The Rise of Partisan Corporate Speech

William Cassidy and Elisabeth Kempf*

PRELIMINARY AND INCOMPLETE

November 24, 2022

Abstract

We develop a novel measure of partisan corporate speech using techniques from natural language processing. Using the entire corpus of tweets from companies listed on the S&P 500, we first establish a large increase in the amount of partisan corporate speech between 2011 and 2022. This increase in partisan speech is disproportionately driven by corporates using speech commonly associated with Democratic politicians; in particular, statements related to climate change as well as diversity, equity, and inclusion. We also explore how intraday stock returns respond to partisan corporate tweets.

*We thank Kai Ding, Niels Gormsen, Jingyu He, Ralph Koijen, Stefan Nagel, and Jesse Shapiro, as well as participants at the UChicago Research Computing Cluster Workshop and the Chicago Booth Asset Pricing Reading Group for helpful comments. We thank Donal O'Cafaigh and Michael Lin for excellent research assistance.

1 Introduction

Recent years have seen many cases of prominent U.S. companies and their CEOs taking a position on social and political issues, including gun laws (Lucas (2019)), voting rights (Gelles and Ross Sorkin (2021)), and racial justice (Hessekiel (2020)). Many of these issues are characterized by a deep partisan divide in Americans' attitudes on these issues (e.g., Pew Research (2019)). However, to date we lack a systematic approach to measure the prevalence of partisan corporate speech. In particular, it is challenging to separate the rise of partisan corporate speech from increased media attention and reporting on these issues.

In this paper, we propose a novel measure of partisan corporate speech, using natural language processing applied to corporate statements shared on social media. We ask three fundamental questions. First, has corporate speech become more partisan over time? Second, what topics do companies discuss when they make partisan statements? Third, how do stock prices respond to partisan corporate speech?

To answer these questions, we collect the corpus of tweets sent by S&P 500 companies with verified Twitter accounts between 2011 and 2021. To detect corporate partisan speech, we measure the degree of similarity in the language used by companies and the language used by Democratic and Republican politicians on social media. Specifically, we estimate multinomial inverse regression (MNIR) on tweets sent by Republican and Democratic politicians, and use the resulting estimates to identify corporate tweets that sound very similar to tweets sent by either Republican or Democratic politicians.

We observe a significant increase in the frequency of partisan corporate speech on Twitter between 2011 and 2021. Prior to 2016, corporate speech on Twitter shared a greater degree of similarity with speech by Republican politicians, mainly due to companies mentioning conservative media outlets, referring to economic indicators, or advocating for fossil fuels. After 2016, we see a strong increase in the amount of Democratic-sounding speech, with peaks around the death of George Floyd and the passage of new voting laws in Georgia. The disproportionate increase in the amount of Democratic corporate speech is particularly pronounced once we remove topics that do not have a direct connection to the operations of the company. Moreover, it is present across almost all industries, including industries with high and low levels of market concentration.

To better understand what companies discuss when they use partisan language, we decompose partisan corporate speech into distinct topics using biterm topic modeling. We find that most of the increase in Democratic-sounding speech is driven by increased discussion of diversity, equity, and inclusion (DEI), climate change, and health issues related to the COVID-19 pandemic. Republican-sounding speech that is not business-related focuses on patriotism and references to conservative media outlets, such as Fox Business and the Wall Street Journal.

Finally, we study how stock prices respond to partisan corporate tweets. Ex ante, the direction of the stock price response is not obvious. On one hand, partisan corporate tweets may be in the financial interests of shareholders, because they could increase loyalty towards the firm in the labor market, product market, or financial markets. On the other hand, such statements may reflect an agency problem between managers and shareholders, with managers acting in their own personal interests. A central challenge associated with studying the financial implications of partisan statements is that their timing may be endogenous. For example, companies may be more likely to issue a statement on a social issue when they have positive financial news to report. We overcome this challenge by exploiting the precise time stamps of the corporate tweets, which allow us to conduct a second-by-second analysis of returns surrounding the tweet. This high-frequency approach reduces concerns about potential confounding events, since these would have to occur within ten minutes of the partisan tweet. An important limitation of our high-frequency approach is that we are only able to capture a very short-term response from investors.

We find positive stock returns around Democratic-sounding tweets and negative returns around Republican-sounding tweets. The average cumulative return over the 20 minutes around the tweet is equal to 1.0 basis points for Democratic tweets and –1.5 basis points for Republican tweets. The difference is statistically significant at the 10% level. Tweets on DEI-related topics most strongly contribute to the overall positive reaction to Democratic tweets. Unsurprisingly, given that we are looking at stock returns over a very short period of time, these effects are economically small. However, these results are nevertheless useful because they suggest the increase in Democratic-sounding corporate speech may be driven by economic considerations rather than by agency problems. The fact that executives in publicly listed U.S. companies are predominantly Republican (Fos et al. (2022)) further supports the interpretation that executives may not be the driving force behind the trend toward more Democratic speech.

Our study contributes to several strands of the literature. First, we contribute to a small but growing literature that studies sociopolitical activism by companies and CEOs. Most of that literature has focused on activism by CEOs. In one of the first attempts to measure the phenomenon, Larcker et al. (2018) use multiple approaches to detect instances of CEO activism, including statements made on Twitter. However, they find that only 11 percent of all S&P 1500 CEOs have active personal Twitter feeds. In contrast, 84% of S&P 500 companies have an active Twitter account during our sample period. Existing studies of investor reactions to corporate and CEO sociopolitical activism have found mixed evidence, with some observing positive stock price reactions at daily frequencies (e.g., Mkrtchyan et al. (2022); Gangopadhyay and HomRoy (2022)) and others observing negative reactions from investors (e.g., Bhagwat et al. (2020)). To the best of our knowledge, our study is the first to apply natural language processing techniques to the entire corpus of tweets from corporate Twitter accounts, as well as to study investor reactions to those tweets at intraday frequencies.

Second, we contribute to a growing literature on the political polarization of corporate America. Fos et al. (2022) show that executive teams have become more politically homogeneous over the past decade. Moreover, a growing number of studies documents how political partisanship shapes individuals' views of the economy and their economic decisions, including in high-stakes, professional environments, such as credit analysts Kempf and Tsoutsoura (2021), asset managers Cassidy and Vorsatz (2021), loan officers Dagostino et al. (2020), and entrepreneurs Engelberg et al. (2021). The results in this paper suggest that U.S. companies are increasingly developing partisan identities, as measured by their speech on social media.

We also contribute to a literature that aims at measuring partisanship via speech. Gentzkow et al. (2019b) study how the speech used by members of Congress has become more polarized over time. Like Gentzkow et al. (2019b), we use MNIR to estimate the probability of using phrases by individuals with different party affiliations.[1] Different than Gentzkow et al. (2019b), we use MNIR for a prediction problem. Our aim is to use MNIR to programatically identify when corporations use speech similar to that of Democrats or Republicans, as opposed to measuring the extent to which speech is polarized across parties. Our approach is therefore more similar to that of Engelberg et al. (2022), who detect partisanship in the speech of financial regulators by identifying partisan phrases in Congressional speech and then observing their usage among regulators. To the best of our knowledge, we are the first to estimate MNIR on tweets in the economics and finance literature.

2 Data and Measure

2.1 Twitter

We begin by collecting the entire corpus of tweets sent by companies listed on the S&P 500 with verified Twitter accounts between 2011 and 2022. We are able to find a verified Twitter account for 632 out of 751 companies (84%).[2] We manually search Twitter for verified accounts with Twitter usernames or handles similar to the name of the firm. If a company has multiple twitter accounts, we map all Twitter accounts to the original firm. For

[1] Gentzkow et al. (2019b), in turn, build on other work in the statistics literature developing computationally feasible methods for estimating MNIR, notably Taddy (2013) and Taddy (2015).

[2] Twitter verifies Twitter accounts for companies and public officials. Once a twitter account is verified, we can be confident that the twitter account actually belongs to the entity that it purports to represent.

example, Alphabet has both the handles "@Google" and "@AlphabetInc". We map both handles to Alphabet. Combined, these 632 companies sent nearly 5 million tweets between 2011 and 2022.

