following Russia’s invasion of Ukraine had more concentrated impacts, following more predictable paths.

These varying transmission patterns reflect the complex nature of risk transmission, influenced by factors such as financial dependence, trade connections, geographic proximity, dependence on natural resources, and historical political relations. For instance, significant increases in Russia-related risk discussions in countries farther from the war in Ukraine, such as Italy and Singapore, underscore the sometimes unexpected nature of risk transmissions across global networks. An examination of the underlying text reveals that Singaporean firms are particularly concerned about market access to and through Russia.) ¹⁴

Taken together, Figure 4 captures the increased global business anxiety in response to the two crises, offering a quantifiable measure of how geopolitical instability can disrupt global economic interactions and alter corporate risk exposures.

6. Earnings Conference Calls: Pitfalls and Lessons Learned

Cheap Talk

Economists typically approach the spoken word with some skepticism, as economic theory emphasizes scenarios in which agents have significant incentives to misrepresent or fabricate

¹⁴We study the global impact of the Brexit and the fallout of the 2011 Fukushima nuclear disaster in (Hassan et al., 2024). We find similar, seemingly surprising patterns for these shocks with noticeable effects in countries that at first sight might appear relatively detached from the event but upon closer inspection have significant connections.

information in “cheap talk” (Crawford and Sobel, 1982). Although we initially shared this view, our findings, and those found in other studies, suggest that discussions in earnings conference calls generally reflect decision makers’ genuine beliefs. Thus, on the whole, it appears sensible to take what call participants say at face value.

First, during earnings calls, management does not have unrestricted freedom to speak. Financial analysts ask questions and there are consequences when managers evade them. Academic research shows that market reactions are immediate to the way managers respond to inquiries, and analysts often follow up with additional questions or expand on queries that were not fully addressed (Hollander et al., 2010; Mayew et al., 2020; Lee, 2016; Bochkay et al., 2020). Fumbling or evading specific questions can lead to immediate stock price reactions. The outright lying or willful withholding of pertinent information is illegal and can have serious consequences. Moreover, our method counts the words of management *and* analysts. Even when managers want to avoid answering questions about a topic, analysts’ questions will count toward the exposure, risk, and sentiment measures.

Second, earnings calls are a repeating game, in which executives know that they will interact with the same analysts and shareholders in the future. Poor performance by executives can significantly harm their reputation (Kim et al., 2023) and future career prospects. In this sense, earnings calls are closer to a setting where participants have an incentive to communicate truthfully. As a case in point, we often see executives volunteering unfavorable information in their presentation, presumably so that the management team can discuss and frame such negative information on their own terms, rather than being forced to disclose in response to probing analyst questions.¹⁵

Some straightforward remedies have proven effective in our research in case concerns remain about the impact of strategic disclosure on textual measures. For example, studies in finance suggest that (unexpected) declines in a firm’s financial performance increase the

¹⁵As with surveys, earnings conference calls reflect the subjective opinions and beliefs of key decision-makers. These viewpoints may diverge from the firm’s objective economic fundamentals. However, even if the executives’ views are biased, they are still likely to drive the firm’s future decision making, serving as a crucial but difficult-to-observe ingredient for many economic models (Flynn and Sastry, 2022).

incentive of managers to distract the attention of shareholders away from reported earnings by misrepresenting information during earnings calls. Therefore, a reasonable check is to evaluate whether estimates on the measured exposure, risk or sentiment change when incorporating controls for recent earnings surprises and recent stock returns patterns in the period leading up to the earnings call ([Hassan et al., 2019](#); [Sautner et al., 2023](#)).¹⁶

Pre-Processing

In contrast to many other text-analysis applications, our experience with earnings-call transcripts has shown that it is generally not helpful to pre-process the text by stemming, lemmatizing, removing punctuation, etc. For example, unlike conventional practice, we find it useful to preserve capitalization (e.g., to distinguish the animal “turkey” from the country “Turkey”). Similarly, since our methods are simple and do not require much computational power, we have not found much benefit in removing stop words, punctuation, or stemming words. Instead, we prefer to keep the original text as intact as possible, and in fact structure our code in such a way that we can, at any time, call up and read the text underlying each data point. This practice has proven extremely valuable when interpreting time variation in our constructed measures: any variation in text-based measures can then be further examined by targeted reading of the underlying text.

Measurement Error

Measures constructed from text are subject to measurement error. The degree to which we need to worry about measurement error depends on the application. For example, aggregating measures from the firm-quarter level to sectors or countries reduces measurement error. Moreover, classical measurement error is generally less of a concern when using panel data for many firms, many time periods, and text-based generated from multiple pages of text.

When we formally test for measurement error in our political risk measures, we find that

¹⁶More ambitious approaches involve trying to deduce from the language (or voice tone) used whether managers are less than fully truthful in their communication (e.g., [Dikolli et al., 2020](#); [Hobson et al., 2012](#)).

about half of the variation at the firm-quarter level is attributable to measurement error. This level of accuracy is similar to typical measures constructed from firm-level accounting data, such as firm-level TFP measures (Bloom et al., 2018). The degree of measurement error likely depends on the topic. For instance, climate change exposure appears to suffer from much less measurement error compared to political risk (Sautner et al., 2023). This may be partly due to how prominently a given topic is discussed on conference calls (Huang et al., 2018).

7. Other Corporate Text Sources: Patents and Job Postings

Patents

Economists have long used patents as a measure of innovation, counting, for example, the number of patents filed in a specific location or the number of citations a given patent garners (Lerner and Seru, 2022; Hall et al., 2005). Recently, several authors have begun to analyze the content of the patent text itself to gain insight into its nature, scope, and impact. Free access to these patent texts is provided by The United States Patent Office (USPTO) along with information about the patent, such as the assigned firm, its inventor, location, and field of invention.

Kelly et al. (2021) use patent text to construct alternative measures of patent importance by analyzing the similarity to patents filed before and after. “Breakthrough” patents are characterized by language distinct from previous patents (low backward similarity) and strongly influencing subsequent patents (high forward similarity).

Similarly, Kalyani (2023) uses the share of new technical terminology in patents to distinguish creative patents from derivative patents. He finds that creative patents—those introducing new technical concepts—are strongly associated with firm- and sector-level TFP growth and abnormally positive stock returns. In contrast, derivative patents, which lack new technical language, do not show these associations.

This approach sheds light on a key puzzle about US economic growth. Despite a steep increase in patents filed per capita, TFP growth has slowed in recent decades. The share of derivative patents in overall patenting activity has increased dramatically, leading to a secular decline in the creativity of patents filed in the United States. This shift toward derivative patents, which lack the innovative content of creative patents, explains the slowdown in TFP growth despite the rise in patent filings.

[Mann and Püttmann \(2018\)](#) use the patent text to classify patents into automation or non-automation inventions. They first read and manually classify a randomly selected sample of patents, and then use this coding to train a (Naive Bayes) supervised learning algorithm to automatically apply this classification to a large number of patents. Using this approach, they calculate a sector-level automation technology index, and, taking as given the sectoral composition of different local economies, evaluate the effects of automation technology on local labor markets.

Combining Text Sources

Several recent papers advance this agenda by integrating the analysis of full patent texts with other text sources to address more complex questions, such as the origins of technology and its effects on the labor market.

For example, [Autor et al. \(2024\)](#) and [Kogan et al. \(2021\)](#) measure complementarities between patenting and occupation using the similarity between the text of patents and the text of task descriptions. [Kogan et al. \(2021\)](#) use the similarity between patent text and task descriptions of routine and nonroutine occupations to identify labor-saving and labor-augmenting inventions. They create a measure of the time-varying exposure to these inventions at the occupation level and evaluate their effects on earnings and employment. Labor-saving inventions reduce employment and earnings, while labor-augmenting inventions have the opposite effect, highlighting the varied impacts of technological advancements on the labor market.

Perhaps the most promising advancement is the integration of online job postings text. Data vendors like Lightcast provide vast databases of hundreds of millions of current and historical online job postings, some dating back to 2007. As job postings have shifted from printed newspapers and classifieds to online platforms, these databases have become more comprehensive and representative of open positions in the US and international labor markets (Hershbein and Kahn, 2018). They offer a variety of information, including firm names, salary levels, job locations, educational requirements, and desired skills. Crucially, they also provide the original text of the job postings, which is particularly valuable for our purposes.