We repeat the same procedure for all members of Congress between 2010 and 2021. There are 155 politicians who served in the Senate and 781 who served in the House of Representatives during this time frame. We are able to match 150 Senators and 721 Representatives to at least one verified Twitter account. When a Congressperson has more than one Twitter account (e.g., an official and a personal one), we use both accounts. Most politicians whom we are not able to match served early in the sample period, before the use of Twitter became ubiquitous among elected officials.

For every Twitter handle we collect, we download the full sample of tweets sent from that Twitter account using the Twitter application programming interface (API). Access to the Twitter API for academic research is granted through Twitter via an application process.

For every tweet, we observe whether the tweet was an original tweet, a retweet, a reply, or a quote tweet. We restrict our sample to original tweets sent by the company itself. We do not download retweets, quote tweets, or replies, because many of these are related to issues concerning customer service and thus less relevant for our exercise. In addition to the text of the tweet, every tweet we download contains the exact date and timestamp of the tweet, as well as a unique tweet ID assigned by Twitter. We also collect metrics designed to measure user engagement with the tweet: the number of times the tweet was retweeted, replied to, or quoted.

The Twitter API also provides information on any other Twitter handles referenced in a tweet, called a "mention," and the hashtags used in the tweet itself. We treat mentions and hashtags as any other bigram. If some hashtags or mentions are used disproportionately by politicians, then our measure will detect this and label their usage as partisan speech.

Table 1, Panel A, provides summary statistics for the sample of corporate tweets by year. The number of unique firms grows over time, as more companies establish Twitter accounts. The average number of tweets per firm is greater than one per day in all years. The distribution of the number of tweets is strongly right-skewed, with the mean being consistently larger than the median. A few firms send a very large number of tweets per day; many of these companies use their Twitter accounts for customer service.

Before constructing a measure of partisan corporate speech, we pre-process the raw text of each tweet in three steps. First, we tokenize each tweet. Tokenization is the process of breaking up a string that is a full sentence into individual tokens. This step effectively removes spaces and punctuation. Second, we remove "stop words"; that is, words that do not substantially contribute to the meaning of the sentence, such as "that" or "the." Third, we stem the remaining words. Stemming maps all words with the same stem, but possibly different suffixes or prefixes, to the same word. For example, both "becoming" and "become" are converted to "becom."

Next, we convert the set of words into n-grams. N-grams are N-length sequences of adjacent words. We use both unigrams and bigrams for different steps of the analysis. Unigrams contain only a single word, whereas bigrams include two words, an example of which is "big data."

2.2 Information on Elected Officials' Demographics

We collect additional demographic and biographical information on the elected officials in our sample by scraping the biographical directory of the United States Congress at https://bioguide.congress.gov. Specifically, we collect information on the official's home state, the highest educational degree attained, and age. To construct a proxy for a Congressperson's ethnicity, we use the python package "ethnicolr," which infers the ethnicity of individuals from their place of birth, state of residence, age, and name.

2.3 Intraday Stock Returns

To measure changes in stock market valuations around tweets, we use second-by-second stock returns based on the Trade and Quote data (TAQ) during a window spanning 10 minutes before and after each tweet. We access TAQ data through the WRDS intraday event study interface. WRDS imposes standard filters on the underlying TAQ data, such as requiring that no more than 20 percent of the underlying prices are missing within a 600-second window around the event. We winsorize cumulative returns at the 5% level.

3 Measure of Partisan Corporate Speech

Our measure of partisan corporate speech is designed to capture how similar the language used in a corporate tweet is to language used by Democratic or Republican politicians. Intuitively, if a corporate tweet uses language that is highly predictive of being used by a Democrat (Republican), then we will label this tweet as Democratic (Republican), respectively. To take this idea to the data, we use multinomial inverse regression (MNIR), a method from natural language processing (NLP) that has also been applied to detect partisan speech in Congress (Gentzkow et al. (2019b)). We first estimate MNIR on tweets sent by Republican and Democratic politicians to find bigrams that are highly associated with usage by either party. We then use the estimated model to detect partisan tweets by corporates.

After estimating MNIR, we also implement topic modeling. We use topic models to group partisan corporate tweets by their subject matter. We train two topic models on the set of corporate tweets that MNIR classifies as very Republican or Democratic, respectively. We describe our method in more detail below.

3.1 Multinomial Inverse Regression

Following the approach in Taddy (2015), we assume that bigram counts (c_{it}) sent by tweeter i at time t are drawn from a multinomial distribution:

$$\boldsymbol{c}_{it} \sim \mathrm{MN}\left(m_{it}, \boldsymbol{q}_t^{P(i)}\left(\boldsymbol{x}_{it}\right)\right). \tag{3.1}$$

There are J total bigrams that the speaker could use. c_{it} is a vector of length J. The j^{th} entry is the number of times that the tweeter uses the j^{th} bigram. There are two arguments to the multinomial distribution $\mathrm{MN}(\cdot)$. m_{it} is the total number of bigrams spoken at time t, referred to as the "verbosity." $\boldsymbol{q}_t^{P(i)}$ is the vector of choice probabilities, also of length J. This vector depends on the covariates of the tweeter at a given point in time, denoted by vector $\boldsymbol{x}_{it}$, as well as on the party affiliation of the tweeter, $P(i) \in \{R, D\}$. We let R and D denote the set of all politician-session pairs for Democratic and Republican politicians, respectively.

MNIR is a bag-of-words model. It disregards the word order or punctuation that human readers use to parse the meaning of sentences. We follow Taddy (2015) in using bigrams as opposed to unigrams to capture some degree of lexical dependence inherent in sentence structure. Using bigrams enables MNIR to distinguish between tweets that use word sequences like "defund police" from tweets that use these two words in completely different parts of the text.

The method described in Taddy (2015) gives a computationally tractable method of estimating the parameters in this multinomial distribution using Poisson regression. The output of this procedure yields the vector of choice probabilities: $\boldsymbol{q}_t^{P(i)}\left(\boldsymbol{x}_{it}\right)$.

We estimate the above model over bigrams used in tweets by Congresspersons with a verified Twitter account between 2011 and 2021. Following Gentzkow et al. (2019b), we analyze speech at the level of the politician–session; i.e., time period t corresponds to a Congressional session. Also similar to the approach in Gentzkow et al. (2019b), we include the control variables year, home region (defined using Census regions), indicators for the highest educational degree attained, age of the speaker, and ethnicity of the speaker. We include these controls to account for demographic variables correlated with both speech and party affiliation.

We estimate MNIR over the set of bigrams used by at least forty distinct speaker-session pairs at least one hundred times. This restriction is imposed because bigrams are sometimes used by chance by only a single party, which can result in a disproportionate number of non-partisan bigrams being spuriously classified as partisan (see Gentzkow et al. (2019b)). We judge that truly partisan phrases should be used relatively frequently and by a broad range of speakers.

We display the ten bigrams most associated with Republican and Democratic speech in Table 2. The list of bigrams is intuitive. Among the Democratic bigrams are those referring to health care, gun violence, climate change, and voting rights. Among the most Republican bigrams are references to illegal immigration, tax reform, and law enforcement.

Next, we compute the posterior probability a listener with a neutral prior would have over an arbitrary politician's party with unknown demographics after hearing a particular bigram. We begin by computing the probability that a Republican politician would use the j^{th} bigram as the average probability that a given Republican politician-session pair uses the j^{th} bigram, averaging across all Republican politician-session pairs:

$$q_j^R = \frac{1}{|R|} \sum_{i \in R} \boldsymbol{q}_t^{P(i)}(\boldsymbol{x}_{it})' \cdot e_j, \tag{3.2}$$

where e_j is a vector of zeros with a single entry of 1 at element j. q_j^D is defined analogously. We then compute the posterior probability that a politician is a Republican after the listener hears the j^{th} bigram, denoted p_j^R, using Bayes rule:

$$p_j^R = \frac{q_j^R}{q_j^R + q_j^D}. \tag{3.3}$$

For bigrams that are not used at least one hundred times by at least forty different Twitter accounts, we set $q_j^R = \frac{1}{2}$.