In a recent paper, Kalyani et al. (2023) intersect the full text of patents, job postings, Wikipedia, and earnings call transcripts to identify new technologies developed in the last four decades and trace their diffusion across jobs in different locations, at different firms, and of different skill levels. Their approach is to (i) identify phrases associated with a newly invented technology, and (ii) characterize this technology and its impact by tracing the diffusion of these phrases through multiple corpora of text.

This approach is versatile and can be adapted to address many research questions. Next, we sketch its outlines. First, we use the full text of US patents between 1976 and 2014 to isolate phrases that appear in multiple patents but did not exist before 1970, thus identifying new terminology specific to influential innovations from the past 40 years. These phrases contain terms that describe new technologies, newly recognized problems (such as “climate change”) or new management terms (such as “performance metrics”). Second, to isolate phrases that describe new technologies, we then search these phrases on Wikipedia and exploit the standardized nature of Wikipedia technology pages to determine those phrases that are primarily associated with pages describing new technologies. This procedure identifies phrases associated with 1,286 unique new technologies developed since 1976.

In the second step, we use this list of new technologies and their associated phrases to identify patents and job postings that mention them. We use patent inventor addresses to determine the locations where each technology was developed and patent application years

to pinpoint when the technology first experienced a significant increase in patent citations.

Essentially, we build a dataset of new technologies by identifying newly appearing phrases in patents. We then link these phrases to patents assigned to firm i with inventor j living in location c in year t .

Finally, we cross-reference this list of technology phrases with the full text of online job postings to identify 51 million jobs advertised between 2010 and 2019 that mention new technologies. These granular data allow us to track the number, location, and skill requirements of job postings associated with each new technology.

Figure 5 illustrates these relations for one of the 1,286 new technologies: AI. Panel A plots a given Core Based Statistical Area’s (CBSA) share of patents mentioning AI prior to 2014 against its share of jobs mentioning AI post-2014. In each case, we normalize by the size of the local labor market so that the “normalized share” of jobs and patents indicates the local over- or under-representation of AI patents and jobs in a given area.

We observe a strong association between AI’s early development and subsequent job creation related to the technology. Places like San Francisco, Boston, Philadelphia, and Austin pioneered AI patents, and later hosted the lion’s share of AI-related jobs.

This pattern holds true for the remaining 1,285 other technologies, where the vast majority of cases show a strong link between early patenting in the technology and subsequent job growth in using and producing that technology. This advantage of pioneer locations is particularly concentrated in *high-skill* jobs related to the new technology and persists for multiple decades (Kalyani et al., 2023).

Panel B illustrates the same relationship at the firm level. Companies like Apple and Alphabet, which pioneered early development of the technology, show high levels of AI patenting before 2014 and subsequent job creation. However, the regression line’s slope is now significantly flatter, indicating positive spillovers from these companies to other local companies using and producing the technology. This finding suggests that the relationship is steeper at the CBSA- than at the firm-level.

Together, this combined textual source approach offers valuable insight into how technology is adopted and spreads. As shown in Figure 5, it reveals patterns of geographic technology distribution, highlighting the importance of innovation hubs. In addition, Bloom et al. (2021) emphasizes the critical role of universities in technological advancement and the symbiotic relationship between academia and industry in driving progress. These analyses provide a foundation for policymakers to develop informed workforce development strategies and address skill gaps by examining labor market responses to evolving technology.

Conclusion

Corporate texts (such as earnings call transcripts, patents, and job postings) offer a direct avenue for researchers and policymakers to access insights from firm executives, financial analysts, and other market participants on contemporary events. Currently, there is no comparable data set that gathers the opinions, preferences, and expectations of crucial corporate decision-makers on various key issues with the same frequency. This valuable text-based information can be gathered not only for a large sample of publicly traded companies in the United States but also for numerous international firms. Using these granular data to improve economic surveillance and forecasting is just one way to take advantage of them. The methods we have outlined can be employed to explore how risk is transmitted between firms, within industries, and across country borders. By referring back to the original text sources, researchers can corroborate their inferences, drawing on textual cues to understand how companies and market participants interpret a shock, what actions they plan to take, and their long-term expectations. Overall, numerous possible new applications for textual methodologies in economic surveillance await.

Tables and Figures


Table 1: Example AI-related excerpts from earnings-call transcripts

Panel A: Risk	
Spotify, 2023 Q1	“I guess on the risk side would be not just for Spotify, but I think for the entire creative ecosystem is obviously the question around copyrights and who owns what copyrights and what the fairway would be to attribute value when you’re doing things in name and likeness situations or inspired by a certain artist, etc.”
News Corp, 2023 Q3	“Candidly, generative AI may pose a challenge to our intellectual property and to the future of journalism. As those who’ve experimented with ChatGPT will be aware, the answers are only as insightful and factual as the source material and are more retrospective than contemporary.”
Panel B: Negative sentiment	
Universal Music Group, 2023 Q1	“Now with respect to the artist community and what we’re hearing, we’ve heard a number of ours expressing their concerns about AI being employed to create content that’s training on their work without consent. And this is really content that’s attempting to hijack their brand, exploit their celebrity for a moment of attention that’s given to the anonymous uploader, not to mention the numerous other very nefarious uses of their content with disrespect and tarnish their reputation.”
Chegg Inc, 2023 Q1	“However, since March we saw a significant spike in student interest in ChatGPT. We now believe it’s having an impact on our new customer growth rate. (...) [I]t’s prudent to be more cautious with our forward outlook.”
Panel C: Positive sentiment	
Alphabet Inc, 2023 Q1	“From Google Lens to multi-search, to visual exploration in Search, immersive view in Maps, Google Translate, to all the language models powering Search today, we have used AI to open up access to knowledge in powerful ways. We’ll continue to incorporate generative AI advances to make Search better in a thoughtful and deliberate way. We’ll be guided by data and years of experience about what people want and our high standards for quality.”
Microsoft Corp, 2023 Q3	“Our hiring business took share for the third consecutive quarter. The excitement around AI is creating new opportunities across every function for marketing, sales and finance to software development and security. LinkedIn is increasingly where people are going to learn, discuss and up-level their skills with more than 100 AI courses.”


Notes: This table shows AI excerpts with risk and uncertainty synonyms (Panel A), negative sentiment (Panel B), and positive sentiment (Panel C) words.

Figure 1: Text-based risk


Panel A: Unconditional risk



Panel B: Brexit risk




Panel C: Risk decomposition




Notes: This figure plots, as a percentage of sentences, unconditional risk (Panel A), Brexit risk (Panel B), and a breakdown of risk (Panel C) related to AI, Brexit, Covid-19, inflation, Russia, supply chain, and trade, for our sample of U.S.-headquartered firms.

Figure 2: Text-based sentiment


Panel A: Event




Panel B: Firm: AI (2023) sentiment



Panel C: Country: Russia (2022) sentiment



Panel D: Sector: Covid-19 (2020) sentiment



Notes: This figure plots event-topic sentiment in Panel A, measured as topic-specific (i.e., Russia (2022), Covid-19 (2020), Brexit (2016), or AI (2023)) sentiment divided by topic-specific exposure; AI (2023) sentiment, in percentage of sentences, by firm in Panel B; Russia (2022) sentiment, in percentage of sentences, by country in Panel C; and Covid-19 (2020) sentiment, in percentage of sentences, by sector (3 digit SIC) in Panel D. The whiskers in Panel A denote 95% confidence intervals around the average.


Table 2: Risk, sentiment, and investment rate

	Investment Rate _{<i>i,t</i>} * 100		
	(1)	(2)	(3)
<i>Unconditional Risk_{i,t}</i>	-0.039*** (0.006)	-0.039*** (0.006)	-0.039*** (0.006)
<i>AI Risk_{i,t}</i>		-0.333*** (0.125)	
<i>AI PositiveSentiment_{i,t}</i>		0.073* (0.041)	
<i>AI NegativeSentiment_{i,t}</i>		-0.002 (0.077)	
<i>Brexit Risk_{i,t}</i> - pre 2019			-0.298** (0.136)
<i>Brexit Risk_{i,t}</i> - post 2019			0.119 (0.326)
<i>Brexit NegativeSentiment_{i,t}</i> - pre 2019			0.095 (0.118)
<i>Brexit NegativeSentiment_{i,t}</i> - post 2019			-0.256* (0.154)
R-squared	0.191	0.191	0.191
N	84,492	84,492	84,492