Finally, in order to obtain a measure of the partisanship of a corporate tweet, we apply the estimates from the MNIR that was estimated on the tweets of Congresspeople to tweets sent by corporations. In this step, the unit of observation is an individual tweet. We calculate the posterior that the corporate sender of tweet k is Republican or Democrat from the expression

$$p_k^R = \frac{\prod_{j \in J^\star} q_j^R}{\prod_{j \in J^\star} q_j^R + \prod_{j \in J^\star} q_j^D}, \tag{3.4}$$

where $J^\star$ denotes the set of bigrams used in the corporate tweet. We refer to variable p_k^R as the "partisan speech index" (PSI) and define a tweet as partisan speech if p_k^R or $p_k^D = 1 - p_k^R$ is sufficiently close to one. For most of our analysis, we use a cutoff of $p_k^R \leq 0.03$ and those with $p_k^R \geq 0.97$. Intuitively, the posterior will be close to zero if a tweet comprises phrases such as the ones in the first column of Table 2. Conversely, it will be close to one if the tweet uses phrases such as "illegal immigration" or "tax reform," which are strongly associated with Republican politicians.

Panels B and C of Table 1 provide summary statistics for the sample of Democratic and Republican tweets, using a threshold of $p_k^R \leq 0.03$ and $p_k^R \geq 0.97$. As expected, partisan tweets constitute a relatively small share of all corporate tweets. The distribution of partisan tweets is also highly right-skewed, with a significantly larger mean than median.

Table 2 lists the most frequently used bigrams by U.S. companies in the set of corprorate tweets with $q_j^R \geq 0.97$ or $q_j^D \geq 0.97$. Some commonly used Democratic partisan bigrams are "climate change," "diversity and inclusion," and "clean energy." Bigrams that are very predictive of being used by Republican politicians and frequently used by companies include "natural gas," "men (and) women," and "missile defense."

3.2 Topic Models

To better understand the content of the tweets that our above method characterizes as partisan, we decompose the subject matter of these tweets into distinct topics using a biterm topic model. Topic models model documents as draws from abstract topics. Topics are probability distributions over words. An example topic could feature a high probability of using the words "trade," "tariff," and "embargo." A reasonable label for such a topic would be trade. An important characteristic of a good topic model is that it is interpretable.

After estimating MNIR, we take two resulting sets of tweets: those with $p_k^R \leq 0.1$ and those with $p_k^D \geq 0.9$. We choose less stringent cut-offs for the purpose of our topic model in order to have a sufficiently large set of partisan tweets to analyze. We then separately train topic models on both sets of tweets. We experimented with using a single topic model on both sets of tweets, but this resulted in less interpretable topics. Moreover, for the sake of computational tractability, we use unigrams instead of bigrams when estimating the topic models, following Yan et al. (2013) and Blei et al. (2003).

We estimate biterm topic models as opposed to the more common approach in the finance literature, which is Latent Dirichlet Allocation (e.g., Bybee et al. (2021), Hansen et al. (2017)). LDA models the words in individual documents as drawn from abstract topics. Unfortunately, LDA performs poorly with short texts, such as tweets, because it requires a substantial amount of text within each document to estimate the parameters of the topic model. Biterm topic models, on the other hand, estimate topics over the entire corpus of tweets. The set of speech within the union of all tweets is large enough such that the topic model parameters can be estimated precisely. Biterm models treat a single tweet as drawn from a single topic, as opposed to many, thus allowing for more precise inference of the tweet topic. Biterm is widely used in the statistics and NLP literature when working with short texts, such as tweets (e.g., Qiang et al. (2022)).

The number of topics in a topic model is a subjective choice of the researcher. We estimate a 75-topic model on the Republican tweets and a 125-topic model on the Democratic tweets. We choose the number of topics through the following procedure: we estimate many topic models, each with multiples of five topics. We started from the topic model with the smallest number of topics and iteratively examined the next smallest topic model. If the next smallest topic model had clearly interpretable topics that were not included in the prior topic model, we continued. If it did not, we stopped. The Republican topic model has fewer topics, largely reflecting that there are fewer distinct topics in the set of Republican-sounding tweets identified by MNIR.

We found that this manual procedure performed better than automated procedures using either the out-of-sample log-likelihood or "perplexity," another measure of topic model fit used in the NLP literature. We found that these two methods resulted in a very large number of topics, many of which were uninterpretable. These drawbacks are well known to the NLP literature and remain an active area of research (Stevens et al. (2012)).

For each tweet, we infer the most important topic for tweet k using a posterior implied by the estimated topic model:

$$\text{Topic Posterior}_{k,n} = \frac{\mathbb{P}\left(\text{Words Drawn from Topic } n\right)}{\sum_{m \in M} \mathbb{P}\left(\text{Words Drawn from Topic } m\right)}. \tag{3.5}$$

In practice, this posterior will be high if the tweet uses words found in Tables A1 or A2. We then say that the tweet belongs to the topic that has the largest posterior probability. Because tweets are short snippets of text and typically refer to a single topic, this "most important" posterior measure does a good job of characterizing the content of individual tweets.

The full results from our biterm topic model estimation are shown in Tables A1 and A2 in the Appendix. The topics are ordered by how frequently they are the most important topic for an individual tweet. We report the five most important unigrams for each topic. For our subsequent analysis, we also manually assign each topic a broader topic label depending on its subject matter.

Whereas topic models are often uninterpretable to a human reader, ours are highly interpretable. The words associated with each topic in Tables A1 and A2 mostly belong to clearly distinguishable groups. We conjecture that this is because of the strong factor structure in partisan speech. Partisan speech, particularly on Twitter, is often issue-specific and thus well-suited for estimation and inference using topic models.

4 Results

4.1 Trends in Partisan Corporate Speech

In Figure 1, we plot histograms of the partisan speech index for all corporate tweets for the years 2011, 2016 and 2021. X-axis values closer to zero (one) indicate corporate language is more similar to that of Democratic (Republican) politicians, respectively.

In 2011, the mass of the distribution is centered around 0.5, indicating that most tweets by corporations do not use very partisan language. If anything, Republican-sounding speech seems to be more prevalent than Democratic-sounding speech. Moving forward in time, we observe a notable shift of the distribution to the left, indicating corporations are using language more similar to that of Democratic politicians. Overall, the distribution in 2021 looks more similar to a uniform distribution than the distribution in 2011, consistent with a rise in partisan corporate speech.

To see the year-by-year evolution in the number of partisan corporate tweets more starkly, we plot the month-by-month counts of tweets using highly Republican- or Democratic-sounding speech, respectively, in Figure 2. Panel A shows the counts of tweets with a PSI value less than 0.03 (blue line) or greater than 0.97 (red line), respectively.[3] In subsequent discussion, we refer to these tweets as "Republican tweets" and "Democratic tweets," respectively.

Figure 2 confirms the findings in Figure 1. We observe a strong upward trend in

[3] We plot counts using alternative posterior cutoffs in Appendix Figure A.4.

the amount of partisan corporate speech. Democratic-sounding corporate speech exploded around the year 2017, reaching unprecedented levels in the last two years of our sample period. For example, there were close to zero Democratic tweets during January 2011. During the later part of our sample period, we see months with nearly 1,000 such tweets.

The time-series plot in Panel A displays significant variation around major events. A visible spike can be observed during June 2020, shortly following the death of George Floyd. An example of a Democratic corporate tweet from this time is the following tweet by Duke Energy on June 26, 2020:

> "The heartbreaking loss of George Floyd's life and the powerful response to it are excruciating reminders of the progress we still need to make in our communities. We're pledging $1 million to nonprofit orgs committed to social justice and racial equity."

MNIR judges this tweet to be highly partisan Democratic speech; it has a PSI-value of approximately 2×10^{-5}. We view the fact that our PSI measure correctly identifies this tweet as Democratic-sounding as a validation of our measurement approach.

The second largest spike for the entire series of Democratic tweets is in March 2021. This is the month in which the state of Georgia passed a high-profile voting law that many perceived as restricting voting rights for political gain. Many Democratic corporate tweets from this month explicitly refer to voting rights and/or to this law specifically. Here is one such example, sent by Salesforce:

> "A person's right to cast their ballot is the foundation of our democracy. Georgia HB 531 would limit trustworthy, safe & equal access to voting by restricting early voting & eliminating provisional ballots. That's why Salesforce opposes HB 531 as it stands. #gapol "

The maximum of the series of Democratic tweets occurs in June 2021. Unlike for the previous two spikes, we were not able to attribute this spike to a one-time event. Many corporate tweets in June 2021, which was Pride month, advocated for LGBTQ rights, indicating an increase in the public celebration of this event by large U.S. corporations.

In the time series of Republican-sounding tweets, we observe fewer pronounced spikes compared to the series of Democratic tweets. One of the months with the largest increase in the number of Republican tweets is November 2018; the month in which the Tax Cuts and Jobs Act (TCJA) was passed. Several Republican-sounding corporate tweets refer explicitly to the TCJA, such as this one sent by CF Industries Holdings, Inc.:

> "#TaxReform is essential to keeping workers, job creators & economy competitive in the 21st-century #TaxReformTuesday"

Panel B of Figure 2 investigates the extensive margin of tweets, in order to understand whether the increase in partisan tweets could be driven by a few companies sending a disproportionate number of partisan tweets. In each month, we plot the percentage of firms with active Twitter accounts that send at least one Democratic or Republican partisan tweet. We see a substantial increase in the percentage of S&P 500 companies sending partisan tweets over time, with the increase again being stronger for Democratic tweets. For example, in 2011, the percentage of companies sending a Democratic partisan tweet is near zero. By 2021, as many as 60% of companies send a Democratic partisan tweet in a single month. The percentage of companies sending Republican partisan tweets also increases around 2018, but remains in the range of 10 to 30% for most of our sample period.

Figure 3 repeats Panel A of Figure 2, separately for each Global Industry Classification Standard (GICS) sector. The strong increase in the number of Democratic partisan tweets is remarkably consistent across a broad range of sectors. Moreover, we do not find very pronounced differences across sectors with high or low market concentration (measured by their Herfindahl Index), as shown in the Appendix. If anything, the increase is somewhat more pronounced for industries with low-HHI, casting doubt on the interpretation that partisan corporate speech may reflect managerial entrenchment.

4.2 Topics

We also implement biterm topic modeling in order to better understand the subject of partisan corporate tweets and how it has changed over time. In the Appendix, we report the full list of topics for Democratic and Republican corporate tweets. Figure 4 presents results from our topic analysis at a higher level of aggregation, where we manually pool similar topics into broader topic categories. For example, the category "Diversity, Equity, and Inclusion" (DEI) in Panel A subsumes topics such as "Diversity and Inclusion," "LGBTQ," and "Black and Hispanic History." In Panel B, the category "Energy/Environment" includes topics such as "Oil and Gas," "Energy Costs," and "Natural Disasters." We report our exact mapping of topics into broader topic categories in Appendix Tables A1 and A2.

Figure 4, Panel A, reports the number of tweets across different topic categories for Democratic-sounding tweets. Many Democratic-sounding tweets are related to DEI, environment/climate change, and healthcare. We see a strong increase in the prevalence of DEI-related tweets after 2017, explaining a large part of the increase in the number of Democratic tweets. We also observe an increase in tweets related to climate change, as well as an increase in the number of health-related tweets around the onset of the Covid-19 pandemic.

Panel B provides the topic breakdown for Republican-sounding tweets. The most common topic categories for Republican sounding-tweets are patriotism/defense, conservative media, energy/environment, and the economy. At the beginning of our sample period, most Republican tweets refer to conservative media outlets, such as Fox Business or the Wall Street Journal (see topic category "Media"). However, over time, companies refer less to these media outlets. Instead, we see an increase in the number of Republican tweets discussing the economy (especially since 2018), as well as issues related to energy (e.g., fossil fuels).

The breakdown of the topics in Figure 4 reveals that some tweets that we identify as partisan have a clear connection to the business of the company (e.g., companies discussing economic indicators or an oil & gas company discussing a pipeline project). Whether a topic is business-related depends not only on the subject but also on the industry of the tweeting firm. For example, if a manufacturer of wind turbines tweets favorably about subsidies for renewable energy, then this statement is directly connected to its core business. However, if a grocery store operator tweets about the same topic, then the link to its core business is less obvious. We therefore define, for each tweet topic, a set of industries whose core business is directly connected to the topic of the tweet (see Appendix Table A2).

Figure 5 reveals that the vast majority of Democratic-sounding tweets are not business related, according to our definition (see Panel A). In contrast, a substantial share of Republican-sounding tweets are related to the business of the company (Panel B), either because these tweets refer to aggregate economic data/conditions, or because they discuss policies that are directly relevant to the company's core business model (e.g., an oil & gas company advocating for a pipeline project).

The results in Figures 2 and 5 illustrate a striking fact. U.S. corporations now discuss partisan topics unrelated to their core business with much greater frequency than previously.

4.3 Intraday Stock Price Reaction

An important remaining question is how investors view partisan corporate speech. Ex ante, the direction of the stock price response to partisan tweets is not obvious. On one hand, partisan corporate tweets may be in the financial interests of shareholders, because they could increase loyalty towards the firm in the labor market, product market, or financial markets. On the other hand, such statements may reflect an agency problem between managers and shareholders, with managers acting in their own personal interests at the expense of the firm's financial value. A central challenge associated with studying the financial implications of partisan statements is that their timing may be endogenous. For example, companies may be more likely to issue a statement on a social issue when they have positive financial news to report. The problem of endogenous timing is very difficult to resolve at daily frequencies.

We overcome this challenge by exploiting the precise time stamps of the corporate tweets, allowing us to conduct a second-by-second analysis of returns in a 20-minute window around the tweet. The fact that we can analyze returns is such a small window helps rule out that our results are driven by confounding events. The key identifying assumption is that other important news for the same company are not systematically released in the ten minutes before or after the tweet is sent.

To study stock returns to partisan corporate tweets, we first restrict the sample to tweets that have a greater chance of being salient to investors: tweets that receive several retweets. We use a threshold of 20 retweets for our main analysis and report robustness tests using alternative cutoffs in the Appendix.

The results from this exercise are displayed in Figure 6. Panel A plots cumulative returns around Democratic tweets. We observe a small, positive reaction to the average Democratic tweet. The average cumulative return in the 20-minute window is 1.0 basis points.

In contrast, the stock price reaction to the average Republican tweet is negative: the cumulative 20-minute return is −1.5 basis points. The difference between the cumulative return around Democratic and Republican tweets is statistically significant at the 10% level. We further find that DEI-related tweets, shown in Panel C, are an important driver of the overall positive reaction to Democratic-sounding tweets, generating cumulative returns of 1.6 basis points on average.

The magnitude of these effects is economically small. We conjecture that the above results mask substantial heterogeneity in the financial impact of partisan corporate tweets. Some tweets are newsworthy and provide new information to market participants, but many tweets do not. Moreover, restricting the window to twenty minutes around the tweet will necessarily miss the effects of announcements that are disseminated to market participants via other channels and then echoed on twitter some time later. Distinguishing between newsworthy and less important tweets is a hard problem and difficult to accomplish purely using NLP.

The positive stock price reaction to Democratic and, especially, DEI-related tweets suggests that the increase in the amount of Democratic corporate speech may reflect economic considerations rather than an agency problem between managers and shareholders. There are a variety of mechanisms through which partisan tweets could affect valuations. First, partisan speech may affect the asset demand of investors. Investors may have non-pecuniary motives for holding stocks that enact policies to address climate change or social equality. Partisan speech may also impact the ability of firms to differentiate themselves in product markets, or to attract highly-skilled labor.