Notes: This table reports estimation results from a regression of investment rate for firm i at time (quarterly) t on risk and sentiment variables. Investment rate is calculated as the ratio of capital expenditures reported by firm i at time t to the stock of property, plant and equipment reported at time $t - 1$. All risk and sentiment variables are in units of percentage of sentences in earnings calls. *Unconditional Risk_{i,t}* is the percentage of sentences with risk synonyms in an earnings call hosted by firm i at time t . *AI (Brexit) Risk* and *Sentiment* is the percentage of sentences in an earnings call with AI (Brexit) keywords mentioned in conjunction with a risk synonym as defined on page 6. *Brexit Risk* - pre 2019 (similarly for - post 2019) is the interaction of *Brexit Risk* with a dummy equal to one for the period before (post) 2019; zero otherwise. Standard errors are clustered by firm. All specifications include controls for sector (2 digit SIC) fixed effects and time (quarter) fixed effects. Only firms with at least 10 earnings calls during 2015-2023.

Figure 4: Risk contagion

(a) Transmission of Russia country risk during the Ukraine war




(b) Transmission of US country risk during the Global Financial Crisis



Notes: This figure plots the transmission of country risk to countries around the world during two crises: (i) the first year of the war in Ukraine (Panel A) and (ii) the first three quarters of the Global Financial Crisis (Panel B). In each panel, the horizontal axis measures the risk transmitted to countries before the respective crisis quarters as the share of sentences about the country (weighted by the country ngram's weight) averaged across firms headquartered in the destination country. The vertical axis measures the same during the crisis quarters 2022 Q1 through 2023 Q1 for Panel A and 2008 Q1 through 2008 Q3 for Panel B.

Figure 5: Machine learning/AI jobs and patents



Notes: In Panel A, figure plots a binned scatter plot of the normalized share of jobs posted in a CBSA between 2015-2019 that mention one of the 'Machine learning/AI' keywords against the normalized share of patents that mention the same keywords filed in a CBSA between 1976-2014. Normalized share of jobs for each CBSA is calculated as $\text{Normalized share jobs}_{c,\tau} = \frac{\text{Share of ML/AI Jobs}_{c,\tau}}{\text{Share of all jobs}_{c,\tau}}$ - the share of jobs in a CBSA posted during 2015-2019 which mention keywords associated with machine learning out of all such jobs in the US divided by the share of jobs posted in the CBSA out of all jobs posted in the US during 2015-2019. Normalized share of patents for each CBSA is calculated as the share of patents in a CBSA filed during 1976-2014 which mention keywords associated with machine learning divided by the share of all patents in the CBSA filed during 1976-2014. Panel B replicates Panel A using firm-level data in place of CBSA level data.