5 Conclusion

We apply new techniques in natural language processing to the entire corpus of tweets sent by S&P 500 companies between 2011 and 2021 to detect partisan speech by corporations. Our measure of partisan corporate speech detects instances when corporations use language similar to that of Republican or Democratic politicians on Twitter. We show that the amount of partisan corporate speech on Twitter has dramatically increased in recent years across all sectors. We further show that the increase is disproportionally driven by speech similar to that used by Democratic politicians; in particular, statements related to climate change as well as diversity, equity, and inclusion. Intraday stock returns respond more positively to Democratic-sounding tweets than to Republican-sounding tweets.

In ongoing work, we seek to understand the determinants of corporate political speech and why companies have dramatically increased their participation in the political process. This increase could plausibly through several channels. Companies may engage in partisan speech to compete in product markets. Conversely, partisan speech may be an attempt to attract young workers or a response to demand from some participants in financial markets. We study each of these questions in turn. We hope these analyses will shed light on what forces are compelling companies to participate in the political system with increased frequency.

References

[1] Allcott, H., Boxell, L., Conway, J., Gentzkow, M., Thaler, M., and Yang, D. (2020). Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. *Journal of Public Economics*, 191:104254.

[2] Bertrand, M. and Kamenica, E. (2018). Coming apart? cultural distances in the united states over time. Technical report, National Bureau of Economic Research.

[3] Bhagwat, Y., Warren, N. L., Beck, J. T., and Watson, G. F. I. (2020). Corporate sociopolitical activism and firm value. *Journal of Marketing*, 84(5):1–12.

[4] Bizjak, J. M., Kalpathy, S. L., Mihov, V. T., and Ren, J. (2021). Ceo political leanings and store-level economic activity during covid-19 crisis: Effects on shareholder value and public health. *Available at SSRN 3674512*.

[5] Blei, D. M., Ng, A. Y., and Jorgan, M. I. (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research*.

[6] Bybee, L., Kelly, B., Manela, A., and Xui, D. (2021). Business news and business cycles. *Available at SSRN 3446225*.

[7] Cassidy, W. and Vorsatz, B. (2021). Partisanship and portfolio choice: Evidence from mutual funds. Working Paper.

[8] Chen, D. L. (2019). Priming ideology: Why do presidential elections affect us judges. *Available at SSRN 2816245*.

[9] Coibion, O., Gorodnichenko, Y., and Weber, M. (2020). Political polarization and expected economic outcomes. Technical report, National Bureau of Economic Research.

[10] Cookson, J. A., Engelberg, J., and Mullins, W. (2021). Does partisanship shape investor beliefs? Evidence from the COVID-19 pandemic. *The Review of Asset Pricing Studies*.

[11] Dagostino, R., Gao, J., and Ma, P. (2020). Partisanship in loan pricing. *Available at SSRN 3701230*.

[12] Di Giuli, A. and Kostovetsky, L. (2014). Are red or blue companies more likely to go green? politics and corporate social responsibility. *Journal of Financial Economics*, 111(1):158–180.

[13] Engelberg, J., Guzman, J., Lu, R., and Mullins, W. (2021). Partisan entrepreneurship. Working Paper.

[14] Engelberg, J., Henriksson, M., Manela, A., and Williams, J. (2022). The partisanship of financial regulators. *Available at SSRN 3481564*.

[15] Fos, V., Kempf, E., and Tsoutsoura, M. (2022). The political polarization of corporate america. *Available at SSRN 3784969*.

[16] Gangopadhyay, S. and HomRoy, S. (2022). Strategic CEO activism in polarized markets. *Available at SSRN 3622605*.

[17] Gelles, D. and Ross Sorkin, A. (2021). Hundreds of companies unite to oppose voting rights limits, but others abstain. *https://www.nytimes.com/2021/04/14/business/ceos-corporate-america-voting-rights.html*.

[18] Gentzkow, M., Kelly, B., and Taddy, M. (2019a). Text as Data. *Journal of Economic Literature*, 57(3):535–574.

[19] Gentzkow, M., Shapiro, J. M., and Taddy, M. (2019b). Measuring Group Differences in High-Dimensional Choices: Method and Application to Congressional Speech. *Econometrica*, 87(4):1307–1340.

[20] Gerber, A. S. and Huber, G. A. (2010). Partisanship, political control, and economic assessments. *American Journal of Political Science*, 54(1):153–173.

[21] Hansen, S., McMahon, M., and Prat, A. (2017). Transparency and deliberation within the fomc: a computational linguisitics approach. *The Quartelry Journal of Economics*, 133:801–870.

[22] Hart, O. and Zingales, L. (2022). The new corporate governance. *University of Chicago Business Law Review*, 1(1).

[23] Hessekiel, D. (2020). Companies taking a public stand in the wake of george floyd's death. *https://www.forbes.com/sites/davidhessekiel/2020/06/04/companies-taking-a-public-stand-in-the-wake-of-george-floyds-death/?sh=1f01639a7214*.

[24] Hutton, I., Jiang, D., and Kumar, A. (2014). Corporate policies of republican managers. *Journal of Financial and Quantitative Analysis*, 49(5-6):1279–1310.

[25] Kempf, E. and Tsoutsoura, M. (2021). Partisan Professionals: Evidence from Credit Rating Analysts. *Journal of Finance*, 76(6):2805–2856.

[26] Larcker, D. F., Miles, S. A., Tayan, B., and Wright-Violich, K. (2018). The double-edged sword of CEO activism. *Stanford Closer Look Series, Corporate Governance Research Initiative*.

[27] Lucas, A. (2019). Chief executives of 145 companies urge senate to pass gun control laws. *https://www.cnbc.com/2019/09/12/chief-executives-of-145-companies-urge-senate-to-pass-gun-control-laws.html*.

[28] Mkrtchyan, A., Sandvik, J., and Zhu, V. (2022). CEO activism and firm value. *Available at SSRN 3699082*.

[29] Pew Research (2019). In a politically polarized era, sharp divides in both partisan coalitions. *https://www.pewresearch.org/politics/2019/12/17/in-a-politically-polarized-era-sharp-divides-in-both-partisan-coalitions/*.

[30] Qiang, J., Qian, Z., Li, Y., Yuan, Y., and Wu, X. (2022). Short text topic modeling techniques, applications, and performance: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 34.

[31] Rabinovich, M. and Blei, D. M. (2014). The Inverse Regression Topic Model. *Proceedings of the 31st International Conference on Machine Learning*, 32:199–207.

[32] Stevens, K., Kegelmeyer, P., Andrzejewski, D., and Buttler, D. (2012). Exploring topic coherence over many models and many topics. In *Proceedings of the 2012 joint conference on empirical methods in natural language processing and computational natural language learning*, pages 952–961.

[33] Taddy, M. (2013). Multinomial Inverse Regression for Text Analysis. *Journal of the American Statistical Association*, 108(503):755–770.

[34] Taddy, M. (2015). Distributed multinomial regression. *The Annals of Applied Statistics*, 9(3):1394–1414.

[35] Yan, X., Guo, J., Lan, Y., and Cheng, X. (2013). A biterm topic model for short texts. *Proceedings of the 22nd international conference on World Wide Web - WWW '13*, pages 1445–1456.

Figure 1
Distribution of PSI-scores for Corporate Tweets

Panel A: 2011 **Panel B: 2016** **Panel C: 2021**

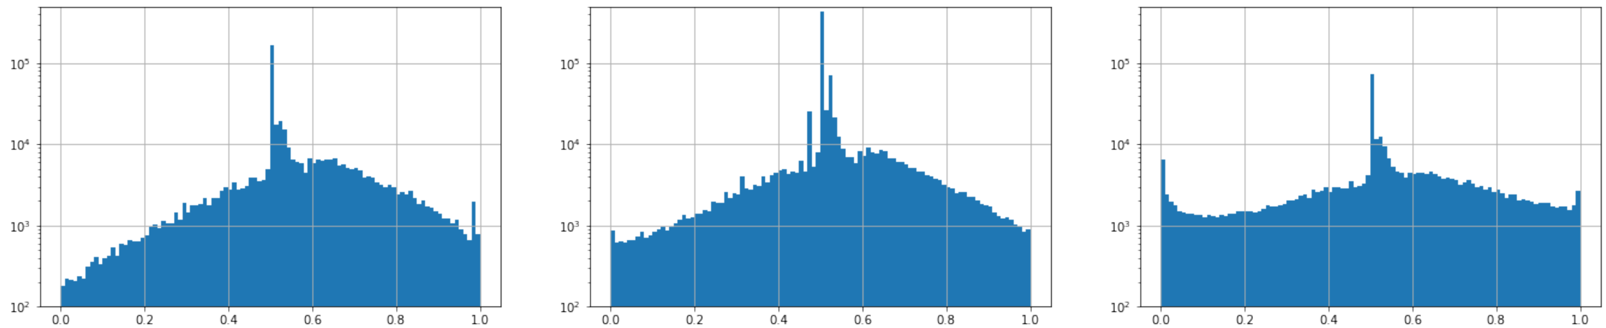

The figure displays the histograms of PSI-scores for corporate tweets sent in 2011, 2016, and 2021, respectively. A PSI-value near zero uses strongly Democratic-sounding language and a PSI-value near one uses strongly Republican-sounding language. The y-axis shows the logged number of tweets with a PSI-value falling within a particular bin.

Figure 2
Time Series of Partisan Corporate Tweets

Panel A: Raw Counts | **Panel B: Extensive Margin**

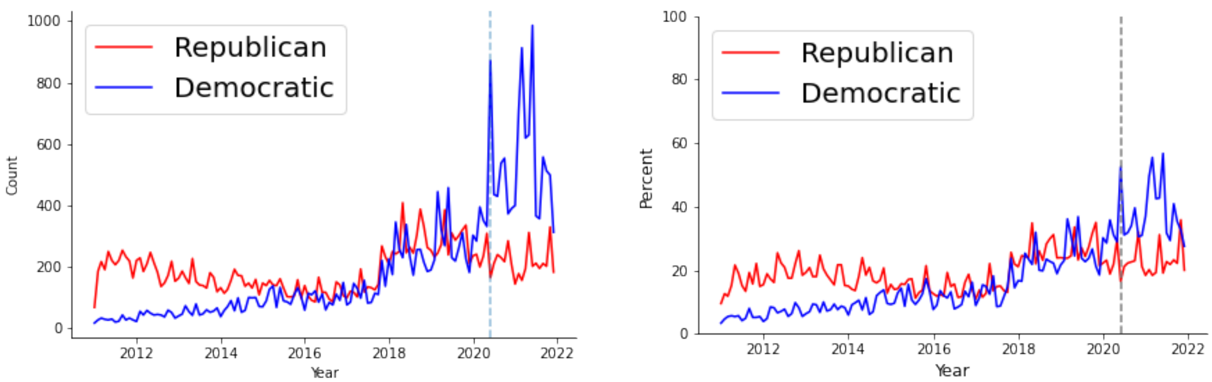

The first panel of this figure plots the raw counts of partisan tweets by calendar month. The second panel displays the ratio of the number companies that send at least one partisan tweet in a given month to the number of companies with a verified Twitter account in that month. In both plots, the blue (red) line corresponds to Democratic (Republican) partisan tweets, respectively. Democratic tweets are tweets with a PSI-value ≤ 0.03 and Republican tweets are tweets with a PSI-value ≥ 0.97. The dashed vertical line corresponds to June 2020.

Figure 3
Partisan Corporate Tweets by GICS Sector

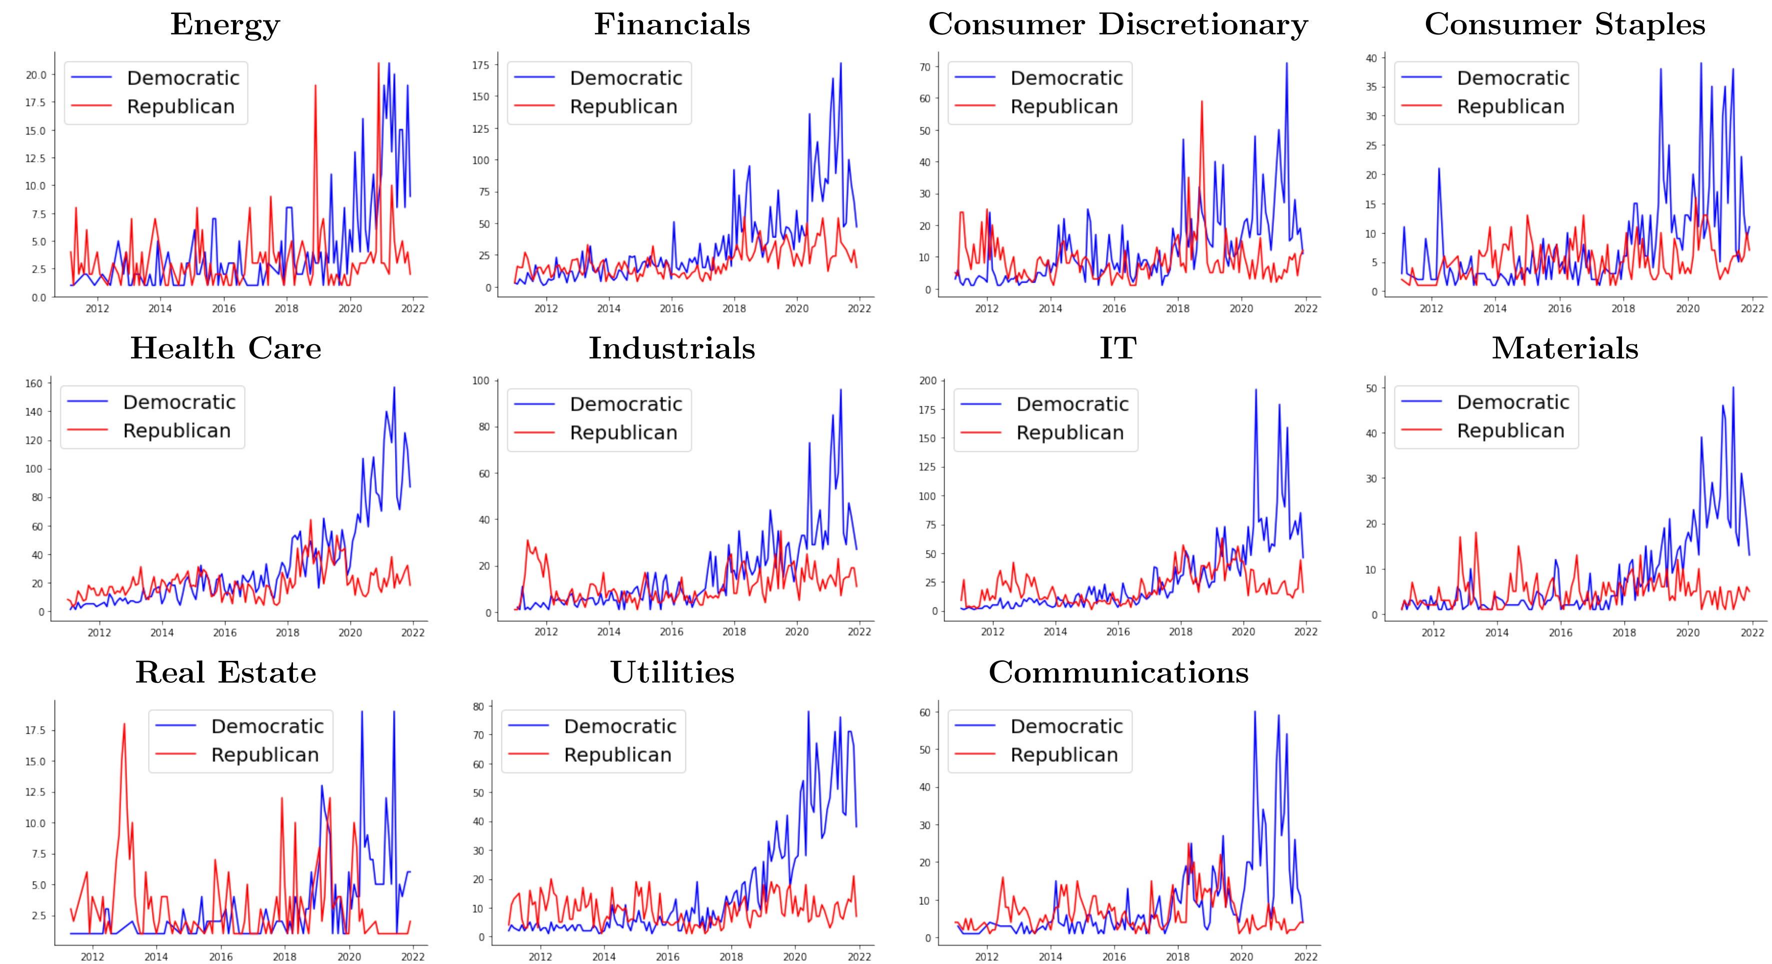

These figures repeat Figure 2 separately for firms operating in a given GICS sector. GICS sectors are obtained from Compustat Annual.