References

- Ahn, J., B. Carton, A. Habib, D. Malacrino, D. Muir, and A. Presbitero (2023, May). Geoeconomic fragmentation and foreign direct investment. In *World Economic Outlook*, Chapter 4, pp. 2–45. International Monetary Fund.
- Albrizio, S., J. Bluedorn, A. Dizioli, C. Koch, and P. Wingender (2023, October). Managing expectations: Inflation and monetary policy. In *World Economic Outlook*, Chapter 2, pp. 49–69. International Monetary Fund.
- Albrizio, S., A. Dizioli, and P. V. Simon (2023, October). Mining the gap: Extracting firms’ inflation expectations from earnings calls. Working Paper 202, International Monetary Fund.
- Alfaro, L. and D. Chor (2023, August). Global supply chains: The looming ‘great reallocation’. Working Paper 24-012, Harvard Business School. in proceedings Jackson Hole Symposium, 2023 (also NBER WP 31661).
- Andersson, M., P. Neves, and C. Nunes (2023, June). Earnings calls: New evidence on corporate profits, investment, and financing conditions. *European Central Bank Economic Bulletin* 4, 56–62.
- Autor, D., C. Chin, A. Salomons, and B. Seegmiller (2024). New frontiers: The origins and content of new work, 1940–2018. *The Quarterly Journal of Economics*, qjae008.
- Bae, J., C.-Y. Hung, and L. van Lent (2023). Mobilizing text as data. *European Accounting Review* 32(5), 1085–1106.
- Benati, L. and P. Surico (2009). VAR analysis and the Great Moderation. *American Economic Review* 99(4), 1636–1652.
- Berger, D., I. Dew-Becker, and S. Giglio (2019, 04). Uncertainty shocks as second-moment news shocks. *The Review of Economic Studies* 87(1), 40–76.
- Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2018). Really uncertain business cycles. *Econometrica* 86(3), 1031–1065.
- Bloom, N., T. A. Hassan, A. Kalyani, J. Lerner, and A. Tahoun (2021). The diffusion of disruptive technologies. *National Bureau of Economic Research w28999*.
- Bochkay, K., J. Hales, and S. Chava (2020). Hyperbole or reality? Investor response to extreme language in earnings conference calls. *The Accounting Review* 95(2), 31–60.
- Cao, S., W. Jiang, B. Yang, and A. Zhang (2023). How to talk when machines are listening: Corporate disclosure in the age of AI. *Review of Financial Studies* 36(9), 3603–3642.
- Correa, R., J. di Giovanni, L. S. Goldberg, and C. Minoiu (2023, November). Trade uncertainty and U.S. bank lending. Staff Reports 1076, Federal Reserve Bank of New York.
- Crawford, V. P. and J. Sobel (1982). Strategic information transmission. *Econometrica* 50(6), 1431–1451.
- Dikolli, S. S., T. Keusch, W. J. Mayew, and T. D. Steffen (2020). CEO behavioral integrity, auditor responses, and firm outcomes. *The Accounting Review* 95(2), 61–88.

- Dixit, A. K. and R. S. Pindyck (1994). *Investment under uncertainty*. Princeton university press.
- Faella, F., B. Javorcik, L. Kitzmüller, V. Langella, and H. Schweiger (2022, November). Global supply chains in turbulence. In *European Bank for Reconstruction and Development Transition Report*, Chapter 3, pp. 59–61. European Bank for Reconstruction and Development.
- Flynn, J. P. and K. Sastry (2022). Attention cycles. *Available at SSRN 3592107*.
- Frankel, R., M. Johnson, and D. Skinner (1999). An empirical examination of conference calls as a voluntary disclosure medium. *Journal of Accounting Research* 37, 133–150.
- Galí, J. and L. Gambetti (2009). On the sources of the Great Moderation. *American Economic Journal: Macroeconomics* 1(1), 26–57.
- Gallemore, J., S. Hollander, M. Jacob, and X. Zheng (2024). Tax policy expectations and investment. *Working paper* <https://ssrn.com/abstract=3775710>.
- Gosselin, M.-A. and T. Taskin (2023, June). What can earnings calls tell us about the output gap and inflation in Canada? Technical report, Bank of Canada.
- Hall, B. H., A. B. Jaffe, and M. Trajtenberg (2005). Market value and patent citations. *The RAND Journal of Economics* 36, 16–38.
- Hassan, T. A., S. Hollander, L. v. Lent, and A. Tahoun (2024). The global impact of Brexit uncertainty. *The Journal of Finance* 79(1), 413–458.
- Hassan, T. A., S. Hollander, L. van Lent, and A. Tahoun (2019). Firm-level political risk: Measurement and effects. *Quarterly Journal of Economics* 134(4), 2135–2202.
- Hassan, T. A., S. Hollander, L. van Lent, and A. Tahoun (2023). Firm-level exposure to epidemic diseases: Covid-19, SARS, and H1N1. *Review of Financial Studies* 34(6), 2409–2449.
- Hassan, T. A., J. Schreger, M. Schwedeler, and A. Tahoun (2023). Sources and transmission of country risk. *Review of Economic Studies*, rdad080.
- Hershbein, B. and L. B. Kahn (2018). Do recessions accelerate routine-biased technological change? Evidence from vacancy postings. *American Economic Review* 108(7), 1737–1772.
- Hobson, J. L., W. J. Mayew, and M. Venkatachalam (2012). Analyzing speech to detect financial misreporting. *Journal of Accounting Research* 50(2), 349–392.
- Hollander, S., M. Pronk, and E. Roelofsen (2010). Does silence speak? An empirical analysis of disclosure choices during conference calls. *Journal of Accounting Research* 48(3), 531–563.
- Huang, A. H., R. Lehavy, A. Y. Zang, and R. Zheng (2018). Analyst information discovery and interpretation roles: A topic modeling approach. *Management Science* 64(6), 2833–2855.
- Jha, M., J. Qian, M. Weber, and B. Yang (2024). ChatGPT and corporate policies. *Working paper Chicago Booth*.