Figure 4
Partisan Corporate Tweets by Time and Topic

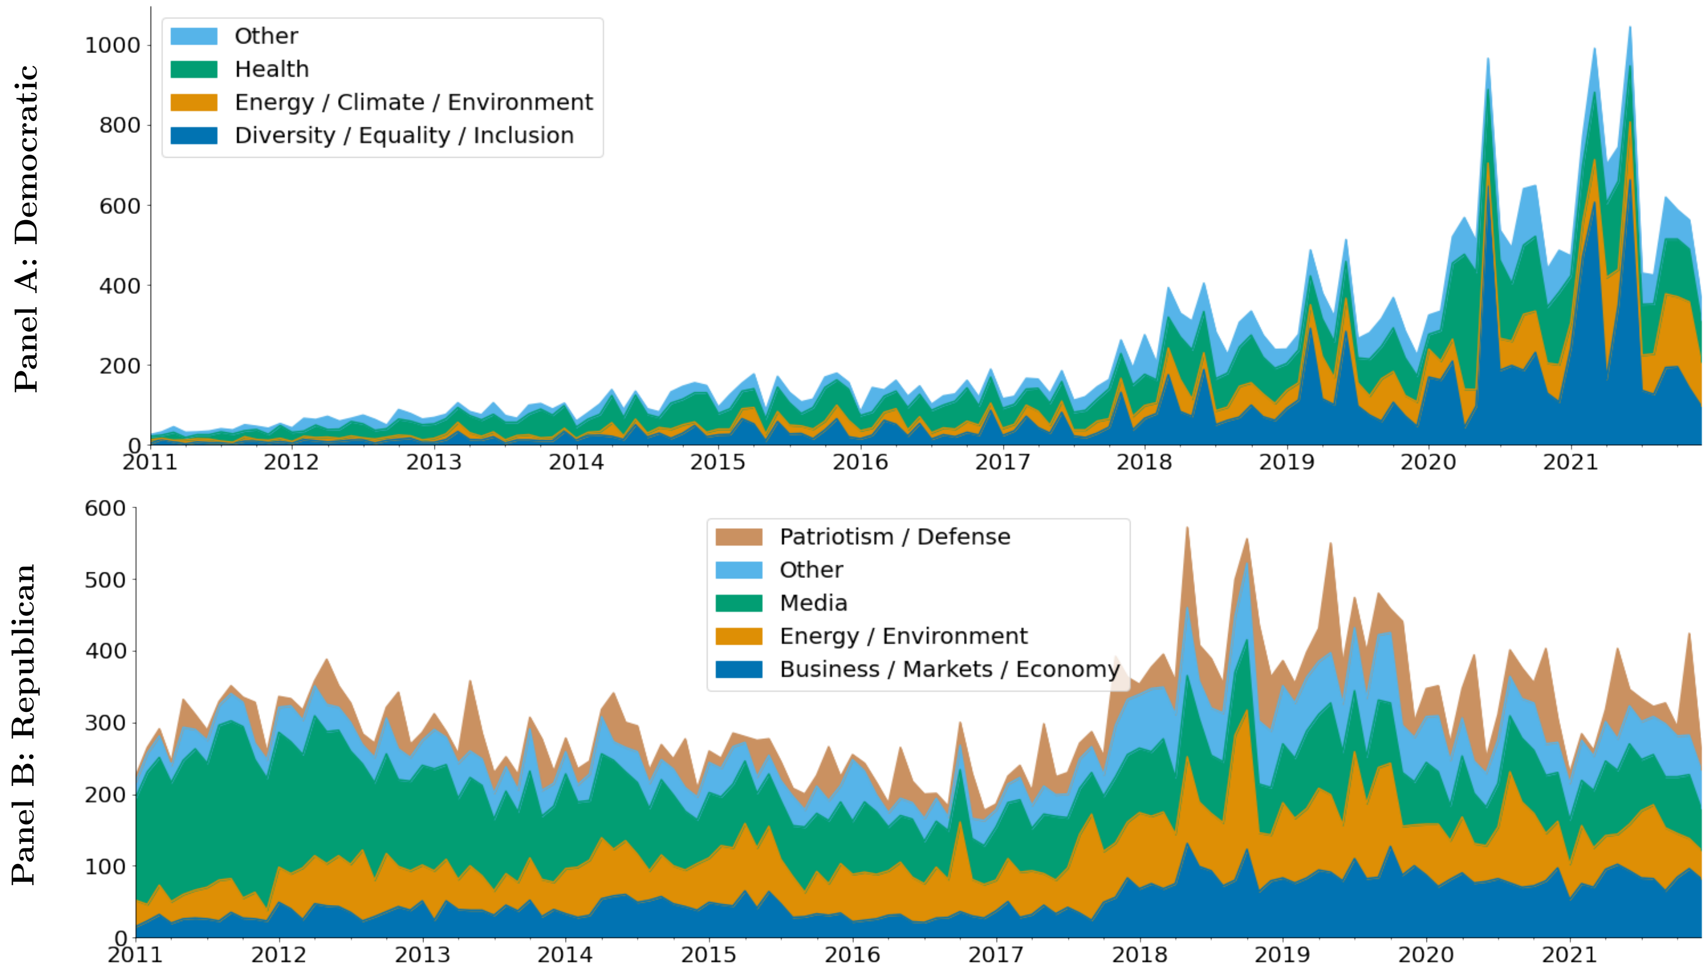

The figure displays the evolution of partisan corporate speech by topic. Panel A shows the number of Democratic tweets broken down by topic. Panel B does the same for Republican tweets. Democratic tweets are tweets with a PSI-value ≤ 0.03 and Republican tweets are tweets with a PSI-value ≥ 0.97. Topics are estimated using a biterm topic model and manually grouped into larger topic categories.

Figure 5
Partisan Corporate Tweets: Business vs. Non-Business Related

Panel A: Democratic Tweets **Panel B: Republican Tweets**

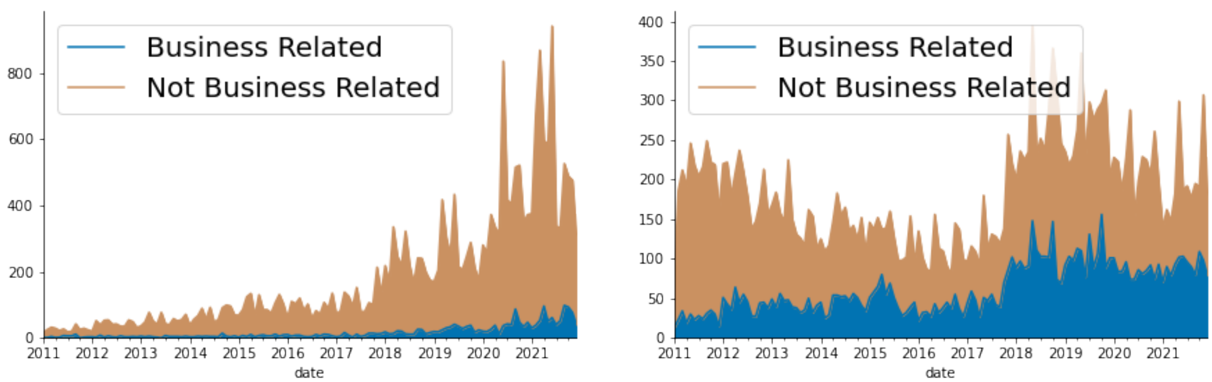

The figure displays the counts of Republican and Democratic partisan tweets that are business-related (blue) or not business-related (brown), respectively. We manually classify which topics are business-related based on the nature of the topic and the industry of the tweeting firms. We list all combinations of topics and industries that are classified as business-related in Appendix Tables A2 and A1.

Figure 6
Intraday Returns

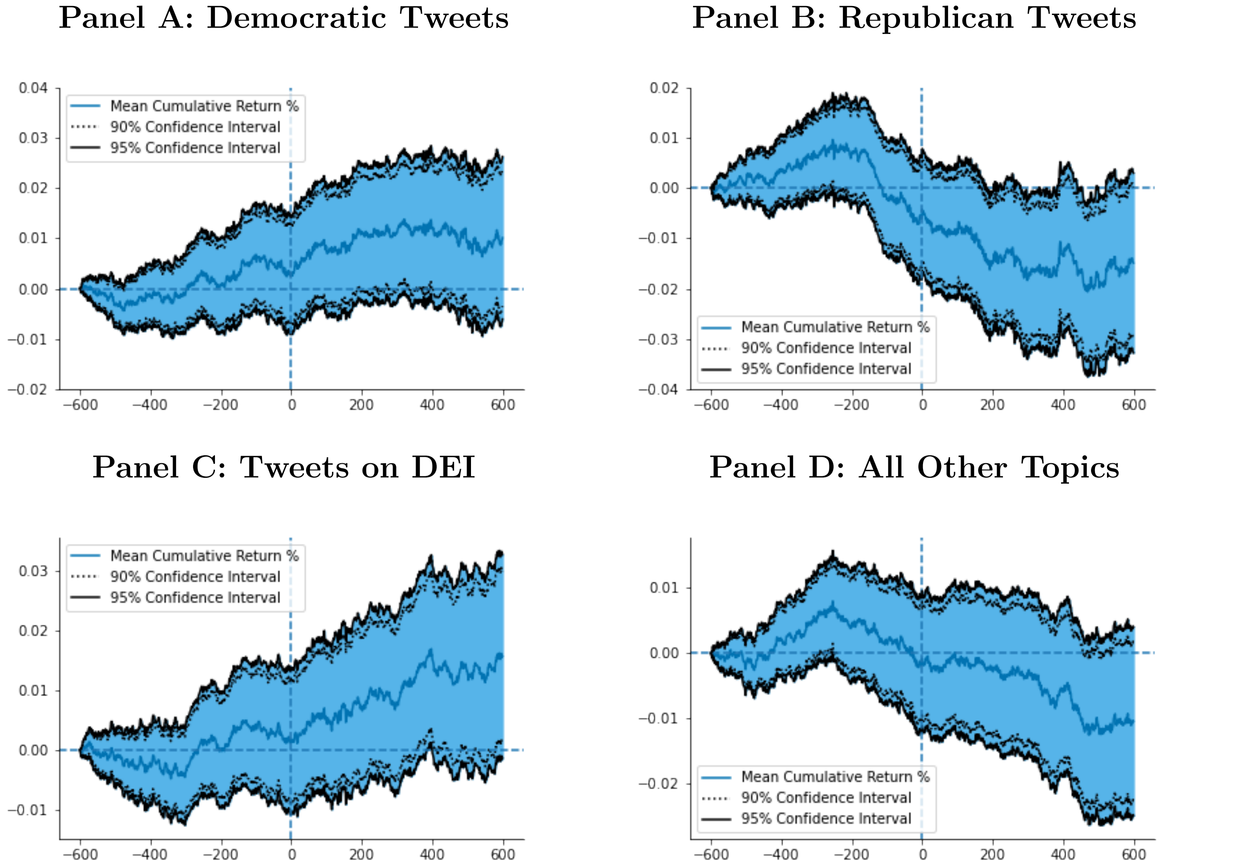

The figure displays cumulative stock returns in a ten-minute window around partisan corporate tweets. The x-axis is measured in event-time seconds and the y-axis is measured in percentage points. Cumulative returns are winsorized at the 5% level. Panel A plots returns around Democratic tweets and Panel B around Republican tweets, using PSI-value cutoffs of 0.03 and 0.97, respectively. Panel C plots returns around tweets falling under the topic of Diversity, Equity, and Inclusion (DEI). Panel B shows returns for all other topics.

Table 1
Corporate Tweets: Summary Statistics

The table reports summary statistics for all tweets sent by firms listed on the S&P 500 via their verified Twitter accounts between 2011 and 2021.

Year:	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
Panel A: Full Sample											
Unique Firms	389	440	461	493	509	524	539	545	551	553	555
Average Tweets Per Firm	701.12	902.75	1032.76	1082.27	1070.53	1355.12	842.29	727.72	639.69	545.96	521.13
Standard Deviation of Tweets Per Firm	1544.81	1662.88	1859.47	2179.83	2456.74	9266.95	2032.02	1844.34	1599.96	1513.98	1717.92
Minimum Number of Tweets	1	1	1	1	1	1	1	1	1	1	1
Median Number of Tweets	310	472	560	629	632	565	475	413	353	290	270
Maximum Number of Tweets	19552	21699	25886	39165	50390	206277	42046	39224	34794	32611	37727
Panel B: Democratic Tweets (PSI-Score $\leq$ 0.03)											
Unique Firms	135	173	201	250	285	276	315	392	426	468	495
Average Tweets Per Firm	3.41	4.02	4.98	5.36	6.68	6.39	7.59	10.99	11.84	16.02	19.44
Standard Deviation of Tweets Per Firm	4.73	5.47	10.31	14.73	24.72	22.1	39.22	51.72	51.18	55.6	64.34
Minimum Number of Tweets	1	1	1	1	1	1	1	1	1	1	1
Median Number of Tweets	2	2	2	2	3	2	3	4	5	6	9
Maximum Number of Tweets	36	51	120	197	400	345	687	993	1024	1123	1327
Panel C: Republican Tweets (PSI-Score $\geq$ 0.97)											
Unique Firms	246	313	333	340	324	336	363	458	453	419	436
Average Tweets Per Firm	11.47	9.63	8.75	8.89	8.45	7.43	7.8	10.71	11.88	12.6	13.14
Standard Deviation of Tweets Per Firm	46.06	27.56	30.84	43.7	51.68	45.28	40.77	57.51	79.55	97.9	119.76
Minimum Number of Tweets	1	1	1	1	1	1	1	1	1	1	1
Median Number of Tweets	3	3	3	3	3	2	3	4	4	4	4
Maximum Number of Tweets	596	306	433	635	919	809	754	1205	1680	1991	2472

Table 2
Most Republican and Democratic Bigrams

Panel A shows the ten bigrams most associated with use by Republican or Democratic politicians on Twitter. Panel B shows the ten most common bigrams among corporate speech that PSI classifies as partisan, using a PSI cutoff of 0.03 and 0.97, respectively.

Panel A: Politician Speech		**Panel B: Corporate Speech**	
Democratic	Republican	Democratic	Republican
health care	god bless	climat chang	top stori
gun violenc	illeg immigr	divers inclus	today via
climat chang	tax reform	join us	stori today
trump administr	ron paul	student loan	natur gas
make sure	law enforc	clean energi	men women
vote right	men women	lgbtq communiti	small busi
million american	joe biden	health care	via wsj
work famili	look forward	histori month	missil defens
preexist condit	small busi	around world	suppli chain
public health	presid realdonaldtrump	sustain futur	defens system