- Kalyani, A. (2023). The creativity decline: Evidence from US patents. <http://dx.doi.org/10.2139/ssrn.4318158>.
- Kalyani, A., N. Bloom, M. Carvalho, T. Hassan, J. Lerner, and A. Tahoun (2023). The diffusion of new technologies. <http://dx.doi.org/10.2139/ssrn.3875827>.
- Kelly, B., D. Papanikolaou, A. Seru, and M. Taddy (2021). Measuring technological innovation over the long run. *American Economic Review: Insights* 3(3), 303–320.
- Kim, A. G., M. Muhn, and V. V. Nikolaev (2023). Bloated disclosures: Can ChatGPT help investors process information? *Chicago Booth Research Paper* (23-07).
- Kogan, L., D. Papanikolaou, L. D. Schmidt, and B. Seegmiller (2021). Technology-skill complementarity and labor displacement: Evidence from linking two centuries of patents with occupations. <http://www.nber.org/papers/w29552>.
- Lee, J. (2016). Can investors detect managers’ lack of spontaneity? Adherence to predetermined scripts during earnings conference calls. *The Accounting Review* 91(1), 229–250.
- Lerner, J. and A. Seru (2022). The use and misuse of patent data: Issues for finance and beyond. *The Review of Financial Studies* 35(6), 2667–2704.
- Loughran, T. and B. McDonald (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *Journal of Finance* 66(1), 35–65.
- Ma, Y. and K. Zimmermann (2023, September). Monetary policy and innovation. Working Paper 31698, National Bureau of Economic Research. 37 pages.
- Mann, K. and L. Püttmann (2018). Benign effects of automation: New evidence from patent texts. *Available at SSRN 2959584*.
- Manning, C. D., P. Raghavan, and H. Schütze (2008). Introduction to information retrieval. *Cambridge: Cambridge University Press*.
- Matsumoto, D., M. Pronk, and E. Roelofsen (2011). What makes conference calls useful? The information content of managers’ presentations and analysts’ discussion sessions. *The Accounting Review* 86, 1383–1414.
- Mayew, W. J., M. Sethuraman, and M. Venkatachalam (2020). Individual analysts’ stock recommendations, earnings forecasts, and the informativeness of conference call question and answer sessions. *The Accounting Review* 95(6), 311–337.
- Min, S. (2023). Clegg shares drop more than 40% after company says ChatGPT is killing its business. *CNBC May 2*.
- Pindyck, R. S. (1991). Irreversibility, uncertainty, and investment. *Journal of Economic Literature* 29, 1110–1148.
- Sautner, Z., L. van Lent, G. Vilkov, and R. Zhang (2023). Firm-level climate change exposure. *Journal of Finance* 78(3), 1449–1498.
- Song, M. and Y.-F. B. Wu (2008). Handbook of research on text and web mining technologies. *Hershey, PA: IGI Global*.

- Staff, I. (2022, October). Global prospects and policies. In *World Economic Outlook*, Chapter 1, pp. 1–49. International Monetary Fund.
- Taskin, T. and F. U. Ruch (2023, June). Global demand and supply sentiment: Evidence from earnings calls. Staff Working Paper 37, Bank of Canada.
- Vafa, K., J. Chen, J. Kleinberg, S. Mullainathan, and A. Rambachan (2024). Evaluating the world model implicit in a generative model. *Working paper*, arxiv.org/abs/2406.03689.